

Random walks on dynamical percolation: mixing times, mean squared displacement and hitting times

Yuval Peres*

Alexandre Stauffer†

Jeffrey E. Steif‡

August 13, 2014

Abstract

We study the behavior of random walk on dynamical percolation. In this model, the edges of a graph G are either open or closed and refresh their status at rate μ while at the same time a random walker moves on G at rate 1 but only along edges which are open. On the d -dimensional torus with side length n , we prove that in the subcritical regime, the mixing times for both the full system and the random walker are n^2/μ up to constants. We also obtain results concerning mean squared displacement and hitting times. Finally, we show that the usual recurrence transience dichotomy for the lattice \mathbb{Z}^d holds for this model as well.

Keywords and phrases. Percolation, dynamical percolation, random walk, mixing times.
MSC 2010 subject classifications. Primary 60K35, 60K37,

1 Introduction

Random walks on finite graphs and networks have been studied for quite some time; see [1]. Here we study random walks on certain randomly evolving graphs. The simplest such evolving graph is given by dynamical percolation. Here one has a graph $G = (V, E)$ and parameters p and μ and one lets 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 $(1 - p)\mu$. We assume $\mu \leq 1$. Let $\{\eta_t\}_{t \geq 0}$ denote the resulting Markov process on $\{0, 1\}^E$ whose stationary distribution is product measure π_p . 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 neighbor (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 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. One motivation for the model is that in real networks, the structure of the network itself can evolve over time; however the time scale at which this might occur is much longer than the time scale for the walker itself. This would correspond to the case $\mu \ll 1$ which is indeed the interesting regime for our results.

*Microsoft Research, Redmond WA, U.S.A. Email: peres@microsoft.com

†University of Bath, Bath, U.K. Email: a.stauffer@bath.ac.uk

‡Chalmers University of Technology and Gothenburg University, Gothenburg, Sweden Email: steif@chalmers.se

We state at this point that all the results in this paper concern only the d -dimensional discrete torus and the lattice \mathbb{Z}^d . Our first result shows that the usual recurrence/transience criterion for ordinary random walk holds for this model as well. We consider Theorem 1.1 and Theorem 3.1 (to be given later on) as minor results. It is certainly possible that these are obtainable by methods in some of the papers being referenced here; however, we find our proofs of these two results quite simple.

Theorem 1.1. *Assume the random walker starts at the origin.*

(i). *If $G = \mathbb{Z}^d$ with d being 1 or 2, then for any $p \in [0, 1]$, $\mu > 0$ and initial bond configuration η_0 , we have that, for any $s_0 \geq 0$, $\mathbf{P} \left(\bigcup_{s \geq s_0} \{X_s = 0\} \right) = 1$.*

(ii). *If $G = \mathbb{Z}^d$ with $d \geq 3$, then for any $p \in (0, 1]$, $\mu > 0$ and initial bond configuration η_0 , we have that*

$$\lim_{t \rightarrow \infty} X_t = \infty \text{ a.s.}$$

We note that when G is finite and has constant degree, one can check that $u \times \pi_p$ is the unique stationary distribution and that the process is reversible; u here is the uniform distribution.

Next, our main theorem gives the mixing time up to constants for $\{M_t\}_{t \geq 0}$ and $\{X_t\}_{t \geq 0}$ on the d -dimensional discrete torus with side length n , denoted by $\mathbb{T}^{d,n}$, in the subcritical regime for percolation, where importantly μ may depend on n .

Let $\|m_1 - m_2\|_{\text{TV}}$ denote the total variation distance between two probability measures m_1 and m_2 , T_{mix} denote the mixing time for a Markov chain and let $p_c(d)$ denote the critical value for percolation on \mathbb{Z}^d . See Section 2 for definitions of all these terms. Next, starting the walker at the origin and taking the initial bond configuration to be distributed according to π_p , let

$$T_{\text{mix}}^{\text{RW}}(\epsilon) := \inf\{t \geq 0: \|\mathcal{L}(X_t) - u\|_{\text{TV}} \leq \epsilon\}.$$

(The superscript RW refers to the fact that we are only considering the walker here rather than the full system.) Below $p_c(\mathbb{Z}^d)$ denotes the critical value for bond percolation on \mathbb{Z}^d and $\theta_d(p)$ denotes the probability that the origin is in an infinite component when the parameter is p ; see Section 2.

Theorem 1.2. (i). *For any $d \geq 1$ and $p \in (0, p_c(\mathbb{Z}^d))$, there exists $C_1 < \infty$ such that, for all n and for all μ , considering the full system $\{M_t\}_{t \geq 0}$ on $\mathbb{T}^{d,n}$, we have*

$$T_{\text{mix}} \leq \frac{C_1 n^2}{\mu}.$$

(ii). *For any $d \geq 1$, $p \in (0, p_c(\mathbb{Z}^d))$ and $\epsilon < 1$, there exist $C_2 > 0$ and $n_0 > 0$ such that, for all $n \geq n_0$ and for all μ , considering the system on $\mathbb{T}^{d,n}$, we have*

$$T_{\text{mix}}^{\text{RW}}(\epsilon) \geq \frac{C_2 n^2}{\mu}.$$

- Remarks 1.3.**
1. (ii) implies that the upper bound in (i) is also a lower bound up to constants.
 2. (i) implies, using (5) in Section 2, that the lower bound in (ii) is also an upper bound up to (ϵ dependent) constants.
 3. The theorem shows that the “mixing time” for the random walk component of the chain is the same as for the full system. However, as this component is not Markovian, there is no well established definition of the mixing time in this context; this is why we write “mixing time”.
 4. Part (ii) becomes a stronger statement when ϵ becomes larger.

One of the key steps in order to prove (ii) of Theorem 1.2 is to prove that the mean squared displacement of the walker is at most linear on the time scale $1/\mu$ uniform in the size of the torus. This result, which is also of independent interest, is presented next. Here and throughout the paper, $\text{dist}(x, y)$ will denote the graph distance between two vertices x and y in a given graph.

Theorem 1.4. *Fix d and $p \in (0, p_c(\mathbb{Z}^d))$. Then there exists $C_{1.4} = C_{1.4}(d, p) < \infty$ so that for all n , for all μ and for all t , if $G = \mathbb{T}^{d, n}$, we have that*

$$\mathbf{E} \left[\text{dist}(X_{\frac{t}{\mu}}, X_0)^2 \right] \leq C_{1.4}(t \vee 1) \quad (1)$$

when we start the full system in stationarity with $u \times \pi_p$.

Remark 1.5. The above inequality is false if the “ $\vee 1$ ” is removed since if $t = \mu$ is very small, then the LHS is not arbitrarily close to 0.

From Theorem 1.4, we can obtain a similar bound for the full lattice \mathbb{Z}^d .

Corollary 1.6. *Fix d and $p \in (0, p_c(\mathbb{Z}^d))$. Then for all μ and for all t , if $G = \mathbb{Z}^d$, we have that*

$$\mathbf{E} \left[\text{dist}(X_{\frac{t}{\mu}}, 0)^2 \right] \leq C_{1.4}(t \vee 1)$$

when we start the system with distribution $\delta_0 \times \pi_p$ and where $C_{1.4}$ comes from Theorem 1.4.

In Theorem 1.2(ii), Theorem 1.4 and Corollary 1.6, it was assumed that the bond configuration was started in stationarity (in Theorem 1.2(ii), this is true since this was incorporated into the definition of $T_{\text{mix}}^{\text{RW}}(\epsilon)$). For Theorem 1.4 and Corollary 1.6, if the initial bond configuration is identically 1, $\mu = \frac{1}{n^{d+2}}$ and $t = \frac{1}{n^{d+1}}$, then the LHS’s of these results grow to ∞ while the RHS’s stay bounded and hence these results no longer hold. The reason for this is that the bonds which the walker might encounter during this time period are unlikely to refresh and so the walker is just doing ordinary random walk on \mathbb{Z}^d or $\mathbb{T}^{d, n}$. For similar reasons, if one takes the initial bond configuration to be identically 1 in the definition of $T_{\text{mix}}^{\text{RW}}(\epsilon)$, then if μ is very small, $T_{\text{mix}}^{\text{RW}}(\epsilon)$ will be of the much smaller order n^2 . However, due to Theorem 1.2(i), one cannot on the other hand make $T_{\text{mix}}^{\text{RW}}(\epsilon)$ larger than order $\frac{n^2}{\mu}$ by choosing an appropriate initial bond configuration.

For general p , we obtain the following lower bounds on the mixing time. This is only of interest in the supercritical and critical cases $p \geq p_c$ since Theorem 1.2(ii) essentially implies this in the subcritical case; one minor difference is that in (i) below, the constants do not depend on p .

Theorem 1.7. (i). *Given $d \geq 1$ and $\epsilon < 1$, there exist $C_1 > 0$ and $n_0 > 0$ such that, for all p , for all $n \geq n_0$ and for all μ , if $G = \mathbb{T}^{d, n}$, then*

$$T_{\text{mix}}^{\text{RW}}(\epsilon) \geq C_1 n^2.$$

(ii). *Given $d \geq 1$, p and $\epsilon < 1 - \theta_d(p)$, there exists $C_2 > 0$ and $n_0 > 0$ such that, for all $n \geq n_0$ and for all for μ , if $G = \mathbb{T}^{d, n}$, then*

$$T_{\text{mix}}^{\text{RW}}(\epsilon) \geq \frac{C_2}{\mu}. \quad (2)$$

In particular, for $\epsilon < 1 - \theta_d(p)$, we get a lower bound for $T_{\text{mix}}^{\text{RW}}(\epsilon)$ of order $\frac{1}{\mu} + n^2$.

Remark 1.8. The lower bound in (ii) holds only for sufficiently small $\epsilon \in (0, 1)$ depending on p . To see this, take $d = 2$ and choose p sufficiently close to 1 so that $\theta_2(p) > .999$. Take n large, $\mu \ll n^{-4}$ and $t = Cn^2$, where C is a large enough constant. Then, with probability going to 1 with n , the giant cluster at time 0 will contain at least .999 fraction of the vertices. Therefore, with probability about .999, the origin will be contained in this giant cluster. By [10], with probability going to 1 with n , this giant cluster will have a mixing time of order n^2 . Therefore, if C is large, a random walk on this giant cluster run for $t = Cn^2$ units of time will be, in total variation, within .0001 of the uniform distribution on this cluster and hence be, in total variation, within .0002 of the uniform distribution u . Since $\mu \ll n^{-4}$, no edges will refresh up to time t with very high probability and hence $\|\mathcal{L}(X_t) - u\|_{\text{TV}} \leq .0002$. Since $t \ll \frac{1}{\mu}$, we see that (2) above is not true for all ϵ but rather only for small ϵ depending on p . This strange dependence of the mixing time on ϵ cannot occur for a Markov process but can only occur here since $\{X_t\}_{t \geq 0}$ is not Markovian.

We mention that heuristics suggest that the lower bound of $\frac{1}{\mu} + n^2$ for the supercritical case should be the correct order.

We now give an analogue of Theorem 1.4 for general p . This is also of interest in itself and as before is a key step in proving Theorem 1.7(i). While it is of course very similar to Theorem 1.4, the fundamental difference between this result and the latter result is that we do not now obtain linear mean squared displacement on the time scale $1/\mu$ as we had before.

Theorem 1.9. *Fix $d \geq 1$. Then there exists $C_{1.9} = C_{1.9}(d)$ so that for all p , for all n , for all μ and for all t , if $G = \mathbb{T}^{d,n}$, then*

$$\mathbf{E} [\text{dist}(X_t, X_0)^2] \leq C_{1.9}t \quad (3)$$

when we start the full system in stationarity with $u \times \pi_p$.

From this, we can obtain, as before, a similar bound on the full lattice \mathbb{Z}^d .

Corollary 1.10. *Fix $d \geq 1$. For all p , for all μ and for all t , if $G = \mathbb{Z}^d$, we have*

$$\mathbf{E} [\text{dist}(X_t, 0)^2] \leq C_{1.9}t \quad (4)$$

when we start the full system with distribution $\delta_0 \times \pi_p$ and where $C_{1.9}$ comes from Theorem 1.9.

Remark 1.11. Theorem 1.9 and Corollary 1.10 are false if we start the bond configuration in an arbitrary configuration. In [9], a subgraph G of \mathbb{Z}^2 is constructed such that if one runs random walk on it, the expected mean squared distance to the origin at time t^2 is much larger than t for large t . If we start the bond configuration in the state G and μ is sufficiently small, then clearly (4) fails for large t provided (for example) that $t^{d+1}\mu = o(1)$. Similarly, (3) fails for such t and large n .

For Markov chains, one is often interested in studying hitting times. For discrete time random walk on the torus $\mathbb{T}^{d,n}$, it is known that the maximum expectation of the time it takes to hit a point from an arbitrary starting point behaves, up to constants, as n^d for $d \geq 3$, $n^2 \log n$ for $d = 2$ and n^2 for $d = 1$. Here we obtain an analogous result for our dynamical model in the subcritical regime.

For $y \in \mathbb{T}^{d,n}$, let $\sigma_y := \inf\{t \geq 0 : X_t = y\}$ be the first hitting time of y . We will always start our system with distribution $\delta_0 \times \pi_p$ (otherwise the results would be substantially different). Finally, we let $H_d(n) := \max\{\mathbf{E}[\sigma_y] : y \in \mathbb{T}^{d,n}\}$ denote the maximum expected hitting time.

Theorem 1.12. *For all $d \geq 1$ and $p \in (0, p_c(\mathbb{Z}^d))$, there exists $C_{1.12} = C_{1.12}(d, p) < \infty$ so that the following hold.*

(i). *For all n and for all $\mu \leq 1$,*

$$\frac{n^2}{C_{1.12}\mu} \leq H_1(n) \leq \frac{C_{1.12}n^2}{\mu}.$$

(ii). For all n and for all $\mu \leq 1$,

$$\frac{n^2 \log n}{C_{1.12} \mu} \leq H_2(n) \leq \frac{C_{1.12} n^2 \log n}{\mu}.$$

(iii). For all $d \geq 3$, for all n and for all $\mu \leq 1$,

$$\frac{n^d}{C_{1.12} \mu} \leq H_d(n) \leq \frac{C_{1.12} n^d}{\mu}.$$

One of the usual methods for obtaining hitting time results (see [23]) is to first develop and then to apply results from electrical networks. However, in our case, where the network itself is evolving in time, this approach does not seem to be applicable. More generally, many of the standard methods for analyzing Markov chains do not seem helpful in studying the case, such as this, where the transition probabilities are evolving in time stochastically.

Previous work.

Studying random walk in a random environment has been done since the early 1970's. In the initial models studied, one chose a random environment which would then be used to give the transition probabilities for a random walk. Once chosen, this environment would be fixed. There are many papers on random walk in random environment, far too many to list here.

After this, one studied random walks in an evolving random environment. The evolving random environment could be of a quite general nature. A sample of papers from this area are [2], [3], [4], [5], [6], [8], [12], [13], [14], [15], [16], [18], [21] and [26].

However, the focus of these papers and the questions addressed in them are of a very different nature than the focus and questions addressed in the present paper. We therefore do not attempt to describe at all the results in these papers.

Organization.

The rest of the paper is organized as follows. In Section 2, various background will be given. In Section 3, Theorem 1.1 is proved as well as a central limit theorem and a general technical lemma which will be used here as well as later on. In Section 4, we prove the mean squared displacement results and the lower bound on the mixing time in the subcritical regime: Theorem 1.4, Corollary 1.6 and Theorem 1.2(ii). In Section 5, we prove the mean squared displacement results and the lower bound on the mixing time in the general case: Theorem 1.9, Corollary 1.10 and Theorem 1.7. In Section 6, the upper bound on the mixing time in the subcritical regime, Theorem 1.2(i), is proved. Finally, in Section 7, Theorem 1.12 is proved. In Section 8, we state an open question.

2 Background

In this section, we provide various background.

Percolation.

In percolation, one has a connected locally finite graph $G = (V, E)$ and a parameter $p \in (0, 1)$. One then declares each edge to be open (state 1) with probability p and closed (state 0) with probability $1 - p$ independently for different edges. Throughout this paper, we write π_p for the corresponding product measure. One then studies the structure of the connected components (clusters) of the resulting subgraph of G consisting of all vertices and all open edges. We will use $\mathbf{P}_{G,p}$ to denote

probabilities when we perform p -percolation on G . If G is infinite, the first question that can be asked is whether an infinite connected component exists (in which case we say *percolation* occurs). Writing \mathcal{C} for this latter event, Kolmogorov's 0-1 law tells us that the probability of \mathcal{C} is, for fixed G and p , either 0 or 1. Since $\mathbf{P}_{G,p}(\mathcal{C})$ is nondecreasing in p , there exists a critical probability $p_c = p_c(G) \in [0, 1]$ such that

$$\mathbf{P}_{G,p}(\mathcal{C}) = \begin{cases} 0 & \text{for } p < p_c \\ 1 & \text{for } p > p_c. \end{cases}$$

For all $x \in V$, let $C(x)$ denote the connected component of x , i.e. the set of vertices having an open path to x . Finally, we let $\theta_d(p) := \mathbf{P}_{\mathbb{Z}^d,p}(|C(0)| = \infty)$ where 0 denotes the origin of \mathbb{Z}^d . See [19] for a comprehensive study of percolation.

Let $\mathbb{T}^{d,n}$ be the d -dimensional discrete torus with vertex set $[0, 1, \dots, n]^d$. This is a transitive graph with n^d vertices. For large n , the behavior of percolation on $\mathbb{T}^{d,n}$ is quite different depending on whether $p > p_c(\mathbb{Z}^d)$ or $p < p_c(\mathbb{Z}^d)$; in this way the finite systems “see” the critical value for the infinite system. In particular, if $p > p_c(\mathbb{Z}^d)$, then with probability going to 1 with n , there will be a unique connected component with size of order n^d (called the *giant cluster*) while for $p < p_c(\mathbb{Z}^d)$, with probability going to 1 with n , all of the connected components will have size of order at most $\log(n)$.

Dynamical Percolation.

This model was introduced in Section 1. Here we mention that the model can equally well be described by having each edge of G independently refresh its state at rate μ and when it refreshes, it chooses to be in state 1 with probability p and in state 0 with probability $1 - p$ independently of everything else. We have already mentioned that for all G, p and μ , the product measure π_p is a stationary reversible probability measure for $\{\eta_t\}_{t \geq 0}$. Dynamical percolation was introduced independently in [20] and by Itai Benjamini. The types of questions that have been asked for this model is whether there exist exceptional times at which the percolation configuration looks markedly different from that at a fixed time. See [27] for a recent survey of the subject. Our focus however in this paper is quite different.

Random walk on Dynamical Percolation.

Random walk on Dynamical Percolation was introduced in Section 1. Throughout this paper, we will assume $\mu \leq 1$. This model is most interesting when $\mu \rightarrow 0$ as the size of the graph gets large. Note that if $\mu = \infty$, then $\{X_t\}_{t \geq 0}$ would simply be ordinary simple random walk on G with time scaled by p and hence would not be interesting. One would similarly expect that if μ is of order 1, the system should behave in various ways like ordinary random walk. (We will see for example that the usual recurrence/transience dichotomy for random walk holds in this model for fixed μ .) This is why $\mu \rightarrow 0$ is the interesting regime.

Mixing times for Markov chains.

We recall the following standard definitions. Given two probability measures m_1 and m_2 on a finite set S , we define the total variation distance between m_1 and m_2 to be

$$\|m_1 - m_2\|_{\text{TV}} := \frac{1}{2} \sum_{s \in S} |m_1(s) - m_2(s)|.$$

If X and Y are random variables, by $\|\mathcal{L}(X) - \mathcal{L}(Y)\|_{\text{TV}}$, we will mean the total variation distance between their laws. There are other equivalent definitions; see [23], Section 4.2. One which we will need is that

$$\|\mathcal{L}(X) - \mathcal{L}(Y)\|_{\text{TV}} = \inf\{\mathbf{P}(X' \neq Y')\}$$

where the infimum is taken over all pairs of random variables (X', Y') defined on the same space where X' has the same distribution as X and Y' has the same distribution as Y .

Given a continuous time finite state irreducible Markov chain with state space S , $t \geq 0$ and $x \in S$, we let $P^t(x, \cdot)$ be the distribution of the chain at time t when started in state x and we let π denote the unique stationary distribution for the chain.

Next, one defines

$$T_{\text{mix}}(\epsilon) := \inf\{t \geq 0: \max_{x \in S} \|P^t(x, \cdot) - \pi\|_{\text{TV}} \leq \epsilon\},$$

in which case the standard definition of the mixing time of a chain, denoted by T_{mix} , is $T_{\text{mix}}(1/4)$. It is well known (see [23], Section 4.5 for the discrete-time analogue) that $\max_x \|P^t(x, \cdot) - \pi\|_{\text{TV}}$ is decreasing in t and that

$$T_{\text{mix}}(\epsilon) \leq \lceil \log_2 \epsilon^{-1} \rceil T_{\text{mix}}. \quad (5)$$

In the theory of mixing times, one typically has a sequence of Markov chains that one is interested in and one studies the limiting behavior of the corresponding sequence of mixing times.

3 Recurrence/transience dichotomy

In this section, we prove Theorem 1.1 as well as a central limit theorem for the process $\{X_t\}$.

Proof of Theorem 1.1. Since d , p and μ are fixed, we drop these superscripts. We first prove this result when the bond configuration starts in state π_p ; at the end of the proof we extend this to a general initial bond configuration.

For this analysis, we let \mathcal{F}_t be the σ -algebra generated by $\{M_s\}_{0 \leq s \leq t}$ as well as, for each edge e , the times before t at which e is refreshed and at which the random walker attempted to cross e . We now define a sequence of sets $\{A_k\}_{k \geq 0}$. Let $A_0 = \emptyset$. For $k \geq 1$, define A_k to be the set of edges of A_{k-1} that did not refresh during $[k-1, k]$ plus the set of edges that the walker attempted to cross during $[k-1, k]$ which did not refresh before time k after the last time in $[k-1, k]$ that the walker attempted to cross it. Note that A_k is measurable with respect to \mathcal{F}_k . Let $\tau_0 = 0$ and, for $k \geq 1$, define

$$\tau_k = \min\{i > \tau_{k-1} : A_i = \emptyset\}.$$

We will see below that for all k , $\tau_k < \infty$ a.s. Note that the random variables $\{\tau_k - \tau_{k-1}\}_{k \geq 1}$ are i.i.d. For $k \geq 1$, let $U_k = X_{\tau_k} - X_{\tau_{k-1}}$. Clearly the $\{U_k\}_{k \geq 1}$ are i.i.d. and hence $\{X_{\tau_n}\}_{n \geq 0}$ is a random walk on \mathbb{Z}^d with step distribution U_1 . It is easy to check that U_1 takes the value 0 as well as any of the $2d$ neighbors of 0 with positive probability and hence the random walk is fully supported, irreducible and aperiodic.

Let J_i denote the number of attempted steps by the random walk during $[i-1, i]$ and $Z_k := \sum_{i=\tau_{k-1}+1}^{\tau_k} J_i$ be the number of attempted steps by the random walk between τ_{k-1} and τ_k . Now clearly $\{J_i\}_{i \geq 1}$ are i.i.d. as are $\{Z_k\}_{k \geq 1}$. A key step is to show that

$$\mathbf{E} [e^{cZ_1}] < \infty \text{ for some } c > 0. \quad (6)$$

(Note that this implies that each τ_k is finite a.s.) Assuming (6) for the moment, we finish the proof of (i) when the bond configuration is started in stationarity. (6) implies, since $\text{dist}(U_1, 0) \leq Z_1$, that $\text{dist}(U_1, 0)$ has an exponential tail and in particular a finite second moment. Since U_1 is obviously symmetric, it therefore necessarily has mean zero. The fact that $\{X_{\tau_n}\}_{n \geq 0}$ is recurrent in one and two dimensions now follows from [22, Theorem 4.1.1]. This proves (i) when the bond configuration is started in stationarity.

If $d \geq 3$, it follows from [22, Theorem 4.1.1] that $\{X_{\tau_n}\}_{n \geq 0}$ is transient and so approaches ∞ a.s. To show (ii), we need to deal with times between τ_k and τ_{k+1} . Fix M , let $B_{\text{dist}}(0, M)$ be the ball around 0 of dist-radius M and let E_k be the event that the random walk returns to $B_{\text{dist}}(0, M)$ during $[\tau_k, \tau_{k+1}]$, we have

$$\sum_{k=0}^{\infty} \mathbf{P}(E_k) \leq \sum_{k=0}^{\infty} \mathbf{P}\left(E_k \mid \text{dist}(X_{\tau_k}, 0) \geq k^{\frac{1}{4d}}\right) + \mathbf{P}\left(\text{dist}(X_{\tau_k}, 0) \leq k^{\frac{1}{4d}}\right).$$

Now $\mathbf{P}\left(E_k \mid \text{dist}(X_{\tau_k}, 0) \geq k^{\frac{1}{4d}}\right) \leq \mathbf{P}\left(Z_1 \geq k^{\frac{1}{4d}} - M\right)$ and since Z_1 has an exponential tail, the first terms are summable. Next, the local central limit theorem (cf. [22, Theorem 2.1.1]) implies that for large k , $\mathbf{P}\left(\text{dist}(X_{\tau_k}, 0) \leq k^{\frac{1}{4d}}\right) \leq \frac{Ck^{\frac{1}{4d}}}{k^{\frac{d}{2}}}$ for some constant C . Since $d \geq 3$, it follows that the second terms are also summable. It now follows from Borel-Cantelli that the walker eventually leaves $B_{\text{dist}}(0, M)$ a.s. and hence by countable additivity (ii) holds when the bond configuration is started in stationarity.

We will now verify (6) by using Proposition 3.2 with $Y_k = |A_k| + 1$ and \mathcal{F}_k being itself. Property (1) is immediate. With J_k as above and R_k being the number of edges of A_{k-1} that were refreshed during $[k-1, k]$, it is easy to see that

$$|A_k| \leq |A_{k-1}| - R_k + J_k.$$

This implies that

$$\mathbf{E}[|A_k| \mid \mathcal{F}_{k-1}] \leq \mathbf{E}[|A_{k-1}| - R_k + J_k \mid \mathcal{F}_{k-1}] = |A_{k-1}|e^{-\mu} + 1.$$

This easily yields that there are positive numbers $a_0 = a_0(\mu)$ and $b_0 = b_0(\mu) < 1$ so that $\mathbf{E}[Y_k \mid \mathcal{F}_{k-1}] \leq b_0 Y_{k-1}$ on the event $Y_{k-1} > a_0$. This verifies property (2). Since properties (3) and (4) are easily checked, we can conclude from Proposition 3.2 at the end of this section that τ_1 has some positive exponential moment. An application of Lemma 3.3 to τ_1 and the J_i 's allows us to conclude that (6) holds. This completes the proof when the bond configuration starts in stationarity.

We now analyze the situation starting from an arbitrary bond configuration. Let E_n be the event that some vertex whose dist-distance to the origin is n has an adjacent edge which does not refresh by time \sqrt{n} . Let H_n be the event that the number of attempted steps by the random walker by time \sqrt{n} is larger than n . It is elementary to check that

$$\sum_n \mathbf{P}[E_n] < \infty \text{ and } \sum_n \mathbf{P}[H_n] < \infty.$$

By Borel-Cantelli, given $\epsilon > 0$ there exists n_0 such that

$$\mathbf{P}[\cap_{n \geq n_0} E_n^c \cap \cap_{n \geq n_0} H_n^c] \geq 1 - \epsilon.$$

Now let η be an arbitrary initial bond configuration and let η^p be the random configuration which is the same as η at edges within distance n_0 of the origin and otherwise is chosen according to π_p .

We claim that by what we have already proved, we can infer that when the initial bond configuration is η^p , the random walker returns to 0 at arbitrarily large times a.s. if d is 1 or 2 and converges to ∞ a.s. if $d \geq 3$. To see this, first observe that by Fubini's Theorem, we can infer from what we have proved that for π_p -a.s. bond configuration, the random walker has the desired behavior a.s. Therefore, since such a random bond configuration takes the same values as η at edges within

distance n_0 of the origin with positive probability, it must be the case that a.s. η^p is such that the random walk has the desired behavior. Using Fubini's Theorem again demonstrates this claim.

One can next couple the random walker when the initial bond configuration is η with the random walker when the initial bond configuration is η^p in the obvious way. They will remain together provided $\cap_{n \geq n_0} E_n^c$ holds and $\cap_{n \geq n_0} H_n^c$ holds for the walks. Therefore the random walker with initial bond configuration η has the claimed behavior with probability $1 - \epsilon$. As ϵ is arbitrary, we are done. \square

We now provide a central limit theorem for the walker.

Theorem 3.1. *Given $d, p \in (0, 1)$ and μ , there exists $\sigma \in (0, \infty)$ so that random walk in dynamical percolation on \mathbb{Z}^d with parameters p and μ started from an arbitrary configuration satisfies*

$$\left\{ \frac{X_{kt}}{\sqrt{k}} \right\}_{t \in [0,1]} \Rightarrow \{B_t\}_{t \in [0,1]}$$

in $C[0, 1]$ as $k \rightarrow \infty$ where $\{B_t\}_{t \in [0,1]}$ is a standard d -dimensional Brownian motion with variance σ^2 . Moreover

$$\sigma^2 = \frac{\text{Var}(U_1^{(1)})}{\mathbf{E}[\tau_1]}$$

where U_1 and τ_1 are given in the proof of Theorem 1.1 and $U_1^{(1)}$ is the first coordinate of U_1 .

Proof. This type of argument is very standard and so we only sketch the proof. We therefore only mention that a tightness argument is needed to prove the convergence for the fixed value $t = 1$. We also assume that the bond configuration is started in stationarity with distribution π_p and that the walker starts at the origin. To deal with general initial bond configurations, the methods in the proof of Theorem 1.1 can easily be adapted.

Now, symmetry considerations give that $\mathbf{E}[U_1^{(1)}] = 0$ and that the different coordinates are uncorrelated. We have also seen that $U_1^{(1)}$ has a finite second moment. The central limit theorem now tells us that

$$\frac{X_{\tau_n}}{\sqrt{n} \text{Var}(U_1^{(1)})^{1/2}} = \frac{\sum_1^n U_i}{\sqrt{n} \text{Var}(U_1^{(1)})^{1/2}} \Rightarrow \mathcal{N} \tag{7}$$

where \mathcal{N} is a standard d -dimensional Gaussian.

We need to show that $\frac{X_k}{\sqrt{k}}$ converges to the appropriate Gaussian. Let $n(k) := \lceil k / \mathbf{E}[\tau_1] \rceil$ and write $\frac{X_k}{\sqrt{k}}$ as

$$\frac{\sqrt{n(k)}}{\sqrt{k}} \left[\frac{X_k - X_{n(k)} \mathbf{E}[\tau_1]}{\sqrt{n(k)}} + \frac{X_{n(k)} \mathbf{E}[\tau_1] - X_{\tau_{n(k)}}}{\sqrt{n(k)}} + \frac{X_{\tau_{n(k)}}}{\sqrt{n(k)}} \right].$$

The first fraction converges to $1/\sqrt{\mathbf{E}[\tau_1]}$. The first fraction in the second factor is easily shown to converge in probability to 0. The weak law of large numbers gives that $\tau_{n(k)}/n(k)$ converges in probability to $\mathbf{E}[\tau_1]$ which easily leads to the second fraction in the second factor converging in probability to 0. Finally, using (7) for the last fraction, we obtain the result. \square

Technical lemma.

We now present the technical lemma which was used in the previous proof and will be used again later on. This result is presumably well known in some form but we provide the proof nevertheless for completeness. As the result is ‘‘obvious’’, the reader might choose to skip the proof.

Proposition 3.2. For all $\alpha \geq 1$, $\delta < 1$, $\epsilon > 0$ and γ , there exist $c_{3.2} = c_{3.2}(\alpha, \delta, \epsilon, \gamma) > 0$ and $C_{3.2} = C_{3.2}(\alpha, \delta, \epsilon, \gamma) < \infty$ with the following property. If $\{Y_i\}_{i \geq 0}$ is a discrete time process taking values in $\{1, 2, \dots\}$ adapted to a filtration $\{\mathcal{F}_i\}_{i \geq 0}$ satisfying

- (1) $Y_0 = 1$,
 - (2) for all i , $\mathbf{E}[Y_{i+1} | \mathcal{F}_i] \leq \delta Y_i$ on $Y_i > \alpha$,
 - (3) for all i , $\mathbf{P}[Y_{i+1} = 1 | \mathcal{F}_i] \geq \epsilon$ on $Y_i \leq \alpha$ and
 - (4) for all i , $\mathbf{E}[Y_{i+1} | \mathcal{F}_i] \leq \gamma$ on $Y_i \leq \alpha$,
- and if $T := \min\{i \geq 1 : Y_i = 1\}$, then

$$\mathbf{E}[e^{c_{3.2}T}] \leq C_{3.2}$$

and so in particular by Jensen's inequality

$$\mathbf{E}[T] \leq \frac{\log C_{3.2}}{c_{3.2}}.$$

To prove this, we will need a slight strengthening of a lemma from [24] which essentially follows the same proof.

Lemma 3.3. Given positive numbers λ and a , there exist $c_{3.3} = c_{3.3}(\lambda, a) > 0$ and $C_{3.3} = C_{3.3}(\lambda, a) < \infty$ so that if $\{X_i\}$ are nonnegative random variables adapted to a filtration $\{\mathcal{G}_i\}_{i \geq 0}$ satisfying

$$\|\mathbf{E}[e^{\lambda X_{i+1}} | \mathcal{G}_i]\|_\infty \leq a \text{ for all } i \geq 0, \quad (8)$$

and M is a nonnegative integer valued random variable satisfying

$$\mathbf{E}[e^{\lambda M}] \leq a, \quad (9)$$

then

$$\mathbf{E}[e^{c_{3.3} \sum_{i=1}^M X_i}] \leq C_{3.3}. \quad (10)$$

(Note that this implies, by Jensen's inequality, that $\mathbf{E}[\sum_{i=1}^M X_i] \leq \frac{\log C_{3.3}}{c_{3.3}}.$)

Proof. Choose $k = k(\lambda, a)$ sufficiently large so that $a < e^{\lambda k}$. We claim there exist $b < 1$ and B , depending only on λ and a , such that for all n

$$\mathbf{P}\left[\sum_{i=1}^M X_i \geq kn\right] \leq Bb^n. \quad (11)$$

To see this, note that this above probability is at most

$$\mathbf{P}[M \geq n] + \mathbf{P}\left[\sum_{i=1}^n X_i \geq kn\right] \leq ae^{-\lambda n} + \frac{\mathbf{E}[e^{\lambda \sum_{i=1}^n X_i}]}{e^{\lambda kn}}$$

by Markov's inequality and (9). Using (8), taking conditional expectations and iterating, one sees that

$$\mathbf{E}[e^{\lambda \sum_{i=1}^n X_i}] \leq a^n$$

and hence the second term is at most $(\frac{a}{e^{\lambda k}})^n$. We can conclude that there are b and B , depending only on λ and a , such that (11) holds. Since b, B and k all depend only on λ and a , it then easily follows from (11) that there are $c_{3.3}$ and $C_{3.3}$ depending only on λ and a so that (10) holds. \square

Proof of Proposition 3.2. Let $U_0 = 0$ and for $k \geq 1$, let $U_k := \min\{i \geq U_{k-1} + 1 : Y_i \leq \alpha\}$. Let $M := \min\{k \geq 0 : Y_{U_k+1} = 1\}$. Property (3) implies that M is stochastically dominated by a geometric random variable with parameter ϵ . Next, clearly

$$T \leq 1 + \sum_{k=1}^M (U_k - U_{k-1}). \quad (12)$$

(Equality does not necessarily hold since it is possible that $Y_{T-1} > \alpha$ in which case $T - 1$ does not correspond to any U_k .) We claim that for all $k \geq 1$

$$\left\| \mathbf{E} \left[\left(\frac{1}{\delta} \right)^{U_k - U_{k-1}} \mid \mathcal{F}_{U_{k-1}} \right] \right\|_{\infty} \leq \frac{\gamma}{\delta}. \quad (13)$$

We write

$$\mathbf{E} \left[\left(\frac{1}{\delta} \right)^{U_k - U_{k-1}} \mid \mathcal{F}_{U_{k-1}} \right] = \mathbf{E} \left[\mathbf{E} \left[\left(\frac{1}{\delta} \right)^{U_k - U_{k-1}} \mid \mathcal{F}_{U_{k-1}+1} \right] \mid \mathcal{F}_{U_{k-1}} \right].$$

Concerning the inner conditional expectation, we claim that

$$\mathbf{E} \left[\left(\frac{1}{\delta} \right)^{U_k - U_{k-1}} \mid \mathcal{F}_{U_{k-1}+1} \right] \leq \frac{Y_{U_{k-1}+1}}{\delta}. \quad (14)$$

Case 1: $Y_{U_{k-1}+1} \leq \alpha$.

In this case, $U_k = U_{k-1} + 1$ and so

$$\mathbf{E} \left[\left(\frac{1}{\delta} \right)^{U_k - U_{k-1}} \mid \mathcal{F}_{U_{k-1}+1} \right] = \frac{1}{\delta}.$$

Case 2: $Y_{U_{k-1}+1} > \alpha$.

In this case, we first make the important observation that property (2) implies that on the event $Y_{U_{k-1}+1} > \alpha$,

$$M_j := \left(\frac{1}{\delta} \right)^{j \wedge (U_k - U_{k-1} - 1)} Y_{(j \wedge (U_k - U_{k-1} - 1)) + U_{k-1} + 1}, \quad j \geq 0$$

is a supermartingale with respect to $\{\mathcal{F}_{j+U_{k-1}+1}\}_{j \geq 0}$. From the theory of nonnegative supermartingales (see [17], Chapter 5), we can let $j \rightarrow \infty$ in the defining inequality of a supermartingale and conclude that

$$\mathbf{E} \left[\left(\frac{1}{\delta} \right)^{U_k - U_{k-1} - 1} Y_{U_k} \mid \mathcal{F}_{U_{k-1}+1} \right] \leq Y_{U_{k-1}+1},$$

It follows that

$$\mathbf{E} \left[\left(\frac{1}{\delta} \right)^{U_k - U_{k-1}} \mid \mathcal{F}_{U_{k-1}+1} \right] \leq \frac{Y_{U_{k-1}+1}}{\delta}. \quad (15)$$

Since the Y_i 's are at least 1, this establishes (14).

Taking the conditional expectation of the two sides of (14) with respect to $\mathcal{F}_{U_{k-1}}$ and using property (4) finally yields (13).

Lastly, Lemma 3.3 (with $X_i = U_i - U_{i-1}$, $\mathcal{G}_i = \mathcal{F}_{U_i}$, $M = M$, $\lambda = \min\{\frac{\epsilon}{2}, \log(\frac{1}{\delta})\}$ and $a = \max\{\frac{\gamma}{\delta}, 2\}$) together with (12), (13) and the fact that M is dominated by a geometric random variable with parameter ϵ gives us the desired conclusion. \square

Remark 3.4. We see that $c_{3.2}$ and $C_{3.2}$ actually depend only on δ, ϵ and γ but not on α .

4 Proofs of the mixing time lower bound in the subcritical case

In this section, we prove Theorem 1.4, Theorem 1.2(ii) and Corollary 1.6. We begin with the proof of Theorem 1.4 as this will be used in the proof of Theorem 1.2(ii).

Proof of Theorem 1.4. Fix d and $p \in (0, p_c(\mathbb{Z}^d))$. Choose $\beta = \beta(d, p)$ sufficiently small so that for all n and μ , the probability that, for $\{\eta_t\}_{t \geq 0}$, a fixed edge e is open at some point in $[0, \beta/\mu]$ is less than $p_c(\mathbb{Z}^d)$. By time scaling, the latter probability is independent of μ (and of course of n).

The main idea of the proof is that by setting β as above we can easily obtain an upper bound for $\mathbf{E} \left[\text{dist}(X_{\frac{\beta}{\mu}}, X_0)^2 \right]$ (cf. Lemma 4.1 below), and then apply a general result for discrete-time Markov chains (Lemma 4.2 below) which allows us to bound $\mathbf{E} \left[\text{dist}(X_{\frac{t}{\mu}}, X_0)^2 \right]$ in terms of $\mathbf{E} \left[\text{dist}(X_{\frac{\beta}{\mu}}, X_0)^2 \right]$. In order to carry out this last step, we introduce the function g_n below that measures distance over the torus in a clean way, and which is bi-Lipschitz so that an upper bound on $\mathbf{E} \left[(g_n(X_{\frac{t}{\mu}}) - g_n(X_0))^2 \right]$ translates to an upper bound on $\mathbf{E} \left[\text{dist}(X_{\frac{t}{\mu}}, X_0)^2 \right]$. We now proceed to the details of the proof.

Let $g_n : \mathbb{T}^{d,n} \rightarrow \mathbb{R}^{2d}$ be given by

$$g_n(x_1, \dots, x_d) := (n \cos(2\pi x_1/n), n \sin(2\pi x_1/n), \dots, n \cos(2\pi x_d/n), n \sin(2\pi x_d/n)). \quad (16)$$

Observe that for fixed d , the functions $\{g_n\}_{n \geq 1}$ are uniformly bi-Lipschitz when $\mathbb{T}^{d,n}$ is equipped with the metric dist and \mathbb{R}^{2d} has its usual metric. Let $C_{\text{Lip}} = C_{\text{Lip}}(d)$ be a uniform bound on the bi-Lipschitz constants.

We need the following two lemmas. The first lemma will be proved afterwards while the second lemma, which is implicitly contained in [7], is stated explicitly in [25]; in fact a strengthening of it yielding a maximal version is proved in [25].

Lemma 4.1. *There exists $C_{4.1} = C_{4.1}(d, p)$ so that for all n , for all μ and for all $s \leq \beta$,*

$$\mathbf{E} \left[\text{dist}(X_{\frac{s}{\mu}}, X_0)^2 \right] \leq C_{4.1}$$

when we start the full system in stationarity.

Lemma 4.2. *Let $\{Y_i\}_{i \in \mathbb{Z}}$ be a discrete time stationary reversible Markov chain with finite state space S and let $h : S \rightarrow \mathbb{R}^m$. Then for each $k \geq 1$,*

$$\mathbf{E} \left[\|h(Y_k) - h(Y_0)\|_{L^2}^2 \right] \leq k \mathbf{E} \left[\|h(Y_1) - h(Y_0)\|_{L^2}^2 \right]$$

where $\|\cdot\|_{L^2}$ denotes the Euclidean norm on \mathbb{R}^m .

We may assume that $\beta \leq 1$. For $t \leq \beta(\leq 1)$, the LHS of (1) is by Lemma 4.1 at most $C_{4.1}$ which is at most $C_{1.4}(t \vee 1)$ if $C_{1.4}$ is taken to be larger than $C_{4.1}$.

On the other hand, if $t \geq \beta$, choose $\ell \in \mathbb{N}$ so that $v := \frac{t}{\ell} \in [\beta/2, \beta]$. Consider the discrete time finite state stationary reversible Markov chain given by

$$Y_k := M_{\frac{k v}{\mu}}, \quad k \in \mathbb{Z}$$

with state space $S := \mathbb{T}^{d,n} \times \{0, 1\}^{E(\mathbb{T}^{d,n})}$. With all the parameters for the chain fixed, let $h_n : S \rightarrow \mathbb{R}^{2d}$ be given by $h_n(x, \eta) := g_n(x)$. Then Lemma 4.1 (with $s = v$) together with the uniform bi-Lipschitz property of the g_n 's implies that

$$\mathbf{E} [\|h_n(Y_1) - h_n(Y_0)\|_{L^2}^2] \leq C_{\text{Lip}}^2 C_{4.1}.$$

We now can apply Lemma 4.2 with $k = \ell$ and obtain

$$\mathbf{E} [\|g_n(X_{\frac{t}{\mu}}) - g_n(X_0)\|_{L^2}^2] \leq C_{\text{Lip}}^2 C_{4.1} \ell.$$

Since $\frac{t}{\ell} \in [\beta/2, \beta]$, we have that $\ell \leq 2t/\beta$. Using this and the bi-Lipschitz property of the g_n 's again, we obtain

$$\mathbf{E} [\text{dist}(X_{\frac{t}{\mu}}, X_0)^2] \leq 2C_{\text{Lip}}^4 C_{4.1} t/\beta.$$

As all of the terms except t on the RHS only depend on d and p , we are done. \square

The proof of Lemma 4.1 requires the following important result concerning subcritical percolation. For $V' \subseteq V$, we let $\text{Diam}(V') := \max\{\text{dist}(x, y) : x, y \in V'\}$ denote the diameter of V' . This following theorem is Theorem 5.4 in [19] in the case of \mathbb{Z}^d . The statement for \mathbb{Z}^d immediately implies the result for $\mathbb{T}^{d,n}$.

Theorem 4.3. *For any $d \geq 1$ and $\alpha \in (0, p_c(\mathbb{Z}^d))$, there exists $C_{4.3} = C_{4.3}(d, \alpha) > 0$ so that for all $r \geq 1$,*

$$\mathbf{P}_{\mathbb{Z}^d, \alpha}(\text{Diam}(C(0)) \geq r) \leq e^{-C_{4.3}r}.$$

The previous line holds with \mathbb{Z}^d replaced by $\mathbb{T}^{d,n}$.

We now give the

Proof of Lemma 4.1. Let $\bar{\eta}$ be the set of edges which are open some time during $[0, \beta/\mu]$. By our choice of β , there exists $p_0 = p_0(d, p) < p_c(\mathbb{Z}^d)$ so that for all n and all μ , the distribution of $\bar{\eta}$ is π_{p_0} .

Letting $C_{\bar{\eta}}(x)$ denote the cluster of x with respect to the bond configuration $\bar{\eta}$, the observation above and Theorem 4.3 implies that there exists a constant $C_{4.1.1} = C_{4.1.1}(d, p)$ so that for all n , for all μ and for all $x \in \mathbb{T}^{d,n}$,

$$\mathbf{E} [(\text{Diam}(C_{\bar{\eta}}(x)))^2] \leq C_{4.1.1}.$$

By independence of X_0 and $\bar{\eta}$, we get

$$\mathbf{E} [(\text{Diam}(C_{\bar{\eta}}(X_0)))^2] \leq C_{4.1.1}. \tag{17}$$

Since the random walker can only move along $\bar{\eta}$ during $[0, \beta/\mu]$, we have that for all $s \leq \beta$, $X_{\frac{s}{\mu}}$ necessarily belongs to $C_{\bar{\eta}}(X_0)$ and hence $\text{dist}(X_{\frac{s}{\mu}}, X_0) \leq \text{Diam}(C_{\bar{\eta}}(X_0))$. The result now follows from (17). \square

We next move to the

Proof of Corollary 1.6. Fix μ and t . By Theorem 1.4 and symmetry, we have that for each n ,

$$\mathbf{E} [\text{dist}(X_{\frac{t}{\mu}}, 0)^2 | \delta_0 \times \pi_p] \leq C_{1.4}(t \vee 1).$$

Clearly $\text{dist}(X_{\frac{t}{\mu}}, 0) | \delta_0 \times \pi_p$ converges in distribution to $\text{dist}(X_{\frac{t}{\mu}}, 0)$ as $n \rightarrow \infty$ where the latter is started in $\delta_0 \times \pi_p$. The result now follows by squaring and applying Fatou's lemma. \square

Proof of Theorem 1.2(ii). Fix $d, p \in (0, p_c(\mathbb{Z}^d))$ and $\epsilon < 1$. It suffices to show that there exists $\delta = \delta(d, p, \epsilon) > 0$ and $n_0 = n_0(d, p, \epsilon) > 0$ so that for $n \geq n_0$ and $s \leq \frac{\delta n^2}{\mu}$

$$\|\mathcal{L}(X_s) - u\|_{\text{TV}} > \epsilon.$$

By symmetry, the distribution of $\text{dist}(X_{\frac{t}{\mu}}, X_0)$ conditioned on $\{X_0 = a\}$ does not depend on a . Hence by Theorem 1.4 and Markov's inequality, we have that for all $\lambda > 0$, for all n , for all μ and for all t

$$\mathbf{P}\left(\text{dist}(X_{\frac{t}{\mu}}, 0) \geq \lambda|\delta_0 \times \pi_p\right) \leq C_{1.4}(t \vee 1)/\lambda^2 \quad (18)$$

where $C_{1.4}$ comes from Theorem 1.4. Next, choose $b = b(d, \epsilon) > 0$ so that

$$(2b)^d < \frac{1 - \epsilon}{2}.$$

We then have that there exists $n_0 = n_0(d, p, \epsilon) > 0$ sufficiently large so that for all $n \geq n_0$ we have that

$$|\{x \in \mathbb{T}^{d,n} : \text{dist}(x, 0) \leq bn\}| \leq \frac{(1 - \epsilon)n^d}{2} \quad (19)$$

and

$$\frac{C_{1.4}}{(bn)^2} < \frac{1 - \epsilon}{2}.$$

Next choose $\delta = \delta(d, p, \epsilon) > 0$ so that

$$\frac{C_{1.4}\delta}{b^2} < \frac{1 - \epsilon}{2}.$$

We now let $n \geq n_0$ and $s \leq \frac{\delta n^2}{\mu}$. Applying (18) with $t = s\mu \leq \delta n^2$ and $\lambda = bn$ yields

$$\mathbf{P}(\text{dist}(X_s, 0) \geq bn|\delta_0 \times \pi_p) \leq \frac{C_{1.4}(\delta n^2 \vee 1)}{(bn)^2} < \frac{1 - \epsilon}{2}. \quad (20)$$

Letting $E_n := \{x \in \mathbb{T}^{d,n} : \text{dist}(x, 0) \leq bn\}$, we have by (20)

$$\mathbf{P}(X_s \in E_n|\delta_0 \times \pi_p) \geq \frac{1 + \epsilon}{2}$$

and by (19), we have

$$u(E_n) \leq \frac{1 - \epsilon}{2}.$$

Hence, by considering the set E_n , it follows that $\|\mathcal{L}(X_s) - u\|_{\text{TV}} \geq \frac{1+\epsilon}{2} - \frac{1-\epsilon}{2} = \epsilon$, completing the proof. \square

We end this section by proving that not only is the mixing time for the full system of order at least n^2/μ but that this is also a lower bound on the relaxation time. Moreover, in proving this, we will only use Lemma 4.1 and do not need to appeal to the so-called Markov type inequality contained in Lemma 4.2.

Proposition 4.4. *For any d and $p \in (0, p_c(\mathbb{Z}^d))$, there exists $C_{4.4} = C_{4.4}(d, p) > 0$ such that, for all n and for all μ , the relaxation time of the full system is at least $\frac{C_{4.4}n^2}{\mu}$.*

Remarks 4.5. While it is a general fact that the mixing time is bounded below by (a universal constant times) the relaxation time, this does not provide an alternative proof of Theorem 1.2(ii) for two reasons. First, in the latter, we have a lower bound for the “mixing time” of the walker (which is stronger than just having a lower bound on the mixing time for the full system) and secondly ϵ in Theorem 1.2(ii) can be taken close to 1 while one could only conclude this for $\epsilon < \frac{1}{2}$ directly from a lower bound on the relaxation time.

Proof. We will obtain an upper bound on the spectral gap by considering the usual Dirichlet form; see Section 13.3 in [23]. Consider the function $f_n : \mathbb{T}^{d,n} \times \{0, 1\}^{E(\mathbb{T}^{d,n})} \rightarrow \mathbb{R}$ given by $f_n(x, \eta) := \text{dist}(x, 0)$ where 0 is the origin in $\mathbb{T}^{d,n}$. Clearly, there exists a constant $C_{4.4.1} = C_{4.4.1}(d) > 0$ such that for all d, p, n and μ ,

$$\text{Var}(f_n) \geq C_{4.4.1} n^2$$

where $\text{Var}(f_n)$ denotes the variance of f_n with respect to the stationary distribution.

Letting β be defined as in the proof of Theorem 1.4, Lemma 4.1 and the triangle inequality imply that

$$\mathbf{E} \left[|f_n(M_{\frac{\beta}{\mu}}) - f_n(M_0)|^2 \right] \leq 4C_{4.1}.$$

Hence

$$\frac{\mathbf{E} \left[|f_n(M_{\frac{\beta}{\mu}}) - f_n(M_0)|^2 \right]}{2\text{Var}(f_n)} \leq \frac{2C_{4.1}}{C_{4.4.1} n^2}.$$

By Section 13.3 in [23], we conclude that the spectral gap for the discrete time process viewed at times $0, \frac{\beta}{\mu}, \frac{2\beta}{\mu}, \dots$ is at most $\frac{2C_{4.1}}{C_{4.4.1} n^2}$. If $-\lambda = -\lambda(d, p, n, \mu)$ is the nonzero eigenvalue of minimum absolute value for the infinitesimal generator of the continuous time process (in which case λ is the spectral gap for the continuous time process), then the spectral gap for the above discrete time process is $1 - e^{-\frac{\lambda\beta}{\mu}}$ and so

$$1 - e^{-\frac{\lambda\beta}{\mu}} \leq \frac{2C_{4.1}}{C_{4.4.1} n^2}.$$

We can conclude that for large n , for any μ , we have that $\frac{\lambda\beta}{\mu} \leq \frac{1}{2}$. Since $1 - e^{-x} \geq x/2$ on $[0, 1]$, we conclude that

$$\frac{\lambda\beta}{2\mu} \leq \frac{2C_{4.1}}{C_{4.4.1} n^2}$$

or

$$\lambda \leq \frac{4C_{4.1}\mu}{\beta C_{4.4.1} n^2}.$$

Since the relaxation time is the reciprocal of the spectral gap, we are done. \square

5 Proofs of the mixing time lower bounds in the general case

In this section, we prove Theorem 1.9, Theorem 1.7 and Corollary 1.10. We begin with the proof of Theorem 1.9 as this will be used in the proof of Theorem 1.7.

Proof of Theorem 1.9. Clearly, $\text{dist}(X_a, X_0)$ is stochastically dominated by a Poisson random variable with parameter a . It follows that

$$\mathbf{E} [\text{dist}(X_s, X_0)^2] \leq s + s^2. \quad (21)$$

(This will be used in the same way that Lemma 4.1 was used.) (21) tells us that (3) holds for $t \leq 1$ if $C_{1.9} \geq 2$.

If $t \geq 1$, choose $\ell \in \mathbb{N}$ so that $v := \frac{t}{\ell} \in [1/2, 1]$. Consider the discrete time finite state stationary reversible Markov chain given by

$$Y_k := M_{kv}, \quad k \in \mathbb{Z}.$$

Letting S and h_n be as in the proof of Theorem 1.4, (21) (with $s = v \leq 1$) together with the uniform bi-Lipschitz property of the g_n 's implies that

$$\mathbf{E} [\|h(Y_1) - h(Y_0)\|_{L^2}^2] \leq 2C_{\text{Lip}}^2.$$

Lemma 4.2 with $k = \ell$ now yields

$$\mathbf{E} [\|g_n(X_t) - g_n(X_0)\|_{L^2}^2] \leq 2C_{\text{Lip}}^2 \ell.$$

Since $\frac{t}{\ell} \in [1/2, 1]$, we have that $\ell \leq 2t$. Using this and the bi-Lipschitz property of the g_n 's again, we obtain

$$\mathbf{E} [\text{dist}(X_t, X_0)^2] \leq 4C_{\text{Lip}}^4 t.$$

Letting $C_{1.9} := 4C_{\text{Lip}}^4 (\geq 2)$, we obtain (3) for $t \geq 1$ as well. \square

Proof of Corollary 1.10. This can be obtained from Theorem 1.9 in the exact same way as Corollary 1.6 was obtained from Theorem 1.4. \square

Proof of Theorem 1.7. (i). One can check that in the same way that Theorem 1.2(ii) is proved using Theorem 1.4, one can use Theorem 1.9 to prove this part. The details are left to the reader.

(ii). Fix $\epsilon < 1 - \theta_d(p)$. Let $\rho = \rho(d, p, \epsilon) := \sqrt{\frac{1 - \theta_d(p) + \epsilon}{2(1 - \theta_d(p))}} \in (0, 1)$. By countable additivity, there exists $\kappa = \kappa(d, p, \epsilon)$ so that $\mathbf{P}_{\mathbb{Z}^d, p}(|C(0)| \leq \kappa) \geq (1 - \theta_d(p))\rho$. For $n > \kappa$, we therefore have that $\mathbf{P}_{\mathbb{T}^{d, n, p}}(|C(0)| \leq \kappa) \geq (1 - \theta_d(p))\rho$. Choose $C_{1.7.2} = C_{1.7.2}(d, p, \epsilon)$ sufficiently small so that $e^{-2d\kappa C_{1.7.2}} \geq \rho$.

Now, let $C_t(0)$ be the cluster of the origin at time t . For any n larger than κ and for any μ , conditioned on $\{|C_0(0)| \leq \kappa\}$, the conditional probability that no edges adjacent to $C_0(0)$ refresh during $[0, \frac{C_{1.7.2}}{\mu}]$ is at least $e^{-2d\kappa C_{1.7.2}}$ which was chosen larger than ρ . If $|C_0(0)| \leq \kappa$ and no edges adjacent to $C_0(0)$ refresh during $[0, \frac{C_{1.7.2}}{\mu}]$, then it is necessarily the case that $\text{dist}(X_{\frac{C_{1.7.2}}{\mu}}, 0) \leq \kappa$. Hence

$$\mathbf{P} \left(\text{dist}(X_{\frac{C_{1.7.2}}{\mu}}, 0) \leq \kappa \mid \delta_0 \times \pi_p \right) \geq (1 - \theta_d(p))\rho^2 = \frac{1 - \theta_d(p) + \epsilon}{2} > \epsilon.$$

Letting $E_n := \{x \in \mathbb{T}^{d, n} : \text{dist}(x, 0) \leq \kappa\}$, we therefore have

$$\mathbf{P} \left(X_{\frac{C_{1.7.2}}{\mu}} \in E_n \mid \delta_0 \times \pi_p \right) > \epsilon.$$

On the other hand, it is clear that $u(E_n)$ goes to 0 as $n \rightarrow \infty$. This demonstrates (2) and completes the proof. \square

We end this section by stating a proposition concerning the relaxation time analogous to Proposition 4.4 which holds for all p . The proof of this proposition follows the proof of Proposition 4.4 in a similar way to how the proof of Theorem 1.9 followed the proof of Theorem 1.4. The details are left to the reader.

Proposition 5.1. *For any d , there exists $C_{5.1} = C_{5.1}(d) > 0$ such that, for all n, p and μ , the relaxation time of the full system is at least $C_{5.1}n^2$.*

6 Proof of the mixing time upper bound in the subcritical case

In this section, we prove Theorem 1.2(i). This section is broken into four subsections. The first sets up the key technique of increasing the state space, the second gives a sketch of the proof, the third provides some percolation preliminaries and the fourth finally gives the proof.

6.1 Increasing the state space in the general case

In this subsection, we fix an arbitrary graph $G = (V, E)$ with constant degree and parameters p and μ and consider the resulting random walk in dynamical percolation which we denote by $\{M_t\}_{t \geq 0} = \{(X_t, \eta_t)\}_{t \geq 0}$. In order to obtain upper bounds on the mixing time, it will be useful to introduce another Markov process which we denote by $\{\tilde{M}_t\}_{t \geq 0} = \{(X_t, \tilde{\eta}_t)\}_{t \geq 0}$ which will incorporate more information than $\{M_t\}_{t \geq 0}$; the extra information will be the set of edges that the random walker has attempted to cross since their last refresh time. The state space for this Markov process will be

$$\Omega := \{(v, \tilde{\eta}) \in V \times \{0, 1, 0^*, 1^*\}^E : \tilde{\eta}(e) \in \{0, 1\} \text{ for each } e \text{ adjacent to } v\}. \quad (22)$$

If we identify 0^* with 0 and 1^* with 1, we want to recover our process $\{M_t\}_{t \geq 0}$. The idea of the possible extra \star for the state of the edge e is that this will indicate that the walker has not touched the endpoints of e since e 's last refresh time. Hence, for such an edge e , whether there is a 1^* or 0^* at e at that time is independent of everything else.

With the above in mind, it should be clear that we should define $\{\tilde{M}_t\}_{t \geq 0}$ as follows. An edge refreshes itself at rate μ . Independent of everything else before the refresh time, the state of the edge after the refresh time will be 1^* with probability p and 0^* with probability $1 - p$ unless the edge is adjacent to the walker at that time. If the latter is the case, then the state of the edge after the refresh time will instead be 1 with probability p and 0 with probability $1 - p$. The random walker will as before choose at rate 1 a uniform neighbor (in the original graph) and move along that edge if the edge is in state 1 and not if the edge is in state 0. (Note that this edge can only be in state 1 or 0 since it is adjacent to the walker.) Finally, when the random walker moves along an edge, the \star 's are removed from all edges which become adjacent to the walker. Clearly, dropping \star 's recovers the original process $\{M_t\}_{t \geq 0}$. We call an edge open if its state is 1^* or 1 and closed otherwise.

In order to exploit the \star -edges, we want that conditioned on (i) the position of the walker, (ii) the collection of \star -edges and (iii) the states of the non- \star -edges, we have no information concerning the states of the \star -edges. This is not necessarily true for all starting distributions. We therefore restrict ourselves to a certain class of distributions. To define this, we first let

$$\Pi : \{0, 1, 0^*, 1^*\}^E \rightarrow \{0, 1, \star\}^E$$

be defined by identifying 1^* and 0^* .

Definition 6.1. *A probability measure on Ω (as defined in (22)) is called good if conditioned on $(v, \Pi(\eta))$, the conditional distribution of η at the \star -edges is i.i.d. 1^* with probability p and 0^* with probability $1 - p$.*

Note that if ν is supported on $V \times \{0, 1\}^E$, then ν is good.

We will let $\{\mathcal{F}_t\}_{t \geq 0}$ be the natural filtration of σ -algebras generated by $\{M_t\}_{t \geq 0}$ which also keeps track of all the refresh times and the attempted steps made by the walker. Note that $\{\tilde{M}_t\}_{t \geq 0}$ is

measurable with respect to this filtration. Next, let $\{\mathcal{F}_t^\star\}_{t \geq 0}$ be the smaller filtration of σ -algebras which is obtained when one does not distinguish 1^\star and 0^\star but is otherwise the same. This filtration will be *critical* for our analysis.

A key property of good distributions, which also indicates the importance of the filtration $\{\mathcal{F}_t^\star\}_{t \geq 0}$, is given in the following obvious lemma, whose proof is left to the reader.

Lemma 6.2. *If the starting distribution for the Markov process $\{\tilde{M}_t\}_{t \geq 0}$ is good, then, for all s , the conditional distribution of \tilde{M}_s given \mathcal{F}_s^\star is good, as is the unconditional distribution of \tilde{M}_s . More generally, if S is a $\{\mathcal{F}_t^\star\}_{t \geq 0}$ stopping time, then the conditional distribution of \tilde{M}_S given \mathcal{F}_S^\star is good, as is the unconditional distribution of \tilde{M}_S .*

6.2 Sketch of proof

In order to make this argument more digestable, we explain here first the outline of the proof. Throughout this and the next subsection, our processes of course depend on d, p, n and μ ; however, we will drop these in the notation throughout which will not cause any problems. We start $\{\tilde{M}_t\}_{t \geq 0}$ with two initial configurations both in $\mathbb{T}^{d,n} \times \{0, 1\}^{E(\mathbb{T}^{d,n})}$; recall that these are necessarily good distributions.

We want to find a coupling of the two processes and a random time T with mean of order at most n^2/μ so that after time T , the two configurations agree. Since $\{M_t\}_{t \geq 0}$ is obtained from $\{\tilde{M}_t\}_{t \geq 0}$ by dropping the \star 's, we will obtain our result.

This coupling will be achieved in three distinct stages.

Stage 1.

In this first phase, we will run the processes independently until they simultaneously reach the set

$$\Omega_{\text{REG}} := \{(x, \tilde{\eta}) \in \Omega : \tilde{\eta}(e) = 0 \text{ for all } e \text{ adjacent to } x \text{ and } \tilde{\eta}(e) \in \{1^\star, 0^\star\} \text{ for all other } e\}. \quad (23)$$

Proposition 6.14 says that this will take at most order $\log n/\mu$ units of time. To prove this, one considers the set of edges

$$A_s := \{e : \tilde{\eta}_s(e) \in \{0, 1\}\} \text{ (i.e., the set of edges without a } \star \text{ at time } s).$$

The hardest step is to show that on the appropriate time scale of order $1/\mu$, the sets A_s tend to decrease in size; this is the content of Proposition 6.11, which relies on comparisons with subcritical percolation. The fact that A_s tends to decrease is intuitive as follows. A fixed proportion of the set A_s will be refreshed during an interval of order $1/\mu$ while the random walker (which is causing A_s to increase by encountering new edges) is somewhat confined even on this time scale since we are in a subcritical setting. Next Lemma 6.13 will tell us that once A_s is relatively small, then the process will enter Ω_{REG} within a time interval of order $1/\mu$ with a fixed positive probability. Proposition 6.11 and Lemma 6.13 will allow us to prove Proposition 6.14.

Stage 2.

At the start of the second stage, the two distributions are the same up to a translation σ . At this point, we look at excursions from Ω_{REG} at discrete times on the time scale $1/\mu$. Proposition 6.11 and Lemma 6.13 will now be used again in conjunction with Proposition 3.2 to show that the number of steps in such an excursion is of order 1 which means order $1/\mu$ in real time; this is stated in Theorem 6.18. The joint distribution of the number of steps in an excursion and the increment of the walker during this excursion is complicated but it has a component of a fixed size where

the excursion is one step and the increment is a lazy simple random walk. Coupling lazy simple random walk on $\mathbb{T}^{d,n}$ takes on order n^2 steps and so we can couple two copies of our process by having them do the exact same thing off of this component of the distribution and doing usual lazy simple random walk coupling on this component. Since this component has a fixed probability, this coupling will couple in order n^2 excursions and hence in order n^2/μ time.

Stage 3.

After this, we can couple the full systems by a color switch.

We carry this all out in detail at the end of this section.

6.3 Some percolation preliminaries

In this subsection, we gather a number of results concerning percolation.

Theorem 6.3. *For any $d \geq 1$ and $\alpha \in (0, p_c(\mathbb{Z}^d))$, there exists $C_{6.3} = C_{6.3}(d, \alpha) > 0$ so that for all $r \geq 2$,*

$$\mathbf{P}_{\mathbb{Z}^d, \alpha}(|C(0)| \geq r) \leq e^{-C_{6.3}r}.$$

The previous line holds with \mathbb{Z}^d replaced by $\mathbb{T}^{d,n}$.

Proof. This is Theorem 6.75 in [19] in the case of \mathbb{Z}^d . Next, Theorem 1 in [11] states that if one has a covering map from a graph G to a graph H , then the size of a vertex component in H is stochastically dominated by the size of the corresponding vertex component in G . (This is stated for site percolation but site percolation is more general than bond percolation.) Since we have a covering map from \mathbb{Z}^d to $\mathbb{T}^{d,n}$, we obtain the result for $\mathbb{T}^{d,n}$ from the result for \mathbb{Z}^d . \square

We collect here some graph theoretic definitions that we will need.

Definition 6.4. *If $V' \subseteq V$, then $E(V')$ will denote the set $\{e \in E : \exists v \in V' \text{ with } v \in e\}$. (It is not required that both endpoints of e are in V' .)*

Definition 6.5. *If $E' \subseteq E$, then $V(E')$ will denote the union of the endpoints of the edges in E' .*

Definition 6.6. *If $V' \subseteq V$, then $\mathcal{N}_k(V') := \{x \in V : \exists v \in V' \text{ with } \text{dist}(x, v) \leq k\}$ will be called the k -neighborhood of V' .*

Definition 6.7. *If $V' \subseteq V$, then $E \setminus V'$ is defined to be those edges in E which have at least one endpoint not in V' .*

Given a set of vertices F of \mathbb{Z}^d or $\mathbb{T}^{d,n}$ and a bond configuration η , let F^η to be the set of vertices reachable from F using open edges in η . If F is a set of vertices, then the configuration η might only be specified for edges in $E \setminus F$ but note that this has no consequence for the definition of F^η . For a set of vertices F , we also let F^α be the random set obtained by choosing $\eta \subseteq E \setminus F$ according to π_α and then taking F^η . We let $F^{\alpha,1} := F^\alpha$ and we also define inductively, for $L \geq 2$, $F^{\alpha,L} := (F^{\alpha,L-1})^\alpha$. It is implicitly assumed here that we use independent randomness in each iteration.

Theorem 6.8. *Fix $d \geq 1$ and $\alpha \in (0, p_c(\mathbb{Z}^d))$. Then for all L , there exists $C_{6.8}(L) = C_{6.8}(d, \alpha, L)$ so that for all finite $F \subseteq \mathbb{Z}^d$ and for all $\ell \geq 1$,*

$$\mathbf{P}(F^{\alpha,L} \not\subseteq \mathcal{N}_\kappa(F)) \leq \frac{L}{2^{\frac{\ell}{\log \ell}}}$$

where $\kappa := \ell C_{6.8}(L) \log(|F| \vee 2)$. In addition, for the case $L = 1$, the $\log \ell$ term can be removed. Finally, the same result holds for $\mathbb{T}^{d,n}$ as well.

Proof. The following proof works for both \mathbb{Z}^d and $\mathbb{T}^{d,n}$. We prove this by induction on L . The case $L = 1$ without the $\log \ell$ term follows easily from Theorem 4.3 and is left to the reader. We now assume the result for $L = 1$ (without the $\log \ell$ term) and for $L - 1$ and prove it for L . It is elementary to check that

$$\{F^{\alpha,L} \not\subseteq \mathcal{N}_\kappa(F)\} \subseteq E_1 \cup (E_2 \cap E_3)$$

where $E_1 := \{F^{\alpha,L-1} \not\subseteq \mathcal{N}_{\kappa_1}(F)\}$, $E_2 := \{F^{\alpha,L-1} \subseteq \mathcal{N}_{\kappa_1}(F)\}$, $E_3 := \{F^{\alpha,L} \not\subseteq \mathcal{N}_{\kappa-\kappa_1}(F^{\alpha,L-1})\}$ and $\kappa_1 := \ell C_{6.8}(L-1) \log(|F| \vee 2)$.

The probability of the first event is, by induction, at most $\frac{L-1}{2^{\log \ell}}$. Note next that when E_2 occurs, it is necessarily the case that

$$|F^{\alpha,L-1}| \leq |F|(2\ell C_{6.8}(L-1) \log(|F| \vee 2) + 1)^d.$$

The latter yields

$$\log(|F^{\alpha,L-1}| \vee 2) \leq \log(|F|) + d \log(2\ell C_{6.8}(L-1) \log(|F| \vee 2) + 1). \quad (24)$$

Now the neighborhood size arising in the event E_3 is

$$\frac{\ell \log(|F| \vee 2)(C_{6.8}(L) - C_{6.8}(L-1))}{C_{6.8}(1) \log(|F^{\alpha,L-1}| \vee 2)} \times C_{6.8}(1) \log(|F^{\alpha,L-1}| \vee 2).$$

By (24), this first factor is at least

$$\frac{\ell \log(|F| \vee 2)(C_{6.8}(L) - C_{6.8}(L-1))}{C_{6.8}(1)(\log(|F|) + d \log(2\ell C_{6.8}(L-1) \log(|F| \vee 2) + 1))}.$$

It is easy to show that given $C_{6.8}(1)$ and $C_{6.8}(L-1)$, one can choose $C_{6.8}(L)$ sufficiently large so that for all F and for all ℓ , this is larger than $\frac{\ell}{\log \ell}$. It now follows from the $L = 1$ case (where no $\log \ell$ term appears) that $\mathbf{P}(E_2 \cap E_3) \leq \frac{1}{2^{\log \ell}}$. Adding this to the first term yields the result. \square

The previous theorem gave bounds on how far F^α (and its higher iterates) could be from F . The next proposition yields bounds on the size of F^α in terms of F . We will only need a bound on the mean which would then be easy to extend to higher iterates.

Theorem 6.9. *Fix $d \geq 1$ and $\alpha \in (0, p_c(\mathbb{Z}^d))$. Then there is a constant $C_{6.9} = C_{6.9}(d, \alpha)$ so that for all finite $F \subseteq \mathbb{Z}^d$, one has*

$$\mathbf{E}|F^\alpha| \leq C_{6.9}|F|.$$

The same result holds for $\mathbb{T}^{d,n}$ as well.

Proof. The following proof works for both \mathbb{Z}^d and $\mathbb{T}^{d,n}$. Theorem 4.3 immediately implies that $\mathbf{E}_{d,\alpha}|C(0)| < \infty$. Note now that $F^\alpha \subseteq \cup_{x \in F} C(x)$ where $C(x)$ is the set of vertices that can be reached from x using the α -open edges in $E \setminus F$ and so $|F^\alpha| \leq \sum_{x \in F} |C(x)|$. This now gives the result with $C_{6.9}(d, \alpha) = \mathbf{E}_{d,\alpha}[|C(0)|]$. \square

6.4 Details of the proof

We now fix d and $p \in (0, p_c(\mathbb{Z}^d))$ for the rest of the argument. We next choose $\epsilon = \epsilon(d, p)$ so that

$$\epsilon < \frac{p_c - p}{4}$$

and we may assume that ϵ is an inverse integer. Since the probability that an edge is refreshed during an interval of length ϵ/μ is $1 - e^{-\epsilon} < \epsilon$, it follows that if, conditioned on \mathcal{F}_s^* , an edge is open at some fixed time $t \geq s$ with conditional probability at most $\frac{p_c + p}{2}$, then the conditional probability given \mathcal{F}_s^* that e is open some time during $[t, t + \frac{\epsilon}{\mu}]$ is at most

$$\frac{p_c + p}{2} + \epsilon < \frac{3}{4}p_c + \frac{1}{4}p.$$

In particular, the unconditional probability that e is open some time during $[t, t + \frac{\epsilon}{\mu}]$ is at most $\frac{3}{4}p_c + \frac{1}{4}p$. For notational convenience, we let

$$p' := \frac{3}{4}p_c + \frac{1}{4}p.$$

Proposition 6.11 is a crucial result showing A_s decreases by a fixed amount on an appropriate time scale. Before doing this, we need the following lemma which yields an a priori bound on the growth rate of $|A_s|$.

Lemma 6.10. *There exists a constant $C_{6.10} = C_{6.10}(d, p)$ so that for all $n, \mu, k \in \mathbb{N}$ and s , if we consider $\{\tilde{M}_t\}_{t \geq 0}$ with an arbitrary good initial distribution, then*

$$\mathbf{E} \left[|A_{s + \frac{k\epsilon}{\mu}}| \mid \mathcal{F}_s^* \right] \leq C_{6.10}^k |A_s| \text{ a.s.} \quad (25)$$

Moreover $C_{6.10}$ can be taken to be

$$4dC_{6.9}(d, p')$$

where the latter constant comes from Theorem 6.9.

Proof. We prove this by induction on k . The main step is $k = 1$. Define a random configuration $\bar{\eta}$ of edges of $E \setminus A_s$ consisting of those edges which are open some time during $[s, s + \epsilon/\mu]$. Note that, by Lemma 6.2 and the way ϵ was chosen, for all μ and n , conditioned on \mathcal{F}_s^* , $\bar{\eta}$ is i.i.d. with density at most p' . The key observation to make is that

$$A_{s + \frac{\epsilon}{\mu}} \subseteq E(V(A_s)\bar{\eta}) \quad (\text{see Definitions 6.4 and 6.5}).$$

It follows from Theorem 6.9 that

$$\mathbf{E} \left[|A_{s + \frac{\epsilon}{\mu}}| \right] \leq C_{6.10} |A_s|$$

where $C_{6.10} = 4dC_{6.9}(d, p')$. For $k \geq 2$, we have

$$\begin{aligned} \mathbf{E} \left[|A_{s + \frac{k\epsilon}{\mu}}| \mid \mathcal{F}_s^* \right] &= \mathbf{E} \left[\mathbf{E} \left[|A_{s + \frac{k\epsilon}{\mu}}| \mid \mathcal{F}_{s + \frac{(k-1)\epsilon}{\mu}}^* \right] \mid \mathcal{F}_s^* \right] \leq \\ &\mathbf{E} \left[C_{6.10} |A_{s + \frac{(k-1)\epsilon}{\mu}}| \mid \mathcal{F}_s^* \right] \leq C_{6.10}^k |A_s| \end{aligned}$$

where the first inequality follows from the $k = 1$ case already proved and the last inequality follows by induction. \square

We now move to

Proposition 6.11. *There exist positive constants $C_{6.11.1} = C_{6.11.1}(d, p)$ and $C_{6.11.2} = C_{6.11.2}(d, p)$ so that for all n and for all μ , if we consider $\{\tilde{M}_t\}_{t \geq 0}$ started with an arbitrary good initial distribution, then we have that*

$$\mathbf{E} \left[|A_{s + \frac{C_{6.11.1}}{\mu}}| \mid \mathcal{F}_s^* \right] \leq \frac{|A_s|}{4} + C_{6.11.2} \log |A_s|. \quad (26)$$

Proof. Given \mathcal{F}_s^* , the conditional probability that an edge e is not refreshed during $[s, s + \frac{k_1}{\mu}]$ or is refreshed and is open at time $s + \frac{k_1}{\mu}$ is $e^{-k_1} + (1 - e^{-k_1})p$. Choose an integer $k_1 = k_1(d, p)$ sufficiently large so that

$$e^{-k_1} + (1 - e^{-k_1})p \leq \frac{pc + p}{2}.$$

Given a sufficiently large integer $k_2 = k_2(d, p)$ to be chosen later, we let $t_1 := s + \frac{k_1}{\mu}$ and $t_2 := t_1 + \frac{k_2}{\mu}$.

Denoting the range of the random walker during the time interval $[u, v]$ by $\mathcal{R}[u, v]$, a key observation is that

$$|A_{s + \frac{k_1 + k_2}{\mu}}| \leq Q + |E(\mathcal{R}[t_1, t_2])| \quad (\text{see Definition 6.4})$$

where Q is the number of edges in A_{t_1} which are not refreshed during $[t_1, t_2]$. The proof would be completed (by taking $C_{6.11.1}$ to be $k_1 + k_2$) if we can choose $k_2 = k_2(d, p)$ and $C_{6.11.2} = C_{6.11.2}(d, p)$ so that for all n and μ ,

$$\mathbf{E}[Q \mid \mathcal{F}_s^*] \leq \frac{|A_s|}{4} \quad (27)$$

and

$$\mathbf{E}[|E(\mathcal{R}[t_1, t_2])| \mid \mathcal{F}_s^*] \leq C_{6.11.2} \log |A_s|. \quad (28)$$

We start by finding k_2 so that (27) holds for all n and μ . For this, we simply choose k_2 so that

$$e^{-k_2} C_{6.10}^{\frac{k_1}{\epsilon}} \leq \frac{1}{4} \quad (29)$$

where $C_{6.10}$ comes from Lemma 6.10 and note that k_2 only depends on d and p .

Clearly $\mathbf{E}[Q \mid \mathcal{F}_{t_1}^*] = e^{-k_2} |A_{t_1}|$ which implies that

$$\mathbf{E}[Q \mid \mathcal{F}_s^*] = \mathbf{E}[\mathbf{E}[Q \mid \mathcal{F}_{t_1}^*] \mid \mathcal{F}_s^*] = \mathbf{E}[e^{-k_2} |A_{t_1}| \mid \mathcal{F}_s^*].$$

Since $t_1 = s + \frac{k_1}{\epsilon} \frac{\epsilon}{\mu}$ (with ϵ being an inverse integer), Lemma 6.10 and (29) imply that this last expression is at most $\frac{|A_s|}{4}$, demonstrating (27).

With k_1 and k_2 now chosen, we want to find $C_{6.11.2} = C_{6.11.2}(d, p)$ so that for all n and μ , (28) holds. Since $|E(\mathcal{R}[t_1, t_2])| \leq 2d|\mathcal{R}[t_1, t_2]|$, it suffices to prove such a bound for $|\mathcal{R}[t_1, t_2]|$ instead.

We first observe that by the way k_1 was chosen, conditioned on \mathcal{F}_s^* , the probability that an arbitrary edge e is open at some fixed time $t \geq t_1$ is at most $\frac{pc+p}{2}$ and hence by our choice of ϵ , we have that for any interval $I := [y, y + \frac{\epsilon}{\mu}] \subseteq [t_1, t_2]$, the conditional probability given \mathcal{F}_s^* that e is open some time during $[y, y + \frac{\epsilon}{\mu}]$ is at most p' . Note that conditioned on \mathcal{F}_s^* , the evolution of the states of the different edges after time s are independent although they are not identically distributed.

We next partition $[t_1, t_2]$ into $D = D(d, p)$ disjoint intervals of length $\frac{\epsilon}{\mu}$. Note importantly that D does not depend on n or μ . It now suffices to show that if $I := [y, y + \frac{\epsilon}{\mu}] \subseteq [t_1, t_2]$ with $y = t_1 + \frac{\ell\epsilon}{\mu}$ and $\ell \in \mathbb{N}$, then

$$\mathbf{E}[|\mathcal{R}[I]| \mid \mathcal{F}_s^*] \leq C \log |A_s| \text{ a.s.} \quad (30)$$

for some C depending only on d and p . To do this, it suffices to show that

$$\left\| \sum_{\ell \geq 1} \mathbf{P} \left[|\mathcal{R}[I]| \geq \ell \frac{4d}{C_{6.3}(d, p')} \log |V(A_s)| \mid \mathcal{F}_s^* \right] \right\|_{\infty} < \infty$$

where $C_{6.3}(d, p')$ comes from Theorem 6.3. Let $\bar{\eta}$ be the set of edges which are open some time during I . By our choice of ϵ , conditioned on \mathcal{F}_s^* , $\bar{\eta}$ is stochastically dominated by an i.i.d. process with density p' . Since $\mathcal{R}[I]$ is necessarily contained inside of a $\bar{\eta}$ -cluster, we have

$$\mathbf{P} \left[|\mathcal{R}[I]| \geq \ell \frac{4d}{C_{6.3}(d, p')} \log |V(A_s)| \mid \mathcal{F}_s^* \right] \leq \mathbf{P} [\text{dist}(X_y, X_s) \geq \ell \Xi \mid \mathcal{F}_s^*] + \quad (31)$$

$$\mathbf{P} \left[\bar{\eta} \text{ contains a cluster of size } \geq \ell \frac{4d}{C_{6.3}(d, p')} \log |V(A_s)| \text{ intersecting } B(X_s, \ell \Xi) \mid \mathcal{F}_s^* \right]$$

where

$$\Xi := 4dC_{6.8}(d, p', D) |V(A_s)| \log |V(A_s)|,$$

where $C_{6.8}(d, p', D)$ comes from Theorem 6.8 and where, as before, $B(v, r)$ is the set of vertices within dist-distance r of v .

It is easy to check that Theorem 6.3 together with a union bound implies that the second terms are summable over ℓ uniform in the conditioning, i.e.

$$\left\| \sum_{\ell \geq 1} \mathbf{P} \left[\bar{\eta} \text{ contains a cluster of size } \geq \ell \frac{4d}{C_{6.3}(d, p')} \log |V(A_s)| \text{ intersecting } B(X_s, \ell \Xi) \mid \mathcal{F}_s^* \right] \right\|_{\infty} < \infty.$$

To deal with the first terms, we partition the interval $[s, y]$ into successive intervals J_1, J_2, \dots, J_L of lengths $\frac{\epsilon}{\mu}$ where $L \leq D$. Let $\bar{\eta}_i$ be the set of edges which are open some time during J_i . The key observation is (see Definition 6.6) that for each $w \geq |V(A_s)| \log |V(A_s)|$

$$\{\text{dist}(X_y, X_s) > 4dw\} \subseteq \{V(A_s)^{\bar{\eta}_1, \dots, \bar{\eta}_L} \not\subseteq \mathcal{N}_{\frac{w}{|V(A_s)|}}(V(A_s))\}. \quad (32)$$

To see this, one first makes the important observation that the random walk path between times s and y is contained in $V(A_s)^{\bar{\eta}_1, \dots, \bar{\eta}_L}$ and a geometric argument shows that $\{\text{dist}(X_y, X_s) > 4dw\}$ implies that $\mathcal{R}[s, y]$ cannot be contained in $\mathcal{N}_{\frac{w}{|V(A_s)|}}(V(A_s))$.

We claim that conditioned on \mathcal{F}_s^* , the set

$$V(A_s)^{\bar{\eta}_1, \dots, \bar{\eta}_L} \text{ is stochastically dominated by } V(A_s)^{p', L}. \quad (33)$$

To see this, one first notes that since we are assuming a good initial distribution, when we condition on \mathcal{F}_s^* , the edges not in A_s , which are the only relevant edges in the construction of $V(A_s)^{\bar{\eta}_1, \dots, \bar{\eta}_L}$, are in stationarity and hence, conditioned on \mathcal{F}_s^* , each $\bar{\eta}_i$, off of the edge set A_s , is stochastically dominated by an i.i.d. process with density p' . Second, one notes that when one further conditions on the sets $V(A_s)^{\bar{\eta}_1}, \dots, V(A_s)^{\bar{\eta}_1, \dots, \bar{\eta}_j}$, this can only stochastically decrease the edges of $\bar{\eta}_{j+1}$ which are relevant at that point. Hence we obtain (33). This stochastic domination implies that

$$\mathbf{P} \left[V(A_s)^{\bar{\eta}_1, \dots, \bar{\eta}_L} \not\subseteq \mathcal{N}_{\frac{w}{|V(A_s)|}}(V(A_s)) \right] \leq \mathbf{P} \left[V(A_s)^{p', L} \not\subseteq \mathcal{N}_{\frac{w}{|V(A_s)|}}(V(A_s)) \right]. \quad (34)$$

Letting $w := \ell C_{6.8}(d, p', D) |V(A_s)| \log |V(A_s)|$, (we may assume $C_{6.8}(d, p', D) \geq 1$) (32), (34) and Theorem 6.8 now imply that the first terms of (31) are also summable over ℓ uniform in the conditioning, i.e.,

$$\left\| \sum_{\ell \geq 1} \mathbf{P} [\text{dist}(X_y, X_s) \geq \ell \Xi \mid \mathcal{F}_s^*] \right\|_\infty < \infty.$$

□

Remark 6.12. If one wants, one could avoid the use of Theorem 6.3 and thereby the need for Theorem 1 from [11]. One could do this by modifying the above proof using only Theorem 4.3 obtaining (26) with the log term replaced by a power of log which would suffice for the rest of the proof.

We are now going to look at our process at integer multiples of $\frac{C_{6.11.1}}{\mu}$ where $C_{6.11.1}$ comes from Proposition 6.11. We therefore let, for integer $k \geq 0$,

$$D_k := (X_{\frac{kC_{6.11.1}}{\mu}}, \tilde{\eta}_{\frac{kC_{6.11.1}}{\mu}}), \quad B_k := A_{\frac{kC_{6.11.1}}{\mu}} \quad \text{and} \quad \mathcal{G}_k := \mathcal{F}_{\frac{kC_{6.11.1}}{\mu}}^*. \quad (35)$$

We will use the set Ω_{REG} (see (23)) as a regenerative set. We therefore define the set of regeneration times along our subsequence of times by

$$\mathcal{I} := \{k \geq 0 : D_k \in \Omega_{\text{REG}}\}. \quad (36)$$

Our next lemma says that when we are not far away from Ω_{REG} , then we have a good chance of entering it.

In the proof of this lemma, if $E' \subseteq E$, then we let $\partial(E')$ denote the set $\{e \notin E' : \exists e' \in E' \text{ with } e \text{ and } e' \text{ adjacent}\}$; this will be called the boundary of E' .

Lemma 6.13. *For all R , there exists $\alpha_{6.13} = \alpha_{6.13}(d, p, R) > 0$ so that for all n and for all μ , if we consider $\{\tilde{M}_t\}_{t \geq 0}$ with an arbitrary good initial distribution, then*

$$\mathbf{P}[j+1 \in \mathcal{I} \mid \mathcal{G}_j] \geq \alpha_{6.13}$$

on the event $|B_j| \leq R$.

Proof. First, let E_1 be the event that all edges of the boundary of B_j , $\partial(B_j)$ are closed at time $\frac{jC_{6.11.1}}{\mu}$, E_2 be the event that no edge of $\partial(B_j)$ refreshes during $[\frac{C_{6.11.1}}{\mu}j, \frac{C_{6.11.1}}{\mu}(j + \frac{1}{2})]$ and E_3 be the event that all edges in B_j are refreshed closed during $[\frac{C_{6.11.1}}{\mu}j, \frac{C_{6.11.1}}{\mu}(j + \frac{1}{2})]$. Observe that if $E_1 \cap E_2 \cap E_3$ occurs, then the edges next to the walker necessarily are closed and belong to either B_j or $\partial(B_j)$. Next, let E_4 be the event that the edges adjacent to $X_{\frac{C_{6.11.1}}{\mu}(j+\frac{1}{2})}$ do not refresh during $[\frac{C_{6.11.1}}{\mu}(j + \frac{1}{2}), \frac{C_{6.11.1}}{\mu}(j + 1)]$ and E_5 be the event that all the edges of B_j and $\partial(B_j)$ except for the edges adjacent to $X_{\frac{C_{6.11.1}}{\mu}(j+\frac{1}{2})}$ refresh during $[\frac{C_{6.11.1}}{\mu}(j + \frac{1}{2}), \frac{C_{6.11.1}}{\mu}(j + 1)]$. It is elementary to check, using the fact that we started with a good initial distribution, that $\mathbf{P}[\cap_{i=1}^5 E_i \mid \mathcal{G}_j]$ is bounded away from 0 on the event $|B_j| \leq R$ uniformly in n and μ and the conditioning and that $\cap_{i=1}^5 E_i \subseteq \{j+1 \in \mathcal{I}\}$. This completes the proof. □

Our next proposition tells us that we can do the first stage of the coupling described in the sketch; namely, two copies of our system will enter Ω_{REG} within $\log n$ many steps and hence within $\log n/\mu$ units of time.

Proposition 6.14. *There exists $C_{6.14} = C_{6.14}(d, p) < \infty$ so that for all n and for all μ , if $\{D_k^1\}_{k \geq 0}$ and $\{D_k^2\}_{k \geq 0}$ are two independent copies of $\{D_k\}_{k \geq 0}$ each starting from arbitrary initial configurations in $\mathbb{T}^{d,n} \times \{0, 1\}^{E(\mathbb{T}^{d,n})}$, and if we set*

$$T_1 := \min\{k \geq 1 : D_k^1 \in \Omega_{\text{REG}} \text{ and } D_k^2 \in \Omega_{\text{REG}}\},$$

then $\mathbf{E}[T_1] \leq C_{6.14} \log n$.

Proof. Let $Z_k := |B_k^1| + |B_k^2|$ using obvious notation. By Proposition 6.11, we obtain that for all n and for all μ ,

$$\mathbf{E}[Z_{k+1} \mid \mathcal{G}_k \times \mathcal{G}_k] \leq \frac{Z_k}{4} + 2C_{6.11.2} \log(Z_k).$$

This immediately gives that there is a constant $C_{6.14.1} = C_{6.14.1}(d, p) < \infty$ so that

$$\mathbf{E}[Z_{k+1} \mid \mathcal{G}_k \times \mathcal{G}_k] \leq \frac{Z_k}{3} \tag{37}$$

on the event $Z_k \geq C_{6.14.1}$. Now, noting that $Z_0 = 2dn^d$ (which is $\geq C_{6.14.1}$ for n large), (37) implies that $3^{k \wedge \tilde{T}} Z_{k \wedge \tilde{T}}$ is a $\mathcal{G}_k \times \mathcal{G}_k$ -supermartingale where

$$\tilde{T} := \min\{k : Z_k \leq C_{6.14.1}\}.$$

From the theory of stopping times for nonnegative supermartingales [17, Chapter 5], we obtain the fact that

$$\mathbf{E}\left[3^{\tilde{T}} Z_{\tilde{T}}\right] \leq Z_0.$$

Since the Z_k 's are always at least 1 and $Z_0 = 2dn^d$, we infer that

$$\mathbf{E}\left[3^{\tilde{T}}\right] \leq 2dn^d$$

which in turn implies, by Jensen's inequality, that

$$\mathbf{E}\left[\tilde{T}\right] \leq \log_3(2dn^d).$$

Since both $|B_{\tilde{T}}^1|$ and $|B_{\tilde{T}}^2|$ are less than $C_{6.14.1}$, using the fact that \tilde{T} is a stopping time and the independence of the two processes, we can conclude from Lemma 6.13 that for all n and μ , the probability that both $D_{\tilde{T}+1}^1$ and $D_{\tilde{T}+1}^2$ are in Ω_{REG} is at least $\alpha_{6.13}(d, p, C_{6.14.1})^2$. If this fails, we start again and wait on average another at most $\log_3(2n^d)$. After a geometric number of trials, we are done. By Wald's Theorem [17, Theorem 4.1.5], this proves the result. \square

We now return to looking at just one copy of our system and study the excursions of $\{\tilde{M}_t\}_{t \geq 0}$ away from Ω_{REG} . Assume now that \tilde{M}_0 has a good distribution supported on Ω_{REG} . We let $\tau_0 = 0$ and, for $j \geq 1$, define

$$\tau_j := \min\{i > \tau_{j-1} : i \in \mathcal{I}\}, \quad U_j := D'_{\tau_j} - D'_{\tau_{j-1}} \tag{38}$$

where D'_k denotes the first coordinate of D_k .

Note that, by Lemma 6.2, for each j , the distribution of the process at time $\frac{\tau_j C_{6.11.1}}{\mu}$ conditioned on \mathcal{G}_{τ_j} is good.

Remark 6.15. It is easy to see that the set of good probability measures supported on Ω_{REG} can be described as follows; one chooses a vertex v at random according to any distribution and then sets the edges adjacent to v to be in state 0 and all other edges are (conditionally) independently chosen to be 1^* with probability p and 0^* with probability $1 - p$.

The following lemma, whose proof is left to the reader, is clear.

Lemma 6.16. *If our initial distribution is a good distribution supported on Ω_{REG} , then $\{(\tau_j - \tau_{j-1}, U_j)\}_{j \geq 1}$ are i.i.d. and moreover, for each $j \geq 1$, given $\mathcal{G}_{\tau_{j-1}}$, the conditional distribution of $(\tau_j - \tau_{j-1}, U_j)$ is the same as (τ_1, U_1) .*

Remark 6.17. Of course $\tau_j - \tau_{j-1}$ and U_j are not independent of each other.

The next theorem tells us that the number of steps in one of our excursions away from Ω_{REG} is of order 1.

Theorem 6.18. *There exists $C_{6.18} = C_{6.18}(d, p) < \infty$ such that, for all n and for all μ , if we start with a good initial distribution supported on Ω_{REG} , then*

$$\mathbf{E}[\tau_1] \leq C_{6.18}.$$

Proof. Let $Y_k := 1 + B'_k + L_k$ where B'_k is the number of edges not adjacent to the walker which are in state 0 or 1 at time $\frac{kC_{6.11.1}}{\mu}$ and where L_k is the number of edges adjacent to the walker which are in state 1 at time $\frac{kC_{6.11.1}}{\mu}$. Note $Y_k = 1$ if and only if $D_k \in \Omega_{\text{REG}}$; hence $Y_0 = 1$ and τ_1 corresponds to the first return of the Y process to 1. We will apply Proposition 3.2 with \mathcal{F}_i there being \mathcal{G}_i . Property (1) trivially holds. Proposition 6.11 easily implies that for $\delta = 1/3$ and for α sufficiently large, but only depending on d and p , Property (2) holds for all n and μ . Next, with α now fixed, Lemma 6.10 implies that there exists γ sufficiently large, but only depending on d and p , so that Property (4) holds for all n and μ . Finally, since α is now fixed, Lemma 6.13 guarantees that property (3) holds for all n and μ for some positive ϵ also only depending on d and p . Proposition 3.2 now yields the result. \square

We now finally have all of the ingredients to give the

Proof of Theorem 1.2(ii). Let $\{\tilde{M}_t^1\}_{t \geq 0}$ and $\{\tilde{M}_t^2\}_{t \geq 0}$ denote two copies of our process $\{\tilde{M}_t\}_{t \geq 0}$, each starting from an arbitrary configuration in $\mathbb{T}^{d,n} \times \{0, 1\}^{E(\mathbb{T}^{d,n})}$. We will find a coupling $(\{\tilde{M}_t^1\}_{t \geq 0}, \{\tilde{M}_t^2\}_{t \geq 0}, T)$ of our two processes and a nonnegative random variable T so that

- (1) $\tilde{M}_t^1 = \tilde{M}_t^2$ for all $t \geq T$ and
- (2) $\mathbf{E}[T] \leq O_{d,p}(1) \frac{n^2}{\mu}$.

From here, it is standard that this gives a bound on the mixing time as follows. If $t \geq 4 \mathbf{E}[T]$, then

$$\|\mathcal{L}(M_t^1) - \mathcal{L}(M_t^2)\|_{\text{TV}} \leq \mathbf{P}[M_t^1 \neq M_t^2] \leq \mathbf{P}[\tilde{M}_t^1 \neq \tilde{M}_t^2] \leq \mathbf{P}[T > t] \leq \frac{1}{4}$$

by Markov's inequality. As outlined earlier, we do this coupling in two separate stages.

Stage 1.

In this first stage we run the two processes independently until both processes simultaneously hit Ω_{REG} at some time T_1 of the form $\frac{kC_{6.11.1}}{\mu}$, $k \in \mathbb{N}$. By Proposition 6.14, this first stage will take in expectation at most $O_{d,p}(1) \frac{\log n}{\mu}$ time. Let $\mathcal{F}_{T_1}^{1,*}$ be the σ -algebra generated by $\{\tilde{M}_t^1\}_{0 \leq t \leq T_1}$ including

all the refresh times in $[0, T_1]$ but where one does not distinguish 1^* and 0^* . Let $\mathcal{F}_{T_1}^{2,*}$ be defined analogously. (A trivial variant of) Lemma 6.2 and the independence of the two processes imply that, conditioned on (1) T_1 , (2) $\mathcal{F}_{T_1}^{1,*}$ and (3) $\mathcal{F}_{T_1}^{2,*}$, we have that these two conditional distributions at time T_1 are good. Also, the two conditional distributions of the walker are trivially degenerate. In view of Remark 6.15, these two conditional distributions will then agree up to a translation σ which translates one walker to the other. Moreover, the two processes at time T_1 are conditionally independent given (1), (2) and (3) above.

Stage 2.

We now condition on (1), (2) and (3) above. Viewing things from time T_1 , we are then in the setting of Lemma 6.16 where we can view our two processes as being independent and having good distributions supported on Ω_{REG} with the walker having a degenerate distribution.

Before completing the second stage, we first recall for the reader how one couples two lazy simple random walks on $\mathbb{T}^{d,n}$. By definition, a lazy simple random walk on $\mathbb{T}^{d,n}$ with probability $1/2$ stays where it is and with probability $1/2$ moves to a neighbor chosen uniformly at random. The usual coupling of two lazy simple random walks on $\mathbb{T}^{d,n}$ is a Markov process on $\mathbb{T}^{d,n} \times \mathbb{T}^{d,n}$ defined as follows. One chooses one of the d coordinates uniformly at random. If the two walkers agree in this coordinate, then they both stay still, both move “right” in this coordinate or both move “left” in this coordinate with respective probabilities $1/2$, $1/4$ and $1/4$. If the two walkers disagree in this coordinate, then one of the walkers, chosen at random, stays still while the other one moves “right” or “left” in this coordinate, each with probability $1/2$. It is easy to check that this is a coupling of the two lazy simple random walks. It is known (see Section 5.3 in [23]) that the expected time until the two walkers meet is at most $O_d(1)n^2$.

Denote the joint distribution of (τ_1, U_1) by $\nu_{d,p,n,\mu}$; this is a probability measure on $\mathbb{N} \times \mathbb{T}^{d,n}$. It is easy to show, along the same lines as the proof of Lemma 6.13, that for any d and p , there is a $\gamma = \gamma(d, p) > 0$ so that for any n and μ ,

$$\nu_{d,p,n,\mu} = \gamma(\delta_1 \times \nu_{\text{LSRW}}) + (1 - \gamma)m_{d,p,n,\mu} \quad (39)$$

where $\delta_1 \times \nu_{\text{LSRW}}$ is the probability measure on $\mathbb{N} \times \mathbb{T}^{d,n}$ where the first coordinate is always 1 and the second coordinate is a step of a lazy simple random walk on $\mathbb{T}^{d,n}$ and where $m_{d,p,n,\mu}$ is some probability measure on $\mathbb{N} \times \mathbb{T}^{d,n}$.

We will first couple the random sequence $\{(\tau_j - \tau_{j-1}, U_j)\}_{j \geq 1}$ for the two systems and denote the coupled variables by $\{((\tau_j^1 - \tau_{j-1}^1, U_j^1), (\tau_j^2 - \tau_{j-1}^2, U_j^2))\}_{j \geq 1}$. We will do this in such a way that $\tau_j^1 - \tau_{j-1}^1 = \tau_j^2 - \tau_{j-1}^2$ for all j . Given $\{((\tau_j^1 - \tau_{j-1}^1, U_j^1), (\tau_j^2 - \tau_{j-1}^2, U_j^2))\}_{1 \leq j \leq K}$ so that $\tau_j^1 - \tau_{j-1}^1 = \tau_j^2 - \tau_{j-1}^2$ for all $1 \leq j \leq K$, we know where the two walkers are located at time

$$\frac{C_{6.11.1}}{\mu} \sum_{j=1}^K (\tau_j^1 - \tau_{j-1}^1) \quad (= \frac{C_{6.11.1}}{\mu} \tau_K^1 = \frac{C_{6.11.1}}{\mu} \tau_K^2 = \frac{C_{6.11.1}}{\mu} \sum_{j=1}^K (\tau_j^2 - \tau_{j-1}^2)).$$

We now define $((\tau_{K+1}^1 - \tau_K^1, U_{K+1}^1), (\tau_{K+1}^2 - \tau_K^2, U_{K+1}^2))$ as follows. With probability $1 - \gamma$, one chooses an element from $\mathbb{N} \times \mathbb{T}^{d,n}$ according to distribution $m_{d,p,n,\mu}$ and uses it for both systems. With probability γ , one takes the first coordinate to be 1 for both systems and does the coupled lazy simple random walk (described above) for the second coordinates of the two systems. (This coupling is reminiscent of a coupling due to D. Ornstein.) Note that this is a coupling and with it, we have that $\tau_j^1 - \tau_{j-1}^1 = \tau_j^2 - \tau_{j-1}^2$ for all j and $(\tau_j^1 - \tau_{j-1}^1, U_j^1) = (\tau_j^2 - \tau_{j-1}^2, U_j^2)$ for all large j depending on ω .

By (39), the fact that γ is uniformly bounded away from 0, the fact that the standard coupling of two lazy simple random walks on $\mathbb{T}^{d,n}$ couples in expected time at most $O_d(1)n^2$ and Wald's Theorem [17, Theorem 4.1.5], we have that the expected number of steps, denoted by N , in the above coupling until the two walkers meet is at most $O_{d,p}(1)n^2$. Using Theorem 6.18, another application of Wald's Theorem tells us that the expected value of

$$T_2 := \frac{C_{6.11.1}}{\mu} \sum_{j=1}^N (\tau_j^1 - \tau_{j-1}^1)$$

is at most $O_{d,p}(1)\frac{n^2}{\mu}$. We now let

$$T := T_1 + T_2$$

and note that we have $\mathbf{E}[T] \leq O_{d,p}(1)\frac{n^2}{\mu}$.

Stage 3.

By construction, we have that $\{T_1, \mathcal{F}_{T_1}^{1,*}\}$ is independent of the random variables $\{(\tau_j^1 - \tau_{j-1}^1, U_j^1)\}_{j \geq 1}$ from the second stage and similarly, we have that $\{T_1, \mathcal{F}_{T_1}^{2,*}\}$ is independent of the random variables $\{(\tau_j^2 - \tau_{j-1}^2, U_j^2)\}_{j \geq 1}$ from the second stage. It also follows by construction that the conditional distribution of $\{\tilde{M}_t^1\}_{t \geq T}$ conditioned on $\{T_1, \mathcal{F}_{T_1}^{1,*}, \{(\tau_j^1 - \tau_{j-1}^1, U_j^1)\}_{j \geq 1}, T\}$ is the same as the conditional distribution of $\{\tilde{M}_t^2\}_{t \geq 0}$ conditioned on $\{T_1, \mathcal{F}_{T_1}^{2,*}, \{(\tau_j^2 - \tau_{j-1}^2, U_j^2)\}_{j \geq 1}, T\}$. This implies that we have a coupling of the desired form and completes the proof. \square

7 Hitting time results

This section is devoted to proving Theorem 1.12. We will use a number of the results which were derived in Section 6. We first need some lemmas which might be of independent interest.

Our first lemma follows easily from the usual local central limit theorem for lazy simple random walk on \mathbb{Z}^d and therefore no proof is given.

Lemma 7.1. *Fix $d \geq 1$ and let P_n^k be the k -step transition probability for lazy simple random walk on $\mathbb{T}^{d,n}$. Then there exists a constant $C(d)$ such that for all n and k*

$$\sup_{x,y \in \mathbb{T}^{d,n}} P_n^k(x,y) \leq C(d) \left(\frac{1}{k^{d/2}} \vee \frac{1}{n^d} \right) \quad (40)$$

and in addition for any $\alpha > 0$, there is a constant $C(d, \alpha)$ such that for all n

$$\inf_{x,y \in \mathbb{T}^{d,n}, k \geq \alpha n^2} P_n^k(x,y) \geq \frac{1}{C(d, \alpha) n^d} \quad (41)$$

and such that for all k

$$\inf_{n,x,y \in \mathbb{T}^{d,n}, \alpha(\text{dist}(x,y))^2 \leq k} P_n^k(x,y) \geq \frac{1}{C(d, \alpha) k^{d/2}}. \quad (42)$$

Lemma 7.2. *Fix $\gamma > 0$ and $d \geq 1$. Then there exists a constant $C_{7.2} = C_{7.2}(\gamma, d)$ so that if $S^{(n)}$ is any discrete time random walk on $\mathbb{T}^{d,n}$ whose step distribution ν_n satisfies*

$$\nu_n = \gamma \nu_{\text{LSRW}} + (1 - \gamma) m_n$$

where ν_{LSRW} is the distribution of a step of a lazy simple random walk on $\mathbb{T}^{d,n}$ and where m_n is some probability measure on $\mathbb{T}^{d,n}$, then the following hold.

(i). If $d = 1$, then for all n ,

$$\max\{\mathbf{E}[\sigma_y] : y \in \mathbb{T}^{1,n}\} \leq C_{7.2}n^2.$$

(ii). If $d = 2$, then for all n ,

$$\max\{\mathbf{E}[\sigma_y] : y \in \mathbb{T}^{2,n}\} \leq C_{7.2}n^2 \log n.$$

(iii). If $d \geq 3$, then for all n ,

$$\max\{\mathbf{E}[\sigma_y] : y \in \mathbb{T}^{d,n}\} \leq C_{7.2}n^d.$$

Proof. Letting $P_{S^{(n)}}^k$ be the k -step transition probability for $S^{(n)}$, we claim that there exists a constant $C(d, \gamma)$ so that for all n and k

$$\sup_{x, y \in \mathbb{T}^{d,n}} P_{S^{(n)}}^k(x, y) \leq C(d, \gamma) \left(\frac{1}{k^{d/2}} \vee \frac{1}{n^d} \right) \quad (43)$$

and so that for all n

$$\inf_{x, y \in \mathbb{T}^{d,n}, k \geq n^2/2} P_{S^{(n)}}^k(x, y) \geq \frac{1}{C(d, \gamma)n^d}. \quad (44)$$

We argue (43). Let $Y(n, k)$ be the number of times that $S^{(n)}$ uses the lazy simple random walk component up to time k and observe that $Y(n, k)$ has a binomial distribution with parameters k and γ . Therefore, the probability of $\{Y(n, k) \leq k\gamma/2\}$ is exponentially small in k . On the other hand, if $Y(n, k) \geq k\gamma/2$, then we can bound the probability of being at some y at time k by conditioning on the steps up to time k taken by $S^{(n)}$ when it doesn't use the lazy simple random walk component and then using (40) for the remaining steps. Together, this yields

$$P_{S^{(n)}}^k(x, y) \leq e^{-c(\gamma)k} + C(d) \left(\left(\frac{2}{k\gamma} \right)^{d/2} \vee \frac{1}{n^d} \right)$$

which is bounded above by $C(d, \gamma) \left(\frac{1}{k^{d/2}} \vee \frac{1}{n^d} \right)$ for some constant $C(d, \gamma)$, obtaining (43). (44) is handled similarly by using (41) instead.

Next, given a general Markov chain $X = \{X_n\}$, an element y in the state space and $k \geq 1$, we let $L_X(k, y) = L(k, y) := \sum_{i=1}^k I_{\{X_i=y\}}$ be the local time at y up to time k .

Case (i): $d = 1$. By using (44) and summing over $k \in \{n^2/2, \dots, n^2\}$ and then using (43) and summing over $k \in \{1, \dots, n^2\}$, we obtain

$$\frac{n}{2C(1, \gamma)} \leq \min_y \mathbf{E} [L_{S^{(n)}}(n^2, y)] \leq \max_y \mathbf{E} [L_{S^{(n)}}(n^2, y)] \leq 2C(1, \gamma)n. \quad (45)$$

The inequalities in (45) easily yield that for any y ,

$$\mathbf{E} [(L_{S^{(n)}}(n^2, y))^2] \leq 8C^2(1, \gamma)n^2 \leq 32C^4(1, \gamma)(\mathbf{E} [L_{S^{(n)}}(n^2, y)])^2.$$

The Cauchy-Schwarz inequality allows us to conclude that for any y ,

$$\mathbf{P} [L_{S^{(n)}}(n^2, y) > 0] \geq \frac{1}{32C^4(1, \gamma)}.$$

By symmetry, starting from any state, we hit y after n^2 steps with probability at least $\frac{1}{32C^4(1,\gamma)}$. This immediately gives an upper bound of the desired form.

Case (ii): $d = 2$. Proceed exactly as in the case $d = 1$ except we consider the local time up to time $n^2 \log n$ and sum (44) and (43) over $k \in \{n^2/2, \dots, n^2 \log n\}$ and $k \in \{1, \dots, n^2 \log n\}$ respectively. The right and left most terms of what will become (45) are then of order $\log n$ in this case.

Case (iii): $d \geq 3$. Proceed exactly as in the case $d = 1$ except we consider the local time up to time n^d and sum (44) and (43) over $k \in \{n^d/2, \dots, n^d\}$ and $k \in \{1, \dots, n^d\}$ respectively. The right and left most terms of what will become (45) are then of order 1 in this case. \square

Lemma 7.3. *Fix $\gamma > 0, R > 0$, and $d \geq 1$. Then there exists a constant $C_{7.3} = C_{7.3}(\gamma, R, d)$ so that if $S^{(n)}$ is any discrete time random walk on $\mathbb{T}^{d,n}$ whose step distribution ν_n satisfies*

$$\nu_n = \gamma \nu_{\text{LSRW}} + (1 - \gamma) m_n$$

where ν_{LSRW} is the distribution of a step of a lazy simple random walk on $\mathbb{T}^{d,n}$ and where m_n is some probability measure on $\mathbb{T}^{d,n}$ and

$$\mathbf{E} \left[\text{dist}(S_1^{(n)}, 0)^2 \right] \leq R, \quad (46)$$

then it follows that for any k ,

$$\inf_{n,x,y \in \mathbb{T}^{d,n}: \text{dist}(x,y)^2 \leq k} P_{S^{(n)}}^k(x,y) \geq \frac{C_{7.3}}{k^{d/2}}. \quad (47)$$

Proof. We just sketch the proof. This can be proved in a similar fashion as (43) but now using (42) as follows. By the uniform variance assumption (46), if we add up the steps taken up to time k when we don't use the lazy simple random walk component, then, with fixed probability, this is not much larger than \sqrt{k} (times a constant depending on R). Since $\text{dist}(x,y) \leq \sqrt{k}$, we won't be more than order \sqrt{k} away from y . Also, with a fixed probability, the number of times we use the lazy simple random walk component up to time k is at least $k\gamma/2$. By (42), these latter steps will bring us to y with probability at least a constant times $\frac{1}{k^{d/2}}$. \square

Since we are always starting in $\delta_0 \times \pi_p$, we may consider all the edges to have a \star , as in Section 6. We let $\tau_0 = 0$ and, for $j \geq 1$, let, as in (38), $\tau_j := \min\{i > \tau_{j-1} : i \in \mathcal{I}\}$ (recall (35) and (36)). We have that $\{\tau_j - \tau_{j-1}\}_{j \geq 1}$ are independent and all of these, except τ_1 , have the same distribution as those introduced in (38). τ_1 has a different distribution since we are not starting in Ω_{REG} . As it will be simpler to look at the real time corresponding to τ_k , we let $\tilde{\tau}_k := \frac{C_{6.11.1} \tau_k}{\mu}$. Recall, as introduced in the proof of Proposition 6.11, that $\mathcal{R}[u, v]$ denotes the range of the random walker during the time interval $[u, v]$.

Lemma 7.4. *For all d and for all $p \in (0, p_c(\mathbb{Z}^d))$, there are finite constants $C_{7.4.1} = C_{7.4.1}(d, p)$ and $C_{7.4.2} = C_{7.4.2}(d, p)$ so that for all n , for all μ and for all $i \geq 1$,*

$$\mathbf{E} \left[e^{C_{7.4.1} |\mathcal{R}[\tilde{\tau}_i, \tilde{\tau}_{i+1}]|} \right] \leq C_{7.4.2}.$$

In particular, for some constant $C_{7.4.3} = C_{7.4.3}(d, p)$, $\mathbf{E} [|\mathcal{R}[\tilde{\tau}_i, \tilde{\tau}_{i+1}]|] \leq C_{7.4.3}$ for all n , for all μ and for all $i \geq 1$.

Proof. We prove this for $i = 1$. For $i \geq 2$, the distributions are of course the same while one can modify the argument for $i = 0$. We say that a family of random variables $\{X_\alpha\}$ has *uniform exponential tails* if there are constants c_1 and c_2 so that $\mathbf{E}[e^{c_1 X_\alpha}] \leq c_2$ for all α . The proof of Theorem 6.18, where Proposition 3.2 is critically invoked, actually shows that the family of random variables $\tau_2 - \tau_1$ as we vary n and μ has uniform exponential tails. (One can similarly argue that this is also true for τ_1 .) Note next that

$$|\mathcal{R}[\tilde{\tau}_1, \tilde{\tau}_2]| \leq \sum_{k=0}^{\tau_2 - \tau_1 - 1} \left| \mathcal{R} \left[\frac{\tau_1 C_{6.11.1}}{\mu} + \frac{k C_{6.11.1}}{\mu}, \frac{\tau_1 C_{6.11.1}}{\mu} + \frac{(k+1) C_{6.11.1}}{\mu} \right] \right|. \quad (48)$$

For $\beta = \beta(d, p)$ sufficiently small, we have that for all n and μ , the probability that, for $\{\eta_t\}_{t \geq 0}$ started in stationarity, a fixed edge e is open at some point in $[0, \beta/\mu]$ is less than $p_c(\mathbb{Z}^d)$. It follows from Theorem 6.3 that started from stationarity, the family of random variables $|\mathcal{R}[s, s + \frac{\beta}{\mu}]|$ as n and μ vary has uniform exponential tails. Since the distribution of $|\mathcal{R}[\frac{\tau_1 C_{6.11.1}}{\mu} + s, \frac{\tau_1 C_{6.11.1}}{\mu} + s + \frac{\beta}{\mu}]|$ is just the distribution of $|\mathcal{R}[u, u + \frac{\beta}{\mu}]|$ conditioned on an event whose probability is uniformly bounded away from 0 (namely, that the edges next to the walker are closed at the appropriate time), we have that the family of random variables $|\mathcal{R}[\frac{\tau_1 C_{6.11.1}}{\mu} + s, \frac{\tau_1 C_{6.11.1}}{\mu} + s + \frac{\beta}{\mu}]|$ as we vary n , μ and s has uniform exponential tails. Since each of the summands in (48) is a sum of a uniformly bounded number of such random variables, the family of summands also has uniform exponential tails. Finally, this together with the fact that the variables $\tau_2 - \tau_1$ has uniform exponential tails (as n and μ vary) allows us to invoke Lemma 3.3 to conclude that the family of random variables $\{|\mathcal{R}[\tilde{\tau}_1, \tilde{\tau}_2]|\}$ has uniform exponential tails (as n and μ vary). This is exactly the claim of the lemma. \square

Proof of Theorem 1.12. It is easy to show that for any fixed n , one can find constants so that the upper and lower bounds hold for all μ . Hence we need to only consider large n below. We first establish the lower bounds. Fix d and p with $p \in (0, p_c(\mathbb{Z}^d))$. As usual, we choose $\beta = \beta(d, p)$ sufficiently small, so that for all n and μ , the probability that, for $\{\eta_t\}_{t \geq 0}$ started at stationarity, a fixed edge e is open at some point in $[0, \beta/\mu]$ is less than $p_c(\mathbb{Z}^d)$. Since $u \times \pi_p$ is stationary, it easily follows that starting from $\delta_0 \times \pi_p$, the distribution of the cluster of the random walker at any time t has the same distribution as an ordinary cluster at the origin. It follows, as we saw in the proof of Lemma 7.4, that for any t , $E|\mathcal{R}[t, t + \beta/\mu]| \leq C_{1.12.1}$ where $C_{1.12.1}$ only depends on d and p .

We now move to the dimension dependent parts of the argument, going in order of increasing difficulty.

Case (iii): $d \geq 3$: lower bound.

It follows from the above that for any N ,

$$\mathbf{E}[|\mathcal{R}[0, N/\mu]|] \leq C_{1.12.1} N/\beta.$$

Letting $N := \frac{\beta n^d}{3C_{1.12.1}}$, we have

$$\mathbf{E} \left[\left| \mathcal{R} \left[0, \frac{\beta n^d}{3C_{1.12.1}\mu} \right] \right| \right] \leq n^d/3.$$

On the other hand, Markov's inequality implies that starting from $\delta_0 \times \pi_p$, we have that for any n and d and any $y \in \mathbb{T}^{d,n}$,

$$\mathbf{P}[\sigma_y \geq 2H_d(n)] \leq 1/2$$

from which it follows that

$$\mathbf{E}[|\mathcal{R}[0, 2H_d(n)]|] \geq n^d/2.$$

We conclude that $H_d(n) \geq \frac{\beta n^d}{6C_{1.12.1}\mu}$ as desired.

Case (i): $d = 1$: lower bound.

Let y have maximum distance from 0 in $\mathbb{T}^{1,n}$. Observe that

$$\{\sigma_y < \frac{\alpha n^2}{\mu}\} \subseteq E_1 \cup E_2$$

where

$$E_1 := \left\{ \max \left\{ \text{dist}(X_{\frac{\beta k}{\mu}}, 0) : 1 \leq k \leq \frac{\alpha n^2}{\beta} \right\} \geq \frac{n}{8} \right\}$$

and

$$E_2 := \cup_{k=1}^{\frac{\alpha n^2}{\beta}} \{|\mathcal{R}[\beta(k-1)/\mu, \beta k/\mu]| \geq n/8\}.$$

By Markov's inequality,

$$\mathbf{P}[E_1] \leq \mathbf{E} \left[\max \left\{ \text{dist}^2(X_{\frac{\beta k}{\mu}}, 0) : 1 \leq k \leq \frac{\alpha n^2}{\beta} \right\} \right] \frac{64}{n^2}.$$

By the proof of Theorem 1.4 but using the maximal version of Lemma 4.2 mentioned earlier instead, we obtain that the latter term is at most $O_{d,p}(1)\alpha$. Therefore, if $\alpha = \alpha(d, p)$ is taken sufficiently small, then $\mathbf{P}[E_1] \leq 1/4$. Since $|\mathcal{R}[s, s + \frac{\beta}{\mu}]|$ has uniform exponential tails as n and μ vary, it follows that $\mathbf{P}[E_2] \rightarrow 0$ as $n \rightarrow \infty$. We conclude that for large n , $\mathbf{P}[\sigma_y < \frac{\alpha n^2}{\mu}] \leq 1/2$ and hence $\mathbf{E}[\sigma_y] \geq \frac{\alpha n^2}{2\mu}$ yielding the desired lower bound.

Case (ii): $d = 2$: lower bound.

Note that we have that for all k ,

$$\tilde{\tau}_k \geq \frac{k C_{6.11.1}}{\mu}. \quad (49)$$

By Lemma 7.4, there exists R such that the second moments of $|\mathcal{R}[\tilde{\tau}_i, \tilde{\tau}_{i+1}]|$ are bounded uniformly (in n and μ) by R . We also let γ be such that the distribution of $X_{\tilde{\tau}_{i+1}} - X_{\tilde{\tau}_i}$ contains, for all n and μ , $\gamma\nu_{\text{LSRW}}$.

For simplicity, let $C_{1.12.2} := \frac{C_{7.3}(\gamma, R, d)}{16C_{7.4.3}}$. If there exists $y \in \mathbb{T}^{d,n}$ such that

$$\mathbf{P}[\sigma_y \leq \tilde{\tau}_{C_{1.12.2} n^2 \log n}] \leq 3/4, \quad (50)$$

then, by (49), we have

$$\mathbf{P} \left[\sigma_y \leq \frac{C_{6.11.1} C_{1.12.2} n^2 \log n}{\mu} \right] \leq 3/4.$$

This would lead to

$$\mathbf{E}[\sigma_y] \geq \frac{C_{6.11.1} C_{1.12.2} n^2 \log n}{4\mu}$$

and we would be done. Hence we may assume that for all $y \in \mathbb{T}^{d,n}$

$$\mathbf{P}[\sigma_y \leq \tilde{\tau}_{C_{1.12.2} n^2 \log n}] \geq 3/4. \quad (51)$$

We will prove below that under this assumption, for all large n , we have that

$$\mathbf{E} [|\mathcal{R}[0, \tilde{\tau}_{C_{1.12.2} n^2 \log n}]|] \leq \frac{n^2}{4}. \quad (52)$$

On the other hand, by summing (51) over y , one immediately obtains

$$\mathbf{E} \left[|\mathcal{R}[0, \tilde{\tau}_{C_{1.12.2} n^2 \log n}]| \right] \geq \frac{3n^2}{4}. \quad (53)$$

This contradiction tells us that in fact (51) cannot be true for all y and so we would have our desired lower bound on the hitting time.

We are now left to prove that (51) implies (52). First we abbreviate the event $\{\sigma_y \leq \tilde{\tau}_{C_{1.12.2} n^2 \log n}\}$ by U_y . Next, for $i \geq 1$, let $Y_i := \mathcal{R}[\tilde{\tau}_{i-1}, \tilde{\tau}_i]$. Clearly for all $y \in \mathbb{T}^{d,n}$, we have that

$$\mathbf{P}[U_y] \leq \frac{\mathbf{E} \left[\sum_{i=1}^{2C_{1.12.2} n^2 \log n} I_{\{y \in Y_i\}} \right]}{\mathbf{E} \left[\sum_{i=1}^{2C_{1.12.2} n^2 \log n} I_{\{y \in Y_i\}} \mid U_y \right]}.$$

We will argue further down that for all y

$$\mathbf{E} \left[\sum_{i=1}^{2C_{1.12.2} n^2 \log n} I_{\{y \in Y_i\}} \mid U_y \right] \geq C_{7.3}(\gamma, R, d) \log n / 2. \quad (54)$$

Assuming this for the moment, we plug this into the previous inequality, sum over $y \in \mathbb{T}^{d,n}$, use Lemma 7.4 and conclude (52). This leaves us with proving (54).

Let $G_i := \{|\mathcal{R}[\tilde{\tau}_{i-1}, \tilde{\tau}_i]| \leq (\log n)^2\}$ and $G = \cap_{i=1}^{n^3} G_i$. Using Lemma 7.4, one easily sees that

$$\lim_{n \rightarrow \infty} \mathbf{P}[G] = 1. \quad (55)$$

Note that (55) and (51) imply that

$$\mathbf{P}[G \mid U_y] = 1 - o(1).$$

Let $T := \min\{i : y \in Y_i\}$. By writing $\mathbf{P}[G \mid U_y]$ as

$$\sum_{j=1}^{C_{1.12.2} n^2 \log n} \mathbf{P}[G \mid T = j] \mathbf{P}[T = j \mid U_y],$$

we see that there must be a set $J_n \subseteq \{1, 2, \dots, C_{1.12.2} n^2 \log n\}$ so that

$$\mathbf{P}[T \in J_n \mid U_y] = 1 - o(1) \text{ and } \inf_{j \in J_n} \mathbf{P}[G \mid T = j] = 1 - o(1).$$

These imply that

$$\begin{aligned} \mathbf{E} \left[\sum_{i=1}^{2C_{1.12.2} n^2 \log n} I_{\{y \in Y_i\}} \mid U_y \right] &= \sum_j \mathbf{E} \left[\sum_{i=1}^{2C_{1.12.2} n^2 \log n} I_{\{y \in Y_i\}} \mid T = j \right] \mathbf{P}[T = j \mid U_y] \geq \\ &(1 - o(1)) \min_{j \in J_n} \mathbf{E} \left[\sum_{i=1}^{2C_{1.12.2} n^2 \log n} I_{\{y \in Y_i\}} \mid T = j \right] \geq \\ &(1 - o(1)) \min_{j \in J_n} \mathbf{P}[G_j \mid T = j] \mathbf{E} \left[\sum_{i=1}^{2C_{1.12.2} n^2 \log n} I_{\{y \in Y_i\}} \mid G_j, T = j \right] \geq \end{aligned}$$

$$(1 - o(1)) \min_{j \in J_n} \mathbf{E} \left[\sum_{i=1}^{2C_{1.12.2} n^2 \log n} I_{\{y \in Y_i\}} \mid G_j, T = j \right] \geq$$

$$(1 - o(1)) \min \left\{ \sum_{i=j+(\log n)^4}^{j+n} \mathbf{P} [X_{\tilde{\tau}_i} = y \mid G_j, T = j] : j \in \{1, 2, \dots, C_{1.12.2} n^2 \log n\} \right\}.$$

Now, recall R was chosen so that the i.i.d. increments of $X_{\tilde{\tau}_j}, X_{\tilde{\tau}_{j+1}}, \dots$ have variances bounded by R , uniformly in n and μ . In addition, these increments are independent of the event $\{G_j, T = j\}$. The occurrence of G_j implies that $\text{dist}(y, X_{\tilde{\tau}_j}) \leq (\log n)^2$. It then follows from Lemma 7.3 that for $i \in \{j + (\log n)^4, \dots, j + n\}$, $\mathbf{P} [X_{\tilde{\tau}_i} = y \mid G_j, T = j] \geq \frac{C_{7.3}(\gamma, R, d)}{i-j}$. Summing over our set of i yields (54).

We now move to the upper bounds which can be obtained simultaneously for all dimensions with an application of Lemma 7.2 as follows.

Letting $U_k := X_{\tilde{\tau}_k}$ as in (38), we have that $\{U_k - U_{k-1}\}_{k \geq 1}$ are independent with all, except the first, having the same distribution. (Recall d and n are suppressed in the notation.) We will apply Lemma 7.2 to the random walk on $\mathbb{T}^{d,n}$ whose steps have distribution $\{U_2 - U_1\}$. The key hypothesis of this lemma holds as indicated in (39). Letting $\hat{\sigma}_y$ denote the hitting time for our induced discrete time system, Lemma 7.2 yields that $\mathbf{E} [\hat{\sigma}_y]$ is bounded above by $O_{d,p}(1)n^2$, $O_{d,p}(1)n^2 \log n$ and $O_{d,p}(1)n^d$ respectively in dimensions 1, 2 and ≥ 3 . (Note that while the constant in Lemma 7.2 depends on γ , our γ here only depends on d and p .)

Note next that

$$\sigma_y \leq \sum_{i=1}^{\hat{\sigma}_y} \tilde{\tau}_i - \tilde{\tau}_{i-1}$$

By Wald's lemma [17, Theorem 4.1.5], we obtain

$$\mathbf{E} [\sigma_y] \leq \frac{C_{6.11.1}}{\mu} \mathbf{E} [\tau_2 - \tau_1] \mathbf{E} [\hat{\sigma}_y].$$

We cheated a little since τ_1 has a different distribution than the other $\tau_i - \tau_{i-1}$'s; however, as mentioned earlier, the proof of Theorem 6.18 easily yields that $\mathbf{E} [\tau_1]$ is also bounded by a constant depending only on d and p and so this can be ignored in the above. Finally, Theorem 6.18 and our derived bound on $\mathbf{E} [\hat{\sigma}_y]$ yields the desired upper bound on $\mathbf{E} [\sigma_y]$. \square

8 Open questions

In Theorem 3.1, the diffusion constant σ^2 depends on d, p and μ . It is easy to see that $\lim_{\mu \rightarrow \infty} \sigma^2(d, p, \mu) = p$ since in the large μ limit, this just corresponds to random walk time scaled by p . Also Corollary 1.6 implies that for $p < p_c(d)$, we have that $\liminf_{\mu \rightarrow 0} \frac{\sigma^2(d, p, \mu)}{\mu} < \infty$. One can also show that the corresponding lim sup is positive in this case as well.

1. Show that $\sigma^2(d, p, \mu)$ is increasing in both p and μ .
2. Show that even $\frac{\sigma^2(d, p, \mu)}{\mu}$ is increasing in μ .

Acknowledgements

A major part of this work was carried out when the second and third authors were visiting (at different times) Microsoft Research at Redmond, WA, and they thank the Theory Group for its hospitality and for creating a stimulating research environment. JES also acknowledges the support of the Swedish Research Council and the Knut and Alice Wallenberg Foundation.

References

- [1] D. Aldous and J. Fill. Reversible Markov chains and random walks on graphs. Available at <http://www.stat.berkeley.edu/~aldous/RWG/book.html>.
- [2] S. Andres. Invariance principle for the random conductance model with dynamic bounded conductances, 2012. arXiv:1202.0803.
- [3] L. Avena. Symmetric exclusion as a model of non-elliptic dynamical random conductances. *Electron. Commun. Probab.*, 17:1–8, 2012.
- [4] L. Avena, F. den Hollander, and F. Redig. Large deviation principle for one-dimensional random walk in dynamic random environment: attractive spin-flips and simple symmetric exclusion. *Markov Process. Related Fields*, 16:139–168, 2010.
- [5] L. Avena, F. den Hollander, and F. Redig. Law of large numbers for a class of random walks in dynamic random environments. *Electron. J. Probab.*, 16:587–617, 2011.
- [6] L. Avena, R. dos Santos, and F. Völlering. Transient random walk in symmetric exclusion: limit theorems and an Einstein relation. *ALEA, Latin American Journal of Probability and Mathematical Statistics*, 10(2):693–709, 2013.
- [7] K. Ball. Markov chains, Riesz transforms and Lipschitz maps. *Geom. Funct. Anal.*, 2:137–172, 1992.
- [8] A. Bandyopadhyay and O. Zeitouni. Random walk in dynamic Markovian random environment. *ALEA, Latin American journal of probability and mathematical statistics*, 1:205–224, 2006.
- [9] M.T. Barlow and E.A. Perkins. Symmetric Markov chains in \mathbf{Z}^d : how fast can they move? *Probab. Theory Related Fields*, 82:95–108, 1989.
- [10] I. Benjamini and E. Mossel. On the mixing time of a simple random walk on the super critical percolation cluster. *Probab. Theory Related Fields*, 125:408–420, 2003.
- [11] I. Benjamini and O. Schramm. Percolation beyond \mathbf{Z}^d , many questions and a few answers. *Electron. Comm. Probab.*, 1:no. 8, 71–82 (electronic), 1996.
- [12] O. Blondel. Tracer diffusion at low temperature in kinetically constrained models, 2013. arXiv:1306.6500.
- [13] C. Boldrighini, R.A. Minlos, and A. Pellegrinotti. Almost-sure central limit theorem for a Markov model of random walk in dynamical random environment. *Probability Theory and Related Fields*, 109:245–273, 1997.

- [14] D. Dolgopyat, G. Keller, and C. Liverani. Random walk in Markovian environment. *The Annals of Probability*, 36:1676–1710, 2008.
- [15] D. Dolgopyat and C. Liverani. Random walk in deterministically changing environment. *ALEA, Latin American journal of probability and mathematical statistics*, 4:89–116, 2008.
- [16] D. Dolgopyat and C. Liverani. Non-perturbative approach to random walk in Markovian environment. *Electronic Communication in Probability*, 14:245–251, 2009.
- [17] R. Durrett. *Probability Theory: Theory and Examples*. Cambridge University Press, 4th edition, 2010.
- [18] D. Figueiredo, P. Nain, B. Ribeiro, E. de Souza e Silva, and D. Towsley. Characterizing continuous time random walks on time varying graphs. In *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE joint international conference on Measurement and Modeling of Computer Systems*, 2012.
- [19] G. Grimmett. *Percolation*. Springer-Verlag, 2nd edition, 1999.
- [20] O. Häggström, Y. Peres, and J.E. Steif. Dynamical percolation. *Annales de l’Institut Henri Poincaré Probabilités et Statistiques*, 33:497–528, 1997.
- [21] F. Hollander, R. dos Santos, and V. Sidoravicius. Law of large numbers for non-elliptic random walks in dynamic random environments. *Stochastic Process. Appl.*, 123:156–190, 2013.
- [22] G.F. Lawler and V. Limic. *Random walk: a modern introduction*. Cambridge University Press, 2nd edition, 2010.
- [23] D.A. Levin, Y. Peres, and E.L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2009.
- [24] Şerban Nacu and Yuval Peres. Fast simulation of new coins from old. *Ann. Appl. Probab.*, 15(1A):93–115, 2005.
- [25] A. Naor, Y. Peres, O. Schramm, and S. Sheffield. Markov chains in smooth Banach spaces and Gromov-hyperbolic metric spaces. *Duke Mathematical Journal*, 134:165–197, 2006.
- [26] F. Redig and F. Völlering. Random walks in dynamic random environments: a transference principle. *The Annals of Probability*, 41(5):3157–3180, 2013.
- [27] J.E. Steif. A survey of dynamical percolation. In *Fractal Geometry and Stochastics IV*, pages 147–174. Birkhäuser Verlag, 2009.

Yuval Peres
 Microsoft Research
 One Microsoft Way
 Redmond, WA
 98052-6399
 USA
 peres@microsoft.com
<http://research.microsoft.com/en-us/um/people/peres/>

Alexandre Stauffer
 Department of Mathematical Sciences

University of Bath
Claverton Down, Bath
BA2 7AY, U.K.
a.stauffer@bath.ac.uk
<http://people.bath.ac.uk/ados20/>

Jeffrey E. Steif
Mathematical Sciences
Chalmers University of Technology
and
Mathematical Sciences
Göteborg University
SE-41296 Gothenburg, Sweden
steif@math.chalmers.se
<http://www.math.chalmers.se/~steif/>