Annales de l'Institut Henri Poincaré - Probabilités et Statistiques
2020, Vol. 56, No. 4, 2745–2773
https://doi.org/10.1214/20-AIHP1057

ANNALES DE L'INSTITUT HENRI POINCARÉ
PROBABILITÉS ET STATISTIQUES

Cutoff for random walk on dynamical Erdős–Rényi graph

Perla Sousi* and Sam Thomas†

Statistical Laboratory, University of Cambridge, UK. E-mail: *p.sousi@statslab.cam.ac.uk; †s.m.thomas@statslab.cam.ac.uk

Received 20 March 2019; revised 30 October 2019; accepted 6 March 2020

Abstract. We consider dynamical percolation on the complete graph K_n, where each edge refreshes its state at rate $\mu \ll 1/n$, and is then declared open with probability $p = \lambda/n$ where $\lambda > 1$. We study a random walk on this dynamical environment which jumps at rate $1/n$ along every open edge. We show that the mixing time of the full system exhibits cutoff at $\frac{3}{2}\log n/\mu$. We do this by showing that the random walk component mixes faster than the environment process; along the way, we control the time it takes for the walk to become isolated.

Résumé. Nous considérons le modèle de percolation dynamique sur le graphe complet K_n, où chaque arête réactualise son état au taux $\mu \ll 1/n$, et est ensuite déclarée ouverte avec probabilité $p = \lambda/n$, où $\lambda > 1$. Nous étudions une marche aléatoire sur cet environnement dynamique qui saute à taux $1/n$ à travers chaque arête ouverte. Nous montrons que le temps de mélange de tout ce processus a un cutoff au temps $\frac{3}{2}\log n/\mu$. Nous l'obtenons en montrant que la composante marche aléatoire mélange plus vite que le processus d'environnement; au passage nous contrôlons le temps que met la marche avant d'être isolée.

MSC2020 subject classifications: 05C81; 60J27; 60K35; 60K37

Keywords: Dynamical percolation; Erdős–Rényi; Random walk; Mixing times; Coupling

1. Introduction

In this paper we consider a random walk on a dynamically evolving random graph. Fix an underlying (undirected) graph $G = (V, E)$. Write $n = |V|$. The dynamics of the graph are that of *dynamical percolation*: edges refresh independently at rate μ; upon refreshing, the edge is declared open with probability p and closed with probability $1 - p$. We denote the state of the graph at time t by $\eta_t \in \{0, 1\}^E$: 0 corresponds to a closed edge and 1 to an open edge. The location of the random walker at time t is denoted $X_t \in V$: it moves at rate 1; when its exponential clock rings, it chooses uniformly at random a neighbour v of its current location, x say, and jumps from x to v if and only if the edge connecting x and v is open (at this time), otherwise it remains in place.

Write π_{RW} for the invariant distribution of the nearest-neighbour simple random walk on G (ie the degree-biased distribution) and π_p for the product measure on $\{0, 1\}^E$ with density p. The full process, (X, η), is reversible with invariant distribution $\pi = \pi_{\mathrm{RW}} \times \pi_p$.

We emphasise that the pair $(X_t, \eta_t)_{t \geq 0}$ is Markovian, as is just the graph process $(\eta_t)_{t \geq 0}$, while the location of the walker alone $(X_t)_{t \geq 0}$ is not: indeed, its transitions depend on the current graph. (Note that $(\eta_t)_{t \geq 0}$ is a biased simple random walk on the hypercube $\{0, 1\}^E$.) For all our results, we take $p = \lambda/n$, and emphasise that λ is a fixed constant, while n and $\mu = \mu_n$ vary.

This model was introduced by Peres, Stauffer and Steif in [21]. They used the torus $\mathbb{Z}_n^d$ (with d fixed) as their underlying graph; in this paper we use the complete graph K_n as our underlying graph. Hence from now on we take

$$V = \{1, \dots, n\} \quad \text{and} \quad E = \big\{(i, j) \mid i, j \in \{1, \dots, n\}, i \neq j\big\}.$$

Percolation on the complete graph gives precisely the Erdős–Rényi graph, and hence the name 'dynamical Erdős–Rényi'; we denote the measure of an Erdős–Rényi graph by π_{ER}, and note that in this case $\pi_p = \pi_{\mathrm{ER}}$. Also, we denote the uniform measure on $\{1, \dots, n\}$ by π_U, and note that in this case $\pi_{\mathrm{RW}} = \pi_U$.

This research was supported by the EPSRC PS by EP/R022615/1 and ST by Doctoral Training Grant #1885554.

Taking $p = \lambda/n$, for a constant λ, corresponds to the sparse regime for percolation, in which the expected degree of a vertex is order 1 (ie independent of n in the limit). Since the majority of the degrees are order 1, the walker takes steps on the timescale n (the majority of the time). We consider $\mu \ll 1/n$ so that the walk takes a large number of steps before seeing a local update to the graph; for bounded degree graphs, we could take $\mu \ll 1$. Our proofs actually require very slightly more, namely a polylogarithmic factor: we consider $\mu \ll (\log n)^{-\alpha}/n$ for a fixed $\alpha > 0$; no attempt has been made to optimise this parameter.

As in [21], we look at the ε-mixing time of the *full system* (X, η):

$$t_{\text{mix}}(\varepsilon) = \inf\Big\{ t \geq 0 \,|\, \max_{x_0,\eta_0} \big\| \mathbb{P}_{x_0,\eta_0}\big((X_t, \eta_t) \in \cdot\big) - \pi \big\|_{\text{TV}} \leq \varepsilon \Big\}.$$

When the leading order term of $t_{\text{mix}}(\varepsilon)$ is independent of ε, we say that there is *cutoff*. The *cutoff window* is given by the order of $t_{\text{mix}}(\varepsilon) - t_{\text{mix}}(1-\varepsilon)$, which will depend on ε.

In order to state our mixing result, we must first define some notation for iterated logarithm:

set $\log_{(1)} n = \log n$ and define inductively $\log_{(m+1)} n = \log(\log_{(m)} n)$ for $m \geq 1$.

Our main result considers the *supercritical regime* of percolation, ie has $p = \lambda/n$ where $\lambda > 1$ is a constant, and states that the full system (X, η) then exhibits cutoff at time $\frac{3}{2}(\log n)/\mu$ with cutoff window of smaller order than $(\log_{(M)} n)/\mu$ for all M.

Theorem 1.1 (Cutoff for Full System). *For all $\lambda > 1$, all $\varepsilon \in (0, 1)$, all $M \in \mathbb{N}$ and all n sufficiently large, for $p = \lambda/n$ and $\mu \leq (\log n)^{-20}/n$, we have*

$$\left| t_{\text{mix}}(\varepsilon) - \frac{3}{2}\big(\log n\big)/\mu \right| \leq (\log_{(M)} n)/\mu.$$

We also consider the 'mixing' of the *random walk component*:

$$t^{\text{RW}}_{\text{mix}}(\varepsilon, \eta_0) = \sup\Big\{ t \geq 0 \,|\, \max_{x_0} \big\| \mathbb{P}_{x_0,\eta_0}(X_t \in \cdot) - \pi_U \big\|_{\text{TV}} \geq \varepsilon \Big\} \quad \text{for } \eta_0 \in \{0, 1\}^E.$$

Since X is not a Markov chain, we do not have a priori that the total variation distance from uniform is decreasing, hence we do not define the mixing time to be 'the first time the total variation distance is below ε', but rather 'the last time the total variation distance is above ε'. (Of course, for a Markov chain, these notions are the same.) Note that, trivially by projection, $t^{\text{RW}}_{\text{mix}}(\varepsilon, \eta_0) \leq t_{\text{mix}}(\varepsilon)$ for all ε and all η_0. We show that 'the walk mixes faster than the environment' when the initial environment is 'typical', in the following precise sense.

Theorem 1.2 (Mixing of Random Walk). *For all $\lambda > 1$, all $\varepsilon \in (0, 1)$, all $M \in \mathbb{N}$ and all n sufficiently large, for $p = \lambda/n$ there exists a subset $H \subseteq \{0, 1\}^E$ with $\pi_{\text{ER}}(H) = 1 - o(1)$ so that, for all $\eta_0 \in H$, for $\mu \leq (\log n)^{-20}/n$, we have*

$$t^{\text{RW}}_{\text{mix}}(\varepsilon, \eta_0) \leq \log_{(M)} n/\mu;$$

ie, if $\eta_0 \sim \pi_{\text{ER}}$ then, for all $\varepsilon \in (0, 1)$, all $M \in \mathbb{N}$ and all n sufficiently large, we have

$$t^{\text{RW}}_{\text{mix}}(\varepsilon, \eta_0) \leq \log_{(M)} n/\mu \quad \textit{with probability } 1 - o(1).$$

Remark. From this intuitively it is clear why we get cutoff. First, let the environment mix. This is just a (biased) random walk on the hypercube $\{0, 1\}^N$ where $N = \binom{n}{2}$, and has cutoff at time $\frac{1}{2}\log(N/p)/\mu = (\frac{3}{2}\log n + \mathcal{O}(1))/\mu$. We prove this in Proposition 5.2; cf. [17, Example 12.19], where the unbiased case is considered. At this time, the graph is 'approximately' Erdős–Rényi, and so likely to be in the set H (from Theorem 1.2). Finally we let the walk mix. This takes time little-o of the mixing of the environment. This is indeed the heuristic that we use, but one has to be careful due to correlations between the walk and environment.

Above we are allowed to choose X_0 dependent on η_0 (and vice versa). In Section 6, we consider drawing η_0 according to π_{ER}, and then choosing X_0 *independently* of η_0. By symmetry, we may assume $X_0 = 1$. We then look at the mixing time of the walk on this (evolving) graph:

$$t^{\text{RW}}_{\text{mix}}(\varepsilon) = \sup\big\{ t \geq 0 \,|\, \big\| \mathbb{P}_{1,\text{ER}}(X_t \in \cdot) - \pi_U \big\|_{\text{TV}} \geq \varepsilon \big\},$$

where $\mathbb{P}_{x_0,\mathrm{ER}}(\cdot) = \sum_{\eta_0} \mathbb{P}_{x_0,\eta_0}(\cdot)\pi_{\mathrm{ER}}(\eta_0)$ averages the initial environment with respect to the Erdős–Rényi measure. We prove a sharp, up to constants, result on this mixing time. The result does not require us to consider the supercritical regime, ie $\lambda > 1$, but allows any $\lambda \in (0, \infty)$, including the critical case $\lambda = 1$. In Section 6, we prove the following.

Theorem 1.3. *For all $\lambda \in (0, \infty)$, there exists a constant C so that, for all $\varepsilon \in (0, 1)$ and all n sufficiently large, $p = \lambda/n$, we have the following bounds on the mixing time*:

$$t_{\mathrm{mix}}^{\mathrm{RW}}(\varepsilon) \ge \frac{1}{\mu} \cdot \frac{1}{2\lambda} \log(1/\varepsilon) \quad \text{if } \varepsilon \in \left(0, e^{-3\lambda} \wedge 1\right) \text{ for any } \mu;$$

$$t_{\mathrm{mix}}^{\mathrm{RW}}(\varepsilon) \le \frac{1}{\mu} \cdot C \log(1/\varepsilon) \quad \text{if } \varepsilon \in \left(0, \frac{1}{4}\right) \text{ when } \mu \le \frac{2}{3}(1+\lambda)^{-1}/n.$$

Note that once all edges of the graph have been refreshed, the graph has the Erdős–Rényi measure, and is independent of η_0. This is the coupon-collector problem, and takes time concentrated at $\log N \approx 2\log n$. At first glance, then, it appears that the idea of the above remark along with Theorem 1.3 can be applied to easily give pre-cutoff at $2\log n$. However, this is not the case: to prove the above statement, it is crucial that X_0 is chosen *independently* of η_0; we then exploit symmetry. At time $2\log n$, the walk and environment are correlated, and so the argument does not apply; hence the need for Theorem 1.2.

The above statement *suggests* that the upper bound in Theorem 1.2 is probably not sharp. However, the main interest in Theorem 1.2 compared with Theorem 1.1 is not the specific upper bound, but the fact that 'most' η_0 have $t_{\mathrm{mix}}^{\mathrm{RW}}(\varepsilon, \eta_0) \ll t_{\mathrm{mix}}(\varepsilon)$, ie that the walk mixes faster than the full system. This is what allows us to show cutoff.

Remark. For the rest of the paper, with the exception of Section 6, we assume that $\mu \le (\log n)^{-20}/n$; we shall not repeat this in the statement of every theorem.

The model of dynamical percolation (without the random walker) was introduced by Häggström, Peres and Steif in [13]. The model with the random walk was then introduced by Peres, Stauffer and Steif in [21], with underlying graph the torus $\mathbb{Z}_n^d$. They considered the *subcritical regime*, ie $p < p_c(\mathbb{Z}^d)$, the critical probability for bond percolation on $\mathbb{Z}^d$, and obtained the correct order for the mixing of the full system, showing further that the order of the mixing of the walk is the same order as the mixing of the full system (in contrast to our model). While they give a very complete picture of the subcritical mixing, key to their proofs is that the percolation clusters are all small, and that there is no giant component (taking up a constant proportion of the vertices).

The *supercritical regime*, ie $p > p_c(\mathbb{Z}^d)$ where such a giant does exist, was then considered by Peres, Sousi and Steif in [20]. They considered the 'quenched' case, where a 'typical' environment process $\{\eta_t\}_{t\ge 0}$ is fixed in advance, and the walker walks on this. They obtained the correct order for the mixing of the walk, up to polylogarithmic factors, but only in the regime where $\theta(p) > \frac{1}{2}$, ie the probability that the component at 0 in $\mathbb{Z}^d$ is infinite is greater than $\frac{1}{2}$. The case $\theta(p) \le \frac{1}{2}$ remains open.

Avena, Güldaş, van der Hofstad and den Hollander in [2,3] studied the mixing time of the non-backtracking random walk on a *dynamical configuration model*. The *configuration model* generates a random graph with a prescribed degree sequence, and the dynamics at every time step 'rewire' uniformly at random a given proportion of the edges: this rewiring involves cutting edges into two half-edges and then randomly repairing the half-edges.

It is straightforward to see that in our model when the walker first crosses a refreshed edge (which is a randomised stopping time) it is then (almost) uniform. The authors of [2,3] considered an analogous time for the nearest-neighbour simple random walk, namely the first time the walk crosses a rewired edge, and showed that the distribution of the walk at this time is (almost) its invariant distribution. In both cases, however, these times are not sufficient to show mixing: it is not the case that the walk 'remains close to invariant'; for example, there is significant probability that the walk will cross back over the same edge to its previous location. It is possible that a more refined analysis of a related stopping time—eg the first time the walk crosses a rewired edge and then 'escapes', not recrossing this edge again (for a long time)—would work. This approach is not taken in [2,3], though, and is left open. Rather, to resolve this the authors consider the non-backtracking random walk which, along with the locally tree-like structure of the configuration model, removes this 'crossing back' issue. They then show sharp asymptotics for the mixing time of this non-backtracking random walk, using the aforementioned stopping time.

In contrast to the above examples, in our work we show cutoff for the full process (X, η) for the entire supercritical regime, ie consider $p = \lambda/n$ for any constant $\lambda > 1$, and obtain the correct order of the mixing of the walk, up to an iterated log factor. Furthermore, our methods adapt immediately to the subcritical regime, ie $p = \lambda/n$ with $\lambda < 1$. We have not considered the details for this, but with a few concentration results on the structure of a subcritical Erdős–Rényi graph, a similar mixing result will follow.

Fountoulakis and Reed in [10] and Benjamini, Kozma and Wormald in [4] studied the mixing time of the nearest-neighbour simple random walk on the giant component of a supercritical Erdős–Rényi random graph, without any graph dynamics: they prove that the mixing time is order exactly $(\log n)^2$. Fountoulakis and Reed carefully studied the ratio between the size of the edge boundary of a set and the set itself, using a variation of the Lovász-Kannan integral, which they developed in [9]. Benjamini, Kozma and Wormald used a more geometric approach, defining a stripping process to analyse the (2-)*core* and the *kernel* of the graph; they show the kernel is a (type of) expander, and describe the *decorations* attached to the kernel.

The two works above consider mixing from the worst-case starting point. Berestycki, Lubetzky, Peres and Sly in [5] consider mixing when the starting point is chosen according to the invariant distribution. They show then that the mixing time is actually order $\log n$, obtaining the correct constant and also showing cutoff; contrast this with order $(\log n)^2$ for the worst-case.

2. Outline of proof and preliminaries

2.1. *Outline of proofs*

We now give a brief, informal outline of the proofs of the main results. First consider the following scenario: suppose a walker is isolated at vertex u (ie the walk is at the vertex u which is an isolated vertex in the current graph), and suppose it becomes non-isolated by the edge (u, v) opening, where v was isolated immediately before (u, v) opened. Now the pair $\{u, v\}$ is a component of the graph. Because $\mu \ll 1/n$, we see that the walker takes a large number of steps before this edge closes. If it closes before any other edge incident to $\{u, v\}$ opens (which has order 1 probability, by counting open/closed edges), then the walker is approximately uniformly distributed on $\{u, v\}$. So it has approximately 'done a lazy simple random walker step'.

This motivates the following coupling. First wait for the two walkers to be simultaneously isolated *in the same environment* of two full systems. Then to couple we want to imitate the standard coupling of the lazy simple random walk on the complete graph; we do this by considering the event that when the walkers become non-isolated they connect to a vertex that was isolated immediately prior. We give a precise definition of the coupling that we use in Section 5.

In order to find the time it takes for two walkers to be simultaneously isolated in the same environment, we first consider how long it takes one walker to become isolated. To find this time, we observe that a walker can only become isolated if it is at a degree 1 vertex and this vertex becomes isolated prior to the walk's leaving it. This motivates looking at the rate at which degree 1 vertices are hit. To do this, we compare the number of degree 1 vertices hit by a walker on the dynamic graph and the same quantity for a walker on a static graph; we then apply a Chernoff-style bound due to Gillman [12]. We give the precise details of this in Section 4.

A key element in studying the isolation time is to control how long the walk remains in the giant once it has entered. We show that since the graph updates slowly, as $\mu n \ll (\log n)^{-19}$, the walk does not see updates to the graph for some while; this time is long enough for the walk to become approximately uniform on the giant prior to seeing a change. We can then use structure results on the Erdős–Rényi graph to see how 'near the core' of the giant the walk is.

We then use the fact that an Erdős–Rényi graph with one vertex conditioned to be isolated has the distribution of an Erdős–Rényi graph on $n-1$ vertices (with edge-probability p) union an isolated vertex: this allows us to say that 'conditioning on one walker's being isolated has almost no affect on the other walker', which will allow us to treat the walkers as almost independent. We give the precise details of this in Section 4.2.

2.2. *Notation and terminology*

For functions f and g, we write $f(n) \lesssim g(n)$, or $f(n) = \mathcal{O}(g(n))$, if there exists a positive constant C so that $f(n) \le Cg(n)$ for all n; write $f(n) \gtrsim g(n)$, or $f(n) = \Omega(g(n))$, if $g(n) \lesssim f(n)$. We write $f(n) \asymp g(n)$, or $f(n) = \Theta(g(n))$, if we have both $f(n) \lesssim g(n)$ and $g(n) \lesssim f(n)$. Write $f(n) \ll g(n)$, or $f(n) = o(g(n))$, if $f(n)/g(n) \to 0$ as $n \to \infty$; write $f(n) \gg g(n)$, or $f(n) = \omega(g(n))$, if $g(n) \ll f(n)$.

For random variables X and Y, we write $X \preccurlyeq Y$ if Y stochastically dominates X (from above), ie if $\mathbb{P}(X \ge z) \le \mathbb{P}(Y \ge z)$ for all z; write $X \succcurlyeq Y$ if $Y \preccurlyeq X$.

For a real number $r > 0$, we write $\mathcal{E}(r)$ for the exponential random variable with rate r.

For real numbers α and β, we write $\alpha \wedge \beta = \min\{\alpha, \beta\}$.

2.3. *Good graphs and Erdős–Rényi structure results*

In this section we state some results on Erdős–Rényi graphs. Since a realisation can be *any* graph, we want to describe explicitly what we shall mean by a *good* graph. We now define some notation for graph properties. For the moment, let c_* and C_* be any two constants.

Notation 2.1. Let $n \in \mathbb{N}$; in the following, we suppress the n-dependence. For a graph $G = (V, E)$ with $V = \{1, \dots, n\}$, we use the following notation.

(i) For $x \in V$, write $d(x)$ for the degree of x (in G).
(ii) Write $\mathcal{G}$ for the (set of vertices in the) largest component, and call it the *giant*; if there is a tie, choose the option that includes the smallest labelled vertex.
(iii) For $x \in V$, call an edge a *removal edge* for x if its removal breaks the component of x in two, leaving x in the smaller component (breaking ties as above). Write $\mathcal{R}(x)$ for the set of removal edges for x, and write $R(x) = |\mathcal{R}(x)|$.
(iv) For $M \in \mathbb{N}$, write $\mathcal{W}^M$ for the vertices of the giant with at most $C_* \log_{(M)} n$ removal edges, ie

$$\mathcal{W}^M = \{x \in \mathcal{G} \mid R(x) \le C_* \log_{(M)} n\}.$$

(v) Write γ for the spectral gap and Φ for the isoperimetric constant, ie

$$\Phi = \min\left\{ \frac{|\partial S|}{d(S)} \,\middle|\, S \subseteq V, S \neq \varnothing, d(S) \le |E| \right\},$$

where $d(S) = \sum_{x \in S} d(x)$, $\partial S = \{(x, y) \in E \mid x \in S, y \notin S\}$; also write γ for the spectral gap of the transition matrix of the nearest-neighbour simple random walk on G (ie gap between 1 and the second largest eigenvalue).

If the graph $G = \zeta$, then we use subscript ζ, eg writing $\mathcal{G}_\zeta$ for the giant of ζ. If the graph $G = \eta_t$, then we add a subscript t, eg writing $\mathcal{G}_t$, $R_t(x)$ or $d_t(x)$.

Let ω_* be any function of n. We now define what we mean by a *good graph*.

Definition 2.2 (Good Graph). Let $n \in \mathbb{N}$; in the following, we suppress the n-dependence. We say that a graph $G = (V, E)$ with $V = \{1, \dots, n\}$ is *good*, and write $G \in \mathscr{G}$, if it has a unique component $\mathcal{G}$ with $|\mathcal{G}| \ge C_* \log n$, which we call the *giant*, which satisfies the following properties.

(i) *Size.* We have $|\mathcal{G}| \ge c_* n$.
(ii) *Maximum degree.* The maximum degree of $\mathcal{G}$ is at most $C_* \log n$.
(iii) *Number of edges.* There are at most $C_* n$ edges in $\mathcal{G}$.
(iv) *Number of degree 1 vertices.* The number of degree 1 vertices in $\mathcal{G}$ is at least $c_* n$.
(v) *Removal edges.* We have $R(x) \le C_* \log n$ for all $x \in \mathcal{G}$.
(vi) *Vertices far from the core.* For all $2 \le M \le \omega_*(n)$, the proportion of vertices x of $\mathcal{G}$ with $R(x) \ge C_* \log_{(M)} n$ is at most $(\log_{(M-1)} n)^{-4}$, ie $|\mathcal{G} \setminus \mathcal{W}^M|/|\mathcal{G}| \le (\log_{(M-1)} n)^{-4}$.
(vii) *Expansion properties.* We have $\Phi_{\mathcal{G}} \ge c_*(\log n)^{-2}$ and $\gamma_{\mathcal{G}} \ge c_*(\log n)^{-4}$.

This concept of *good* depends on the choice of constants c_* and C_* and function ω_*. The next proposition says that we can choose these parameters suitably so that an Erdős–Rényi graph is overwhelmingly likely to be good. Write $\mathscr{G}_t = \{\eta_t \in \mathscr{G}\}$, and also

$$\mathscr{G}[s, t] = \{\eta_u \in \mathscr{G} \ \forall u \in [s, t]\} = \bigcap_{s \le u \le t} \mathscr{G}_u.$$

Proposition 2.3. *There exist positive constants c_* and C_* and a function $\omega_*(n) \to \infty$ so that, for $G \sim \pi_{\mathrm{ER}}$, we have*

$$\mathbb{P}(G \notin \mathscr{G}) = \mathcal{O}(n^{-9}) \quad \text{and} \quad \mathbb{P}\left(\frac{1}{n}\big|\{x \in V \mid d(x) = 0\}\big| \le c_*\right) = \mathcal{O}(n^{-9}).$$

The claims in this proposition are fairly standard, but usually in the literature the proved decay rate is only $o(1)$, whereas we desire the quantitative $\mathcal{O}(n^{-9})$. We omit the technical details of this proof, but they can be found in the appendix of [23]. The proof follows standard arguments and results, particularly from [1,6–8,11,14,16,18,19,22,24,25].

For the remainder of the paper, we select (ω_*, c_*, C_*) as guaranteed by this proposition and fix them permanently; whenever ω_*, c_* or C_* is written below, it will refer to these constants.

We now consider our graph dynamics. We want not only the starting graph to be good, but we want it to remain good for a long time.

Definition 2.4. We make the following definitions:

$$\begin{aligned} \mathscr{H} &= \{\eta_0 \in \{0,1\}^E \mid \mathbb{P}_{\eta_0}(\mathscr{G}[0,1/\mu]^c) \le n^{-1}\}; \\ \mathscr{H}[s,t] &= \{\eta_u \in \mathscr{H}\ \forall u \in [s,t]\}; \\ H &= \{\eta_0 \in \{0,1\}^E \mid \mathbb{P}_{\eta_0}(\mathscr{H}[0,n/\mu]^c) \le n^{-1}\}. \end{aligned}$$

Further, if we are considering two environment processes, η and ξ say, then we (abuse notation slightly and) use the same notation, eg

$$\mathscr{H}[s,t] = \{\eta_u, \xi_u \in \mathscr{H}\ \forall u \in [s,t]\}.$$

We have the following result on 'how good' an Erdős–Rényi graph is.

Proposition 2.5. *For all $t \le n/\mu$, if $\eta_0 \sim \pi_{\mathrm{ER}}$, then we have*

$$\mathbb{P}(\mathscr{H}[0,t]^c) = \mathcal{O}(n^{-3}), \quad \textit{and hence} \quad \pi_{\mathrm{ER}}(H) = 1 - o(1).$$

Moreover, these still hold if we add the condition that at least a proportion c_* of the vertices are isolated to the definition of a good graph.

We state a large deviations result on the Poisson distribution, which we use on a number of occasions throughout the paper.

Lemma 2.6. *We have the following bounds, valid for all $\lambda > 0$ and all $\varepsilon \in (0,1)$:*

$$\begin{aligned} &\mathbb{P}(\mathrm{Po}(\lambda) \ge (1+\varepsilon)\lambda) \le \exp\left(-\frac{1}{2}\lambda\varepsilon^2\left(1 - \frac{1}{3}\varepsilon\right)\right); \\ &\mathbb{P}(\mathrm{Po}(\lambda) \le (1-\varepsilon)\lambda) \le \exp\left(-\frac{1}{2}\lambda\varepsilon^2\right). \end{aligned}$$

Proof of Proposition 2.5. Since π_{ER} is the invariant distribution for the environment, using the concentration of the Poisson distribution and Proposition 2.3 we find that

$$\mathbb{P}(\mathscr{G}[0,1/\mu]^c) = \mathbb{P}(\exists t \le 1/\mu \text{ s.t. } \eta_t \notin \mathscr{G}) \le n^2 \cdot \mathcal{O}(n^{-9}) + \exp\left(-\frac{1}{12}n^2\right) = \mathcal{O}(n^{-7}).$$

We now restrict the implicit sum in $\mathbb{P}$ from $\eta_0 \in \{0,1\}^E$ to $\eta_0 \in \mathscr{H}^c$:

$$\begin{aligned} \mathbb{P}(\mathscr{G}[0,1/\mu]^c) &= \sum_{\eta_0 \in \{0,1\}^E} \mathbb{P}_{\eta_0}(\mathscr{G}[0,1/\mu]^c)\pi_{\mathrm{ER}}(\eta_0) \\ &\ge \sum_{\eta_0 \in \mathscr{H}^c} \mathbb{P}_{\eta_0}(\mathscr{G}[0,1/\mu]^c)\pi_{\mathrm{ER}}(\eta_0) \ge n^{-1} \cdot \pi_{\mathrm{ER}}(\mathscr{H}^c). \end{aligned}$$

Hence we deduce that $\pi_{\mathrm{ER}}(\mathscr{H}^c) = \mathcal{O}(n^{-6})$. Repeating the same argument, we find that

$$\pi_{\mathrm{ER}}(H^c) = \mathcal{O}(n^{-2}).$$

□

3. Hitting and exit times of the giant

In this section we study the hitting time of the giant, and how long the walk remains in the giant given that it starts there. Write the following for the hitting and exit times of the giant:

$$\tau_{\mathcal{G}} = \inf\{t \geq 0 \mid X_t \in \mathcal{G}_t\} \quad \text{and} \quad \tau'_{\mathcal{G}} = \inf\{t \geq 0 \mid X_t \notin \mathcal{G}_t\}.$$

Lemma 3.1 (Hitting the Giant). *There exists a positive constant c so that, for all n sufficiently large and all* (x_0, η_0), *we have*

$$\mathbb{P}_{x_0,\eta_0}(\tau_{\mathcal{G}} \leq 1/\mu) \geq c - \mathbb{P}_{\eta_0}\big(\mathscr{G}[0, 1/\mu]^c\big).$$

Proof. Write $\mathcal{U}$ for the first time an edge incident to the walker refreshes and opens. By the memoryless property, $\mathcal{U} \sim \mathcal{E}(\lambda\mu(1 - 1/n))$. When such an edge opens, it connects to the giant with probability $|\mathcal{G}_{\mathcal{U}}|/(n-1)$. Write $\theta_t = |\mathcal{G}_t|/n$. We then have

$$\begin{aligned}
\mathbb{P}_{x_0,\eta_0}(\tau_{\mathcal{G}} \leq 1/\mu) &\geq \mathbb{P}_{x_0,\eta_0}(\tau_{\mathcal{G}} \leq \mathcal{U}, \mathcal{U} \leq 1/\mu, \theta_{\mathcal{U}^-} \geq c_*) \\
&\geq \mathbb{P}_{x_0,\eta_0}(\tau_{\mathcal{G}} \leq \mathcal{U} \mid \mathcal{U} \leq 1/\mu, \theta_{\mathcal{U}^-} \geq c_*) \cdot \mathbb{P}_{x_0,\eta_0}(\mathcal{U} \leq 1/\mu, \theta_t \geq c_* \ \forall t \leq 1/\mu) \\
&\geq c_*\big(\mathbb{P}_{x_0,\eta_0}(\mathcal{U} \leq 1/\mu) - \mathbb{P}_{\eta_0}\big(\{\theta_t \geq c_* \ \forall t \leq 1/\mu\}^c\big)\big).
\end{aligned}$$

The proof is completed by noting that $c_* \leq 1$ and $U \sim \mathcal{E}(\lambda\mu(1 - 1/n))$ so $\mathbb{P}_{x_0,\eta_0}(\mathcal{U} \leq 1/\mu) \asymp 1$. □

We now consider how long the walker remains in the giant once it enters. Recall the definition of R and $\mathcal{W}^M$ from Section 2.3. Since the number of removal edges satisfies $R(x) \leq C_* \log n$ for all $x \in \mathcal{G}$ when the graph is good, while the graph is good a trivial bound is $\tau'_{\mathcal{G}} \succcurlyeq \mathcal{E}(C_*\mu \log n)$. We do a more careful analysis which shows that, for all $M \in \mathbb{N}$, 'most of the time' $X_t \in \mathcal{W}_t^M$, ie satisfies $R(X_t) \leq C_* \log_{(M)} n$; this is because $|\mathcal{W}^M|/|\mathcal{G}| = 1 - o(1)$ for a good graph. The precise statement that we prove is as follows.

Proposition 3.2 (Exit Time from the Giant). *There exists a constant C so that, for all* $M \in \mathbb{N}$, *all n sufficiently large, all t with* $(\log n)^{-5} \leq \mu t \leq \frac{1}{10}$ *and all* (x_0, η_0) *with* $\eta_0 \in \mathscr{H}$ *and* $x_0 \in \mathcal{G}_0$, *we have*

$$\mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} \leq t\big) \leq C\mu t \log_{(M)} n.$$

Since $\mu n \ll 1$, the 'majority of the time' the walker takes a step before any edge incident to its location changes state. This motivates looking at a random walker moving on a *static* graph, ie one without graph-dynamics. We call such a walk the *static walk*, and the original walk (on the dynamic graph) the *dynamic walk*; we denote them by $\tilde{X}$ and X, respectively.

Consider starting the two walks together. Observe that until the static walk encounters an edge that is in a different state to its original, the two walks have the same distribution, and hence we can couple them to be the same (until this time) as follows. Give X and $\tilde{X}$ the same jump clock. When this clock rings, at time t say, both walks choose the same vertex; $\tilde{X}$ performs the jump if and only if the connecting edge is present in η_0, while X performs the jump if and only if the connecting edge is present in η_t. We call this the *static-dynamic coupling*.

We define the *set of edges seen by the walker* (in an interval $[s, t]$) as the set of all edges (open or closed) that are incident to the walker at some time (in an interval $[s, t]$). Until X sees an edge which is in a different state to its original, we can couple it with $\tilde{X}$, as described above.

Label the edges of the (complete) graph $e_1, \dots, e_N$, where $N = \binom{n}{2}$, in an arbitrary ordering (eg lexicographically). At time t, write $\mathscr{O}_t = \{o_t^1, o_t^2, \dots\}$ for the (ordered) set of open edges in η_t and $\mathscr{C}_t = \{c_t^1, c_t^2, \dots\}$ for the set of closed ones. We say that an edge is a *bridge* for a component if its removal splits the component into two (disconnected) parts. Recall also from Notation 2.1 the definition of a removal edge, and of $\mathcal{R}$.

Fix $\rho = (\log n)^{11}$. We now define an edge-set process $\mathcal{P} = (\mathcal{P}_t)_{t \geq 0}$, with $\mathcal{P}_t \subseteq E$ for all t.

Definition 3.3 (Set Processes). We define the edge-set process $\mathcal{P}$ inductively. Throughout the definition, we assume that the graph is good, ie we define $(\mathcal{P}_s)_{s \leq t}$ on the event $\mathscr{G}[0, t]$; recall the definition of good and $\mathscr{G}[0, t]$ from Definition 2.2 and the display after it.

Suppose we have defined the set process $\mathcal{P}$ up until time t, ie have defined $(\mathcal{P}_s)_{s \in [0,t)}$. Let V' be the most recent 'update time' for the process $(\mathcal{P}_s)_{s \in [0,t)}$ in the following sense:

$$V' = \sup\big\{s \in [0, t) \mid \text{ an edge of } \mathcal{P}_s \text{ changes state at time } s\big\}.$$

(If $t = 0$, then we take $V' = 0$. What follows is also used for the base case of the induction.) For $r \geq 0$, let $\mathcal{A}'_{r+t}$ (respectively $\mathcal{B}'_{r+t}$) be the set of open (respectively closed) edges seen by X in $[V', t+r]$. For $s \geq 0$, write $\mathcal{R}_s = \mathcal{R}_s(X_s)$. If $|\mathcal{A}'_{r+t} \cup \mathcal{R}_{r+t}| \leq \rho$ and $|\mathcal{B}'_{r+t}| \leq \rho n$, then set

$$\mathcal{A}_{r+t} = \mathcal{A}'_{r+t} \cup \mathcal{R}_{r+t} \cup \{o^1_{r+t}, \ldots, o^{a_{r+t}}_{r+t}\} \quad \text{and} \quad \mathcal{B}_{r+t} = \mathcal{B}'_{r+t} \cup \{c^1_{r+t}, \ldots, c^{b_{r+t}}_{r+t}\}, \tag{3.1}$$

where a_{r+t} and b_{r+t} are so that $|\mathcal{A}_{r+t}| = \rho$ and $|\mathcal{B}_{r+t}| = \rho n$; otherwise, set

$$\mathcal{A}_{r+t} = \{o^1_{r+t}, \ldots, o^{\rho}_{r+t}\} \quad \text{and} \quad \mathcal{B}_{r+t} = \{c^1_{r+t}, \ldots, c^{\rho n}_{r+t}\}. \tag{3.2}$$

Let V be the first update time for the process $(\mathcal{A}_{r+t} \cup \mathcal{B}_{r+t})_{r>0}$:

$$V = \inf\{r \geq 0 \mid \text{ an edge of } \mathcal{A}_{r+t} \cup \mathcal{B}_{r+t} \text{ changes state at time } r+t\} + t.$$

Define $\mathcal{P}_{r+t} = \mathcal{A}_{r+t} \cup \mathcal{B}_{r+t}$ for $r \in [0, V)$.

If at time s we use (3.1) to define $\mathcal{A}_s$ and $\mathcal{B}_s$, then we say that the set definitions *succeeded* at time s, and write $\mathcal{S}_s$ for this event; if we used (3.2), then we say they have *failed*.

Finally, for $t > 0$ define the event $\mathcal{S}[0, t) = (\bigcap_{s<t} \mathcal{S}_s) \cap \mathscr{G}[0, t)$.

Definition 3.4. Set $S = n(\log n)^8$ and $U_0 = 0$. We say that $\mathcal{P}$ *updates* (at time t) when an edge in the set $\mathcal{P}$ refreshes and changes state (at time t). We inductively define the sequence $U_1, U_2, \ldots$: for all $k \geq 1$, let V_k be the first time after U_{k-1} that $\mathcal{P}$ updates; set $U_k = V_k \wedge (U_{k-1} + S)$.

Work on the event $\mathcal{S}[0, t)$, and fix $s < t$. Then $\mathcal{A}_s$ is a collection of open edges of size ρ and $\mathcal{B}_s$ is a collection of closed edges of size ρn. Hence the set $\mathcal{P}_s = \mathcal{A}_s \cup \mathcal{B}_s$ updates at rate $\kappa\mu$ where

$$\kappa = (1-p)\rho + p\rho n = (1 - \lambda/n)\rho + \lambda\rho = (1 + \lambda - \lambda/n)\rho; \quad \text{note that } \kappa \asymp \rho = (\log n)^{11}. \tag{3.3}$$

By the memoryless property, $U_k - U_{k-1} \sim^{\text{iid}} \mathcal{E}(\kappa\mu) \wedge S$. Observe also that the walk may only leave the giant when the set $\mathcal{P}$ updates, in particular only at one of the times $U_1, U_2, \ldots$, but note that not all of these times are caused by updates: some are caused because of the threshold S.

We first look at the probability of the event $\{\tau'_{\mathcal{G}} \leq U_k\} \cap \mathcal{S}[0, U_k)$.

Lemma 3.5. *For all n, all k and all (x_0, η_0) with $x_0 \in \mathcal{G}_0$, we have*

$$\mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} \leq U_k, \mathcal{S}[0, U_k)\big) \leq k \max_{\eta_0, x_0 \in \mathcal{G}_0} \mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} = U_1, \mathcal{S}[0, U_1)\big).$$

Proof. Consider any (x_0, η_0) satisfying $x_0 \in \mathcal{G}_0$. By the union bound and the strong Markov property (applied at time U_{j-1} for the j-th term of the sum), we find that

$$\begin{aligned}
&\mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} \leq U_k, \mathcal{S}[0, U_k)\big) \\
&\quad \leq \sum_{j=1}^{k} \mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} = U_j, \mathcal{S}[0, U_j)\big) \\
&\quad \leq \max_{\eta_0, x_0 \in \mathcal{G}_0} \mathbb{P}\big(\tau'_{\mathcal{G}} = U_1, \mathcal{S}[0, U_1)\big) \cdot \sum_{j=1}^{k} \mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} > U_{j-1}, \mathcal{S}[0, U_{j-1})\big).
\end{aligned}$$

Upper bounding the sum by k completes the proof. □

We now determine the *uniform* mixing time of the static walk on a good giant. For a Markov chain Z with transition matrix P and invariant distribution π, the *uniform mixing time* is

$$t_{\text{unif}}(\varepsilon, Z) = \inf\Big\{t \geq 0 \mid \max_{x,y} \big|1 - p_t(x, y)/\pi(y)\big| \leq \varepsilon\Big\}.$$

Lemma 3.6. *Let G be a good graph, and let $\mathcal{G}$ be its (unique) giant. Consider the static walk, denoted $\tilde{X}$, on the giant. Write $t_{\mathrm{unif}}(\varepsilon, \tilde{X})$ for the ε-uniform mixing time of $\tilde{X}$ (on $\mathcal{G}$). Then*

$$t_{\mathrm{unif}}\left(\frac{1}{8}, \tilde{X}\right) \lesssim n(\log n)^6.$$

Proof. To prove this lemma, we compare $\tilde{X}$ with a 'sped-up' version. Consider a walk Z on the giant $\mathcal{G}$ of a good graph ζ. Write $m = |\mathcal{G}|$; so $m \asymp n$. Write $d(z)$ for the degree of z in ζ, and d^* for the maximum degree; note that $d^* \le C_* \log n$. Associate to a vertex $z \in \mathcal{G}$ the following set:

$$\mathcal{V}_z = \mathcal{N}_z \cup \{1, \dots, k_z\} \setminus \{z\} \quad \text{where } k_z \text{ is such that } V_z = |\mathcal{V}_z| = 2C_* \log n,$$

where $\mathcal{N}_z = \{z' \in \mathcal{G} \mid \zeta(z, z') = 1\}$ is the (open) neighbourhood of z (in ζ). (This is possible since $d^* \le C_* \log n$.) Give Z a rate 1 jump clock: when this clock rings, if Z is at z then a vertex z' is chosen uniformly at random from $\mathcal{V}_z$ and Z moves (from z) to z' if and only if the edge (z, z') is open, ie $\zeta(z, z') = 1$. Observe that Z is the same as the static walk $\tilde{X}$, except that it is sped-up by a factor $n/(2C_* \log n)$. Hence the mixing times are in ratio $n/(2C_* \log n)$, for both total variation and uniform mixing. We now calculate the uniform mixing time $t_{\mathrm{mix}}(\frac{1}{8}, Z)$.

Since $V_z = |\mathcal{V}_z| = 2C_* \log n \ge 2d^*$ for all $z \in \mathcal{G}$, we see that the chain Z is *lazy* in the sense that if we discretise by its rate 1 jump clock then the resulting discrete-time chain is *lazy*, ie $p(z, z) \ge \frac{1}{2}$ for all $z \in \mathcal{G}$. Moreover,

$$\pi_Z(z) = 1/|\mathcal{G}| = 1/m \quad \text{and} \quad p(z, z') = \frac{1}{2C_* \log n} \mathbf{1}(\zeta(z, z') = 1).$$

Hence Z is reversible. It is then known that

$$t_{\mathrm{unif}}(\varepsilon, Z) \lesssim \Phi_*^{-2}\big(\log(1/\pi_{\min}) + \log(1/\varepsilon)\big),$$

where $\Phi_* = \inf\{\Phi_S \mid \pi_Z(S) \le \frac{1}{2}\}$ and $\Phi_S = \sum_{x \in A, y \in B} \pi_Z(x) p_Z(x, y)/\pi_Z(S)$; for a proof of this, see [15]. For any set $S \subseteq \mathcal{G}$, we have

$$\Phi_S = \frac{1}{2C_* \log n} \cdot \frac{|\partial S|}{|S|} \ge \frac{1}{2C_* \log n} \cdot \frac{|\partial S|}{d(S)} = \frac{1}{2C_* \log n} \cdot \Phi'_S,$$

where the prime (′) denotes that we are considering the corresponding quantity for the nearest-neighbour discrete-time random walk. But we know that $\Phi'_* \gtrsim (\log n)^{-2}$ since the graph is good, and hence $\Phi_* \gtrsim (\log n)^{-3}$. Hence

$$t_{\mathrm{unif}}\left(\frac{1}{8}, Z\right) \lesssim (\log n)^7, \quad \text{and hence} \quad t_{\mathrm{unif}}\left(\frac{1}{8}, \tilde{X}\right) \lesssim n(\log n)^6. \qquad \square$$

We now use this mixing of the static walk along with our static-dynamic coupling to determine where the dynamic walk is at the update times of $\mathcal{P}$.

Lemma 3.7. *There exists a constant C, so that, for all M, all n sufficiently large, we have*

$$\max_{\eta_0, x_0 \in \mathcal{G}_0} \mathbb{P}_{x_0, \eta_0}\big(\tau'_{\mathcal{G}} = U_1, \mathcal{S}[0, U_1)\big) \le C \mu S \log_{(M)} n.$$

Proof. For this whole proof, we only consider the first update time U_1; as such, we drop the 1 from the subscript, just writing U. Also, we write $\tilde{X}$ for the static walk on η_0 (as above).

For the walk to leave the giant, we need the time U to be triggered by an update to $\mathcal{P}$, ie we need $U < S$. If this is the case, then the walk leaves the giant if and only if the update was caused by the closing of one of the removal edges which, given $R_{U^-}(X_U)$, has probability $R_{U^-}(X_U)/\kappa$; write $R_U = R_{U^-}(X_U)$. Hence

$$\mathbb{P}_{x_0, \eta_0}\big(\tau'_{\mathcal{G}} = U, \mathcal{S}[0, U)\big) = \frac{1}{\kappa} \mathbb{E}_{x_0, \eta_0}\big(R_U \mathbf{1}(U < S) \mathbf{1}(\mathscr{G}[0, U))\big). \tag{3.4}$$

We now set $T = n(\log n)^7$; so $S = T \log n \gg T$. We decompose according to $\{U < T\}$ or $\{U \ge T\}$. When $U < T$ we use the trivial bound $R_U \le C_* \log n$ (which holds whenever the graph is good):

$$\mathbb{E}_{x_0, \eta_0}\big(R_U \mathbf{1}(U < T) \mathbf{1}(\mathscr{G}[0, U))\big) \le C_* \log n \cdot \mathbb{P}_{x_0, \eta_0}\big(U < T, \mathscr{G}[0, U)\big).$$

Since $\mathcal{W}_0^{k+1} \subseteq \mathcal{W}_0^k$, for all M and all n sufficiently large, we have

$$\begin{aligned}
&\mathbb{E}_{x_0,\eta_0}\big(R_U \mathbf{1}(T \le U < S)\mathbf{1}\big(\mathscr{G}[0,U)\big)\big) \\
&\quad = \mathbb{E}_{x_0,\eta_0}\big(R_U \mathbf{1}(T \le U < S)\mathbf{1}\big(X_U \in \mathcal{W}_0^M\big)\mathbf{1}\big(\mathscr{G}[0,U)\big)\big) \\
&\qquad + \sum_{k=1}^{M-1} \mathbb{E}_{x_0,\eta_0}\big(R_U \mathbf{1}(T \le U < S)\mathbf{1}\big(X_U \in \mathcal{W}_0^k \setminus \mathcal{W}_0^{k+1}\big)\mathbf{1}\big(\mathscr{G}[0,U)\big)\big).
\end{aligned} \tag{3.5}$$

When $X_U \in \mathcal{W}_0^k$, we have (by definition) $R_U \le C_* \log_{(k)} n$. Hence we have

$$\begin{aligned}
&\mathbb{E}_{x_0,\eta_0}\big(R_U \mathbf{1}(T \le U < S)\mathbf{1}\big(\mathscr{G}[0,U)\big)\big) \\
&\quad \le C_* \mathbb{P}\big(U < S, \mathscr{G}[0,U)\big) \cdot \left(\log_{(M)} n + \sum_{k=1}^{M-1} \log_{(k)} n \cdot \mathbb{P}\big(X_U \notin \mathcal{W}_0^{k+1} \mid T \le U < S, \mathscr{G}[0,U)\big)\right)).
\end{aligned}$$

What is crucial is that, on the event that the graph is good, the update *times* are independent of the evolution of the walk: since $\mathcal{P}$ always, regardless of the number of edges seen by the walker, contains precisely ρ open edges and ρn closed edges, the update rate is always $\kappa\mu$. Thus an equivalent way of realising $(\mathcal{P}_t)_{t\in[0,U)}$ is the following. Define the processes $(\mathcal{A}_r)_{r\ge 0}$ and $(\mathcal{B}_r)_{r\ge 0}$ as in (3.1), (3.2), taking $t = 0$. Then sample independently $V \sim \mathcal{E}(\kappa\mu)$. At time V with probability $q = (1-p)/(1-p+\lambda)$ choose an edge uniformly at random from $\mathcal{A}_V$ and change its state from open to closed, and with probability $1-q$ choose an edge uniformly at random from $\mathcal{B}_V$ and change its state from closed to open. Then set $\mathcal{P}_r = \mathcal{A}_r \cup \mathcal{B}_r$ for all $r \in [0, V)$. Finally, set $U = V \wedge S$.

It remains to calculate this final probability, of $X_U \notin \mathcal{W}_0^{k+1}$. We want to couple X_U with $\tilde{X}_U$, as we can then apply the (uniform) mixing result Lemma 3.6 to obtain good control over its location. However, we can only do this under certain conditions; sufficient conditions are that none of the edges of $\mathcal{P}_{U-}$ have changed *throughout the entire interval* $[0, U)$. (Note that an edge could change state *before* it is added to the set process $\mathcal{P}$.) Write $\mathcal{C}$ for this sufficient condition. Then

$$\begin{aligned}
&\mathbb{P}_{x_0,\eta_0}\big(X_U \notin \mathcal{W}_0^{k+1} \mid T \le U < S, \mathscr{G}[0,U)\big) \\
&\quad \le \mathbb{P}_{x_0,\eta_0}\big(\tilde{X}_U \notin \mathcal{W}_0^{k+1} \mid T \le U < S\big) + \mathbb{P}_{x_0,\eta_0}\big(\mathcal{C}^c \mid T \le U < S, \mathscr{G}[0,U)\big),
\end{aligned}$$

since conditioning on $\mathscr{G}[0, U)$ has no effect on the static walk.

Since $T \gg n(\log n)^6$, which is the uniform mixing time of the static walk on a good giant (Lemma 3.6), if $U \ge T$ then $\tilde{X}_U$ has (uniformly) mixed and so, since the invariant distribution of the static walk is uniform (on the giant), for all $k \le M$, we have

$$\mathbb{P}_{x_0,\eta_0}\big(\tilde{X}_U \notin \mathcal{W}_0^{k+1} \mid T \le U < S\big) \le \frac{3}{2}\big|\mathcal{G}_0 \setminus \mathcal{W}_0^{k+1}\big|/|\mathcal{G}_0| \le \frac{3}{2}(\log_{(k)} n)^{-4},$$

with the final inequality holding by definition of a good graph; here we have used crucially that the path $(\tilde{X}_t)_{t\ge 0}$ is independent of U. Also, since the update rate of $\mathcal{P}$ is always $\kappa\mu \asymp \mu(\log n)^{11}$ and we run for time $U \le S = n(\log n)^8$, we find that

$$\mathbb{P}_{x_0,\eta_0}\big(\mathcal{C}^c \mid T \le U < S, \mathscr{G}[0,U)\big) \le \frac{\mathbb{P}(\mathcal{E}(\kappa\mu) < S)}{\mathbb{P}_{x_0,\eta_0}(T \le U < S, \mathscr{G}[0,U))} \le \kappa S\mu \cdot \big(1 + o(1)\big) \ll (\log_{(k)} n)^{-4},$$

by the assumption that $\mu \le (\log n)^{-20}/n$ and the fact that $\eta_0 \in \mathscr{H}$. Together, these give

$$\mathbb{P}\big(X_U \notin \mathcal{W}_0^{k+1} \mid T \le U < S, \mathscr{G}[0,U)\big) \le 2(\log_{(k)} n)^{-4}. \tag{3.6}$$

Also, for any $s \ge 0$, we have

$$\mathbb{P}_{x_0,\eta_0}\big(U < s, \mathscr{G}[0,U)\big) \le \mathbb{P}\big(\mathcal{E}(\kappa\mu) \le s\big) = 1 - e^{-\kappa\mu s} \le \kappa\mu s.$$

Hence combining these inequalities, for all M and all n for sufficiently large, we have

$$\mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} = U_1, \mathcal{S}[0,U_1)\big) \le 2C_*\mu S \log_{(M)} n. \qquad \square$$

Let K be the (random) index given by $U_K \le t < U_{K+1}$. Note that

$$\mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} \le t, K \le k-1, \mathcal{S}[0,U_k]\big) \le \mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} \le U_k, \mathcal{S}[0,U_k]\big),$$

by monotonicity of $t \mapsto \{\tau'_{\mathcal{G}} \le t\}$, along with the fact that $t < U_{K+1}$. Hence

$$\mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} \le t\big) \le \mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} \le U_k, \mathcal{S}[0,U_k]\big) + \mathbb{P}_{x_0,\eta_0}\big(K \ge k, \mathcal{S}[0,U_k]\big) + \mathbb{P}_{x_0,\eta_0}\big(\mathcal{S}[0,U_k]^c, \mathscr{G}[0,U_k]\big) + \mathbb{P}_{x_0,\eta_0}\big(\mathscr{G}[0,U_k]^c\big).$$

We have already dealt with the first term in the previous lemmas; we now just need to show that the three 'remainder' terms are sufficiently small. We do this now.

Lemma 3.8. *For all n sufficiently large, all t with $\mu t \ge (\log n)^{-5}$ and all (x_0,η_0) with $x_0 \in \mathcal{G}_0$, for $k = \lceil 5t/S \rceil$, we have*

$$\mathbb{P}_{x_0,\eta_0}\big(K \ge k, \mathcal{S}[0,U_k]\big) \le n^{-5}.$$

Proof. If we were to not have the thresholding by S, then we would have $K \sim \mathrm{Po}(\kappa\mu t)$. However, we do have the thresholding. Set $\tilde{U}_0 = 0$, and inductively define $\tilde{U}_j$, for $j = 1, 2, \dots$, by

$$\tilde{U}_j - \tilde{U}_{j-1} = S \cdot \mathbf{1}(U_j - U_{j-1} = S), \quad \text{and set} \quad \tilde{K} = \inf\{k \ge 0 \mid \tilde{U}_k \le t < \tilde{U}_{k+1}\}.$$

We have $\tilde{U}_j \le U_j$ for all $j \ge 0$, and thus $\tilde{K} \ge K$.

Recall that when the set definitions succeed, $\{U_j - U_{j-1}\}_{j\ge1}$ is a collection of iid random variables, and are independent of the starting point (x_0,η_0). Recalling from Definition 3.4 that $S = n(\log n)^8$ and from (3.3) that $\kappa \asymp (\log n)^{11}$, note that

$$\mathbb{P}\big(\mathcal{E}(\kappa\mu) \ge S\big) = e^{-\kappa\mu S} = 1 - o(1) \ge \frac{1}{2},$$

by the assumption $\mu n \ll (\log n)^{-19}$. Also let us write $k' = \lceil t/S \rceil$; since $\mu t \ge (\log n)^{-5}$ and $\mu n \ll (\log n)^{-14}$, we have $t/S \gg \log n$, and so $k' \gg 1$ and $k' \le 2t/S$. Then, on the event that the set definitions succeed, we have $\tilde{K} \preccurlyeq \mathrm{Po}(4t/S)$, since $\mathbb{P}(\mathcal{E}(\kappa\mu) \ge S) \ge \frac{1}{2}$. Hence

$$\mathbb{P}_{x_0,\eta_0}\big(K \ge k, \mathcal{S}[0,U_k]\big) \le \mathbb{P}\big(\mathrm{Po}(4t/S) \ge 5t/S\big) \le \exp\left(-\frac{1}{10}t/S\right),$$

by Poisson concentration. Since $t/S \gg \log n$, we deduce our lemma. □

Lemma 3.9. *For all n sufficiently large, all k and all (x_0,η_0) with $x_0 \in \mathcal{G}_0$, we have*

$$\mathbb{P}_{x_0,\eta_0}\big(\mathcal{S}[0,U_k]^c, \mathscr{G}[0,U_k]\big) \le k \cdot \exp\left(-\frac{1}{3}C_*(\log n)^9\right).$$

Proof. By the union bound, we have

$$\mathbb{P}_{x_0,\eta_0}\big(\mathcal{S}[0,U_k]^c, \mathscr{G}[0,U_k]\big) \le k \cdot \max_{x_0,\eta_0} \mathbb{P}_{x_0,\eta_0}\big(\mathcal{S}[0,U_1]^c, \mathscr{G}[0,U_1]\big).$$

Since we work on the event that the graph is good, we have at most $C_* \log n$ removal edges for each vertex; we also have that there are $\Omega(n)$ open edges and $\Omega(n^2)$ closed edges. Hence the only part that can 'go wrong' in the definitions is if the number of open or closed edges seen since the last update is too high. However, the maximum degree is at most $C_* \log n$ and we walk for a time at most $S = n(\log n)^8$, so Poisson concentration will tell us that we do not see too many.

Consider a static walk $\tilde{X}$ on a good graph η_0, starting from $x_0 \in \mathcal{G}_0$ and run for a time $S = n(\log n)^8$. Let α be the number of open edges seen in this time, and β the number of closed. Write N for the number of steps taken; by Poisson thinning, we have $N \preccurlyeq \mathrm{Po}(C_* S \log n/n)$, and $S \log n/n = (\log n)^9$. On the event $N \le 2C_*(\log n)^9$, we have $\alpha \le 2C_*^2(\log n)^{10} \ll \rho$ and $\beta \le 2C_* n(\log n)^9 \le \rho n$, as required for the set definitions to succeed. Hence

$$\mathbb{P}_{x_0,\eta_0}\big(\mathcal{S}[0,U_1]^c, \mathscr{G}[0,U_1]\big) \le \mathbb{P}\big(\mathrm{Po}\big(C_*(\log n)^9\big) > 2C_*(\log n)^9\big) \le \exp\left(-\frac{1}{3}C_*(\log n)^9\right),$$

by Poisson concentration. The result now follows from the union bound given above. □

Corollary 3.10. *For all n sufficiently large, all t with $\mu t \ge (\log n)^{-5}$ and all (x_0, η_0) with $x_0 \in \mathcal{G}_0$, for $\mu \ge n^{-8}$ and $k = \lceil 5t/S \rceil$, we have*

$$\mathbb{P}_{x_0,\eta_0}\big(\mathcal{S}[0, U_k)^c, \mathscr{G}[0, U_k)\big) \le \mu t \cdot n^{-5}.$$

Proof. As in Lemma 3.8 for $\lceil t/S \rceil$, we have $k \le 6t/S$. Hence

$$k \cdot \exp\left(-\frac{1}{3} C_* (\log n)^9\right) \le 6S^{-1} \cdot \mu t \cdot n^8 \exp\left(-\frac{1}{3} C_* (\log n)^9\right). \le \mu t \cdot n^{-5}. \qquad \square$$

We now have all the ingredients to prove Proposition 3.2 for the case $\mu \ge n^{-8}$.

Proof of Proposition 3.2 (*when* $\mu \ge n^{-8}$). Fix M. Combining the above results, we have, for $k = \lceil 5t/S \rceil$, recalling that $U_k \le kS \le 6t$ for the times t we are considering, that

$$\begin{aligned}\mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} \le t\big) &\le \mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} \le U_k, \mathscr{G}[0, U_k)\big) + \mathbb{P}_{x_0,\eta_0}\big(K \ge \lceil 5t/S \rceil, \mathcal{S}[0, U_k)\big) \\ &\quad + \mathbb{P}_{x_0,\eta_0}\big(\mathcal{S}[0, U_k)^c, \mathscr{G}[0, U_k)\big) + \mathbb{P}_{x_0,\eta_0}\big(\mathscr{G}[0, 6t]^c\big) \\ &\le C\mu t \log_{(M)} n + n^{-5} + \mu t n^{-5} + n^{-1} \le 2C\mu t \log_{(M)} n \end{aligned} \tag{3.7}$$

since $\eta_0 \in \mathscr{H}$, $\mu \ge n^{-8}$ and $(\log n)^{-5} \le \mu t \le \frac{1}{10}$. $\square$

It remains to prove the proposition in the case $\mu \le n^{-8}$. In this case, 'almost always' the static walk mixes on the entire giant before any of the graph even refreshes; this will make this proof easier. The general idea will be very similar, particularly to Lemma 3.7.

Proof of Proposition 3.2 (*when* $\mu \le n^{-8}$). Fix M. For this proof, let $U_1, U_2, \dots$ be the refresh times of the graph; let $U_0 = 0$. Note then that $U_j - U_{j-1} \sim^{\text{iid}} \mathcal{E}(\mu N)$ where $N = \binom{n}{2} \le n^2$.

We are now interested in the probability that $\tau'_{\mathcal{G}} = U_1$; as previously, drop the subscript 1. We have $U \sim \mathcal{E}(\mu N)$, independent of X.

Suppose $\eta_0 \in \mathscr{G}$ and $x_0 \in \mathcal{G}_0$. Then, similarly to in (3.4), we have

$$\mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} = U, \mathscr{G}[0, U)\big) = \frac{1}{N} \mathbb{E}_{x_0,\eta_0}\big(R_U \mathbf{1}\big(U \ge n^2\big)\mathbf{1}\big(\mathscr{G}[0, U)\big)\big) + \mathbb{P}_{x_0,\eta_0}\big(U \le n^{-2}\big).$$

We know that $\mathbb{P}(U \le n^2) = \mathbb{P}(\mathcal{E}(\mu N) \le n^2) \le \mu n^2 N \le n^{-4}$. Similarly to in (3.5), we have

$$\begin{aligned}&\mathbb{E}_{x_0,\eta_0}\big(R_U \mathbf{1}\big(U \ge n^2\big)\mathbf{1}\big(\mathscr{G}[0, U)\big)\big) \\ &\quad = \mathbb{E}_{x_0,\eta_0}\big(R_U \mathbf{1}\big(U \ge n^2\big)\mathbf{1}\big(X_U \in \mathcal{W}_0^M\big)\mathbf{1}\big(\mathscr{G}[0, U)\big)\big) \\ &\qquad + \sum_{k=1}^{M-1} \mathbb{E}_{x_0,\eta_0}\big(R_U \mathbf{1}\big(U \ge n^2\big)\mathbf{1}\big(X_U \in \mathcal{W}_0^k \setminus \mathcal{W}_0^{k+1}\big)\mathbf{1}\big(\mathscr{G}[0, U)\big)\big).\end{aligned}$$

When $X_U \in \mathcal{W}_0^k$, we have (by definition) $R_U \le C_* \log_{(k)} n$. Hence we have

$$\mathbb{E}_{x_0,\eta_0}\big(R_U \mathbf{1}\big(U \ge n^2\big)\mathbf{1}\big(\mathscr{G}[0, U)\big)\big) \le \log_{(M)} n + \sum_{k=1}^{M-1} \log_{(k)} n \cdot \mathbb{P}\big(X_U \notin \mathcal{W}_0^{k+1} \mid U \ge n^2, \mathscr{G}[0, U)\big).$$

Using our static-dynamic coupling, we may couple $X_t = \tilde{X}_t$ for all $t \le U = U_1$, where $\tilde{X}$ is the static walk, since the U_j are the refresh times of the entire graph. Hence, as in (3.6), but without needing to consider the condition $\mathcal{C}$, we use the (uniform) mixing of the static walk to obtain

$$\mathbb{P}_{x_0,\eta_0}\big(X_U \notin \mathcal{W}_0^{k+1} \mid U \ge n^2, \mathscr{G}[0, U)\big) \le 2\big|\mathcal{G}_0 \setminus \mathcal{W}_0^{k+1}\big|/|\mathcal{G}_0| \le 2(\log_{(k)} n)^{-4}.$$

Hence combining these inequalities, for all M and all n sufficiently large, we have

$$\max_{x_0,\eta_0} \mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} = U_1, \mathscr{G}[0, U_1)\big) \le 2C_* \frac{1}{N} \log_{(M)} n.$$

As in Lemma 3.5 (except replacing $\mathscr{G}$ by $\mathcal{S}$), for all $k \in \mathbb{N}$, we have

$$\mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} \le U_k, \mathscr{G}[0, U_k)\big) \le k \max_{x_0,\eta_0} \mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} = U_1, \mathscr{G}[0, U_1)\big) \le 2C_* k \frac{1}{N} \log_{(M)} n.$$

Observe that $U_k \sim \Gamma(k, \mu N)$; let K be the (random) index given by $U_K \le t < U_{K+1}$, and observe then that $K \sim \mathrm{Po}(\mu t N)$. Set $k = \lceil 2\mu t N \rceil$. By the same arguments as used in (3.7) (except without the 'set-definitions' term) we have

$$\begin{aligned}\mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} \le t\big) &\le \mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} \le U_k, \mathscr{G}[0, U_k)\big) + \mathbb{P}(K \ge 2\mu t N) + \mathbb{P}\big(U_k \ge 3k/(\mu N)\big) + \mathbb{P}_{\eta_0}\big(\mathscr{G}[0, 7t]^c\big) \\ &\le 2C_* \mu t \log_{(M)} n + \exp\left(-\frac{1}{3}\mu t N\right) + \exp(-2\mu t N) + \mathbb{P}_{x_0,\eta_0}\big(\mathscr{G}[0, 7t]^c\big).\end{aligned}$$

We now recall that we restricted consideration of t to satisfy $(\log n)^{-5} \le \mu t \le \frac{1}{10}$; note then that $7t \le 1/\mu$. We also consider only $\eta_0 \in \mathscr{H}$; as such, the graph remainder term in the final line above is at most $1/n$. Since $N \asymp n^2$, we see that the first term dominates, leaving us with

$$\mathbb{P}_{x_0,\eta_0}\big(\tau'_{\mathcal{G}} \le t\big) \le C \mu t \log_{(M)} n \quad \text{for a constant } C. \qquad \square$$

4. Isolation times

In this section we prove two main results on isolation times. They will involve, respectively, a single random walker on a dynamical environment and two independent random walkers on the same dynamical environment. When considering just one walk, we write τ_{isol} for the isolation time; eg for a dynamical percolation system (Z, ζ) we write

$$\tau^Z_{\mathrm{isol}} = \inf\big\{t \ge 0 \mid d^\zeta_t(Z_t) = 0\big\}.$$

When the context is clear, we omit the superscript, just writing τ_{isol}; similarly, when the context is clear we write d for the degree, rather than d^ζ. When we consider two walks, X and Y, on the same system, η, we use superscript X or Y to indicate which walk we are referring to: define

$$\tau^X_{\mathrm{isol}} = \inf\big\{t \ge 0 \mid d_t(X_t) = 0\big\} \quad \text{and} \quad \tau^Y_{\mathrm{isol}} = \inf\big\{t \ge 0 \mid d_t(Y_t) = 0\big\}.$$

Recall from Definition 2.4 that the event $\mathscr{H}[0, t]$ guarantees that the graph is good up until time t.

Theorem 4.1 (Single-Walker Isolation Time). *For all $M \in \mathbb{N}$, all n sufficiently large and all pairs (x_0, η_0), we have*

$$\mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} > t, \mathscr{H}[0, t]\big) \le 2 \exp(-\mu t / \log_{(M)} n).$$

Moreover, if $\mu t \ge 3$, then we may remove the pre-factor of 2.

Once we have proved this, we shall be able to use a type of concentration result to prove a bound on the isolation time of two independent random walkers on the same environment. When we are considering this, we write $\mathbb{P}_{x_0,y_0,\eta_0}$ for the measure. For two walks X and Y on the same (dynamical) environment η, write

$$\tau = \inf\big\{t \ge 0 \mid d_t(X_t) = 0 = d_t(Y_t)\big\}.$$

Theorem 4.2 (Dual-Walker Isolation Time). *For all $M \in \mathbb{N}$, all n sufficiently large and all triples (x_0, y_0, η_0), we have*

$$\mathbb{P}_{x_0,y_0,\eta_0}\big(\tau > t, \mathscr{H}[0, t]\big) \le 2 \exp(-\mu t / \log_{(M)} n).$$

4.1. *Single-walker isolation time*

In this section we prove Theorem 4.1 on the isolation time of a walk X on a dynamical environment η. In order to find the isolation time, we wait until the walk joins the giant and then look at becoming isolated from there. It is easier to consider the giant, rather than subcritical components, because we are able to use concentration results on the structure of the giant.

We first state the proposition on isolation from the giant, and then show how to conclude Theorem 4.1 from it; we then prove the proposition to finish. Throughout, M is a positive integer.

Proposition 4.3 (Isolation from the Giant). *There exists a positive constant c so that, for all M, all n sufficiently large and all (x_0, η_0) with $\eta_0 \in \mathscr{H}$ and $x_0 \in \mathcal{G}_0$, we have*

$$\mathbb{P}_{x_0,\eta_0}\left(\tau_{\mathrm{isol}} \le \frac{1}{\mu \log_{(M)} n}\right) \ge c \cdot \frac{1}{\log_{(M)} n}.$$

Proof of Theorem 4.1. Observe that this trivially holds (for all n large enough) if $\mu t \le 3$. By monotonicity, replacing M by $M-1$, it suffices to prove an upper bound of $\exp -c\mu t/\log_{(M)} n$ for a positive constant c when $\mu t \ge 3$.

Fix M. For this proof, rescale time so that $\mu = 1$. We prove this theorem by performing independent experiments. Note that if $x_0 \in \mathcal{G}_0$ then $\tau_{\mathcal{G}} = 0$, and otherwise we apply Lemma 3.1. By direct calculation, we have

$$\begin{aligned}
&\mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} \le 2, \mathscr{H}[0,1]\big)\\
&\quad\ge \mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} \le 2, \tau_{\mathcal{G}} \le 1, \mathscr{H}[0,1]\big)\\
&\quad\ge \mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} - \tau_{\mathcal{G}} \le 1/\log_{(M)} n, \tau_{\mathcal{G}} \le 1, \mathscr{H}[0,1]\big)\\
&\quad= \sum_{x_0',\eta_0'} \mathbb{P}_{x_0',\eta_0'}(\tau_{\mathrm{isol}} \le 1/\log_{(M)} n) \cdot \mathbb{P}_{x_0,\eta_0}\big(X(\tau_{\mathcal{G}}) = x_0', \eta(\tau_{\mathcal{G}}) = \eta_0', \tau_{\mathcal{G}} \le 1, \mathscr{H}[0,1]\big)\\
&\quad\ge c(\log_{(M)} n)^{-1} \cdot \big(\mathbb{P}_{x_0,\eta_0}(\tau_{\mathcal{G}} \le 1) - \mathbb{P}_{x_0,\eta_0}\big(\mathscr{H}[0,1]^c\big)\big),
\end{aligned}$$

for a positive constant c, where for the final inequality we used that on the event $\mathscr{H}[0,1]$ we have $\eta_0' \in \mathscr{H}$, and hence we may apply Proposition 4.3. Now applying Lemma 3.1, we obtain

$$\mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} \le 2, \mathscr{H}[0,1]\big) \ge \frac{1}{2} c c_1/\log_{(M)} n - \mathbb{P}_{x_0,\eta_0}\big(\mathscr{H}[0,1]^c\big),$$

with the positive constant c_1 coming from Lemma 3.1, noting that $\mathbb{P}_{x_0,\eta_0}(\mathscr{G}[0,1]^c) = o(1)$ since $\eta_0 \in \mathscr{H}$. Rearranging this, we obtain, for a positive constant c, that

$$\mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} > 2, \mathscr{H}[0,2]\big) \le \mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} > 2, \mathscr{H}[0,1]\big) \le \exp(-c/\log_{(M)} n).$$

Hence, for any $k \in \mathbb{N}$, applying the strong Markov property ($k-1$ times), we obtain

$$\mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} > 2k, \mathscr{H}[0,2k]\big) \le \max_{x_0',\eta_0'} \mathbb{P}_{x_0',\eta_0'}\big(\tau_{\mathrm{isol}} > 2, \mathscr{H}[0,2]\big)^k \le \exp(-ck/\log_{(M)} n).$$

This completes the proof. □

It remains to prove Proposition 4.3. We do this via a sequence of lemmas, using the following rough methodology. Observe that the *only* way for the walk to become isolated is to be at a degree 1 vertex and for the one open edge to close before any closed incident edges open or the walker leaves the vertex. This motivates looking at the rate at which the walk hits degree 1 vertices.

Since the walk is on a dynamically evolving graph, even though when we require the graph to be good this includes that the giant has a lot of degree 1 vertices, the location of these degree 1 vertices is changing. This makes using averaging properties (like a law of large numbers) difficult. However, since we take steps at rate at least $1/n$ (when non-isolated) and $\mu n \ll 1/\log n$ (the order of the maximum degree), we see that the vast majority of the time the walker takes a step before any edge incident to its location changes state. This motivates looking at the rate at which a walk (with the same walk-dynamics) hits degree 1 vertices *on a static (good) graph*, and then relating this quantity to the relevant quantity for the walk on the dynamic graph. In Section 3 we referred to this as the *static walk* and the original as the *dynamic walk*, denoting them by $\tilde{X}$ and X, respectively.

With this motivation in mind, we first collect some results regarding a walk with our dynamics on a static graph. To do this, we use a Chernoff-style bound on the number of visits to a set, which is due to Gillman [12]. It applies to *discrete-time* random walks. We do not apply it to a discretisation of our continuous chain, but to the jump chain of the walk (on a static graph). We state it in a general form; an even more general form is given in [12, Theorem 2.1].

Theorem 4.4 (Gillman [12]). *Consider the discrete-time random walk on a weighted, connected graph $G = (V, E)$ with any initial distribution. Let π be the unique invariant distribution. Let $A \subseteq V$, and let N_m be the number of visits to A in* *m steps. Write γ for the spectral gap. Then*

$$\mathbb{P}\big(\big|N_m - m\pi(A)\big| \ge R\big) \le 3\pi_{\min}^{-1/2} \exp\biggl(-\frac{1}{20}\gamma R^2/m\biggr) \quad \textit{for any } \varepsilon \in [0, m].$$

We now apply this to a walk on a good (static) giant.

Lemma 4.5. *Consider the discrete-time nearest-neighbour simple random walk on a graph G, and write N_m for the number of visits to the set of degree 1 vertices in m steps. There exists a positive constant c so that, for all n sufficiently large and all $m \ge (\log n)^6$, if the graph is good, ie $G \in \mathscr{G}$, and the walk starts from its giant, then we have*

$$\mathbb{P}(N_m \le cm) \le n^{-1}.$$

Proof. Note that the invariant measure of this walk, which we denote π', is given by $\pi_i = d_i/d_{\mathcal{G}}$, where $d_{\mathcal{G}} = \sum_{i \in \mathcal{G}} d_i$. Since $d_i \ge 1$ for all $i \in \mathcal{G}$, we have $\pi'_{\min} \ge 1/d_{\mathcal{G}}$. Now, trivially we have that $d_{\mathcal{G}} \le d_G$, where $d_G = \sum_{i \in G} d_i$, and $d_G \le 2C_* n$ by Definition 2.2(iii). Hence $1/\pi'_{\min} \le d_{\mathcal{G}} \le 2C_* n$.

Let $A = \{x \in \mathcal{G} \mid d(x) = 1\}$ be the set of degree 1 vertices in the giant. Definition 2.2(iv) tells us that $|A| \ge c_* n$. Thus, since $d(x) \ge 1$ for all $x \in \mathcal{G}$ and $d_{\mathcal{G}} \le 2C_* n$, we have that $\pi'(A) \ge c_*/(2C_*)$; let $c = c_*/(4C_*)$ so that $\pi'(A) \ge 2c$.

Recall from Definition 2.2(vii) that the spectral gap γ of a good giant satisfies $\gamma \ge c_*(\log n)^{-4}$. We now take $R = \frac{1}{2}\pi'(A)m \le m$ in Theorem 4.4 to obtain

$$\mathbb{P}\biggl(\big|N_m - m\pi'(A)\big| \ge \frac{1}{2}\pi'(A)m\biggr) \le 3\sqrt{2c_*} \cdot n^{-1/2} \exp\biggl(-\frac{1}{20}c_*(\log n)^{-4} \cdot \frac{1}{4}\pi'(A)^2 m\biggr).$$

Since $\pi'(A) \ge 2c$, taking $m \ge (\log n)^6$ gives super-polynomial decay, completing the proof. □

We now make rigorous the motivation given at the start of this section in the following lemma.

Lemma 4.6. *There exists a positive constant q so that, for all n sufficiently large and all (x_0, η_0) with $\eta_0 \in \mathscr{G}$ and $x_0 \in \mathcal{G}_0$, for $s = n(\log n)^6$, we have*

$$\mathbb{P}_{x_0,\eta_0}(\tau_{\text{isol}} \le s) \ge q\mu s.$$

Remark 1. Observe that, as a best-case scenario, if the walker were always at a degree 1 vertex until it becomes isolated, then the isolation time would simply be the time it takes for that one edge to close, which is $\mathcal{E}(\mu(1-p))$. Thus, for any (x_0, η_0) with $d_0(x_0) \ne 0$, we have

$$\mathbb{P}_{x_0,\eta_0}(\tau_{\text{isol}} \le s) \le \mu s/(1-p).$$

Hence, for this s, such a result as Lemma 4.6 is best-possible up to constants.

Proof of Lemma 4.6. Fix a pair (x_0, η_0) with $\eta_0 \in \mathscr{G}$ and $x_0 \in \mathcal{G}_0$; for this proof, drop it from the notation, writing $\mathbb{P}(\cdot)$ in place of $\mathbb{P}_{x_0,\eta_0}(\cdot)$.

Lemma 4.5 tells us the rate at which the static walk hits degree 1 vertices (with high probability). In order to transfer this result to our dynamic walk, we define a coupling between the two walks; this was given in Section 3, but we recall it precisely here. Write $\tilde{X}$ for the static walk, walking on the static graph η_0. Set $X_0 = \tilde{X}_0 = x_0$. Give X and $\tilde{X}$ the same jump clock. When the clock rings, at time t say, both walks choose the same vertex; $\tilde{X}$ performs the jump if and only if the edge is present in η_0, while X performs the jump if and only if the edge is present in η_t. We call this the *static-dynamic coupling*.

We now define $\tilde{T}_i$ to be the i-th time that (the static walk) $\tilde{X}$ hits a degree 1 vertex: set $\tilde{T}_0 = \tilde{T}'_0 = 0$ and define inductively, for $i \ge 1$,

$$\tilde{T}_i = \inf\bigl\{t \ge \tilde{T}'_{i-1} | d_t(\tilde{X}_t) = 1\bigr\} \quad \text{and} \quad \tilde{T}'_i = \inf\bigl\{t \ge \tilde{T}_i | \tilde{X}(t) \ne \tilde{X}(\tilde{T}_i)\bigr\}.$$

Since the jump rate of $\tilde{X}$ is always at least $1/n$, by standard Poisson concentration it takes at least $s/(2n)$ steps in time $s - n$ with probability $1 - o(1)$. Along with Lemma 4.5 this says that

$$\mathbb{P}(\tilde{T}_k \le s - n) = 1 - o(1) \quad \text{for } k = s/(8n). \tag{4.1}$$

For $i \geq 1$ define the event that all the (open or closed) edges incident to a vertex that the static walk visited remain in the same state between visits to degree 1 vertices:

$$\mathscr{E}_i = \{\text{neighbourhood of path of static walk did not change in } [\tilde{T}_{i-1}, \tilde{T}_i]\}.$$

Similarly, for $u \geq v \geq 0$ define

$$\mathscr{E}_{u,v} = \{\text{neighbourhood of path of static walk did not change in } [u, v]\}.$$

Note that, by definition, on the event $\{\tilde{T}_k \leq s\}$ we have $\mathscr{E}_1 \cap \cdots \cap \mathscr{E}_k \supseteq \mathscr{E}_{0,s}$.

Write N for the number of steps taken by $\tilde{X}$ in time s. Since $\eta_0 \in \mathscr{G}$, the maximum degree is at most $C_* \log n$. Hence, by Poisson thinning, $N \preccurlyeq L \sim \text{Po}(C_* s \log n/n)$. Let α and β be the total number of open and closed edges, respectively, that are adjacent to the path of the (static) walk by time s. When $N \leq 2C_* s \log n/n$, we have $\alpha \leq 2C_*^2 s (\log n)^2/n$ and $\beta \leq 2C_* s \log n$. Hence

$$\begin{aligned}\mathbb{P}(\mathscr{E}_{0,s}^c) &\leq \mathbb{P}\big(\mathcal{E}\big(\alpha(1-p)\mu + \beta p \mu\big) \leq s, N \leq 2C_* s \log n/n\big) + \mathbb{P}\big(L > 2C_* s \log n/n\big) \\ &\leq Cs^2(\log n)^2 \mu/n + \exp -cs \log n/n,\end{aligned}$$

for positive constants c and C, by Poisson concentration. Hence, since $\mu n \ll (\log n)^{-14}$, we have

$$\mathbb{P}(\mathscr{E}_{0,s}) = 1 - o(1) \text{when} s = n(\log n)^6. \tag{4.2}$$

Note that X can only become isolated when it is at a degree 1 vertex immediately prior. We use our static-dynamic coupling to lower bound:

$$\begin{aligned}\mathbb{P}(\tau_{\text{isol}} \leq s) &\geq \mathbb{P}\left(\bigcup_{i=1}^{k}\left\{\tau_{\text{isol}} \in \big[\tilde{T}_i, \tilde{T}_i'\big), \tilde{T}_i \leq s - n, \tau_{\text{isol}} - \tilde{T}_i \leq n, \bigcap_{j \leq i} \mathscr{E}_j\right\}\right) \\ &= \sum_{i=1}^{k} \mathbb{P}\left(\tau_{\text{isol}} \in \big[\tilde{T}_i, \tilde{T}_i'\big), \tau_{\text{isol}} - \tilde{T}_i \leq n | \tilde{T}_i \leq s - n, \bigcap_{j \leq i} \mathscr{E}_j\right) \cdot \mathbb{P}\left(\tilde{T}_i \leq s - n, \bigcap_{j \leq i} \mathscr{E}_j\right).\end{aligned} \tag{4.3}$$

Using the static-dynamic coupling on the event $\mathscr{E}_1 \cap \cdots \cap \mathscr{E}_i$, we see that if the unique open edge adjacent to X at time $\tilde{T}_i$ closes before anything else opens or X jumps, then X becomes isolated during $[\tilde{T}_i, \tilde{T}_i')$. Writing $\mathcal{E}_1$, $\mathcal{E}_2$ and $\mathcal{E}_3$ for independent exponential random variables, we have

$$\begin{aligned}&\mathbb{P}\left(\tau_{\text{isol}} \in \big[\tilde{T}_i, \tilde{T}_i'\big), \tau_{\text{isol}} - \tilde{T}_i \leq n | \tilde{T}_i \leq s - n, \bigcap_{j \leq j} \mathscr{E}_j\right) \\ &\quad \geq \mathbb{P}\big(\mathcal{E}_1\big((1-p)\mu\big) < \min\big\{\mathcal{E}_2\big(p(n-1)\mu\big), \mathcal{E}_3(1/n)\big\}, \mathcal{E}_1\big((1-p)\mu\big) \leq n\big) \asymp \mu n,\end{aligned}$$

since $\mu n \ll 1$, by comparing rates. Using this in (4.3) along with (4.1) and (4.2) we obtain

$$\begin{aligned}\mathbb{P}(\tau_{\text{isol}} \leq s) &\gtrsim \mu n k \mathbb{P}(\tilde{T}_k \leq s - n, \mathscr{E}_1 \cap \cdots \cap \mathscr{E}_k) \\ &\geq \mu n k\big(1 - \mathbb{P}\big(\mathscr{E}_{0,s}^c\big) - \mathbb{P}(\tilde{T}_k > s - n)\big) \asymp \mu n k.\end{aligned}$$

Since $k = s/(8n)$, this concludes the proof. □

We now use this to prove our isolation result Proposition 4.3.

Proof of Proposition 4.3. In this proof, we use the following shorthand:

$$\mathbb{P}_{\mathcal{G}}(\cdot) = \min_{\eta_0 \in \mathscr{G}, x_0 \in \mathcal{G}_0} \mathbb{P}_{x_0,\eta_0}(\cdot) \quad \text{and} \quad \mathbb{P}^{\mathcal{G}}(\cdot) = \max_{\eta_0 \in \mathscr{G}, x_0 \in \mathcal{G}_0} \mathbb{P}_{x_0,\eta_0}(\cdot).$$

Consider an initial pair (x_0, η_0) with $\eta_0 \in \mathscr{H}$ and $x_0 \in \mathcal{G}_0$. For any $s \in \mathbb{R}$ and $r \in \mathbb{N}$, using the Markov property we have

$$\mathbb{P}_{x_0,\eta_0}(\tau_{\text{isol}} \in \big(sr, s(r+1)\big], \mathscr{G}[0, sr]\big)$$

$$
\begin{aligned}
&\geq \mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} \leq s(r+1) | \tau_{\mathrm{isol}} > sr, X_{sr} \in \mathcal{G}_{sr}, \mathscr{G}[0,sr]\big)\mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} > sr, X_{sr} \in \mathcal{G}_{sr}, \mathscr{G}[0,sr]\big) \\
&\geq \mathbb{P}_{\mathcal{G}}(\tau_{\mathrm{isol}} \leq s) \cdot \big(\mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} > sr, \mathscr{G}[0,sr]\big) - \mathbb{P}_{x_0,\eta_0}(\exists u \leq sr \text{ s.t. } X_u \notin \mathcal{G}_u)\big).
\end{aligned}
$$

Hence we have

$$
\begin{aligned}
&\mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} > s(r+1), \mathscr{G}[0,sr]\big) \\
&= \mathbb{P}_{x_0,\eta_0}\big(\mathscr{G}[0,sr]\big) - \mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} \leq sr, \mathscr{G}[0,sr]\big) - \mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} \in \big(sr, s(r+1)\big], \mathscr{G}[0,sr]\big) \\
&\leq \mathbb{P}_{x_0,\eta_0}\big(\mathscr{G}[0,sr]\big) - \mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} \leq sr, \mathscr{G}[0,sr]\big) \\
&- \mathbb{P}_{\mathcal{G}}(\tau_{\mathrm{isol}} \leq s) \cdot \big(\mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} > sr, \mathscr{G}[0,sr]\big) - \mathbb{P}_{x_0,\eta_0}(\exists u \leq sr \text{ s.t. } X_u \notin \mathcal{G}_u)\big) \\
&= \mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} > sr, \mathscr{G}[0,sr]\big) \cdot \mathbb{P}^{\mathcal{G}}(\tau_{\mathrm{isol}} > s) \\
&+ \mathbb{P}_{\mathcal{G}}(\tau_{\mathrm{isol}} \leq s) \cdot \mathbb{P}_{x_0,\eta_0}(\exists u \leq sr \text{ s.t. } X_u \notin \mathcal{G}_u).
\end{aligned}
$$

Hence, upon iterating, we obtain

$$
\begin{aligned}
&\mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} > sr, \mathscr{G}[0,sr]\big) \\
&\quad \leq \mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} > sr, \mathscr{G}\big[0, s(r-1)\big]\big) \\
&\quad \leq \mathbb{P}^{\mathcal{G}}(\tau_{\mathrm{isol}} > s)^r + r \cdot \mathbb{P}_{\mathcal{G}}(\tau_{\mathrm{isol}} \leq s)\mathbb{P}_{x_0,\eta_0}(\exists u \leq sr \text{ s.t. } X_u \notin \mathcal{G}_u).
\end{aligned} \tag{4.4}
$$

Observe that, by the memoryless property, we have

$$
\mathbb{P}_{\mathcal{G}}(\tau_{\mathrm{isol}} \leq s) \leq \mu s. \tag{4.5}
$$

We now set $s = n(\log n)^6$, $t = \gamma(\log_{(M)} n)^{-1}/\mu$ for a constant γ, to be chosen later, and $r = \lfloor t/s \rfloor$; note then that $\frac{2}{3}t \leq rs \leq t$ as $\mu n \ll (\log n)^{-5}$. Since $\eta_0 \in \mathscr{H}$, we may apply Proposition 3.2 for this t to obtain a constant C so that

$$
\mathbb{P}_{x_0,\eta_0}(\exists u \leq sr \text{ s.t. } X_u \notin \mathcal{G}_u) \leq C\mu t \log_{(M)} n. \tag{4.6}
$$

Using (4.5) and (4.6) along with Lemma 4.6 in (4.4), we find that

$$
\mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} > t, \mathscr{G}[0,t]\big) \leq (1 - q\mu s)^r + C(\mu t)^2 \log_{(M)} n \leq 1 - \frac{1}{3}q\mu t + C(\mu t)^2 \log_{(M)} n,
$$

valid for any (x_0, η_0) with $\eta_0 \in \mathscr{H}$ and $x_0 \in \mathcal{G}_0$. We then take $\gamma = q/(6C)$ and obtain

$$
\mathbb{P}_{x_0,\eta_0}\big(\tau_{\mathrm{isol}} > t, \mathscr{G}[0,t]\big) \leq 1 - \frac{q^2}{36C} \cdot \frac{1}{\log_{(M)} n}.
$$

Since we can take $C \geq 1$ and $q \leq 1$, we then have

$$
\begin{aligned}
&\mathbb{P}_{x_0,\eta_0}\bigg(\tau_{\mathrm{isol}} > \frac{1}{C\mu \log_{(M)} n}, \mathscr{G}\bigg[0, \frac{1}{C\mu \log_{(M)} n}\bigg]\bigg) \\
&\quad \leq \mathbb{P}_{x_0,\eta_0}\bigg(\tau_{\mathrm{isol}} > \frac{q}{3C\mu \log_{(M)} n}, \mathscr{G}\bigg[0, \frac{q}{3C\mu \log_{(M)} n}\bigg]\bigg).
\end{aligned}
$$

Hence there exists a positive constant c so that

$$
\mathbb{P}_{x_0,\eta_0}\bigg(\tau_{\mathrm{isol}} > \frac{1}{\mu \log_{(M)} n}, \mathscr{G}\bigg[0, \frac{1}{\mu \log_{(M)} n}\bigg]\bigg) \leq 1 - c \cdot \frac{1}{\log_{(M)} n}.
$$

Finally, $\eta_0 \in \mathscr{H}$, so $\mathbb{P}_{\eta_0}(\mathscr{G}[0, (\log_{(M)} n)^{-1}/\mu]^c) \leq n^{-1}$, and the result follows. □

Remark. Observe that, as in the remark after Lemma 4.6, for this time-scale the result of Proposition 4.3 is best-possible, up to constants.

4.2. *Dual-walker (joint) isolation time*

In this section we prove Theorem 4.2 on the joint isolation time of two walkers on a single dynamical environment. We start by introducing some more notation. Consider two walks X and Y, which start from x_0 and y_0 respectively, walking *independently* on the same environment η. Let

$$\tau = \inf\{t \geq 0 \mid d_t(X_t) = 0 = d_t(Y_t)\}.$$

Let $\tau_0^X = \tau_0^Y = \hat{\tau}_0^X = 0$, and for $k \geq 1$ define inductively

$$\tau_k^X = \inf\{t \geq \hat{\tau}_{k-1}^X \mid d_t(X_t) = 0\}, \qquad \hat{\tau}_k^X = \inf\{t \geq \tau_k^X \mid d_t(X_t) \neq 0\} \quad \text{and}$$
$$\tau_k^Y = \inf\{t \geq \tau_k^X \mid d_t(Y_t) = 0\}.$$

We prove a result on the joint-isolation time of two walks, X and Y, walking independently on the same (dynamic) environment η. For the probability measure associated to this system (X, Y, η), when it is started from (x_0, y_0, η_0), we write $\mathbb{P}_{x_0,y_0,\eta_0}$.

In order to prove the dual-walker isolation result, we first state two lemmas that we use. We prove the theorem using the lemmas, then prove the lemmas. Throughout, M is a positive integer.

Lemma 4.7. *There exists a positive constant c_1 so that, for all M, all n sufficiently large and all (x_0, y_0, η_0), we have*

$$\mathbb{P}_{x_0,y_0,\eta_0}\left(\tau_1^Y > \hat{\tau}_1^X, \mathscr{H}\left[\tau_1^X, \hat{\tau}_1^X\right]\right) \leq \exp(-c_1/\log_{(M)} n).$$

Lemma 4.8. *There exists a positive constant c_2 so that, for all M, all n sufficiently large and all (x_0, η_0), we have*

$$\mathbb{P}_{x_0,\eta_0}\left(\tau_{K+1}^X > t, \mathscr{H}[0,t]\right) \leq \exp\left(-\frac{2}{3}K\right) \quad \text{when } K = \lfloor c_2\mu t/\log_{(M)} n\rfloor.$$

Proof of Theorem 4.2. By monotonicity, replacing M by $M-1$, it suffices to find a positive constant c so that the probability is upper bounded by $2\exp -c\mu t/(\log_{(M)} n)^2$ for a positive constant c. Hence we may assume that $\mu t \geq (\log_{(M)} n)^2$, as otherwise the result trivially holds.

For any $t \geq 0$ and for $K = \lfloor c_2\mu t/\log_{(M)} n\rfloor$, using Lemma 4.8 we have

$$\mathbb{P}_{x_0,y_0,\eta_0}\left(\tau > t, \mathscr{H}[0,t]\right) \leq \mathbb{P}_{x_0,y_0,\eta_0}\left(\tau > t, \tau_{K+1}^X \leq t, \mathscr{H}[0,t]\right) + \exp\left(-\frac{2}{3}K\right). \tag{4.7}$$

Since $\tau_{K+1}^X \leq t$ implies $\hat{\tau}_K^X \leq t$, on the event $\{\tau_{K+1}^X \leq t\}$ we have $\mathscr{H}[0,t] \subseteq \mathscr{H}[0, \hat{\tau}_K^X]$. We then use the strong Markov property at time $\hat{\tau}_1^X$ to iterate:

$$\begin{aligned}
&\mathbb{P}_{x_0,y_0,\eta_0}\left(\tau > t, \tau_{K+1}^X \leq t, \mathscr{H}[0,t]\right) \\
&\quad\leq \mathbb{P}_{x_0,y_0,\eta_0}\left(\bigcap_{k=1}^{K}\{\tau_k^Y > \hat{\tau}_k^X\}, \mathscr{H}\left[0, \hat{\tau}_K^X\right]\right) \\
&\quad\leq \mathbb{P}_{x_0,y_0,\eta_0}\left(\bigcap_{k=2}^{K}\{\tau_k^Y > \hat{\tau}_k^X\}, \mathscr{H}\left[\hat{\tau}_1^X, \hat{\tau}_K^X\right] | \tau_1^Y > \hat{\tau}_1^X, \mathscr{H}\left[0, \hat{\tau}_1^X\right]\right) \\
&\qquad\times \mathbb{P}_{x_0,y_0,\eta_0}\left(\tau_1^Y > \hat{\tau}_1^X, \mathscr{H}\left[\tau_1^X, \hat{\tau}_1^X\right]\right) \\
&\quad\leq \max_{x_0',y_0',\eta_0'} \mathbb{P}_{x_0',y_0',\eta_0'}\left(\bigcap_{k=1}^{K-1}\{\tau_k^Y > \hat{\tau}_k^X\}, \mathscr{H}\left[0, \hat{\tau}_{K-1}^X\right]\right) \cdot \mathbb{P}_{x_0,y_0,\eta_0}\left(\tau_1^Y > \hat{\tau}_1^X, \mathscr{H}\left[\tau_1^X, \hat{\tau}_1^X\right]\right) \\
&\quad\leq \cdots \leq \max_{x_0',y_0',\eta_0'} \mathbb{P}_{x_0',y_0',\eta_0'}\left(\tau_1^Y > \hat{\tau}_1^X, \mathscr{H}\left[\tau_1^X, \hat{\tau}_1^X\right]\right)^K.
\end{aligned} \tag{4.8}$$

Since $\mu t \ge (\log_{(M)} n)^2$, we have $K \ge \frac{1}{2} c_2 \mu t / \log_{(M)} n$. Using (4.8) and Lemma 4.7 in (4.7), we have

$$\mathbb{P}_{x_0,y_0,\eta_0}\big(\tau > t, \mathscr{H}[0,t]\big) \le \max_{x_0',y_0',\eta_0'} \mathbb{P}_{x_0',y_0',\eta_0'}\big(\tau_1^Y > \hat{\tau}_1^X, \mathscr{H}\big[\tau_1^X, \hat{\tau}_1^X\big]\big)^K + \exp\left(-\frac{2}{3}K\right)$$

$$\le \exp(-c_1 K / \log_{(M)} n) + \exp\left(-\frac{2}{3}K\right) \le 2\exp\left(-\frac{1}{2}c_1 c_2 \mu t / (\log_{(M)} n)^2\right). \qquad \square$$

It remains to prove Lemmas 4.7 and 4.8.

Proof of Lemma 4.7. By the strong Markov property, used in the same way as above, and recalling that τ_1^Y is the first time after τ_1^X that Y becomes isolated, we have

$$\max_{x_0,y_0,\eta_0} \mathbb{P}_{x_0,y_0,\eta_0}\big(\tau_1^Y > \hat{\tau}_1^X, \mathscr{H}\big[\tau_1^X, \hat{\tau}_1^X\big]\big) \le \max_{\substack{x_0,y_0,\eta_0 \\ d_0(x_0)=0}} \mathbb{P}_{x_0,y_0,\eta_0}\big(\tau_1^Y > \hat{\tau}_1^X, \mathscr{H}\big[0, \hat{\tau}_1^X\big]\big). \tag{4.9}$$

For the moment, we emphasise that our underlying graph has n vertices: we do this by using super- and subscript n, eg $\mathbb{P}^n$ and $\mathscr{G}_n$. Recall Theorem 4.1, which says that

$$\mathbb{P}^n_{y_0,\eta_0}\big(\tau^Y_{\text{isol}} > t, \mathscr{H}[0,t]\big) \le \exp(-\mu t / \log_{(M)} n) \quad \text{when } \mu t \ge 3.$$

We wish to bound (the related quantity)

$$\mathbb{P}^n_{x_0,y_0,\eta_0}\big(\tau_1^Y > \hat{\tau}_1^X, \mathscr{H}\big[0, \hat{\tau}_1^X\big] | \hat{\tau}_1^X\big).$$

This is trivially 0 for $x_0 = y_0$; consider $x_0 \ne y_0$. To bound this, we observe that, conditional on the value of $\hat{\tau}_1^X = T$, this is conditioning the vertex x_0 to be isolated until time $\hat{\tau}_1^X = T$; the rest of the graph is unaffected.

Let η^n be a dynamical environment on n vertices, and let $x_0 \in \{1,\dots,n\}$ be a vertex. Define $\tilde{\eta}^n$ by conditioning on the event that the vertex x_0 is isolated until time T. Write $\hat{\eta}^n$ for the restriction of $\tilde{\eta}^n$ to $\{1,\dots,n\} \setminus \{x_0\}$. Observe then that $\hat{\eta}^n \sim \eta^{n-1}$ (where η^{n-1} is a dynamical environment on $n-1$ vertices), up to relabelling of vertices. (In words, this says that if a vertex is conditioned to be isolated, then the rest of the graph behaves as a dynamical environment on $n-1$ vertices.) Note also that two edges do not update at the same time, so we cannot have $\tau_1^Y = \hat{\tau}_1^X$ (since the first requires an edge to close and the second an edge to open).

Hence, $(Y_t \mid t \le \hat{\tau}_1^X)$ is a walk on the environment $\hat{\eta}^n$, which has the distribution of η^{n-1}. Note that Y may still pick the (conditioned to be isolated) vertex x_0 (with probability $1/(n-1)$), in which case it does not move; thus, under this conditioning, $(Y, \hat{\eta}^n)$ is simply a realisation of dynamical percolation on $n-1$ vertices, but with added laziness: when Y's $\mathcal{E}(1)$ clock rings, with probability $1/(n-1)$ it does nothing; with the remaining probability, it performs the usual step.

Note that we can rescale μ to get rid of the laziness of Y. Indeed, the laziness has the effect of changing the walker's clock from rate 1 to rate $1-1/(n-1)$. As such, if we replace μ by $\mu' = \mu(1-1/(n-1))$, then the ratio of the rate edge-clocks to the rate of the walker-clock is μ: we have simply slowed both down. We then speed up everything by a factor $1-1/(n-1)$. We apply previous results with μ replaced by μ'. The restrictions on μ are satisfied by μ' also, since $\mu' \le \mu$.

For all m, define $\mathscr{G}_m'$ by replacing c_* and C_* in Definition 2.2 by $\frac{1}{2}c_*$ and $2C_*$, respectively; define $\mathscr{H}'$ in terms of $\mathscr{G}'$ as in Definition 2.4. We then have that if $\eta_0^n \in \mathscr{G}_n$ and the vertex x_0 is isolated (in η_0^n), then $\tilde{\eta}_0^n = \eta_0^n - \{x_0\}$ defined by removing the vertex x_0 satisfies $\tilde{\eta}_0^n \in \mathscr{G}_{n-1}'$ (for n sufficiently large). Hence we have the following inequality: let (Z,ζ) be a full system, independent of X and Y, on $n-1$ vertices, and start it from $(Z_0,\zeta_0) = (y_0, \hat{\eta}_0^n)$; we then have

$$\mathbb{P}^n_{x_0,y_0,\eta_0^n}\big(\tau_1^Y > \hat{\tau}_1^X, \mathscr{H}_n\big[0,\hat{\tau}_1^X\big] | \hat{\tau}_1^X\big) \le \mathbb{P}^{n-1}_{y_0,\hat{\eta}_0^n}\big(\tau_1^Z > \hat{\tau}_1^X, \mathscr{H}_{n-1}'\big[0,\hat{\tau}_1^X\big] | \hat{\tau}_1^X\big).$$

Note that Theorem 4.1 still holds if we replace $\mathscr{H}$ by $\mathscr{H}'$ in its statement. Combining all the above considerations, applying Theorem 4.1, on the event $\{\mu \hat{\tau}_1^X \ge 3\}$ we have

$$\begin{aligned}
&\mathbb{P}^n_{x_0,y_0,\eta_0^n}\big(\tau_1^Y > \hat{\tau}_1^X, \mathscr{H}_n\big[0,\hat{\tau}_1^X\big] | \hat{\tau}_1^X\big) \\
&\quad \le \mathbb{P}^{n-1}_{y_0,\hat{\eta}_0^n}\big(\tau_1^Z > \hat{\tau}_1^X, \mathscr{H}_{n-1}'\big[0,\hat{\tau}_1^X\big] | \hat{\tau}_1^X\big) \\
&\quad \le \exp\left(-\mu \hat{\tau}_1^X \left(1 - \frac{1}{n}\right) / \log_{(M)}(n-1)\right) \le \exp\left(-\frac{1}{2}\mu \hat{\tau}_1^X / \log_{(M)} n\right).
\end{aligned}$$

We now calculate the unconditioned value. Fix (x_0, y_0, η_0) with $d_0(x_0) = 0$. We have

$$\begin{aligned}
&\mathbb{P}_{x_0,y_0,\eta_0}\big(\tau_1^Y > \hat{\tau}_1^X, \mathscr{H}\big[0, \hat{\tau}_1^X\big] | \hat{\tau}_1^X\big) \\
&\quad \le \mathbb{E}_{x_0,y_0,\eta_0}\big(\mathbb{P}_{x_0,y_0,\eta_0}\big(\tau_1^Y > \hat{\tau}_1^X, \mathscr{H}\big[0, \hat{\tau}_1^X\big] | \hat{\tau}_1^X\big) \cdot \mathbf{1}\big(\hat{\tau}_1^X \ge 3/\mu\big)\big) + \mathbb{P}_{x_0,\eta_0}\big(\hat{\tau}_1^X < 3/\mu\big) \\
&\quad \le \mathbb{P}_{x_0,\eta_0}\big(\hat{\tau}_1^X \ge 3/\mu\big) \cdot \exp\left(-\frac{1}{2}\mu(3/\mu)/\log_{(M)} n\right) + \mathbb{P}_{x_0,\eta_0}\big(\hat{\tau}_1^X < 3/\mu\big) \\
&\quad \le 1 - \mathbb{P}_{x_0,\eta_0}\big(\hat{\tau}_1^X < 3/\mu\big)/\log_{(M)} n \le \exp\big(-\mathbb{P}_{x_0,\eta_0}\big(\hat{\tau}_1^X < 3/\mu\big)/\log_{(M)} n\big),
\end{aligned}$$

where we have used the inequality $\exp -\frac{3}{2}x \le 1 - x$, valid for sufficiently small x. Since $d_0(x_0) = 0$, we have $\hat{\tau}_1^X \sim \mathcal{E}(\lambda\mu(1 - 1/n))$, and hence we have $\mathbb{P}_{x_0,\eta_0}(\hat{\tau}_1^X \ge 3/\mu) \asymp 1$. Substituting this into (4.9) gives the required bound. □

Proof of Lemma 4.8. We may assume that $K \ge 1$, otherwise the result is trivial.

For this lemma we only consider one walker, X; as such, we drop the X superscripts. We define

$$\tau_{\text{isol}}(s) = \inf\big\{t \ge s \mid d_t(X_t) = 0\big\} \quad \text{and} \quad \hat{\tau}_{\text{isol}}(s) = \inf\big\{t \ge \tau_{\text{isol}}(s) \mid d_t(X_t) > 0\big\};$$

also write $\tau_{\text{isol}} = \tau_{\text{isol}}(0)$ and $\hat{\tau}_{\text{isol}} = \hat{\tau}_{\text{isol}}(0)$.

For $k = 0, \dots, 3K$, set $t_k = \frac{t}{3K}$ and $t'_k = t_k + \frac{1}{2}t/(3K)$; also, for $k = 1, \dots, 3K$, write

$$\mathscr{H}_k = \mathscr{H}\big[t_{k-1}, t'_{k-1}\big] \quad \text{and} \quad \mathcal{J}_k = \big\{\hat{\tau}_{\text{isol}}(t_{k-1}) \le t_k\big\}.$$

If $\mathcal{J}_k$ occurs then at some point in the interval $[t_{k-1}, t_k]$ the walk is isolated and at a later point (in the same interval) is not. Observe that we have

$$\begin{aligned}
\{\tau_{K+1} > t\} \cap \mathscr{H}[0, t] &\subseteq \left\{\sum_{k=1}^{3K} \mathbf{1}(\mathcal{J}_k) \le K\right\} \cap \mathscr{H}[0, t] \\
&= \left\{\sum_{k=1}^{3K} \mathbf{1}\big(\mathcal{J}_k^c\big) \ge 2K\right\} \cap \mathscr{H}[0, t] \subseteq \left\{\sum_{k=1}^{3K} \mathbf{1}\big(\mathcal{J}_k^c \cap \mathscr{H}_k\big) \ge 2K\right\}.
\end{aligned}$$

Write $J = \sum_{k=1}^{3K} \mathbf{1}(\mathcal{J}_k^c \cap \mathscr{H}_k)$. Note that by the Markov property we have $J \preccurlyeq \text{Bin}(3K, q)$ where

$$q = \max_{x_0,\eta_0} \mathbb{P}_{x_0,\eta_0}\big(\mathcal{J}_1^c \cap \mathscr{H}_1\big).$$

We shall show, for a suitable constant c_2 in the definition of K, that $q \le \frac{1}{3}$, and then deduce that

$$\mathbb{P}_{x_0,\eta_0}\big(\tau_{K+1} > t, \mathscr{H}[0, t]\big) \le \mathbb{P}\left(\text{Bin}\left(3K, \frac{1}{3}\right) \ge 2K\right) \le \exp\left(-\frac{2}{3}K\right).$$

Observe that we have

$$\left\{\hat{\tau}_{\text{isol}} > \frac{t}{3K}, \tau_{\text{isol}} \le \frac{t}{6K}\right\} \subseteq \left\{\hat{\tau}_{\text{isol}} - \tau_{\text{isol}} > \frac{t}{6K}\right\}.$$

Thus we have, for any (x_0, η_0), that

$$\mathbb{P}_{x_0,\eta_0}\big(\mathcal{J}_1^c \cap \mathscr{H}_1\big) \le \mathbb{P}_{x_0,\eta_0}\left(\hat{\tau}_{\text{isol}} - \tau_{\text{isol}} > \frac{t}{6K}\right) + \mathbb{P}_{x_0,\eta_0}\left(\tau_{\text{isol}} > \frac{t}{6K}, \mathscr{H}_1\right).$$

The first term is simply

$$\mathbb{P}\left(\mathcal{E}\big((n-1)\lambda\mu/n\big) > \frac{t}{6K}\right) \le \mathbb{P}\left(\mathcal{E}(\mu) > \frac{t}{6K}\right) = \exp\left(-\frac{1}{6}\mu t/K\right)$$

since there are $n - 1$ edges that can open, and $\lambda > 1$. Applying Theorem 4.1, we have

$$\mathbb{P}_{x_0,\eta_0}\left(\tau_{\text{isol}} > \frac{t}{6K}, \mathscr{H}_1\right) \le 2\exp\left(-\frac{1}{6}\mu t/(K \log_{(M)} n)\right).$$

Combining these two bounds, we then find that

$$q = \max_{x_0,\eta_0} \mathbb{P}_{x_0,\eta_0}(\mathcal{J}_1^c \cap \mathscr{H}_1) \le \exp\left(-\frac{1}{6}\mu t/K\right) + 2\exp\left(-\frac{1}{6}c_1^*\mu t/(K \log_{(M)} n)\right)$$
$$\le 3\exp\left(-\frac{1}{6}c_1^*\mu t/(K \log_{(M)} n)\right).$$

Hence there exists a positive constant c_2 so that if $K = \lfloor c_2\mu t/\log_{(M)} n\rfloor$ then $q \le \frac{1}{3}$. □

5. Coupling

5.1. *Statement and application of coupling to mixing*

For this section only, we call a graph *good* if it satisfies the conditions of Definition 2.2 and in addition the condition that at least a proportion c_* of its vertices are isolated. Since this is an additional condition, the probability that a graph is good decreases, and hence all our isolation results (from Section 4) still hold with this extra condition. Recall also the definition of H from Definition 2.4, and in particular that $\pi_{\mathrm{ER}}(H) = 1 - o(1)$.

In this section, (X, η) and (Y, ξ) are two realisations of the dynamical percolation system; we shall define a Markovian coupling of the two systems, and find a tail bound on the coupling time. We look first at the case when the environments η and ξ start with $\eta_0 = \xi_0$.

Proposition 5.1 (Coupling Tail Bound). *There exists a Markovian coupling, which we denote by* $\mathbb{P}_{(x_0,\eta_0),(y_0,\eta_0)}$ *when* (X, η) *and* (Y, ξ) *start from* (x_0, η_0) *and* (y_0, ξ_0) *respectively, so that, for all* $M \in \mathbb{N}$, *all n sufficiently large, all t and all* (x_0, y_0, η_0), *we have*

$$\mathbb{P}_{(x_0,\eta_0),(y_0,\eta_0)}\big((X_t, \eta_t) \ne (Y_t, \xi_t)\big) \le 3\exp(-\mu t/\log_{(M)} n).$$

From Proposition 5.1 we are able to deduce the upper bounds in Theorems 1.1 and 1.2.

Proof of Theorem 1.2. Observe that we have

$$\big\|\mathbb{P}_{x_0,\eta_0}(X_t = \cdot) - \mathbb{P}_{y_0,\eta_0}(Y_t = \cdot)\big\|_{\mathrm{TV}} \le \big\|\mathbb{P}_{x_0,\eta_0}(X_t = \cdot, \eta_t = \cdot) - \mathbb{P}_{y_0,\eta_0}(Y_t = \cdot, \xi_t = \cdot)\big\|_{\mathrm{TV}}.$$

While the walk component alone is not a Markov chain, the full system is. It is then standard to upper bound the total variation distance by the tail probability of the coalescence time:

$$\big\|\mathbb{P}_{x_0,\eta_0}(X_t = \cdot, \eta_t = \cdot) - \mathbb{P}_{y_0,\eta_0}(Y_t = \cdot, \xi_t = \cdot)\big\|_{\mathrm{TV}} \le \mathbb{P}_{(x_0,\eta_0),(y_0,\eta_0)}(\tau_{\mathrm{c}} > t).$$

The coupling is coalescent, and so this tail bound is monotone in t. Proposition 5.1 now implies that this is $o(1)$ when $t = \frac{1}{\mu}\log_{(M)} n$, since $\eta_0 \in H$ (and replacing M with $M + 1$ in the proposition).

Now recall that the uniform distribution π_{RW} is invariant for our walk on *any* graph. Hence

$$\mathbb{P}_{\pi_{\mathrm{RW}},\eta_0}(Y_t = \cdot) = \pi_{\mathrm{RW}}$$

for any η_0. Thus we obtain our result: for $t = \frac{1}{\mu}\log_{(M)} n$ and any $\eta_0 \in H$, we have

$$\max_{x_0}\big\|\mathbb{P}_{x_0,\eta_0}(X_t = \cdot) - \pi_{\mathrm{RW}}\big\|_{\mathrm{TV}} \le \max_{x_0,y_0}\big\|\mathbb{P}_{x_0,\eta_0}(X_t = \cdot) - \mathbb{P}_{y_0,\eta_0}(Y_t = \cdot)\big\|_{\mathrm{TV}}$$
$$\le \max_{x_0,y_0}\mathbb{P}_{(x_0,\eta_0),(y_0,\eta_0)}(\tau_{\mathrm{c}} > t) = o(1).$$ □

In order to prove the mixing of the full system (X, η), we also need to know the mixing of the environment by itself. First recall that the environment process is simply a p-biased walk on the hypercube $\{0, 1\}^N$, where $N = \binom{n}{2}$ and $p = \lambda/n$, and where each coordinate refreshes at rate μ. We state the result now, then prove it at the end of the subsection.

Proposition 5.2 (Hypercube Mixing). *Consider the rate-*1 *p-biased random walk on the hypercube* $\{0,1\}^N$, *with* $1/N \ll p \le \frac{1}{2}$; *denote it* $\eta = (\eta_t)_{t\ge 0}$, *and its invariant distribution* π_p. *There is cutoff at* $\frac{1}{2}\log(N/p)$ *with window order* 1: *for all* $\varepsilon \in (0,1)$, *there exists a constant* C_ε *so that*

$$\max_{\eta_0}\bigl\|\mathbb{P}_{\eta_0}(\eta_t = \cdot) - \pi_p\bigr\|_{\mathrm{TV}} \le \varepsilon \quad \text{if } t \ge \frac{1}{2}\log(N/p) + C_\varepsilon,$$
$$\min_{\eta_0}\bigl\|\mathbb{P}_{\eta_0}(\eta_t = \cdot) - \pi_p\bigr\|_{\mathrm{TV}} \ge 1-\varepsilon \quad \text{if } t \le \frac{1}{2}\log(N/p) - C_\varepsilon.$$

Since we work in continuous time, we can apply this directly when the refresh-rate is μ.

Proof of Theorem 1.1. We consider first a lower bound on $t_{\mathrm{mix}}(\varepsilon)$. Observe that, trivially,

$$\bigl\|\mathbb{P}_{x_0,\eta_0}\bigl((X_t,\eta_t)\in\cdot\bigr) - \pi_U \times \pi_{\mathrm{ER}}\bigr\|_{\mathrm{TV}} \ge \bigl\|\mathbb{P}_{\eta_0}(\eta_t \in \cdot) - \pi_{\mathrm{ER}}\bigr\|_{\mathrm{TV}}.$$

Thus it suffices to only show that the environment has not mixed by time t. This follows immediately from the lower bound in Proposition 5.2, since $\frac{1}{2}\log(N/p) = \frac{3}{2}\log n + \Theta(1)$.

Now consider the upper bound. Fix (x_0,η_0) and (y_0,ξ_0). By Chapman–Kolmogorov, we have

$$\mathbb{P}_{x_0,\eta_0}\bigl((X_{s+t},\eta_{s+t})\in\cdot\bigr) = \mathbb{E}_{x_0,\eta_0}\bigl(\mathbb{P}_{X_s,\eta_s}\bigl((X_t,\eta_t)\in\cdot\bigr)\bigr).$$

Hence for any coupling $\mathbb{Q}$ of $\mathbb{P}_{x_0,\eta_0}((X_s,\eta_s)=\cdot)$ and $\mathbb{P}_{y_0,\xi_0}((Y_s,\xi_s)=\cdot)$ we have

$$\begin{aligned}&\bigl\|\mathbb{P}_{x_0,\eta_0}\bigl((X_{s+t},\eta_{s+t})\in\cdot\bigr) - \mathbb{P}_{y_0,\xi_0}\bigl((Y_{s+t},\xi_{s+t})\in\cdot\bigr)\bigr\|_{\mathrm{TV}}\\ &\quad\le \mathbb{Q}(\eta_s \ne \xi_s) + \mathbb{P}(\eta_s \notin H) + \max_{x_0',y_0',\eta_0'\in H}\bigl\|\mathbb{P}_{x_0',\eta_0'}\bigl((X_t,\eta_t)\in\cdot\bigr) - \mathbb{P}_{y_0',\eta_0'}\bigl((Y_t,\xi_t)\in\cdot\bigr)\bigr\|_{\mathrm{TV}}.\end{aligned}$$

In particular, consider the following such coupling $\mathbb{Q}$: fix $s \ge 0$, and couple (η_s,ξ_s) using the optimal coupling when started from (η_0,ξ_0); given η_s (and the fixed η_0), sample X_s conditional on η_s (and η_0); do similarly (and independently) for Y_s with ξ_s (and ξ_0). This then has

$$\mathbb{Q}(\eta_s \ne \xi_s) = \bigl\|\mathbb{P}_{\eta_0}(\eta_s\in\cdot) - \mathbb{P}_{\xi_0}(\xi_s\in\cdot)\bigr\|_{\mathrm{TV}}.$$

Now fix $M \in \mathbb{N}$ and choose s so that $\mu s = \frac{3}{2}\log n + \log_{(M+2)} n$, which has $\mu s \ge \frac{1}{2}\log(N/p) + \log_{(M+2)} n$. Then by the upper bound in Proposition 5.2 and the triangle inequality, we have

$$\mathbb{Q}(\eta_s \ne \xi_s) = o(1) \quad \text{and} \quad \mathbb{P}(\eta_s \notin H) \le \pi_{\mathrm{ER}}\bigl(H^c\bigr) + o(1) = o(1).$$

Since our coupling $\mathbb{P}_{\cdot,\cdot}$ from Proposition 5.1 is Markovian and coalescent, we have

$$\bigl\|\mathbb{P}_{x_0,\eta_0}\bigl((X_t,\eta_t)\in\cdot\bigr) - \mathbb{P}_{y_0,\eta_0}\bigl((Y_t,\xi_t)\in\cdot\bigr)\bigr\|_{\mathrm{TV}} \le \mathbb{P}_{(x_0,\eta_0),(y_0,\eta_0)}(\tau_{\mathrm{c}} > t).$$

Noting the conditions of Proposition 5.1, this implies that

$$\mathbb{P}_{(x_0,\eta_0),(y_0,\xi_0)}(\tau_{\mathrm{c}} > t) \le \varepsilon^2 \quad \text{when } t = \frac{1}{\mu}\log\bigl(3/\varepsilon^2\bigr)\log_{(M+1)} n.$$

Combining these three bounds we obtain, for these s and t, that

$$\bigl\|\mathbb{P}_{x_0,\eta_0}\bigl((X_{s+t},\eta_{s+t})\in\cdot\bigr) - \mathbb{P}_{y_0,\xi_0}\bigl((Y_{s+t},\xi_{s+t})\in\cdot\bigr)\bigr\|_{\mathrm{TV}} \le \varepsilon^2 + o(1) + o(1) \le \varepsilon.$$

Hence for all $\varepsilon \in (0,1)$ we have

$$\mu \cdot t_{\mathrm{mix}}(\varepsilon) \le \frac{3}{2}\log n + \log_{(M)} n.$$

This completes the proof of the upper bound. □

It remains to prove Proposition 5.2.

Proof of Proposition 5.2. We prove the upper bound first. We do this by relating the TV distance to the L_∞ distance. The probability an edge is in the same state $z \in \{0, 1\}$ as initially is exactly

$$e^{-t} + \big(1 - e^{-t}\big)\mathbb{P}\big(\mathrm{Bern}(p) = z\big) = p^z(1-p)^{1-z} + e^{-t}\big(1 - p^z(1-p)^{1-z}\big).$$

Also, it is well-known that for a reversible transition kernel $P = (P_t)_{t\ge 0}$ with invariant distribution π, writing $d_p(t)$ for the p-norm at time t (for $p \in [1, \infty]$), we have

$$d_{\mathrm{TV}}(t) = \frac{1}{2}d_1(t) \le \frac{1}{2}d_2(t) \quad \text{and} \quad d_\infty(2t) = \big(d_2(t)\big)^2 = \max_x P_{2t}(x,x)/\pi(x) - 1;$$

see [17, Exercise 4.5 and Proposition 4.15]. We hence deduce that

$$d_\infty(2t) = \max_{\eta_0} \big\| \mathbb{P}_{\eta_0}(\eta_{2t} \in \cdot) - \pi_{\mathrm{ER}} \big\|_\infty = \max_{\eta_0} \mathbb{P}_{\eta_0}(\eta_{2t} = \eta_0)/\pi_{\mathrm{ER}}(\eta_0) - 1.$$

Calculating this directly, recalling that $N = \binom{n}{2}$ and $p = \lambda/n$ with λ a constant, we see that

$$d_\infty(2t) = \big(1 + e^{-2t}(1/p - 1)\big)^N - 1 \le \exp\big(e^{-2t}N/p\big) - 1.$$

Hence if we set $t = \frac{1}{2}\log(N/p) + \frac{1}{2}C_\varepsilon$, for some large constant C_ε, then we obtain

$$d_\infty(2t) \le \exp(1/C_\varepsilon) - 1 \le 2/C_\varepsilon.$$

Finally we deduce that $d_{\mathrm{TV}}(t) \le \frac{1}{2}\sqrt{d_\infty(2t)} \le 1/\sqrt{2C_\varepsilon}$, proving the upper bound.

We now pursue the lower bound. For this, we consider the statistic $N_t = \sum_{e=1}^N \mathbf{1}(\eta_t(e) = 1)$, ie the number of open edges at time t. Observe that $N_t \sim \mathrm{Bin}(N, p)$ when $\eta_0 \sim \pi_{\mathrm{ER}}$. Consider starting η_0 from the all-1 state, which we denote $\mathbf{1} \in \{0, 1\}^N$. Then write

$$q_t = e^{-t} + \big(1 - e^{-t}\big)p = p + e^{-t}(1-p),$$

and observe that $N_t \sim \mathrm{Bin}(N, q_t)$ when $\eta_0 = \mathbf{1}$. Now define the set

$$A_t = \left\{ \zeta \in \{0,1\}^N \mid \sum_{e=1}^N \mathbf{1}\big(\zeta(e) = 1\big) \ge \frac{1}{2}(p + q_t)N \right\}.$$

This will be our distinguishing statistic/set. Recall that

$$\mathbb{E}\big(\mathrm{Bin}(N, r)\big) = Nr \quad \text{and} \quad \mathbb{V}\mathrm{ar}\big(\mathrm{Bin}(N, r)\big) = Nr(1-r) \le Nr.$$

Take $t = \frac{1}{2}\log(N/p) - \frac{1}{2}\log C_\varepsilon$, for some large constant C_ε. Note that $q_t - p \ge \frac{1}{2}\sqrt{C_\varepsilon p/N}$; also $q_t \le 2p$ since $\sqrt{p/N} \ll p$. (This is where we use the condition $p \gg 1/N$.) Hence, by Chebyshev,

$$\mathbb{P}_{\mathbf{1}}(\eta_t \notin A) \le \mathbb{P}\left(\big|\mathrm{Bin}(N, q_t) - q_t N\big| \ge \frac{1}{2}(q_t - p)N \right) \le \frac{4Nq_t}{(q_t - p)^2 N^2} \le \frac{50p}{(Cp/N)\cdot N} = \frac{50}{C_\varepsilon};$$

similarly, $\mathbb{P}_{\pi_{\mathrm{ER}}}(\eta_t \in A) \le 50/C_\varepsilon$. Hence $d_{\mathrm{TV}}(t) \ge 1 - 100/C_\varepsilon$, proving the lower bound. □

It remains to prove Proposition 5.1. To prove this, we carefully define a coupling, and use the result on dual-walker isolation, Theorem 4.2, that we proved in the previous section.

5.2. *Coupling description and proof of tail bound*

Below, we write (x, y) for the *undirected* edge with endpoints x and y; in particular, $(x, y) = (y, x)$. We only use the coupling below once the environments have coupled and the two walks have subsequently become then jointly isolated. We now define the coupling.

Definition 5.3. Suppose that (X, η) and (Y, ξ) are in the states (x, η_0) and (y, ξ_0), respectively. Assume that $\eta_0 = \xi_0$ and both x and y are isolated vertices in the environment $\eta_0 = \xi_0$. Let η evolve in the standard way. Couple ξ to η as follows. Suppose that edge (u, v) refreshes in η:

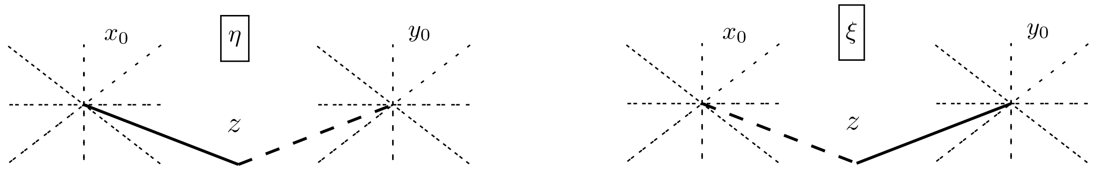

Fig. 1. The dotted lines represent the closed edges incident to x and y (recall that they are both isolated initially). The full line indicates opening an edge to another vertex z; the dashed line indicates leaving it closed. The dotting/dashing is reversed in ξ compared with η.

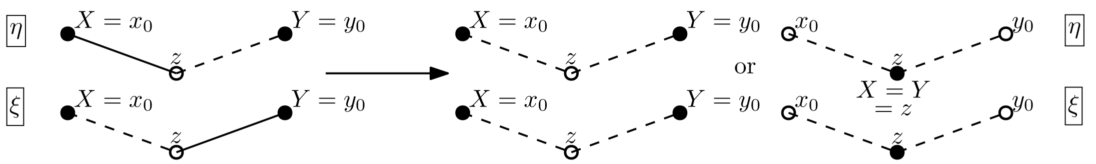

Fig. 2. The full line indicates an open edge; the dashed line indicates a closed edge. The walkers move along the open edges, moving together. On the left-hand side the filled dots represent where the walkers start; on the right-hand side the filled dots represent where the walkers end; the empty circles indicate empty sites.

- if $\{u, v\} \cap \{x, y\} = \varnothing$ or $\{u, v\} = \{x, y\}$, then perform the same update to (u, v) in ξ as in η;
- if $u = x$ and $v \notin \{x, y\}$, then perform the same update to (y, v) in ξ as to (u, v) in η;
- if $u \notin \{x, y\}$ and $v = y$, then perform the same update to (u, x) in ξ as to (u, v) in η.

This corresponds to a relabelling of x and y in ξ. [See Figure 1 for an illustration.]

While the environments are run like this, the environment η from the perspective of the walk X looks exactly the same as ξ from the perspective of Y, modulo the label difference $x-y$. This allows us to couple X and Y, modulo the relabelling. [See Figure 2 for an illustration.] So at every time, we have one of the following three situations:

- both X and Y are at some vertex $z \notin \{x, y\}$;
- X is at x and Y is at y;
- X is at y and Y is at x.

Observe that this defines a genuine Markovian coupling. When the systems (X, η) and (Y, ξ) start from (x, η_0) and (y, ξ_0), respectively, we denote this coupling $\mathbb{P}_{(x,\eta_0),(y,\xi_0)}$.

We now describe how to couple two processes (X, η) and (Y, ξ), when the environments are initially the same, but the walks are not necessarily isolated. (This is the setup of Proposition 5.1.) In the below algorithm, we define a time τ_c at which (X, η) and (Y, ξ) agree.

(i) Run the environments together (in the natural coupling, without any relabelling) and the walks independently until the two walks are jointly isolated, ie until time

$$\tau_0 = \inf\{t \geq 0 \mid d^{\eta}_{\tau_0}(X_{\tau_0}) = 0 = d^{\xi}_{\tau_0}(Y_{\tau_0})\}.$$

Note that $\eta_{\tau_0} = \xi_{\tau_0}$. Write $x = X_{\tau_0}$ and $y = Y_{\tau_0}$.

(ii) Set $k = 1$. Use the coupling from Definition 5.3: run until x becomes non-isolated in η (and hence y becomes non-isolated in ξ), and then on until the first time after that both x and y are isolated (in both η and ξ); call this first time σ_k and the final time τ_k. That is, set

$$\sigma_k = \inf\{t \geq \tau_0 \mid d^{\eta}_t(x) > 0\} \quad \text{and} \quad \tau_k = \inf\{t \geq \sigma_k \mid d^{\eta}_t(x) = 0 = d^{\eta}_t(y)\};$$

by the relabelling of the coupling, in the above definition we could swap $(X, \eta) \leftrightarrow (Y, \xi)$ and the times would be the same.

(iii) If $X_{\tau_k} \notin \{x, y\}$, then $X_{\tau_k} = Y_{\tau_k}$ (and vice versa). In this case, we have successfully coalesced the full processes. We then stop, setting $\tau_c = \tau_k$. (Also set $K = k$.)

Otherwise, we have $\{X_{\tau_k}, Y_{\tau_k}\} = \{x, y\}$. Then the walks are not at the same vertex, but the environments are in the same state and the walks are jointly isolated. Hence we can increment $k \to k + 1$ and return to Step (ii). By symmetry, assume $X_{\tau_k} = x$ and $Y_{\tau_k} = y$.

This means that, once we have the walks simultaneously isolated in the same environment, we can control their evolution very carefully. It will be straightforward to see that the probability that Step (ii) 'succeeds', ie ends with $X_{\tau_k} = Y_{\tau_k}$ is $\frac{1}{2} + o(1)$; hence we run Step (ii) an order 1 number of times to couple. It remains to analyse how long Steps (i) and (ii) take (Step (iii) only makes definitions, and so takes no time to 'run'). Step (i) is given precisely by Theorem 4.2.

Lemma 5.4. *The probability Step* (*ii*) *ends with* $X_{\tau_k} = Y_{\tau_k}$ *is at least* $\frac{1}{3}$, *and hence* $K \preccurlyeq \mathrm{Geo}(\frac{1}{3})$.

The intuition behind this lemma is as follows. Consider just X and the vertex $x = X_{\tau_0}$. In the (random) time interval between when x has degree 1 and when it becomes isolated, the component of x has size at least 2, and moreover $1/\mu \gg n$, so the walk takes a large number of steps if it is in this component, and so has probability $\frac{1}{2} + o(1)$ of being at x when x becomes isolated. This heuristic is made rigorous in the proof of Lemma 6.5, where an analogous claim is considered. From this we deduce that $K \preccurlyeq \mathrm{Geo}(\frac{1}{3})$.

We still need to find the distribution of $\tau_1 - \tau_0$. (Note that $(\tau_k - \tau_{k-1})_{k\geq 1}$ are iid.)

Lemma 5.5. *There exists a constant* C *so that, for all* $L \geq 1$ *and all* n *sufficiently large, we have*

$$\mathbb{P}(\tau_1 - \tau_0 > CL/\mu) \leq e^{-L}.$$

The intuition behind this lemma is as follows. There are $2n-3$ edges incident to $\{x, y\}$, in the complete graph; in equilibrium, in expectation $p(2n-3) \approx 2\lambda$ will be open. The number of edges open is well approximated by a birth and death chain with birth rate $\lambda\mu$ and death rate μ. For this chain, the return time to 0 has mean order $1/\mu$ and an exponential tail. This heuristic is made rigorous in the proof of Lemma 6.3, where a similar claim is considered – there only one vertex is considered, and some random number of edges are open initially; the same argument applies here. To go from $\tau_1 - \tau_0$ to $\tau_k - \tau_0 = \sum_{\ell=1}^{k} (\tau_\ell - \tau_{\ell-1})$, we use the simple fact that if a random variable has an exponential tail, then so does a sum. This can be proved by applying Chernoff; cf. the proof of Lemma 6.4. (Note that these statements and proofs do not require the graph to be 'good'.)

From these two lemmas, we immediately get the following corollary.

Corollary 5.6. *There exists a constant* C *so that, for all* $L \geq 1$ *and all* n *sufficiently large, we have*

$$\mathbb{P}(\tau_c - \tau_0 > CL/\mu) \leq e^{-L}.$$

Now that we know bounds on both τ_0 and $\tau_c - \tau_0$, given by Theorem 4.2 and Corollary 5.6 respectively, the bound on τ_c in Proposition 5.1 follows immediately. (For the application of Theorem 4.2, recall from Definition 2.4 that $\mathbb{P}_{\eta_0}(\mathscr{H}[0, n/\mu]^c \leq n^{-1}) \leq n^{-1}$ for all $\eta_0 \in H$.) Note also that we may assume $\mu t \geq 3 \log_{(M)} n$, else the claim holds trivially.

6. Invariant initial environment

In this section we prove Theorem 1.3, which concerns the case where we draw η_0 according to π_{ER} and set $X_0 = 1$; the reader should recall the precise statement. Throughout this entire section we consider the measure $\mathbb{P}_{1,\mathrm{ER}}(\cdot)$; for ease of notation, we drop the subscript and just write $\mathbb{P}(\cdot)$.

Proof of Theorem 1.3 (*lower bound*). Suppose the walk starts from an isolated vertex: it cannot have mixed before an incident edge opens. We make this idea precise and rigorous. We have

$$\mathbb{P}\big(d_0(1) = 0\big) = (1-p)^n = e^{-\lambda}\big(1 - o(1)\big).$$

Let τ be the first time an edge incident to $X_0 = 1$ opens. By counting edges and their respective rates, we see that $\tau \sim \mathcal{E}(\lambda\mu(1 - 1/n)) \succcurlyeq \mathcal{E}(\lambda\mu)$. Let $T = \lambda\mu t$, and observe that

$$\mathbb{P}\big(\tau > t \mid d_0(1) = 0\big) \geq \mathbb{P}\big(\mathcal{E}(\lambda\mu) > T/(\lambda\mu)\big) = e^{-T}.$$

On the event $\{d_0(1) = 0\} \cap \{\tau > t\}$, we have $X_t = 1$ (in fact $X_s = 1$ for all $s \leq t$), and hence

$$\big\|\mathbb{P}(X_t \in \cdot) - \pi_U\big\|_{\mathrm{TV}} \geq \mathbb{P}(X_t = 1) - \pi_U(1) \geq \mathbb{P}\big(d_0(1) = 0, \tau > t\big) - \pi_U(1) \geq e^{-(T+\lambda)} - \frac{1}{n}.$$

We desire T so that $\|\mathbb{P}(X_t \in \cdot) - \pi_U\|_{\mathrm{TV}} \geq \varepsilon$. By the above, we may take $T = -\log(\varepsilon + 1/n) - \lambda$. If $\varepsilon < e^{-3\lambda}$, then $T \geq \frac{1}{2}\log(1/\varepsilon)$. This proves the lower bound, as $t = T/(\lambda\mu)$. □

The aim of the remainder of this section is to prove the upper bound in Theorem 1.3; herein we assume that $\mu \leq \frac{2}{3}(1+\lambda)^{-1}/n$. To this end, let τ be the first time our initial vertex is isolated and the walk is not there, ie

$$\tau = \inf\{t \geq 0 \mid d_t(X_0) = 0,\, X_t \neq X_0\}.$$

The *idea* is that at time τ we are nearly uniform *and* have lost information about where we started, and so our total variation does not become large in the future. We show this rigorously.

Proposition 6.1. *For all n and all t, we have*

$$\|\mathbb{P}(X_t \in \cdot) - \pi_U\|_{\mathrm{TV}} \leq \mathbb{P}\big(X_t = 1, \tau \leq t \wedge (n/\mu)\big) + \mathbb{P}\big(\tau > t \wedge (n/\mu)\big) + \frac{1}{n}.$$

Proof. Note that by construction and the symmetry of the graph, at all times $t \geq 0$ we must have that $\mathbb{P}(X_t = x)$ is constant over $x \in V \setminus \{1\} = \{2, \dots, n\}$: define $\rho_t = \mathbb{P}(X_t = 1)$; then $\mathbb{P}(X_t = x) = (1-\rho_t)/(n-1)$ for all $x \in \{2, \dots, n\}$. We then have

$$2\|\mathbb{P}(X_t \in \cdot) - \pi_U\|_{\mathrm{TV}} = (n-1)\left|\frac{1-\rho_t}{n-1} - \frac{1}{n}\right| + \left|\rho_t - \frac{1}{n}\right| = 2\left|\rho_t - \frac{1}{n}\right| \leq 2\left(\rho_t + \frac{1}{n}\right).$$

Decomposing according to the event $\{\tau > t \wedge (n/\mu)\}$ completes the proof. □

Proposition 6.2. *There exists a constant C so that, for all $K \geq 2$ and all n sufficiently large, we have*

$$\mathbb{P}(\tau > CK/\mu) \leq e^{-K}.$$

Before we prove this, we define some preliminary notation, and then state three claims. First, let $\sigma_0 = \sigma_1' = 0$, and let σ_1 be the first time the initial vertex, 1, becomes isolated, ie

$$\sigma_1 = \inf\{t \geq 0 \mid d_t(1) = 0\},$$

and for $i \geq 1$ define inductively

$$\sigma_{i+1}' = \inf\{t \geq \sigma_i \mid d_t(1) > 0\}, \qquad \sigma_{i+1} = \inf\{t \geq \sigma_{i+1}' \mid d_t(1) = 0\} \quad \text{and}$$
$$\sigma_i'' = \inf\{t \leq \sigma_i \mid d_s(1) = 1\ \forall s \in [t, \sigma_i)\}.$$

In words, σ_i is the i-th time the vertex 1 becomes isolated, σ_i' is the first time after this that it becomes non-isolated and $[\sigma_i'', \sigma_i)$ is the interval in which it is degree 1 immediately before becoming isolated for the i-th time. By the memoryless property, $\sigma_i' - \sigma_{i-1} \sim^{\mathrm{iid}} \mathcal{E}(\lambda\mu(1-1/n))$.

Define $\tau_i := \sigma_i - \sigma_{i-1}$ for $i \geq 1$; then τ_1 is the time it takes to become isolated initially, and, for $i \geq 2$, τ_i is the time between the $(i-1)$-st and i-th times we become isolated. Note that the random variables $\{\tau_i\}_{i \geq 2}$ are all independent and identically distributed.

We now state three lemmas which we use to deduce Proposition 6.2.

Lemma 6.3. *There exists a constant C so that, for all $K \geq 1$ and all n sufficiently large, we have*

$$\mathbb{P}(\tau_1 > CK/\mu) \leq e^{-K}.$$

Lemma 6.4. *There exists a constant C so that, for all $K \geq 2$ and all n sufficiently large, we have*

$$\mathbb{P}\left(\sum_{i=2}^{K} \tau_i > CK/\mu\right) \leq e^{-K}.$$

Lemma 6.5. *For all n sufficiently large and all $i \geq 1$, we have*

$$\mathbb{P}(X_{\sigma_i} = 1 \mid X_{\sigma_j} = 1 \; \forall j < i) \leq \frac{2}{3}.$$

We now show how to conclude our tail bounds on τ from these three lemmas.

Proof of Proposition 6.2. Consider an integer $K \geq 2$. Lemma 6.5 tells us that

$$\mathbb{P}(X_{\sigma_i} = 1 \; \forall i \leq K) = \prod_{i=1}^{K} \mathbb{P}(X_{\sigma_i} = 1 \mid X_{\sigma_j} = 1 \; \forall j < i) \leq (2/3)^K.$$

Combining this with Lemma 6.3 and Lemma 6.4 tells us that

$$\mathbb{P}\big(\sigma_K \leq C' K/\mu, \exists k \leq K \text{ s.t. } X_{\sigma_k} \neq 1\big) \geq 1 - e^{-K} - e^{-K} - (3/2)^{-K}$$

for a suitably large constant C'. From this we deduce our claim. □

To complete the proof of our tail bound, it remains only to prove our three lemmas; we do this at the end of the section. For now, we turn to upper bounding $\mathbb{P}(X_t = 1, \tau \leq t \wedge (n/\mu))$.

Lemma 6.6. *There exists a constant C so that, for all n sufficiently large and all t, we have*

$$\mathbb{P}\big(X_t = 1, \tau \leq t \wedge (n/\mu)\big) \leq C/n.$$

Proof. Write $\mathcal{I}_t$ for the set of isolated vertices at time t. Write $s = t \wedge (n/\mu)$. First we lower bound the number of isolated vertices at time τ on the event $\{\tau \leq s\}$. From Proposition 2.5,

$$\mathbb{P}\left(\big|\mathcal{I}_\tau \setminus \{X_\tau\}\big| \leq \frac{1}{2} c_* n, \tau \leq s\right) = \mathcal{O}\big(n^{-2}\big).$$

By the symmetry of the complete graph, we must have that $\mathbb{P}(X_t = x \mid \mathcal{F}_\tau)$ is constant over $x \in \mathcal{I}_\tau \setminus \{X_\tau\}$ on the event $\{\tau \leq s\}$; let ξ_t be this (random) value. (ξ_t is an $\mathcal{F}_\tau$-measurable random variable.) Now, by construction of τ, we have that $X_0 = 1 \in \mathcal{I}_\tau \setminus \{X_\tau\}$. This says that

$$\xi_t = \mathbb{P}(X_t = 1 \mid \mathcal{F}_\tau)\mathbf{1}(\tau \leq s) \quad \text{and hence} \quad \mathbb{P}(X_t = 1, \tau \leq s) = \mathbb{E}(\xi_t).$$

It remains to bound $\mathbb{E}(\xi_t)$, which we now do. Note that we have

$$1 \geq \mathbb{P}(\tau \leq s) \geq \mathbb{E}\big(\big|\mathcal{I}_\tau \setminus \{X_\tau\}\big| \cdot \mathbb{P}(X_t = 1 \mid \mathcal{F}_\tau)\mathbf{1}(\tau \leq s)\big).$$

Letting $A = \{|\mathcal{I}_\tau \setminus \{X_\tau\}| \geq \frac{1}{2} c_* n\}$, we have $\mathbb{P}(A^c, \tau \leq s) = \mathcal{O}(n^{-2})$, as above. Hence

$$\begin{aligned} 1 &\geq \mathbb{E}\big(\big|\mathcal{I}_\tau \setminus \{X_\tau\}\big| \cdot \mathbb{P}(X_t = 1 \mid \mathcal{F}_\tau)\mathbf{1}(\tau \leq s)\mathbf{1}(A)\big) \\ &\geq \frac{1}{2} c_* n \mathbb{E}\big(\xi_t \mathbf{1}(A)\big) \geq \frac{1}{2} c_* n \big(\mathbb{E}(\xi_t) - \mathbb{P}\big(A^c, \tau \leq s\big)\big). \end{aligned}$$

Rearranging completes the proof:

$$\mathbb{P}(X_t = 1, \tau \leq s) = \mathbb{E}(\xi_t) \leq \left(\frac{1}{2} c_* n\right)^{-1} + \mathcal{O}\big(n^{-2}\big) \leq 3 c_*^{-1}/n.$$

□

We can now give the proof of the upper bound in Theorem 1.3.

Proof of Theorem 1.3 (*upper bound*). Lemma 6.6 says that, for all t, we have

$$\mathbb{P}\big(X_t = 1, \tau \leq t \wedge (n/\mu)\big) \leq C'/n,$$

for a constant C'. Hence we have

$$\big\|\mathbb{P}(X_t \in \cdot) - \pi_U\big\|_{\mathrm{TV}} \le \mathbb{P}\big(\tau > t \wedge (n/\mu)\big) + \big(C'+1\big)/n.$$

Observe that this upper bound is (weakly) monotone-decreasing in t, and Proposition 6.2 gives us a constant C so that

$$\big\|\mathbb{P}(X_t \in \cdot) - \pi_U\big\|_{\mathrm{TV}} \le \varepsilon^2 + \big(C'+1\big)/n \le \varepsilon \quad \text{when } t = 2C\log(1/\varepsilon)/\mu.$$

Hence we deduce that $t_{\mathrm{mix}}(\varepsilon) \le 2C\log(1/\varepsilon)/\mu$. □

Proof of Lemma 6.3. Write $d_t = d_t(1)$. Also rescale time by μ, so as to remove the μ factors from the workings. Observe that the jump-rates of d are as follows:

$$\begin{aligned} k &\to k+1 \quad \text{at rate } q_+(k) = (n-1-k)p = \big(\lambda - \lambda(1+k)/n\big); \\ k &\to k-1 \quad \text{at rate } q_-(k) = k(1-p) = (k - \lambda k/n). \end{aligned}$$

Let $q(k) = q_+(k) + q_-(k)$, and observe that $q(k) \ge q(0) \ge 1$ for all $k \ge 0$. We now couple d with an auxiliary process d', which has rate-1 jumps. Above 3λ, d' has probability $\frac{2}{3}$ of going up and $\frac{1}{3}$ of going down; below 3λ, it has the same probabilities as d, ie $q_+(k)/q(k)$ for up and $q_-(k)/q(k)$ for down. Set $d'_0 = d_0$, and write τ'_1 for the hitting time of 0 by d'. We then have $\tau_1 \preccurlyeq \tau'_1$.

Note that once d' reaches $\lceil 3\lambda \rceil$, it moves directly to 0 (in $\lceil 3\lambda \rceil$ steps) with probability bounded away from 0. The hitting time of $\lceil 3\lambda \rceil$ is that of a biased simple random walk. Since $d_0 \sim \mathrm{Bin}(n-1, \lambda/n)$, we may assume that $d_0 \le CK$ for some sufficiently large constant C with a penalty e^{-K} to the probability. Given this, we see that the hitting time of $\lceil 3\lambda \rceil$ has mean $\Theta(1)$ and an exponential tail. Once d' hits $\lceil 3\lambda \rceil$, we perform a geometric number of excursions, the length of which have an exponential tail. Hence τ'_1 has mean $\Theta(1)$ and an exponential tail. □

Proof of Lemma 6.4. Again, drop the μ factors. Note that $\tau_i \preccurlyeq \tau_1$, and τ_1 has mean $\Theta(1)$ with an exponential tail. Since the τ_i are independent, we then apply the Chernoff bound to a sum of K independent τ_1 random variables to deduce the lemma. □

Proof of Lemma 6.5. Fix $i \ge 1$. For $t \in (\sigma''_i, \sigma_i)$ we have that $d_t(1) = 1$; write x_i for the neighbour of 1 in the interval (σ''_i, σ_i). Note that all the σ-times depend only on the environment, not also on the walk. We describe a coupling between X and an auxiliary walk X' which is confined to the pair $\{1, x_i\}$. The coupling will have the property that

$$\mathbb{P}(X_{\sigma_i} = 1 \mid X_{\sigma_j} = 1\ \forall j < i) \le \mathbb{P}\big(X'_{\sigma_i} = 1 \mid X_{\sigma_j} = 1\ \forall j < i\big).$$

In particular, X' will be the usual simple random walk on $\{1, x_i\}$, jumping at rate $1/(n-1)$. Thus we shall see that the probability on the right-hand side is 'approximately' $\frac{1}{2}$.

We now explicitly define the coupling. Start X' from 1. If both X and X' are at 1, then move them together; if both X and X' are at x_i and X chooses vertex 1 to jump to, then move them together; otherwise let them evolve independently. Observe that, wherever X is at time σ''_i, we always have for $t \in [\sigma''_i, \sigma_i]$ that $X_t = 1$ implies $X'_t = 1$. Hence our desired inequality holds.

Observe that X' is at 1 if it has taken an even number of steps (and at x_i if odd). Hence

$$\mathbb{P}\big(X'_{\sigma_i} = 1 \mid X_{\sigma_j} = 1\ \forall j < i\big) = \mathbb{P}\big(\mathrm{Po}(r) \text{ is even}\big) = \frac{1}{2}\big(1 + \mathbb{E}\big(e^{-2r}\big)\big) \quad \text{where } r = \big(\sigma_i - \sigma''_i\big)/(n-1).$$

We have $\sigma_i - \sigma''_i \sim \mathcal{E}((\lambda + 1 - 3p)\mu)$ by counting edges and rates, and so the lemma follows since

$$\mathbb{E}\big(e^{-2r}\big) = \frac{(\lambda+1-3p)\mu}{(\lambda+1-3p)\mu + 2/(n-1)} \le \frac{1}{2}(\lambda+1)\mu n \le \frac{1}{3}.$$ □

Acknowledgements

The authors would like to thank Andrew Swan for many helpful discussions at the start of this project.

References

[1] M. Abdullah, C. Cooper and A. Frieze. Cover time of a random graph with given degree sequence. *Discrete Math.* **312** (21) (2012) 3146–3163. MR2957935 https://doi.org/10.1016/j.disc.2012.07.006

[2] L. Avena, H. Güldas, R. van der Hofstad and F. den Hollander. Mixing times of random walks on dynamic configuration models. *Ann. Appl. Probab.* **28** (4) (2018) 1977–2002. MR3843821 https://doi.org/10.1214/17-AAP1289

[3] L. Avena, H. Güldas, R. van der Hofstad and F. den Hollander. *Random Walks on Dynamic Configuration Models: A Trichotomy. Stochastic Processes and Their Applications*, 2018. MR3985565 https://doi.org/10.1016/j.spa.2018.09.010

[4] I. Benjamini, G. Kozma and N. Wormald. The mixing time of the giant component of a random graph. *Random Structures Algorithms* **45** (3) (2014) 383–407. MR3252922 https://doi.org/10.1002/rsa.20539

[5] N. Berestycki, E. Lubetzky, Y. Peres and A. Sly. Random walks on the random graph. *Ann. Probab.* **46** (1) (2018) 456–490. MR3758735 https://doi.org/10.1214/17-AOP1189

[6] B. Bollobas. A probabilistic proof of an asymptotic formula for the number of labelled regular graphs. *European J. Combin.* **1** (4) (1980) 311–316. MR0595929 https://doi.org/10.1016/S0195-6698(80)80030-8

[7] B. Bollobas. *Random Graphs*, 2nd edition. Cambridge University Press, Cambridge, 2001. MR1864966 https://doi.org/10.1017/CBO9780511814068

[8] J. Ding, E. Lubetzky and Y. Peres. Anatomy of the giant component: The strictly supercritical regime. *European J. Combin.* **35** (2014) 155–168. MR3090494 https://doi.org/10.1016/j.ejc.2013.06.004

[9] N. Fountoulakis and B. A. Reed. Faster mixing and small bottlenecks. *Probab. Theory Related Fields* **137** (3–4) (2007) 475–486. MR2278465 https://doi.org/10.1007/s00440-006-0003-8

[10] N. Fountoulakis and B. A. Reed. The evolution of the mixing rate of a simple random walk on the giant component of a random graph. *Random Structures Algorithms* **33** (1) (2008) 68–86. MR2428978 https://doi.org/10.1002/rsa.20210

[11] A. Frieze and M. Karonski. *Introduction to Random Graphs*. Cambridge University Press, Cambridge, 2016. MR3675279 https://doi.org/10.1017/CBO9781316339831

[12] D. Gillman. A Chernoff bound for random walks on expander graphs. *SIAM J. Comput.* **27** (4) (1998) 1203–1220. MR1621958 https://doi.org/10.1137/S0097539794268765

[13] O. Haggstrom, Y. Peres and J. E. Steif. Dynamical percolation. *Ann. Inst. Henri Poincaré Probab. Stat.* **33** (4) (1997) 497–528. MR1465800 https://doi.org/10.1016/S0246-0203(97)80103-3

[14] M. Jerrum and A. Sinclair. Approximate counting, uniform generation and rapidly mixing Markov chains. *Inform. and Comput.* **82** (1) (1989) 93–133. MR1003059 https://doi.org/10.1016/0890-5401(89)90067-9

[15] M. Jerrum and A. Sinclair. Approximating the permanent. *SIAM J. Comput.* **18** (6) (1989) 1149–1178. MR1025467 https://doi.org/10.1137/0218077

[16] G. F. Lawler and A. D. Sokal. Bounds on the L^2 spectrum for Markov chains and Markov processes: A generalization of Cheeger's inequality. *Trans. Amer. Math. Soc.* **309** (2) (1988) 557–580. MR0930082 https://doi.org/10.2307/2000925

[17] D. A. Levin, Y. Peres and E. L. Wilmer. *Markov Chains and Mixing Times*, 2nd edition. American Mathematical Society, Providence, RI, USA, 2017. MR3726904

[18] C. McDiarmid *On the Method of Bounded Differences. Surveys in Combinatorics, 1989*. *London Math. Soc. Lecture Note Ser.*, 148–188. Cambridge University Press, Cambridge, 1989. MR1036755

[19] B. D. McKay and N. C. Wormald. Asymptotic enumeration by degree sequence of graphs with degrees $o(n^l l^2)$. *Combinatorica* **11** (4) (1991) 369–382. MR1137769 https://doi.org/10.1007/BF01275671

[20] Y. Peres, P. Sousi and J. E. Steif. Mixing time for random walk on supercritical dynamical percolation. *Probab. Theory Related Fields* **176** (3–4) (2020) 809–849. MR4087484 https://doi.org/10.1007/s00440-019-00927-z

[21] Y. Peres, A. Stauffer and J. E. Steif. Random walks on dynamical percolation: Mixing times, mean squared displacement and hitting times. *Probab. Theory Related Fields* **162** (3–4) (2015) 487–530. MR3383336 https://doi.org/10.1007/s00440-014-0578-4

[22] B. Pittel, J. Spencer and N. Wormald. Sudden emergence of a giant k-core in a random graph. *J. Combin. Theory Ser. B* **67** (1) (1996) 111–151. MR1385386 https://doi.org/10.1006/jctb.1996.0036

[23] P. Sousi and S. Thomas. Cutoff for random walk on dynamical Erdős–Rényi graph, 2018. Available at arXiv:1807.04719.

[24] R. van der Hofstad. *Random Graphs and Complex Networks*. Cambridge University Press, Cambridge, 2017.

[25] L. Warnke. On the method of typical bounded differences. *Combin. Probab. Comput.* **25** (2) (2016) 269–299. MR3455677 https://doi.org/10.1017/S0963548315000103