



HAL
open science

CUTOFF FOR NON-BACKTRACKING RANDOM WALKS ON SPARSE RANDOM GRAPHS

Anna Ben-Hamou, Justin Salez

► **To cite this version:**

Anna Ben-Hamou, Justin Salez. CUTOFF FOR NON-BACKTRACKING RANDOM WALKS ON SPARSE RANDOM GRAPHS. *Annals of Probability*, 2017, 10.1214/16-AOP1100 . hal-01141192

HAL Id: hal-01141192

<https://hal.science/hal-01141192>

Submitted on 10 Apr 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

CUTOFF FOR NON-BACKTRACKING RANDOM WALKS ON SPARSE RANDOM GRAPHS

ANNA BEN-HAMOU AND JUSTIN SALEZ

ABSTRACT. A finite ergodic Markov chain is said to exhibit cutoff if its distance to stationarity remains close to 1 over a certain number of iterations and then abruptly drops to near 0 on a much shorter time scale. Discovered in the context of card shuffling (Aldous-Diaconis, 1986), this phenomenon is now believed to be rather typical among fast mixing Markov chains. Yet, establishing it rigorously often requires a challengingly detailed understanding of the underlying chain. Here we consider non-backtracking random walks on random graphs with a given degree sequence. Under a general sparsity condition, we establish the cutoff phenomenon, determine its precise window, and prove that the (suitably rescaled) cutoff profile approaches a remarkably simple, universal shape.

1. INTRODUCTION




FIGURE 1. Distance to stationarity along time for the NBRW on a random graph with 10^6 degree 3-vertices and 10^6 degree 4-vertices

1.1. **Setting.** Given a finite set V and a function $\deg: V \rightarrow \{2, 3, \dots\}$ such that

$$(1) \quad N := \sum_{v \in V} \deg(v)$$

is even, we construct a graph G with vertex set V and degrees $(\deg(v))_{v \in V}$ as follows. We form a set \mathcal{X} by “attaching” $\deg(v)$ *half-edges* to each vertex $v \in V$:

$$\mathcal{X} := \{(v, i) : v \in V, 1 \leq i \leq \deg(v)\}.$$

We then simply choose a *pairing* π on \mathcal{X} (i.e., an involution without fixed points), and interpret every pair of matched half-edges $\{x, \pi(x)\}$ as an edge between the corresponding vertices. Loops and multiple edges are allowed.




FIGURE 2. A set of half-edges \mathcal{X} , a pairing π and the resulting graph G

The *non-backtracking random walk* (NBRW) on the graph $G = G(\pi)$ is a discrete-time Markov chain with state space \mathcal{X} and transition matrix

$$P(x, y) = \begin{cases} \frac{1}{\deg(\pi(x))} & \text{if } y \text{ is a neighbour of } \pi(x) \\ 0 & \text{otherwise.} \end{cases}$$

In this definition and throughout the paper, two half-edges $x = (u, i)$ and $y = (v, j)$ are called *neighbours* if $u = v$ and $i \neq j$, and we let $\deg(x) := \deg(u) - 1$ denote the number of neighbours of the half-edge $x = (u, i)$. In words, the chain moves at every step from the current state x to a uniformly chosen neighbour of $\pi(x)$.




FIGURE 3. The non-backtracking moves from x (in red)

Note that the matrix P is symmetric with respect to π : for all $x, y \in \mathcal{X}$,

$$(2) \quad P(\pi(y), \pi(x)) = P(x, y).$$

In particular, P is doubly stochastic: the uniform law on \mathcal{X} is invariant for the chain. The *worst-case total-variation distance to equilibrium* at time $t \in \mathbb{N}$ is

$$(3) \quad \mathcal{D}(t) := \max_{x \in \mathcal{X}} \left\{ \frac{1}{2} \sum_{y \in \mathcal{X}} \left| P^t(x, y) - \frac{1}{N} \right| \right\}.$$

\mathcal{D} is non-increasing, and the number of transitions that have to be made before it falls below a given threshold $0 < \varepsilon < 1$ is known as the *mixing time*:

$$t_{\text{MX}}(\varepsilon) := \inf \{t \in \mathbb{N} : \mathcal{D}(t) < \varepsilon\}.$$

1.2. Result. The present paper is concerned with the typical profile of the function $t \mapsto \mathcal{D}(t)$ under the so-called *configuration model*, i.e. when the pairing π is chosen uniformly at random among the $(N-1)!!$ possible pairings of \mathcal{X} . In order to study large-size asymptotics, we let the vertex set V and degree function $\text{deg}: V \rightarrow \mathbb{N}$ depend on an implicit parameter $n \in \mathbb{N}$, which we omit from the notation for convenience. The same convention applies to all related quantities, such as N or \mathcal{X} . All asymptotic statements are understood as $n \rightarrow \infty$. Our interest is in the *sparse regime*, where the number N of half-edges diverges at a much faster rate than the maximum degree. Specifically, we assume that

$$(4) \quad \Delta := \max_{v \in V} \text{deg}(v) = N^{o(1)}.$$

As the behaviour of the NBRW at degree-2 vertices is deterministic, we will also assume, without much loss of generality, that

$$(5) \quad \min_{v \in V} \text{deg}(v) \geq 3.$$

Remarkably enough, the asymptotics in this regime depend on the degrees through two simple statistics: the mean logarithmic degree of an half-edge

$$(6) \quad \mu := \frac{1}{N} \sum_{v \in V} \text{deg}(v) \log(\text{deg}(v) - 1),$$

and the corresponding variance

$$(7) \quad \sigma^2 := \frac{1}{N} \sum_{v \in V} \text{deg}(v) \{\log(\text{deg}(v) - 1) - \mu\}^2.$$

We will also need some control on the third absolute moment:

$$(8) \quad \varrho := \frac{1}{N} \sum_{v \in V} \text{deg}(v) |\log(\text{deg}(v) - 1) - \mu|^3.$$

It might help the reader to think of μ, σ and ϱ as being fixed, or bounded away from 0 and ∞ . However, we only impose the following (much weaker) condition:

$$(9) \quad \frac{\sigma^2}{\mu^3} \gg \frac{(\log \log N)^2}{\log N} \quad \text{and} \quad \frac{\sigma^3}{\varrho \sqrt{\mu}} \gg \frac{1}{\sqrt{\log N}}.$$

Our main result states that on most graphs with degrees $(\deg(v))_{v \in V}$, the NBRW exhibits a remarkable behaviour, visible on Figure 1 and known as a *cutoff*: the distance to equilibrium remains close to 1 for a rather long time, roughly

$$(10) \quad t_\star := \frac{\log N}{\mu},$$

and then abruptly drops to nearly 0 over a much shorter time scale¹, of order

$$(11) \quad \omega_\star := \sqrt{\frac{\sigma^2 \log N}{\mu^3}}.$$

Moreover, the cutoff shape inside this window approaches a surprisingly simple function $\Phi: \mathbb{R} \rightarrow [0, 1]$, namely the tail distribution of the standard normal:

$$\Phi(\lambda) := \frac{1}{2\pi} \int_\lambda^\infty e^{-\frac{u^2}{2}} du.$$

It is remarkable that this limit shape does not depend at all on the precise degrees.

Theorem 1.1 (Cutoff for the NBRW on sparse graphs). *For every $0 < \varepsilon < 1$,*

$$\frac{t_{\text{MIX}}(\varepsilon) - t_\star}{\omega_\star} \xrightarrow{\mathbb{P}} \Phi^{-1}(\varepsilon).$$

Equivalently, for $t = t_\star + \lambda\omega_\star + o(\omega_\star)$ with $\lambda \in \mathbb{R}$ fixed, we have $\mathcal{D}(t) \xrightarrow{\mathbb{P}} \Phi(\lambda)$.

1.3. Comments. It is interesting to compare this with the d -regular case (i.e., $\deg: V \rightarrow \mathbb{N}$ constant equal to d) studied by Lubetzky and Sly [27]: by a remarkably precise path counting argument, they establish cutoff within constantly many steps around $t_\star = \log N / \log(d - 1)$. To appreciate the effect of heterogeneous degrees, recall that μ and σ are the mean and variance of $\log D$, where D is the degree of a uniformly sampled half-edge. Now, by Jensen's Inequality,

$$t_\star \geq \frac{\log N}{\log \mathbb{E}[D]},$$

and the less concentrated D , the larger the gap. The right-hand side is a well-known characteristic length in G , namely the typical inter-point distance (see e.g., [35]). One notable effect of heterogeneous degrees is thus that the mixing time becomes significantly larger than the natural graph distance. A heuristic explanation is as follows: in the regular case, all paths of length t between two points are equally likely for the NBRW, and mixing occurs as soon as t is large enough for many such paths to exist. In the non-regular case however, different paths have very different weights, and most of them actually have a negligible chance of being seen by the walk. Consequently, one has to make t larger in order to see paths with a "reasonable" weight. Even more remarkable is the impact of heterogeneous degrees on the cutoff width ω_\star , which satisfies $\omega_\star \gg \log \log N$

¹The fact that $\omega_\star \ll t_\star$ follows from condition (4).

against $\omega_\star = \Theta(1)$ in the regular case. Finally, the gaussian limit shape Φ itself is specific to the non-regular case and is directly related to the fluctuations of degrees along a typical trajectory of the NBRW.

Remark 1.2 (Simple graphs). *A classical result by Janson [23] asserts that the graph produced by the configuration model is simple (no loops or multiple edges) with probability asymptotically bounded away from 0, as long as*

$$(12) \quad \sum_{v \in V} \deg(v)^2 = \mathcal{O}(N).$$

*Moreover, conditionally on being simple, it is uniformly distributed over all simple graphs with degrees $(\deg(v))_{v \in V}$. Thus, every property which holds **whp** under the configuration model also holds **whp** for the uniform simple graph model. In particular, under (12), the conclusion of Theorem 1.1 extends to simple graphs.*

Remark 1.3 (IID degrees). *A common setting consists in generating an infinite IID degree sequence $(\deg(v))_{v \in \mathbb{N}}$ from some fixed degree distribution Q and then restricting it to the index set $V = \{1, \dots, n\}$ for each $n \geq 1$. Let D denote a random integer with distribution Q . Assuming that*

$$\mathbb{P}(D \leq 2) = 0, \quad \text{Var}(D) > 0, \quad \text{and} \quad \mathbb{E}[e^{\theta D}] < \infty \text{ for some } \theta > 0,$$

ensures that the conditions (4), (5) and (9) hold almost surely. Thus, Theorem 1.1 applies with the parameters μ, σ and N now being random. But the latter clearly concentrate around their deterministic counterparts, in the following sense:

$$\begin{aligned} N &= n\mathbb{E}[D] + O_{\mathbb{P}}\left(n^{\frac{1}{2}}\right) \\ \mu &= \mu_\star + O_{\mathbb{P}}\left(n^{-\frac{1}{2}}\right) \quad \text{with} \quad \mu_\star = \mathbb{E}[D \log(D-1)]/\mathbb{E}[D] \\ \sigma &= \sigma_\star + O_{\mathbb{P}}\left(n^{-\frac{1}{2}}\right) \quad \text{with} \quad \sigma_\star^2 = \mathbb{E}\left[D \{\log(D-1) - \mu_\star\}^2\right]/\mathbb{E}[D]. \end{aligned}$$

Those error terms are small enough to allow one to substitute $n, \mu_\star, \sigma_\star$ for N, μ, σ without affecting the convergence stated in Theorem 1.1.

1.4. Related work. The first instances of the cutoff phenomenon were discovered in the early 80's by Diaconis and Shahshahani [15] and Aldous [2], in the context of card shuffling: given a certain procedure for shuffling a deck of cards, there exists a quite precise number of shuffles slightly below which the deck is far from being mixed, and slightly above which it is almost completely mixed. The term *cutoff* and the general formalization appeared shortly after, in the seminal paper by Aldous and Diaconis [3]. Since then, this remarkable behaviour has been identified in a variety of other contexts, see e.g., Diaconis [14], Chen and Saloff-Coste [10], or the survey by Saloff-Coste [34] for random walks on finite groups.

Interacting particle systems in statistical mechanics provide a rich class of dynamics displaying cutoff. One emblematic example is the *stochastic Ising model* at high enough temperature, for which the cutoff phenomenon has been established

successively on the complete graph (Levin et al. [26]), on lattices (Ding et al. [16], Lubetzky and Sly [28]), and finally on any sequence of graphs (Lubetzky and Sly [29]). Other examples include the *Potts model* (Cuff et al. [13]), the *East process* (Ganguly et al. [20]), or the *Simple Exclusion process* on the cycle (Lacoin [24]).

The problem of singling out abstract conditions under which the cutoff phenomenon occurs, without necessarily pinpointing its precise location, has drawn considerable attention. In 2004, Peres [31] proposed a simple spectral criterion for reversible chains, known as the *product condition*. Although counter-examples have quickly been constructed (see Levin et al. [25, Chapter 18] and Chen and Saloff-Coste [10, Section 6]), the condition is widely believed to be sufficient for “most” chains. This has already been verified for birth-and-death chains (Ding et al. [17]) and, more generally, for random walks on trees (Basu et al. [5]). The latter result relies on a promising characterization of cutoff in terms of the concentration of hitting times of “worst” (in some sense) sets. See also Oliveira [30], Peres and Sousi [32], Griffiths et al. [21] and Hermon [22].

Many natural families of Markov chains are now believed to exhibit cutoff. Yet, establishing this phenomenon rigorously requires a very detailed understanding of the underlying chain, and often constitutes a challenging task even in situations with a high degree of symmetry. The historical case of random walks on the symmetric group for example, is still far from being completely understood: see Saloff-Coste [34] for a list of open problems, and Berestycki and Sengul [6] for a recent proof of one of them.

Understanding the mixing properties of random walks on sparse random graphs constitutes an important theoretical problem, with applications in a wide variety of contexts (see e.g., the survey by Cooper [11]). A classical result of Broder and Shamir [8] states that random d -regular graphs with d fixed are *expanders* with high probability (see also Friedman [19]). In particular, the simple random walk (SRW) on such graphs satisfies the product condition, and should therefore exhibit cutoff. This long-standing conjecture was confirmed only recently in an impressive work by Lubetzky and Sly [27], who also determined the precise cutoff window and profile. Their result is actually derived from the analysis of the NBRW itself, via a clever transfer argument. Interestingly, the mixing time of the SRW is $d/(d-2)$ times larger than that of the NBRW. This confirms the practical advantage of NBRW over SRW for efficient network sampling and exploration, and complements a well-known spectral comparison for regular expanders due to Alon et al. [4], as well as a recent result by Cooper and Frieze [12] on the cover time of random regular graphs. For other ways of speeding up random walks, see Cooper [11].

In the non-regular case however, the tight correspondence between the SRW and the NBRW breaks down, and there seems to be no direct way of transferring our main result to the SRW. We note that the latter should exhibit cutoff since the product condition holds, as can be seen from the fact that the *conductance* of sparse random graphs with a given degree sequence remains bounded away from 0

(see Abdullah et al. [1]). Confirming this constitutes a challenging open problem. In particular, it would be interesting to see whether the SRW still mixes faster than the NBRW. To the best of our knowledge, no precise conjectural expression for the mixing time of the SRW has been put forward².

2. PROOF OUTLINE

The proof of Theorem 1.1 is divided into two (unequal) halves: for

$$(13) \quad t = t_\star + \lambda w_\star + o(w_\star),$$

with $\lambda \in \mathbb{R}$ fixed, we will show that

$$(14) \quad \mathbb{E}[\mathcal{D}(t)] \geq \Phi(\lambda) - o(1),$$

$$(15) \quad \mathcal{D}(t) \leq \Phi(\lambda) + o_{\mathbb{P}}(1).$$

The lower bound (14) is proved in Section 3. The difficult part is the upper bound (15), due to the maximization over all possible initial states. Starting from state $x \in \mathcal{X}$, the distance to stationarity is

$$\begin{aligned} \mathcal{D}_x(t) &= \sum_{y \in \mathcal{X}} \left(\frac{1}{N} - P^t(x, y) \right)_+ \\ &= \sum_{y \in \mathcal{X}} \left(\frac{1}{N} - P^t(x, \pi(y)) \right)_+, \end{aligned}$$

since π is an involution. Using the symmetry (2), we may re-write $P^t(x, \pi(y))$ as

$$(16) \quad P^t(x, \pi(y)) = \sum_{(u,v) \in \mathcal{X} \times \mathcal{X}} P^{t/2}(x, u) P^{t/2}(y, v) \mathbb{1}_{\{\pi(u)=v\}}.$$

As a first approximation, let us assume that the balls of radius $t/2$ around x and y consist of disjoint trees, in agreement with Figure 2. This is made rigorous by a particular exposure process described in Section 5. The *weight* $\mathbf{w}(u) := P^{t/2}(x, u)$ (resp. $\mathbf{w}(v) := P^{t/2}(y, v)$) can then be unambiguously written as the inverse product of degrees along the unique path from x to u (resp. y to v).

A second approximation consists in eliminating those paths whose weight exceeds some given threshold $\theta > 0$ (the correct choice turns out to be $\theta \approx \frac{1}{N}$):

$$P^t(x, \pi(y)) \approx \sum_{u,v} \mathbf{w}(u) \mathbf{w}(v) \mathbb{1}_{\mathbf{w}(u) \mathbf{w}(v) \leq \theta} \mathbb{1}_{\{\pi(u)=v\}}.$$

Conditionally on the two trees of height $t/2$, this is a weighted sum of weakly dependent Bernoulli variables, and the large-weight truncation should prevent it from deviating largely from its expectation. We make this argument rigorous in Section 6, using Stein's method of exchangeable pairs. Provided the exposure process did

²During the finalization of the manuscript, a solution to this problem has been announced [7].




FIGURE 4. The tree-approximation

not reveal too many pairs of matched half-edges, the conditional expectation of $\mathbb{1}_{\{\pi(u)=v\}}$ remains close to $1/N$, and we obtain the new approximation

$$NP^t(x, \pi(y)) \approx \sum_{u,v} \mathbf{w}(u)\mathbf{w}(v) \mathbb{1}_{\mathbf{w}(u)\mathbf{w}(v) \leq \theta}.$$

Now, the right-hand side corresponds to the quenched probability (conditionally on the graph) that the product of the weights seen by two independent NBRW of length $t/2$, one starting from x and the other from y , does not exceed θ . The last step consists in approximating those trajectories by independent uniform samples from \mathcal{X} , which we denote by X_1^*, \dots, X_t^* . We obtain

$$\begin{aligned} \sum_{u,v} \mathbf{w}(u)\mathbf{w}(v) \mathbb{1}_{\mathbf{w}(u)\mathbf{w}(v) \leq \theta} &\approx \mathbb{P} \left[\frac{1}{\deg(X_1^*)} \cdots \frac{1}{\deg(X_t^*)} \leq \theta \right] \\ &\approx \mathbb{P} \left[\frac{\sum_{k=1}^t (\mu - \log \deg(X_k^*))}{\sigma\sqrt{t}} \leq \frac{\mu t + \log \theta}{\sigma\sqrt{t}} \right] \\ &\approx 1 - \Phi(\lambda), \end{aligned}$$

by the central limit theorem, since $\theta \approx 1/N$ and $t \approx t_* + \lambda\omega_*$. Consequently,

$$\sum_y \left(\frac{1}{N} - P^t(x, \pi(y)) \right)_+ \approx \Phi(\lambda),$$

as desired. This argument is made rigorous in Sections 7, 8 and 9.

3. THE LOWER BOUND

Fix $t \geq 1$, two states $x, y \in \mathcal{X}$, and a parameter $\theta \in (0, 1)$. Let $P_\theta^t(x, y)$ denote the contribution to $P^t(x, y)$ from paths having weight less than θ . Note that

$P_\theta^t(x, y) < P^t(x, y)$ if and only if some path of length t from x to y has weight larger than θ , implying in particular that $P^t(x, y) > \theta$. Thus,

$$\frac{1}{N} - P_\theta^t(x, y) \leq \left(\frac{1}{N} - P^t(x, y) \right)_+ + \frac{1}{N} \mathbf{1}_{P^t(x, y) > \theta}.$$

Summing over all $y \in \mathcal{X}$ and observing that there can not be more than $1/\theta$ half-edges $y \in \mathcal{X}$ satisfying $P^t(x, y) > \theta$, we obtain

$$1 - \sum_{y \in \mathcal{X}} P_\theta^t(x, y) \leq \mathcal{D}_x(t) + \frac{1}{\theta N}.$$

Now, the left-hand side is the quenched probability (i.e., conditional on the underlying pairing) that a NBRW $\{X_k\}_{0 \leq k \leq t}$ starting at x satisfies $\prod_{k=1}^t \frac{1}{\deg(X_k)} > \theta$. Taking expectation w.r.t. the pairing, we arrive at

$$(17) \quad \mathbb{P} \left(\prod_{k=1}^t \frac{1}{\deg(X_k)} > \theta \right) \leq \mathbb{E}[\mathcal{D}_x(t)] + \frac{1}{\theta N},$$

where the average is now over both the NBRW and the pairing (annealed law). A useful property of the uniform pairing is that it can be constructed sequentially, the pairs being revealed along the way, as we need them. We exploit this degree of freedom to generate the walk $\{X_k\}_{k \geq 0}$ and the pairing simultaneously, as follows. Initially, all half-edges are unpaired and $X_0 = x$; then at each time $k \geq 1$,

- (i) if X_{k-1} is unpaired, we pair it with a uniformly chosen other unpaired half-edge; otherwise, $\pi(X_{k-1})$ is already defined and no new pair is formed.
- (ii) in both cases, we let X_k be a uniformly chosen neighbour of $\pi(X_{k-1})$.

The sequence $\{X_k\}_{k \geq 0}$ is then exactly distributed according to the annealed law. Now, if we sample uniformly from \mathcal{X} instead of restricting the random choice made at (i) to unpaired half-edges, the uniform neighbour chosen at step (ii) also has the uniform law on \mathcal{X} . This creates a coupling between the process $\{X_k\}_{k \geq 1}$ and a sequence $\{X_k^*\}_{k \geq 1}$ of IID samples from \mathcal{X} , valid until the first time T where the uniformly chosen half-edge or its uniformly chosen neighbour is already paired. As there are less than $2k$ paired half-edges by step k , a crude union-bound yields

$$\mathbb{P}(T \leq t) \leq \frac{2t^2}{N}.$$

Consequently,

$$(18) \quad \left| \mathbb{P} \left(\prod_{k=1}^t \frac{1}{\deg(X_k)} > \theta \right) - \mathbb{P} \left(\prod_{k=1}^t \frac{1}{\deg(X_k^*)} > \theta \right) \right| \leq \frac{2t^2}{N}.$$

On the other hand, since $\{X_1^*, \dots, X_t^*\}$ are IID, Berry-Esseen's inequality implies

$$(19) \quad \left| \mathbb{P} \left(\prod_{k=1}^t \frac{1}{\deg(X_k^*)} > \theta \right) - \Phi \left(\frac{\mu t + \log \theta}{\sigma \sqrt{t}} \right) \right| \leq \frac{\varrho}{\sigma^3 \sqrt{t}}.$$

We may now combine (17), (18), (19) to obtain

$$\mathbb{E}[\mathcal{D}_x(t)] \geq \Phi\left(\frac{\mu t + \log \theta}{\sigma\sqrt{t}}\right) - \frac{1}{\theta N} - \frac{2t^2}{N} - \frac{\varrho}{\sigma^3\sqrt{t}}.$$

With t as in (13) and $\theta = (\log N)/N$, the right-hand side is $\Phi(\lambda) + o(1)$, thanks to our assumptions on μ, σ, ϱ . This establishes the lower bound (14).

4. THE UPPER-BOUND

Following Lubetzky and Sly [27], we call $x \in \mathcal{X}$ a *root* (written $x \in \mathcal{R}$) if the (directed) ball of radius h centered at x (denoted by \mathcal{B}_x) is a tree, where

$$(20) \quad h := \left\lfloor \frac{\log N}{10 \log \Delta} \wedge \log \log N \right\rfloor.$$

Note that $1 \ll h \ll \omega_*$ by assumptions (4) and (9). The first proposition below shows that we may restrict our attention to paths between roots. The second proposition provides a good control on such paths.

Proposition 4.1 (Roots are quickly reached).

$$\max_{x \in \mathcal{X}} P^h(x, \mathcal{X} \setminus \mathcal{R}) \xrightarrow{\mathbb{P}} 0.$$

Proposition 4.2 (Roots are well inter-connected). *For t as in (13),*

$$\min_{x \in \mathcal{R}} \min_{y \in \mathcal{R} \setminus \mathcal{B}_x} P^t(x, \pi(y)) \geq \frac{1 - \Phi(\lambda) - o_{\mathbb{P}}(1)}{N}.$$

Let us first see how those results imply the upper-bound (15). Observe that

$$\mathcal{D}(t+h) \leq \max_{x \in \mathcal{X}} P^h(x, \mathcal{X} \setminus \mathcal{R}) + \max_{x \in \mathcal{R}} \mathcal{D}_x(t).$$

The first term is $o_{\mathbb{P}}(1)$ by Proposition 4.1. For the second one, we write

$$\mathcal{D}_x(t) = \sum_{y \in \mathcal{R} \setminus \mathcal{B}_x} \left(\frac{1}{N} - P^t(x, \pi(y)) \right)_+ + \sum_{y \in \mathcal{B}_x \cup (\mathcal{X} \setminus \mathcal{R})} \left(\frac{1}{N} - P^t(x, \pi(y)) \right)_+.$$

Proposition 4.2 ensures that the first sum is bounded by $\Phi(\lambda) + o_{\mathbb{P}}(1)$ uniformly in $x \in \mathcal{R}$. To see that the second sum is $o_{\mathbb{P}}(1)$ uniformly in $x \in \mathcal{R}$, it suffices to bound its summands by $1/N$ and observe that $|\mathcal{B}_x| \leq \Delta^h = o(N)$ by (20), while

$$|\mathcal{X} \setminus \mathcal{R}| = \sum_{x \in \mathcal{X}} P^h(x, \mathcal{X} \setminus \mathcal{R}),$$

(P is doubly stochastic), which is $o_{\mathbb{P}}(N)$ uniformly in \mathcal{X} by Proposition 4.1.

Proof of Proposition 4.1. Define $r := \left\lfloor \frac{\log N}{5 \log \Delta} \right\rfloor$ and fix $x \in \mathcal{X}$. The ball of radius r around x can be generated sequentially, its half-edges being paired one after the other with uniformly chosen other unpaired half-edges, until the whole ball has been paired. Observe that at most $k = \frac{\Delta((\Delta-1)^r - 1)}{\Delta-2}$ pairs are formed. Moreover,

for each of them, the number of unpaired half-edges having an already paired neighbour is at most $\Delta(\Delta-1)^r$ and hence the conditional chance of hitting such an half-edge (thereby creating a cycle) is at most $p = \frac{\Delta(\Delta-1)^r-1}{N-2k-1}$. Thus, the probability that more than one cycle is found is at most

$$(kp)^2 = O\left(\frac{(\Delta-1)^{4r}}{N^2}\right) = o\left(\frac{1}{N}\right).$$

Summing over all $x \in \mathcal{X}$ (union bound), we obtain that with high probability, no ball of radius r in $G(\pi)$ contains more than one cycle.

To conclude the proof, we now fix a pairing π with the above property, and we prove that the NBRW on $G(\pi)$ starting from any $x \in \mathcal{X}$ satisfies

$$(21) \quad \mathbb{P}(X_t \text{ is not a root}) \leq 2^{1-t},$$

for all $t \leq r-h$. The claim is trivial if the ball of radius r around x is acyclic. Otherwise, it contains a single cycle \mathcal{C} , by assumption. Write $d(z, \mathcal{C})$ for the minimum length of a non-backtracking path from x to some $z \in \mathcal{C}$. The non-backtracking property ensures that if $d(X_t, \mathcal{C}) < d(X_{t+1}, \mathcal{C})$ for some $t < r-h$, then $X_{t+1}, X_{t+2}, \dots, X_{r-h}$ are all roots. By (5), the conditional chance that $d(X_{t+1}, \mathcal{C}) = d(X_t, \mathcal{C}) + 1$ given the past is at least $1/2$ (unless $d(X_t, \mathcal{C}) = 1$, which can only happen once). This shows (21). We then specialize to $t = h$. \square

5. THE EXPOSURE PROCESS

The remainder of the paper is devoted to the proof of Proposition 4.2. Fix two distinct half-edges $x, y \in \mathcal{X}$. We describe a two-stage procedure that generates a uniform pairing on \mathcal{X} together with a rooted forest \mathfrak{F} keeping track of certain paths from x and y . Initially, all half-edges are unpaired and \mathfrak{F} is reduced to its two roots, x and y . We then iterate the following three steps:

1. An unpaired half-edge $z \in \mathfrak{F}$ is selected according to some rule (see below).
2. z is paired with a uniformly chosen other unpaired half-edge z' .
3. If neither z' nor any of its neighbours was already in \mathfrak{F} , then all neighbours of z' become children of z in the forest \mathfrak{F} .

The *exploration stage* stops when no unpaired half-edge is compatible with the selection rule. We then complete the pairing by matching all the remaining unpaired half-edges uniformly at random: this is the *completion stage*.

The condition in step 3 ensures that \mathfrak{F} remains a forest: any $z \in \mathfrak{F}$ determines a unique sequence (z_0, \dots, z_h) in \mathfrak{F} such that z_0 is a root, z_i is a child of z_{i-1} for each $1 \leq i \leq h$, and $z_h = z$. We shall naturally refer to h and z_0 as the *height* and *root* of z , respectively. We also define the *weight* of z as

$$\mathbf{w}(z) := \prod_{i=1}^h \frac{1}{\deg(z_i)}.$$

Note that this quantity is the quenched probability that the sequence (z_0, \dots, z_h) is realized by a NBRW on G starting from z_0 . In particular,

$$(22) \quad \mathbf{w}(z) \leq P^h(z_0, z).$$

Our rule for step 1 consists in selecting an half-edge with maximal weight³ among all unpaired $z \in \mathfrak{F}$ with height $\mathbf{h}(z) < t/2$ and weight $\mathbf{w}(z) > w_{\min}$, where

$$w_{\min} := N^{-\frac{2}{3}}.$$

The only role of this parameter is to limitate the number of pairs formed during the exploration stage. As outlined in Section 2, we shall be interested in

$$\mathfrak{W} := \sum_{(u,v) \in \mathcal{H}_x \times \mathcal{H}_y} \mathbf{w}(u)\mathbf{w}(v)\mathbf{1}_{\mathbf{w}(u)\mathbf{w}(v) \leq \theta},$$

where \mathcal{H}_x (resp. \mathcal{H}_y) denotes the set of unpaired half-edges with height $\frac{t}{2}$ and root x (resp. y) in \mathfrak{F} at the end of the exploration stage, and where

$$(23) \quad \theta := \frac{1}{N(\log N)^2}.$$

Write $\overline{\mathfrak{W}}$ for the quantity obtained by replacing \leq with $>$ in \mathfrak{W} , so that

$$\begin{aligned} \mathfrak{W} + \overline{\mathfrak{W}} &= \sum_{(u,v) \in \mathcal{H}_x \times \mathcal{H}_y} \mathbf{w}(u)\mathbf{w}(v) \\ &\geq \sum_{z \in \mathcal{H}_x \cup \mathcal{H}_y} \mathbf{w}(z) - 1, \end{aligned}$$

thanks to the inequality $ab \geq a + b - 1$ for $a, b \in [0, 1]$. Now, let \mathfrak{U} denote the set of unpaired half-edges in \mathfrak{F} . By construction, at the end of the exploration stage, each $z \in \mathfrak{U}$ must have height $t/2$ or weight less than w_{\min} , so that

$$\mathfrak{W} + \overline{\mathfrak{W}} \geq \sum_{z \in \mathfrak{U}} \mathbf{w}(z) - \sum_{z \in \mathfrak{F}} \mathbf{w}(z)\mathbf{1}_{(\mathbf{w}(z) < w_{\min})} - 1.$$

Therefore, Proposition 4.2 follows from the following four technical lemmas.

Lemma 5.1. *For every $\varepsilon > 0$,*

$$\mathbb{P}\left(P^t(x, \pi(y)) \leq \frac{\mathfrak{W} - \varepsilon}{N}\right) = o\left(\frac{1}{N^2}\right).$$

Lemma 5.2. *For every $\varepsilon > 0$,*

$$\mathbb{P}\left(\sum_{z \in \mathfrak{F}} \mathbf{w}(z)\mathbf{1}_{(\mathbf{w}(z) < w_{\min})} > \varepsilon\right) = o\left(\frac{1}{N^2}\right).$$

Lemma 5.3. *For every $\varepsilon > 0$,*

$$\mathbb{P}\left(\overline{\mathfrak{W}} > \Phi(\lambda) + \varepsilon\right) = o\left(\frac{1}{N^2}\right).$$

³For definiteness, let us say that we use the lexicographic order on \mathcal{X} to break ties.

Lemma 5.4. *For every $\varepsilon > 0$,*

$$\mathbb{P} \left(x \in \mathcal{R}, y \in \mathcal{R} \setminus \mathcal{B}_x, \sum_{z \in \mathcal{U}} \mathbf{w}(z) < 2 - \varepsilon \right) = o \left(\frac{1}{N^2} \right).$$

6. PROOF OF LEMMA 5.1

Combining the representation (16) with the observation (22) yields

$$P^t(x, \pi(y)) \geq \sum_{(u,v) \in \mathcal{H}_x \times \mathcal{H}_y} \mathbf{w}(u)\mathbf{w}(v) \mathbf{1}_{\mathbf{w}(u)\mathbf{w}(v) \leq \theta} \mathbf{1}_{\pi(u)=v}.$$

The right-hand side can be interpreted as the weight of the uniform pairing chosen at the completion stage, provided we define the weight of a pair (u, v) as

$$(24) \quad \mathbf{w}(u)\mathbf{w}(v) \mathbf{1}_{u \in \mathcal{H}_x} \mathbf{1}_{v \in \mathcal{H}_y} \mathbf{1}_{\mathbf{w}(u)\mathbf{w}(v) \leq \theta}.$$

Lemma 5.1 now follows from the following general concentration inequality, which we apply conditionally on the exploration stage, with \mathcal{I} being the set of half-edges that did not get paired, and weights being given by (24).

Lemma 6.1. *Let \mathcal{I} be an even set, $\{w_{i,j}\}_{(i,j) \in \mathcal{I} \times \mathcal{I}}$ an array of non-negative weights, and π a uniform random pairing on \mathcal{I} . Then for all $a > 0$,*

$$\mathbb{P} \left(\sum_{i \in \mathcal{I}} w_{i, \pi(i)} \leq m - a \right) \leq \exp \left\{ -\frac{a^2}{4\theta m} \right\},$$

where $m = \frac{1}{|\mathcal{I}|-1} \sum_{i \in \mathcal{I}} \sum_{j \neq i} w_{i,j}$ and $\theta = \max_{i \neq j} (w_{i,j} + w_{j,i})$.

Note that in our case, $m = \frac{\mathfrak{W}}{|\mathcal{I}|-1}$. Lemma 5.1 follows easily by taking $a = \frac{\varepsilon}{|\mathcal{I}|-1}$ and observing that $|\mathcal{I}| - 1 \leq N$ and $\mathfrak{W} \leq 1$.

Proof. We exploit the following concentration result for Stein pairs due to Chatterjee [9] (see also Ross [33, Theorem 7.4]): let Y, Y' be bounded variables satisfying

- (i) $(Y, Y') \stackrel{d}{=} (Y', Y)$;
- (ii) $\mathbb{E}[Y' - Y | Y] = -\lambda Y$;
- (iii) $\mathbb{E}[(Y' - Y)^2 | Y] \leq \lambda(bY + c)$,

for some constants $\lambda \in (0, 1)$ and $b, c \geq 0$. Then for all $a > 0$,

$$\mathbb{P}(Y \leq -a) \leq \exp \left\{ -\frac{a^2}{c} \right\} \quad \text{and} \quad \mathbb{P}(Y \geq a) \leq \exp \left\{ -\frac{a^2}{ab + c} \right\}.$$

We shall only use the first inequality. Consider the centered variable

$$Y := \sum_{i \in \mathcal{I}} w_{i, \pi(i)} - m,$$

and let Y' be the corresponding quantity for the pairing π' obtained from π by performing a random switch: two indices i, j are sampled uniformly at random

from \mathcal{I} without replacement, and the pairs $\{i, \pi(i)\}$, $\{j, \pi(j)\}$ are replaced with the pairs $\{i, j\}$, $\{\pi(i), \pi(j)\}$. This changes the weight by exactly

$$(25) \quad \Delta_{i,j} := w_{i,j} + w_{j,i} + w_{\pi(i),\pi(j)} + w_{\pi(j),\pi(i)} - w_{i,\pi(i)} - w_{\pi(i),i} - w_{j,\pi(j)} - w_{\pi(j),j}.$$

It is not hard to see that $(\pi, \pi') \stackrel{d}{=} (\pi', \pi)$, so that (i) holds. Moreover,

$$\begin{aligned} \mathbb{E}[Y' - Y|\pi] &= \frac{1}{|\mathcal{I}|(|\mathcal{I}| - 1)} \sum_{i \in \mathcal{I}} \sum_{j \neq i} \Delta_{i,j} \\ &= \frac{4}{|\mathcal{I}|(|\mathcal{I}| - 1)} \sum_{i \in \mathcal{I}} \sum_{j \neq i} w_{i,j} - \frac{4}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} w_{i,\pi(i)} \\ &= -\frac{4}{|\mathcal{I}|} Y. \end{aligned}$$

Regarding the square $\Delta_{i,j}^2 = |\Delta_{i,j}| |\Delta_{i,j}|$, we may bound the first copy of $|\Delta_{i,j}|$ by 2θ and the second by changing all minus signs to plus signs in (25), yielding

$$\begin{aligned} \mathbb{E}[(Y' - Y)^2|\pi] &= \frac{1}{|\mathcal{I}|(|\mathcal{I}| - 1)} \sum_{i \in \mathcal{I}} \sum_{j \neq i} \Delta_{i,j}^2 \\ &\leq \frac{8\theta}{|\mathcal{I}|(|\mathcal{I}| - 1)} \sum_{i \in \mathcal{I}} \sum_{j \neq i} w_{i,j} + \frac{8\theta}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} w_{i,\pi(i)} \\ &= \frac{8\theta}{|\mathcal{I}|} (2m + Y). \end{aligned}$$

Note that taking conditional expectation with respect to Y does not affect the right-hand side. Thus, (ii) and (iii) hold with $\lambda = \frac{4}{|\mathcal{I}|}$, $b = 2\theta$ and $c = 4m\theta$. \square

7. PROOF OF LEMMA 5.2

We may fix $z_0 \in \{x, y\}$ and restrict our attention to the halved sum

$$Z := \sum_{z \in \mathfrak{X}} \mathbf{w}(z) \mathbf{1}_{(\mathbf{w}(z) < w_{\min})} \mathbf{1}_{(z \text{ has root } z_0)}.$$

Consider $m = \lceil \log N \rceil$ independent NBRWs on $G(\pi)$ starting at z_0 , each being killed as soon as its weight falls below w_{\min} , and write A for the event that their trajectories form a tree of height less than $t/2$. Clearly, $\mathbb{P}(A|\pi) \geq Z^m$. Taking expectation and using Markov inequality, we deduce that

$$\mathbb{P}(Z > \varepsilon) \leq \frac{\mathbb{P}(A)}{\varepsilon^m},$$

where the average is now taken over both the walks and the pairing. Recalling that $m = \lceil \log N \rceil$, it is more than enough to establish that $\mathbb{P}(A) = (o(1))^m$. To do so, we generate the m killed NBRWs one after this other, revealing the underlying pairs along the way, as described in Section 3. Given that the first $\ell - 1$ walks

form a tree of height less than $t/2$, the conditional chance that the ℓ^{th} walk also fulfils the requirement is $o(1)$, uniformly in $1 \leq \ell \leq m$. Indeed,

- either its weight falls below $\eta = (1/\log N)^2$ before it ever reaches an unpaired half-edge: thanks to the tree structure, there are at most $\ell - 1 < m$ possible trajectories to follow, so the chance is less than

$$m\eta = o(1).$$

- or the remainder of its trajectory after the first unpaired half-edge has weight less than $\Delta w_{\text{MIN}}/\eta$: this part consists of at most $t/2$ half-edges which can be coupled with uniform samples from \mathcal{X} for a total-variation cost of mt^2/N , as in Section 3. Thus, the conditional chance is at most

$$\frac{mt^2}{N} + \mathbb{P}\left(\prod_{k=1}^{t/2} \deg(X_k^*) \geq \frac{\eta}{\Delta w_{\text{MIN}}}\right) = o(1),$$

by Chebychev's inequality, since $\log\left(\frac{\eta}{\Delta w_{\text{MIN}}}\right) - \frac{mt}{2} \gg \sigma\sqrt{\frac{t}{2}}$.

8. PROOF OF LEMMA 5.3

Set $m = \lceil \log N \rceil$. On $G(\pi)$, let $X^{(1)}, \dots, X^{(m)}$ and $Y^{(1)}, \dots, Y^{(m)}$ be $2m$ independent NBRWs of length $t/2$ starting at x and y respectively. Let B denote the event that their trajectories form two disjoint trees and that for all $1 \leq k \leq m$,

$$\prod_{\ell=1}^{t/2} \frac{1}{\deg(X_\ell^{(k)})} \prod_{\ell=1}^{t/2} \frac{1}{\deg(Y_\ell^{(k)})} > \theta.$$

Then clearly, $\mathbb{P}(B|\pi) \geq \overline{\mathfrak{W}}^m$. Averaging w.r.t. the pairing π , we see that

$$\mathbb{P}(\overline{\mathfrak{W}} > \Phi(\lambda) + \varepsilon) \leq \frac{\mathbb{P}(B)}{(\Phi(\lambda) + \varepsilon)^m}.$$

Thus, it is enough to establish that $\mathbb{P}(B) \leq (\Phi(\lambda) + o(1))^m$. We do so by generating the $2m$ walks $X^{(1)}, Y^{(1)}, \dots, X^{(m)}, Y^{(m)}$ one after the other along with the underlying pairing, as above. Given that $X^{(1)}, Y^{(1)}, \dots, X^{(\ell-1)}, Y^{(\ell-1)}$ already satisfy the desired property, the conditional chance that $X^{(\ell)}, Y^{(\ell)}$ also does is at most $\Phi(\lambda) + o(1)$, uniformly in $1 \leq \ell \leq m$. Indeed,

- either one of the two walks attains length $s = \lceil 2 \log \log N \rceil$ before reaching an unpaired half-edge: there are at most $\ell - 1 < m$ possible trajectories to follow for each walk, so the conditional chance is at most

$$2m2^{-s} = o(1).$$

- or at least $t - 2s$ unpaired half-edges are encountered, and the product of their degrees falls below $\frac{1}{\theta}$ with conditional probability at most

$$\frac{4mt^2}{N} + \mathbb{P}\left(\prod_{k=1}^{t-2s} \deg(X_k^*) < \frac{1}{\theta}\right) = \Phi(\lambda) + o(1),$$

by the same coupling as above and Berry-Essen's inequality (19).

9. PROOF OF LEMMA 5.4

We denote by τ the (random) number of pairs formed during the exploration stage. For $k \geq 0$, we let \mathfrak{U}_k denote the set of unpaired half-edges in the forest after $k \wedge \tau$ pairs have been formed, and we consider the random variable

$$W_k := \sum_{z \in \mathfrak{U}_k} \mathbf{w}(z).$$

Initially $W_0 = 2$, and this quantity either stays constant or decreases at each stage, depending on whether the condition appearing in step 3 is satisfied or not. More precisely, denoting by z_k (resp. z'_k) the half-edge selected at step 1 (resp. chosen at step 2) of the k^{th} pair, we have for all $k \geq 1$,

$$W_k = W_{k-1} - \mathbf{1}_{(k \leq \tau)} \left(\mathbf{w}(z_k) \mathbf{1}_{(z'_k \in \mathfrak{U}_{k-1}^+)} + \mathbf{w}(z'_k) \mathbf{1}_{(z'_k \in \mathfrak{U}_{k-1})} \right),$$

where \mathfrak{U}_{k-1}^+ is the union of \mathfrak{U}_{k-1} and the set of unpaired neighbours of the roots. Now, let $\{\mathcal{G}_k\}_{k \geq 0}$ be the natural filtration associated with the exploration stage. Note that τ is a stopping time, that $\mathbf{w}(z_k)$ is \mathcal{G}_{k-1} -measurable, and that the conditional law of z'_k given \mathcal{G}_{k-1} is uniform on $\mathcal{X} \setminus \{z_1, \dots, z_k, z'_1, \dots, z'_{k-1}\}$. Thus,

$$\begin{aligned} \mathbb{E}[W_k - W_{k-1} | \mathcal{G}_{k-1}] &= -\mathbf{1}_{(k \leq \tau)} \frac{\mathbf{w}(z_k)(|\mathfrak{U}_{k-1}^+| - 2) + W_{k-1}}{N - 2k + 1}. \\ \mathbb{E}[(W_k - W_{k-1})^2 | \mathcal{G}_{k-1}] &= \mathbf{1}_{(k \leq \tau)} \frac{\mathbf{w}(z_k)^2(|\mathfrak{U}_{k-1}^+| - 4) + 2\mathbf{w}(z_k)W_{k-1} + \sum_{z \in \mathfrak{U}_{k-1}} \mathbf{w}(z)^2}{N - 2k + 1}. \end{aligned}$$

To bound those quantities, observe that each half-edge in \mathfrak{U}_{k-1} has weight at least $\frac{\mathbf{w}(z_k)}{\Delta}$ because its parent has been selected at an earlier iteration and our selection rule ensures that the quantity $\mathbf{w}(z_k)$ is non-increasing with k . Consequently,

$$|\mathfrak{U}_{k-1}| \frac{\mathbf{w}(z_k)}{\Delta} \leq \sum_{z \in \mathfrak{U}_{k-1}} \mathbf{w}(z) \leq 2.$$

Combining this with the bound $|\mathfrak{U}_{k-1}^+| \leq |\mathfrak{U}_{k-1}| + 2\Delta$, we arrive at

$$\begin{aligned} \mathbb{E}[W_k - W_{k-1} | \mathcal{G}_{k-1}] &\geq -\mathbf{1}_{(k \leq \tau)} \frac{4\Delta}{N - 2k + 1} \\ \mathbb{E}[(W_k - W_{k-1})^2 | \mathcal{G}_{k-1}] &\leq \mathbf{1}_{(k \leq \tau)} \frac{4\Delta \mathbf{w}(z_k) + 2}{N - 2k + 1}. \end{aligned}$$

Now recall that $\mathbf{w}(z_k) \geq w_{\text{MIN}}$ and $\mathbf{h}(z_k) < \frac{t}{2}$ as per our selection rule, implying

$$w_{\text{MIN}} \tau \leq \sum_{k \geq 1} \mathbf{w}(z_k) \mathbf{1}_{(\tau \geq k)} \leq \sum_{z \in \mathfrak{F}} \mathbf{w}(z) \mathbf{1}_{(\mathbf{h}(z) < \frac{t}{2})} \leq t.$$

The right-most inequality follows from the fact that the total weight at a given height in \mathfrak{F} is at most 2 (the total weight being preserved from a parent to its children, if any). We conclude that

$$\begin{aligned} \sum_{k=1}^{\tau} \mathbb{E}[W_k - W_{k-1} | \mathcal{G}_{k-1}] &\geq -\frac{4\Delta t}{w_{\min}N - 2t} := -m \\ \sum_{k=1}^{\tau} \mathbb{E}[(W_k - W_{k-1})^2 | \mathcal{G}_{k-1}] &\leq \frac{4\Delta t w_{\min} + 2t}{Nw_{\min} - 2t} := v. \end{aligned}$$

Now, fix $\varepsilon > 0$ and consider the martingale $\{M_k\}_{k \geq 0}$ defined by $M_0 = 0$ and

$$M_k := \sum_{i=1}^k \left\{ (W_{i-1} - W_i) \wedge \varepsilon - \mathbb{E}[(W_{i-1} - W_i) \wedge \varepsilon | \mathcal{G}_{i-1}] \right\}.$$

Then the increments of $\{M_k\}_{k \geq 0}$ are bounded by ε by construction, and the above computation guarantees that almost-surely,

$$\sum_{k=1}^{\tau} \mathbb{E}[(M_k - M_{k-1})^2 | \mathcal{G}_{k-1}] \leq v = N^{-\frac{1}{3} + o(1)}.$$

Thus, the martingale version of Bennett's inequality due to Freedman [18] yields

$$(26) \quad \mathbb{P}(M_{\tau} > 7\varepsilon) \leq \left(\frac{ev}{v + 7\varepsilon^2} \right)^7 = N^{-\frac{7}{3} + o(1)}.$$

But on the event $\{x \in \mathcal{R}, y \in \mathcal{R} \setminus \mathcal{B}_x\}$, all paths from the set $\{x, y\}$ to itself must have length at least h , and since $h \rightarrow \infty$, we must have asymptotically

$$\begin{aligned} \{x \in \mathcal{R}, y \in \mathcal{R} \setminus \mathcal{B}_x\} &\subseteq \left\{ \max_k (W_{k-1} - W_k) \leq \varepsilon \right\} \\ &\subseteq \{W_0 - W_{\tau} \leq M_{\tau} + m\} \end{aligned}$$

With (26), this proves Lemma 5.4 since $W_0 - W_{\tau} = 2 - \sum_{z \in \mathcal{U}} \mathbf{w}(z)$ and $m = o(1)$.

REFERENCES

- [1] M. Abdullah, C. Cooper, and A. Frieze. Cover time of a random graph with given degree sequence. *Discrete Mathematics*, 312(21):3146–3163, 2012.
- [2] D. Aldous. Random walks on finite groups and rapidly mixing Markov chains. In *Seminar on probability, XVII*, volume 986 of *Lecture Notes in Math.*, pages 243–297. Springer, Berlin, 1983.
- [3] D. Aldous and P. Diaconis. Shuffling cards and stopping times. *American Mathematical Monthly*, pages 333–348, 1986.
- [4] N. Alon, I. Benjamini, E. Lubetzky, and S. Sodin. Non-backtracking random walks mix faster. *Communications in Contemporary Mathematics*, 9(04):585–603, 2007.
- [5] R. Basu, J. Hermon, and Y. Peres. Characterization of cutoff for reversible markov chains. 2015.

- [6] N. Berestycki and B. Sengul. Cutoff for conjugacy-invariant random walks on the permutation group. *arXiv preprint arXiv:1410.4800*, 2014.
- [7] N. Berestycki, E. Lubetzky, Y. Peres, and A. Sly. Random walks on the random graph. *ArXiv e-prints*, 2015.
- [8] A. Broder and E. Shamir. On the second eigenvalue of random regular graphs. In *Foundations of Computer Science, 1987., 28th Annual Symposium on*, pages 286–294, Oct 1987.
- [9] S. Chatterjee. Stein’s method for concentration inequalities. *Probability theory and related fields*, 138(1):305–321, 2007.
- [10] G.-Y. Chen and L. Saloff-Coste. The cutoff phenomenon for ergodic markov processes. *Electronic Journal of Probability*, 13(3):26–78, 2008.
- [11] C. Cooper. Random walks, interacting particles, dynamic networks: Randomness can be helpful. In *Structural Information and Communication Complexity*, pages 1–14. 2011.
- [12] C. Cooper and A. Frieze. Vacant sets and vacant nets: Component structures induced by a random walk. *arXiv preprint arXiv:1404.4403*, 2014.
- [13] P. Cuff, J. Ding, O. Luidor, E. Lubetzky, Y. Peres, and A. Sly. Glauber dynamics for the mean-field potts model. *Journal of Statistical Physics*, 149(3):432–477, 2012.
- [14] P. Diaconis. The cutoff phenomenon in finite Markov chains. *Proc. Nat. Acad. Sci. U.S.A.*, 93(4):1659–1664, 1996.
- [15] P. Diaconis and M. Shahshahani. Generating a random permutation with random transpositions. *Probability Theory and Related Fields*, 57(2):159–179, 1981.
- [16] J. Ding, E. Lubetzky, and Y. Peres. The mixing time evolution of glauber dynamics for the mean-field ising model. *Communications in Mathematical Physics*, 289(2):725–764, 2009.
- [17] J. Ding, E. Lubetzky, and Y. Peres. Total variation cutoff in birth-and-death chains. *Probability theory and related fields*, 146(1-2):61–85, 2010.
- [18] D. A. Freedman. On tail probabilities for martingales. *the Annals of Probability*, pages 100–118, 1975.
- [19] J. Friedman. *A proof of Alon’s second eigenvalue conjecture and related problems*. American Mathematical Soc., 2008.
- [20] S. Ganguly, E. Lubetzky, and F. Martinelli. Cutoff for the east process. *arXiv preprint arXiv:1312.7863*, 2013.
- [21] S. Griffiths, R. Kang, R. Oliveira, and V. Patel. Tight inequalities among set hitting times in markov chains. *Proceedings of the American Mathematical Society*, 142(9):3285–3298, 2014.
- [22] J. Hermon. A technical report on hitting times, mixing and cutoff. *arXiv preprint arXiv:1501.01869*, 2015.
- [23] S. Janson. The probability that a random multigraph is simple. *Combinatorics, Probability and Computing*, 18(1-2):205–225, 2009.

- [24] H. Lacoïn. The Cutoff profile for the Simple-Exclusion process on the cycle. *arXiv preprint arXiv:1502.00952*, 2015.
- [25] D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov chains and mixing times*. American Mathematical Soc., 2009.
- [26] D. A. Levin, M. J. Luczak, and Y. Peres. Glauber dynamics for the mean-field ising model: cut-off, critical power law, and metastability. *Probability Theory and Related Fields*, 146(1-2):223–265, 2010.
- [27] E. Lubetzky and A. Sly. Cutoff phenomena for random walks on random regular graphs. *Duke Mathematical Journal*, 153(3):475–510, 2010.
- [28] E. Lubetzky and A. Sly. Cutoff for general spin systems with arbitrary boundary conditions. *Communications on Pure and Applied Mathematics*, 67(6):982–1027, 2014.
- [29] E. Lubetzky and A. Sly. Universality of cutoff for the ising model. *arXiv preprint arXiv:1407.1761*, 2014.
- [30] R. I. Oliveira. Mixing and hitting times for finite markov chains. *Electron. J. Probab*, 17(70):1–12, 2012.
- [31] Y. Peres. American institute of mathematics (AIM) research workshop “sharp thresholds for mixing times” (Palo Alto, December 2004). *Summary available at <http://www.aimath.org/WWN/mixingtimes>*.
- [32] Y. Peres and P. Sousi. Mixing times are hitting times of large sets. *Journal of Theoretical Probability*, pages 1–32, 2013.
- [33] N. Ross. Fundamentals of Stein’s method. *Probab. Surv*, 8:210–293, 2011.
- [34] L. Saloff-Coste. Random walks on finite groups. In *Probability on discrete structures*, pages 263–346. Springer, 2004.
- [35] R. van der Hofstad, G. Hooghiemstra, and P. Van Mieghem. Distances in random graphs with finite variance degrees. *Random Structures Algorithms*, 27(1):76–123, 2005. ISSN 1042-9832.