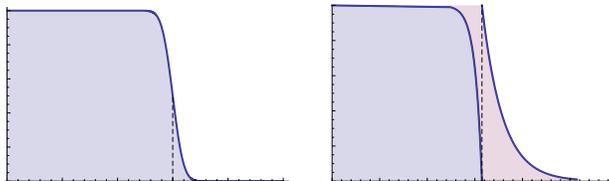


CUTOFF PHENOMENA FOR RANDOM WALKS ON RANDOM REGULAR GRAPHS

EYAL LUBETZKY AND ALLAN SLY

ABSTRACT. The cutoff phenomenon describes a sharp transition in the convergence of a family of ergodic finite Markov chains to equilibrium. Many natural families of chains are believed to exhibit cutoff, and yet establishing this fact is often extremely challenging. An important such family of chains is the random walk on $\mathcal{G}(n, d)$, a random d -regular graph on n vertices. It is well known that almost every such graph for $d \geq 3$ is an expander, and even essentially Ramanujan, implying a mixing-time of $O(\log n)$. According to a conjecture of Peres, the simple random walk on $\mathcal{G}(n, d)$ for such d should then exhibit cutoff **whp**. As a special case of this, Durrett conjectured that the mixing time of the lazy random walk on a random 3-regular graph is **whp** $(6 + o(1)) \log_2 n$.

In this work we confirm the above conjectures, and establish cutoff in total-variation, its location and its optimal window, both for simple and for non-backtracking random walks on $\mathcal{G}(n, d)$. Namely, for any fixed $d \geq 3$, the *simple* random walk on $\mathcal{G}(n, d)$ **whp** has cutoff at $\frac{d}{d-2} \log_{d-1} n$ with window order $\sqrt{\log n}$. Surprisingly, the *non-backtracking* random walk on $\mathcal{G}(n, d)$ **whp** has cutoff already at $\log_{d-1} n$ with *constant* window order. We further extend these results to $\mathcal{G}(n, d)$ for any $d = n^{o(1)}$ that grows with n (beyond which the mixing time is $O(1)$), where we establish concentration of the mixing time on one of two consecutive integers.



1. INTRODUCTION

A finite ergodic Markov chain is said to exhibit *cutoff* if its distance from the stationary measure drops abruptly, over a negligible time period known as the *cutoff window*, from near its maximum to near 0. That is, one has to run the Markov chain until the cutoff point in order for it to even slightly mix, and yet running it any further would be essentially redundant.

Let (X_t) be an aperiodic irreducible Markov chain on a finite state space Ω with transition kernel $P(x, y)$ and stationary distribution π . The worst-case total-variation distance to stationarity at time t is defined by

$$d(t) \triangleq \max_{x \in \Omega} \|\mathbb{P}_x(X_t \in \cdot) - \pi\|_{\text{TV}},$$

where \mathbb{P}_x denotes the probability given $X_0 = x$, and where $\|\mu - \nu\|_{\text{TV}}$, the *total-variation distance* of two distributions μ, ν on Ω , is given by

$$\|\mu - \nu\|_{\text{TV}} \triangleq \sup_{A \subset \Omega} |\mu(A) - \nu(A)| = \frac{1}{2} \sum_{x \in \Omega} |\mu(x) - \nu(x)| .$$

We define $t_{\text{MIX}}(\varepsilon)$, the total-variation *mixing-time* of (X_t) for $0 < \varepsilon < 1$, as

$$t_{\text{MIX}}(\varepsilon) \triangleq \min \{t : d(t) < \varepsilon\} .$$

Next, let $(X_t^{(n)})$ be a family of such chains, each with its corresponding worst-case total-variation distance from stationarity $d_n(t)$, its mixing-times $t_{\text{MIX}}^{(n)}$, etc. We say that this family of chains exhibits *cutoff* at time $t_{\text{MIX}}^{(n)}(\frac{1}{4})$ iff the following sharp transition in its convergence to stationarity occurs:

$$\lim_{n \rightarrow \infty} t_{\text{MIX}}^{(n)}(\varepsilon) / t_{\text{MIX}}^{(n)}(1 - \varepsilon) = 1 \quad \text{for any } 0 < \varepsilon < 1 . \quad (1.1)$$

The rate of convergence in (1.1) is addressed by the following: A sequence $w_n = o(t_{\text{MIX}}^{(n)}(\frac{1}{4}))$ is called a *cutoff window* for the family of chains $(X_t^{(n)})$ if for any $\varepsilon > 0$ there exists some $c_\varepsilon > 0$ such that for all n ,

$$t_{\text{MIX}}^{(n)}(\varepsilon) - t_{\text{MIX}}^{(n)}(1 - \varepsilon) \leq c_\varepsilon w_n . \quad (1.2)$$

That is, there is cutoff at time $t_n = t_{\text{MIX}}^{(n)}(\frac{1}{4})$ with window w_n if and only if

$$t_{\text{MIX}}^{(n)}(s) = (1 + O(w_n)) t_n = (1 + o(1)) t_n \quad \text{for any fixed } 0 < s < 1 ,$$

or equivalently, cutoff at time t_n with window w_n occurs if and only if

$$\begin{cases} \lim_{\lambda \rightarrow \infty} \liminf_{n \rightarrow \infty} d_n(t_n - \lambda w_n) = 1 , \\ \lim_{\lambda \rightarrow \infty} \limsup_{n \rightarrow \infty} d_n(t_n + \lambda w_n) = 0 . \end{cases}$$

Although many natural families of chains are believed to exhibit cutoff, determining that cutoff occurs proves to be an extremely challenging task even for fairly simple chains, as it often requires the full understanding of the delicate behavior of these chains around the mixing threshold. Before reviewing some of the related work in this area, as well as the conjectures that our work addresses, we state a few of our main results.

The focus of this paper is on random walks on a random regular graph, namely on $G \sim \mathcal{G}(n, d)$, a graph uniformly distributed over the set of all d -regular graphs on n vertices, for $d \geq 3$ and n large. This important class of random graphs has been extensively studied, among other reasons due to the remarkable expansion properties of its typical instance. One useful implication of these expansion properties is the rapid mixing of the corresponding *simple random walk* (SRW), the chain whose states are the vertices of the graph, and moves at each step to a uniformly chosen neighbor. Namely, the SRW on such a graph has a mixing time of $O(\log n)$ *with high probability* (**whp**), that is, with probability tending to 1 as $n \rightarrow \infty$.

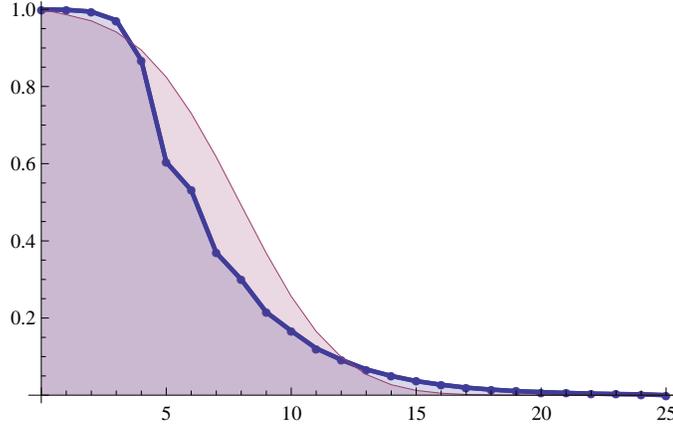


FIGURE 1. Distance from stationarity along time for the SRW on a random 6-regular graph on $n = 5000$ vertices.

Our first result establishes both cutoff and its optimal window for the SRW on a typical instance of $\mathcal{G}(n, d)$ for any $d \geq 3$ fixed. As we later describe, this settles conjectures of Durrett [17] and Peres [24] in the affirmative.

Theorem 1. *Let $G \sim \mathcal{G}(n, d)$ be a random regular graph for $d \geq 3$ fixed. Then **whp**, the simple random walk on G exhibits cutoff at $\frac{d}{d-2} \log_{d-1} n$ with a window of order $\sqrt{\log n}$. Furthermore, for any fixed $0 < s < 1$, the worst case total-variation mixing time **whp** satisfies*

$$t_{\text{MIX}}(s) = \frac{d}{d-2} \log_{d-1} n - (\Lambda + o(1)) \Phi^{-1}(s) \sqrt{\log_{d-1} n},$$

where $\Lambda = \frac{2\sqrt{d(d-1)}}{(d-2)^{3/2}}$ and Φ is the c.d.f. of the standard normal.

The essence of the cutoff for the SRW on a typical $G \sim \mathcal{G}(n, d)$ lies in the behavior of its counterpart, the non-backtracking random walk (NBRW), that does not traverse the same edge twice in a row (formally defined soon). Curiously, this chain also exhibits cutoff on $\mathcal{G}(n, d)$ **whp**, only this time the cutoff window is *constant*: (1.2) holds for $w_n = 1$ and c_ε logarithmic in $1/\varepsilon$:

Theorem 2. *Let $G \sim \mathcal{G}(n, d)$ be a random regular graph for $d \geq 3$ fixed. Then **whp**, the non-backtracking random walk on G has cutoff at $\log_{d-1}(dn)$ with a constant-size window. More precisely, for any fixed $\varepsilon > 0$, the worst case total-variation mixing time **whp** satisfies*

$$\begin{aligned} t_{\text{MIX}}(1 - \varepsilon) &\geq \lceil \log_{d-1}(dn) \rceil - \lceil \log_{d-1}(1/\varepsilon) \rceil, \\ t_{\text{MIX}}(\varepsilon) &\leq \lceil \log_{d-1}(dn) \rceil + 3 \lceil \log_{d-1}(1/\varepsilon) \rceil + 4. \end{aligned}$$

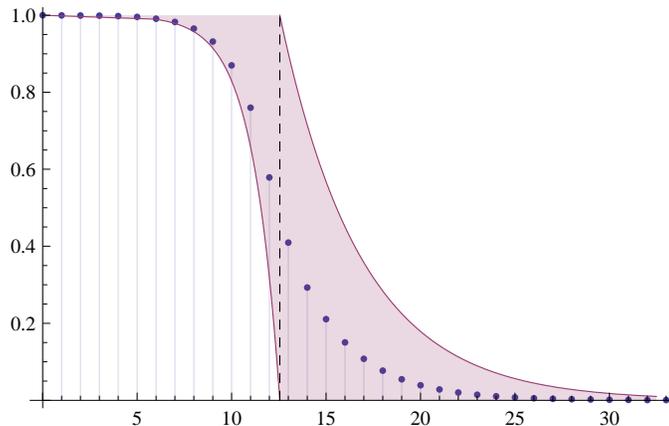


FIGURE 2. Distance from stationarity along time for the NBRW on a random 3-regular graph on $n = 2000$ vertices. Red curves represent a $(4 \log_{d-1}(1/\varepsilon))$ -wide cutoff window.

To gain insight to the above behaviors of the SRW and NBRW on a typical instance of $\mathcal{G}(n, d)$, note that **whp**, the random d -regular graph is locally-tree-like, its diameter is $(1 + o(1)) \log_{d-1} n$ and this is also the distance between a typical pair of vertices. In a d -regular tree, the height of a SRW, started at the root, is analogous to a biased 1-dimensional random walk with speed $(d-2)/d$. Hence, the time it takes this walk to reach height $\log_{d-1} n$ is concentrated around $\frac{d}{d-2} \log_{d-1} n$ with a standard deviation of order $\sqrt{\log n}$. Our results establish that at this time, the walk on $\mathcal{G}(n, d)$ is mixed. One of the keys to showing this is estimating the number of simple paths of length just beyond $\log_{d-1} n$ between most pairs of vertices (see Lemma 3.5 for a more precise statement). In comparison, as the NBRW started at the root of a tree is forbidden from backtracking up, it reaches height $\log_{d-1} n$ after precisely $\log_{d-1} n$ steps, hence the sharper cutoff window.

Establishing the above theorems requires a careful analysis of the local geometry around typical pairs of vertices, via a Poissonization argument. Namely, we show that the number of edges between certain neighborhoods of two prescribed vertices is roughly Poisson. Similar arguments then allow us to formulate analogous results for the case of regular graphs of high degree, that is, $\mathcal{G}(n, d)$ where d is allowed to tend to ∞ with n , up to $n^{o(1)}$.

1.1. Related work. The cutoff phenomenon was first identified for the case of random transpositions on the symmetric group in [14], and for the case of the riffle-shuffle and random walks on the hypercube in [2]. In their seminal paper [3] from 1985, Aldous and Diaconis established cutoff (and coined the term) for the top-in-at-random card shuffling process. See [13]

and [12] for more on the cutoff phenomenon, as well as [27] for a survey of this phenomenon for random walks on finite groups.

Unfortunately, there are relatively few examples where cutoff has been rigorously shown, whereas many important chains are conjectured to exhibit cutoff. Indeed, merely deciding whether a given family of finite Markov chains exhibits cutoff or not (without pinpointing the precise cutoff location) can already be a formidable task (see [13] for more on this problem).

In 2004, Peres [24] proposed the condition $t_{\text{MIX}}(\frac{1}{4}) \cdot \mathbf{gap} \rightarrow \infty$ as a cutoff criterion, where \mathbf{gap} is the spectral gap of the chain (i.e., $\mathbf{gap} \triangleq 1 - \lambda$ where λ is the largest nontrivial eigenvalue of the transition kernel). While this “product-condition” is indeed necessary for cutoff in a family of reversible chains, there are known examples where this condition holds yet there is no cutoff (see [12, Section 6]). Nevertheless, Peres conjectured that for many natural chains the product-condition does imply total-variation cutoff (e.g., this was recently verified in [15] for the class of birth-and-death chains).

An important family of chains, mentioned in this context in [24], is SRWs on transitive “expander” graphs of fixed degree d (graphs where the second eigenvalue of the adjacency matrix is bounded away from d). Chen and Saloff-Coste [12] verified that such chains exhibit cutoff when measuring the convergence to equilibrium via other (less common) norms, and mentioned the remaining open problem of proving total-variation cutoff.

On the other hand, it is well known that almost every d -regular graph for $d \geq 3$ is an expander (see [9], and also [25] for an analogous statement under a closely related combinatorial definition of expansion). In fact, it was shown by Friedman [18] that the second eigenvalue of the adjacency matrix of $G \sim \mathcal{G}(n, d)$ for $d \geq 3$ is **whp** $2\sqrt{d-1} + o(1)$, essentially as far from d as possible. Thus, random regular graphs are a valuable tool for constructing sparse expander graphs, and furthermore, for any fixed $d \geq 3$, any statement that holds **whp** for $\mathcal{G}(n, d)$ also holds for almost every d -regular expander. See, [11],[21] and also [28] for more on the thoroughly studied model $\mathcal{G}(n, d)$.

By the above, it follows that for any fixed $d \geq 3$, the mixing time of the SRW on $G \sim \mathcal{G}(n, d)$ is typically $O(\log n)$, whereas its \mathbf{gap} is bounded away from 0. Hence, if we consider the SRW on graphs $\{G_n \sim \mathcal{G}(n, d)\}$ for some fixed $d \geq 3$, then the product-condition typically holds, and according to the above conjecture of Peres, these chains should exhibit cutoff **whp**.

A special case of this was conjectured by Durrett, following his work with Berestycki [8] studying the SRW on a random 3-regular graph $G \sim \mathcal{G}(n, 3)$. They showed that at time $c \log_2 n$ the distance of the walk from its starting point is asymptotically $(\frac{c}{3} \wedge 1) \log_2 n$. This implies a lower bound of $3 \log_2 n$

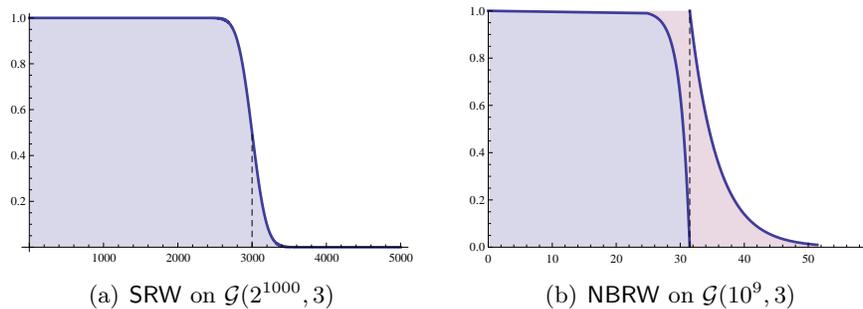


FIGURE 3. Estimates on the total-variation distance from stationarity for SRWs and NBRWs on large 3-regular graphs. (a) Asymptotic behavior of t_{MIX} established by Theorem 1. (b) Lower and upper bounds according to Theorem 2.

for the asymptotic mixing time of random 3-regular graphs, and in particular, an asymptotic lower bound of $6 \log_2 n$ for the *lazy* random walk (the lazy version of a chain with transition kernel P is the chain whose transition kernel is $\frac{1}{2}(P + I)$, i.e., in each step it stays in place with probability $\frac{1}{2}$, and otherwise it follows the rule of the original chain). In [17], Durrett conjectured that this latter bound is tight:

Conjecture (Durrett [17, Conjecture 6.3.5]). *The mixing time for the lazy random walk on the random 3-regular graph is asymptotically $6 \log_2 n$.*

Theorem 1 stated above confirms these conjectures of Peres and Durrett (one can readily infer an upper bound on the mixing time of the lazy random walk from Theorem 1). Not only does this theorem establish cutoff and its location for the SRW on $\mathcal{G}(n, d)$ (an analogous result immediately holds for the lazy walk), but it also determines the second order term in $t_{\text{MIX}}(s)$ for any $0 < s < 1$ (the term corresponding to the cutoff window of order $\sqrt{\log n}$).

The SRW on $\mathcal{G}(n, d)$ for $d = \lfloor (\log n)^a \rfloor$ and $a \geq 2$ fixed, starting from v_1 (not worst-case), was studied by Hildebrand [20]. He showed that in this case there is cutoff **whp** at $(1+o(1)) \log_d n$, and asked whether this also holds for $a < 2$. As we soon show, the answer to this question is positive, even from worst-case starting point and after replacing the $o(1)$ by an additive 2. To describe this result, we must first discuss the NBRW in further detail.

1.2. Cutoff for the SRW and NBRW. While the SRW of a graph is a Markov chain on its vertices, the NBRW has the set of directed edges (i.e., each edge appears in both orientations) as its state space: it moves from an edge (x, y) to a uniformly chosen edge (y, z) with $z \neq x$. However, in most applications for NBRWs on regular graphs (see, e.g., [7] and the

references therein), one often considers the projection of this chain onto the currently visited vertex (i.e., $(x, y) \mapsto y$), as it also converges to the uniform distribution on the vertices, and can thus be compared to the SRW.

In [5] the authors compare the SRW and this projection of the NBRW on regular expander graphs, showing that the NBRW has a faster *mixing rate* (see [22] for the definition of this spectral parameter, which for the SRW coincides with the largest nontrivial eigenvalue in absolute value). However, it was not clear how this spectral data actually translates into a direct comparison of the corresponding mixing times.

Theorems 1 and 2, as a bi-product, enable us to directly compare the mixing times of the SRW and NBRW (not only its projection onto the vertices). Namely, we obtain that the NBRW indeed *mixes faster* than the SRW on almost every d -regular graph, by a factor of $d/(d-2)$. Surprisingly, the delicate result stated in Theorem 2 also shows that once we omit the “noise” created by the backtracking possibility of the SRW, we are able to pinpoint the cutoff location up to $O(1)$ (see [19] for an example of such an $O(1)$ cutoff window related to random walks on the symmetric group).

Recalling that the cutoff window in Theorem 2 had the form $\log_{d-1}(1/\varepsilon)$, one may wonder what the effect of large degrees would be. Our results extend to the case of large d , all the way up to $d = n^{o(1)}$, beyond which the mixing time is constant (see, e.g., [16]) hence there is no point in discussing cutoff. The cutoff window indeed *vanishes* as $d \rightarrow \infty$, and the entire mixing transition occurs within merely two steps of the chain:

Theorem 3. *Let $G \sim \mathcal{G}(n, d)$ where $d = n^{o(1)}$ tends to ∞ with n . Then **whp**, for any fixed $0 < s < 1$, the worst case total-variation mixing time of the non-backtracking random walk on G **whp** satisfies*

$$t_{\text{MIX}}(s) \in \{ \lceil \log_{d-1}(dn) \rceil, \lceil \log_{d-1}(dn) \rceil + 1 \} .$$

*That is, the NBRW on G has cutoff **whp** within two steps of the chain.*

As a corollary, the relation between NBRWs and SRWs directly implies an analogous statement for the SRW on regular graphs of large degree. Here, the cutoff window becomes $\sqrt{(1/d) \log_d n}$ (compared to $\sqrt{\log n}$ for d fixed), and if $\frac{\log n}{\log \log n} = o(d)$ then the walk completely coincides with the NBRW.

Corollary 4. *Let $G \sim \mathcal{G}(n, d)$ where $d = n^{o(1)}$ tends to ∞ with n . Then **whp**, the SRW on G has cutoff at $\frac{d}{d-2} \log_{d-1} n$ with a window of $\sqrt{\frac{\log n}{d \log d}}$. Furthermore, if $\frac{d \log \log n}{\log n} \rightarrow \infty$, then for any fixed $0 < s < 1$, the worst case total-variation mixing time of the SRW on G **whp** satisfies*

$$t_{\text{MIX}}(s) \in \{ \lceil \log_{d-1}(dn) \rceil, \lceil \log_{d-1}(dn) \rceil + 1 \} .$$

In particular, this answers the above question of Hildebrand (the case of $d = \lfloor (\log n)^a \rfloor$ for any $a > 0$ fixed) in the affirmative, even from a worst starting position. Furthermore, instead of a multiplicative $1 + o(1)$, the cutoff point is determined up to an additive 2 if $a \geq 1$.

1.3. Random walks on the hypercube. As mentioned above, one of the original examples of cutoff was for the lazy random walk on the hypercube Q_m , which was shown by Aldous [2] to exhibit cutoff at $\frac{1}{2}m \log m$. When compared to the SRW on $\mathcal{G}(2^m, m)$, guaranteed by Corollary 4 to have cutoff **whp** at $(\log 2 + o(1))m / \log m$ (in this setting, $d = \log_2 n$ has $\frac{d \log \log n}{\log n} \rightarrow \infty$), this demonstrates the slower than typical mixing of the hypercube.

1.4. Organization. The rest of the paper is organized as follows. Section 2 contains several preliminary facts on random regular graphs. In Sections 3 and 4 we prove the main theorems, Theorems 1 and 2 resp., and in Section 5 we extend these proofs to the case of d large.

2. PRELIMINARIES

Let $G = (V, E)$, and let \bar{E} denote the set of directed edges (i.e., \bar{E} contains both orientations of every edge in E). Throughout the paper, we will use x, y, \dots for vertices in V , as opposed to \bar{x}, \bar{y}, \dots for directed edges in \bar{E} .

2.1. The configuration model. This model, introduced by Bollobás [10] and sometimes also referred to as the *pairing model*, provides a convenient method of both constructing and analyzing a random regular graph. We next briefly review some of the properties of this model which we will need for our arguments (see [11],[21] and [28, Section 2] for further information).

Given d and n with dn even, a d -regular (multi-)graph on n vertices is constructed via the configuration model as follows. Each vertex is identified with d distinct points, and a random perfect matching of all these dn points is then produced. The resulting multi-graph is obtained by collapsing every d -tuple into its corresponding vertex (possibly introducing loops or multiple edges). Let SIMPLE denote the event that the outcome is a simple graph.

It can easily be verified that, on the event SIMPLE, the resulting graph is uniformly distributed over $\mathcal{G}(n, d)$. Crucially, for any fixed d ,

$$\mathbb{P}(\text{SIMPLE}) = (1 + o(1)) \exp\left(\frac{1 - d^2}{4}\right), \quad (2.1)$$

where the $o(1)$ -term tends to 0 as $n \rightarrow \infty$. In particular, as this probability is uniformly bounded away from 0, any event that holds **whp** for multi-graphs constructed via the configuration model, also holds **whp** for $\mathcal{G}(n, d)$.

In fact, the statement in equation (2.1) was extended to any $d = o(n^{1/3})$ by McKay [23]. Although the asymptotical behavior of this probability was

thereafter determined for even larger values of d (see [28] for additional information), in this work we are only concerned with the case $d = n^{o(1)}$, and hence this result will suffice for our purposes.

A highly useful property of the configuration model is the following: we can expose the “pairings” sequentially, that is, given a vertex, we reveal the d neighbors of its corresponding points one by one, and so on. This allows us to “explore our way” into the graph, while constantly maintaining the uniform distribution over the pairings of the remaining unmatched points.

2.2. Neighborhoods and tree excess. We need the following definitions with respect to a given graph $G = (V, E)$. Let $\text{dist}(u, v) = \text{dist}_G(u, v)$ denote the distance between two vertices $u, v \in V$ in this graph. For any vertex $u \in V$ and integer t , the t -radius neighborhood of u , denoted by $B_t(u)$, and its (vertex) boundary $\partial B_t(u)$, are defined as

$$B_t(u) \triangleq \{v \in V : \text{dist}(u, v) \leq t\}, \quad \partial B_t(u) \triangleq B_t(u) \setminus B_{t-1}(u). \quad (2.2)$$

The abbreviated form B_t will be used whenever the identity of u becomes clear from the context. The *tree excess* of B_t , denoted by $\text{tx}(B_t)$, is the maximum number of edges that can be deleted from the induced subgraph on B_t while keeping it connected (i.e., the number of extra edges in that induced subgraph beyond $|B_t| - 1$).

The next lemma demonstrates the well known locally-tree-like properties of a typical $G \sim \mathcal{G}(n, d)$ for any fixed $d \geq 3$. Its proof follows from a standard and straightforward application of the above mentioned “exploration process” for the configuration model.

Lemma 2.1. *Let $G \sim \mathcal{G}(n, d)$ for some fixed $d \geq 3$, and let $t = \lfloor \frac{1}{5} \log_{d-1} n \rfloor$. Then **whp**, $\text{tx}(B_t(u)) \leq 1$ for all $u \in V(G)$.*

Proof. Choose $u \in V$ uniformly at random, and consider the process where the neighborhood of u is sequentially exposed level by level, according to the configuration model. When pairing the vertices of level i (and establishing level $i + 1$) for some $i \geq 0$, we are matching

$$m_i \leq d \vee (d - 1)|\partial B_i|$$

points among a pool of $(1 - o(1))dn$ yet unpaired points. For $1 \leq k \leq m_i$, let $\mathcal{F}_{i,k}$ denote the σ -field generated by the process of sequentially exposing pairings up to the k -th unmatched point in ∂B_i . Further let $A_{i,k}$ denote the event that the newly exposed pair of the k -th unmatched point in ∂B_i already belongs to some vertex in B_{i+1} . Clearly,

$$\mathbb{P}(A_{i,k} \mid \mathcal{F}_{i,k}) \leq \frac{(m_i - k) + (d - 1)(k - 1)}{(1 - o(1))dn} \leq \frac{(d - 1)m_i}{(1 - o(1))dn} \leq \frac{m_i}{n} \quad (2.3)$$

(where the last inequality holds for a sufficiently large n), and hence the number of events $\{A_{i,k} : 1 \leq k \leq m_i\}$ that occur is stochastically dominated by a binomial random variable with parameters $\text{Bin}(m_i, m_i/n)$. (We say that μ stochastically dominates ν , denoted by $\mu \succeq \nu$, if $\int f d\mu \geq \int f d\nu$ for every bounded increasing function f .) Moreover, since $m_i \leq d(d-1)^i$ for any $0 \leq i \leq t$, it follows that $\sum_{i=0}^{t-1} m_i \leq d(d-1)^t$, and the number of occurrences in the entire set of events $\{A_{i,k} : i < t\}$ can be stochastically dominated as follows:

$$\sum_{i=0}^{t-1} \sum_{k=1}^{m_i} \mathbf{1}_{A_{i,k}} \preceq \text{Bin} \left(d(d-1)^t, \frac{d(d-1)^{t-1}}{n} \right). \quad (2.4)$$

Notice that, by definition, the number of such events that occur is exactly the tree excess of $B_t(u)$. We thus obtain that

$$\mathbb{P}(\text{tx}(B_t) \geq 2) \leq O \left(\binom{d(d-1)^t}{2} \frac{d^2(d-1)^{2(t-1)}}{n^2} \right) = O \left(n^{-6/5} \right),$$

where the last equality is by the assumption on t . Taking a union bound over all vertices $u \in V$ completes the proof. \blacksquare

When proving cutoff for the NBRW in Section 4, we will be dealing with directed edges rather than vertices. The t -radius neighborhood of a directed edge \bar{x} , denoted by $B_t(\bar{x})$, and its boundary $\partial B_t(\bar{x})$, then consist of directed edges, and are defined analogously to (2.2) (with $\text{dist}(\bar{x}, \bar{y})$ measuring the shortest non-backtracking walk from \bar{x} to \bar{y} ; note that $\text{dist}(\cdot, \cdot)$ is not necessarily symmetric). The tree excess $\text{tx}(B_t(\bar{x}))$ in this case will refer to the undirected underlying graph induced on $B_t(\bar{x})$.

2.3. The cover tree of a regular graph. Let $G = (V, E)$ be a d -regular graph and $u \in V$ be some given vertex in G . The *cover tree of G at u* is a mapping $\varphi : \mathcal{T} \rightarrow V$, where \mathcal{T} is a d -regular tree with root ρ , and the following holds:

$$\begin{cases} \varphi(\rho) = u, \\ N_G(\varphi(x)) = \{\varphi(y) : y \in N_{\mathcal{T}}(x)\} \text{ for any } x \in \mathcal{T}, \end{cases} \quad (2.5)$$

where $N_H(u) = \{v \in V(H) : \text{dist}_H(u, v) = 1\}$ (i.e., $\partial B_1(v)$ for the graph H). That is, the root of \mathcal{T} is mapped to u , and φ respects 1-radius neighborhoods.

The following two simple facts will be useful later on. First, there is a one-to-one correspondence between non-backtracking paths in G starting from u and non-backtracking paths in \mathcal{T} starting from ρ . Second, if X_t is a simple random walk on \mathcal{T} , then $\varphi(X_t)$ is a simple random walk on G .

3. CUTOFF FOR THE SIMPLE RANDOM WALK

In this section, we prove Theorem 1, which establishes cutoff for the SRW on a typical random d -regular graph for any fixed $d \geq 3$. Throughout this section, let $d \geq 3$ be some fixed integer, and consider some $G \sim \mathcal{G}(n, d)$.

We need the following definition concerning the locally tree-like geometry.

Definition 3.1 (K -root). We say that a vertex $u \in V$ is a K -root if and only if the induced subgraph on $B_K(u)$ is a tree, that is, $\mathfrak{tx}(B_K(u)) = 0$.

Recalling Lemma 2.1, **whp** every vertex in our graph $G \sim \mathcal{G}(n, d)$ has a tree excess of at most 1 in its $\lfloor \frac{1}{5} \log_{d-1} n \rfloor$ -radius neighborhood. The next simple lemma shows that in such a graph (in fact, a weaker assumption suffices), a “burn-in” period of $\Theta(\log \log n)$ steps allows the SRW from the worst-case starting position to reposition itself in a typically “nice” vertex.

Lemma 3.2. *Let $K = \lfloor \log_{d-1} \log n \rfloor$, and suppose that every $u \in V$ has $\mathfrak{tx}(B_{5K}(u)) \leq 1$. Then for any $u \in V$, the SRW of length $4K$ from (u, v) ends at a K -root with probability $1 - o(1)$. In particular, there are $n - o(n)$ vertices in G that are K -roots.*

Proof. If $\mathfrak{tx}(B_{5K}(u)) = 0$ then the induced subgraph on B_{5K} is a tree and the result is immediate.

If $\mathfrak{tx}(B_{5K}(u)) = 1$ then the induced subgraph on B_{5K} is cycle C , with disjoint trees rooted on each of its vertices. Let X_t denote the position of the random walk at time t , and let $\rho_t = \text{dist}(X_t, C)$, that is, the length of the shortest path between C and X_t in G .

If the random walk is on the cycle then in the next step it either leaves C with probability $\frac{d-2}{d}$, or remains on C with probability $\frac{2}{d}$. Alternatively, if the random walk is not on C , then it moves one step closer to C with probability $\frac{1}{d}$ and one step further away with probability $\frac{d-1}{d}$. Either way,

$$\mathbb{E}[\rho_{t+1} - \rho_t \mid X_t] = \frac{d-2}{d}.$$

Therefore, $\rho_t - \frac{(d-2)t}{d}$ is a martingale, and the Azuma-Hoeffding inequality (cf., e.g., [6]) ensures that

$$\mathbb{P} \left(\left| \rho_{4K} - \rho_0 - \frac{4K(d-2)}{d} \right| > \frac{K}{3} \right) \leq \exp \left(\frac{-K}{72 \left(1 + \frac{d-2}{d}\right)^2} \right) = o(1).$$

We deduce that, **whp**, $\rho_{4K} \geq \frac{4K(d-2)}{d} - \frac{K}{3} \geq K$ and hence X_{4K} is a K -root.

To obtain the statement on the number of K -roots in G , suppose we start from a uniformly chosen vertex. Clearly, the random walk at time $4K$ is also uniform, thus the probability that a uniformly chosen vertex is not a K -root is $o(1)$, as required. \blacksquare

The following lemma demonstrates the control over the local geometry around a K -root with $K = \Theta(\log \log n)$.

Lemma 3.3. *Set $R = \lfloor \frac{4}{7} \log_{d-1} n \rfloor$ and $K = \lfloor \log_{d-1} \log n \rfloor$. With high probability, every K -root u satisfies*

$$|\partial B_t(u)| \geq (1 - o(1))d(d-1)^{t-1} \text{ for all } t < R .$$

Proof. Let u be a uniformly chosen vertex; expose its K -neighborhood, and assume that it is indeed a K -root. Following the notation from the proof of Lemma 2.1 we let $A_{i,k}$ be the event that, in the process of sequentially matching points, the newly exposed pair of the k -th unmatched point in ∂B_i belongs to a vertex already in B_{i+1} . Further recall that, by (2.3) and the discussion thereafter, the number of events $\{A_{i,k} : 0 \leq i < R\}$ that occur is stochastically dominated by a binomial variable with parameters $\text{Bin}\left(d(d-1)^R, \frac{d(d-1)^{R-1}}{n}\right)$. Since the expectation of this random variable is

$$d^2(d-1)^{2R-1}/n \leq O(n^{1/7}) ,$$

the number of events $A_{i,k}$ with $0 \leq i < R$ that occur is less than $n^{1/6}$ (with room to spare) with probability at least $1 - \exp(-\Omega(n^{1/6}))$.

Each event $A_{i,k}$ reduces the number of leaves in level $i+1$ by at most 2 and so reduces the number of leaves in level $t > i$ by at most $2(d-1)^{t-i-1}$ vertices. It follows that for each $0 \leq t < R$,

$$|\partial B_t| \geq d(d-1)^{t-1} - \sum_{i < t} \sum_k \mathbf{1}_{A_{i,k}} 2(d-1)^{t-i-1} . \quad (3.1)$$

Set $L = \lfloor \frac{1}{5} \log_{d-1} n \rfloor$. As u is a K -root, no events of the form $A_{i,k}$ with $i < K$ occur, and the number of events $A_{i,k}$ which occur with $i < L$ is exactly $\mathbf{tx}(B_L(u))$, giving

$$\sum_{i < L} \sum_k \mathbf{1}_{A_{i,k}} 2(d-1)^{t-i-1} \leq 2(d-1)^{t-K-1} \mathbf{tx}(B_L(u)) .$$

Furthermore, by the above discussion on the number of events $\{A_{i,k}\}$ that occur, we deduce that with probability at least $1 - \exp(-\Omega(n^{1/6}))$

$$\sum_{i=L}^{t-1} \sum_k \mathbf{1}_{A_{i,k}} 2(d-1)^{t-i-1} \leq 2(d-1)^{t-L-1} n^{1/6} = o((d-1)^t) .$$

Plugging the above in (3.1) we get that with probability $1 - \exp(-\Omega(n^{1/6}))$,

$$|\partial B_t| \geq (1 - o(1))d(d-1)^{t-1} - 2(d-1)^{t-K} \mathbf{tx}(B_L(u)) , \quad (3.2)$$

and a union bound implies that (3.2) holds for all K -roots u and all $t < R$ except with probability $\exp(-\Omega(n^{1/6}))$.

Finally, Lemma 2.1 asserts that **whp** every u satisfies $\mathbf{tx}(B_L(u)) \leq 1$. Hence, **whp**, every K -root u satisfies $|\partial B_t| \geq (1 - o(1))d(d-1)^{t-1}$ for all $0 \leq t \leq R$, as required. \blacksquare

Let $\partial B_t^*(u)$ denote the set of vertices in $\partial B_t(u)$ with a single (simple) path of length t to u . We next wish to establish an estimate for the typical number of such vertices, intersected with some other neighborhood $B_{t'}(v)$.

Lemma 3.4. *Let $K = \lfloor \log_{d-1} \log n \rfloor$ and $R = \lfloor \frac{4}{7} \log_{d-1} n \rfloor$. With high probability, any two K -roots u and v with $\text{dist}(u, v) > 2K$ satisfy*

$$|\partial B_t^*(u) \setminus B_{t+1}(v)| = (1 - o(1))d(d-1)^{t-1} \text{ for all } t < R - 1 .$$

Proof. The proof follows the same arguments as the proof of Lemma 3.3, except now we begin with two randomly chosen vertices u, v . Expose $B_K(u)$ and $B_K(v)$, at which point we may assume that both u and v are K -roots, and that $\text{dist}(u, v) > 2K$. Next, we sequentially expand the layers

$$\partial \tilde{B}_i \triangleq \{w \in V : \text{dist}(w, \{u, v\}) = i\} \text{ for } K < i \leq R .$$

By the above assumption on u and v , we have

$$|\partial \tilde{B}_K| = 2d(d-1)^{K-1} .$$

Repeating essentially the same calculations as those appearing in the proof of Lemma 3.3 now shows that with probability $1 - \exp(-\Omega(n^{1/6}))$,

$$|\partial \tilde{B}_t| = (2 - o(1))d(d-1)^{t-1} \text{ for all } t \leq R , \tag{3.3}$$

thus **whp**, the above holds for all pairs of K -roots u, v with $\text{dist}(u, v) > 2K$.

We claim that the statement of the lemma follows directly from (3.3). To see this, assume that (3.3) indeed holds for u, v as above, and let $t < R - 1$. Clearly, at most $d(d-1)^{t-1}$ of the vertices in $\partial \tilde{B}_t$ belong to $\partial B_t(v)$, hence

$$|\partial B_t(u) \setminus B_t(v)| = (1 - o(1))d(d-1)^{t-1} ,$$

and similarly,

$$|\partial B_{t+1}(v) \setminus B_{t+1}(u)| = (1 - o(1))d(d-1)^t .$$

Therefore,

$$\begin{aligned} |\partial B_t(u) \cap B_t(v)| &= o(d(d-1)^{t-1}) , \\ |\partial B_{t+1}(v) \cap B_{t+1}(u)| &= o(d(d-1)^t) , \end{aligned}$$

and altogether we obtain that

$$\begin{aligned} |\partial B_t(u) \cap B_{t+1}(v)| &\leq |\partial B_t(u) \cap B_t(v)| + |B_t(u) \cap \partial B_{t+1}(v)| \\ &= o(d(d-1)^t) . \end{aligned}$$

Since there are at most $d(d-1)^t$ paths of length t from u to $\partial B_t(u)$, and since $|\partial B_t(u)| = (1 - o(1))d(d-1)^{t-1}$, it then follows that

$$|\partial B_t(u) \setminus \partial B_t^*(u)| = o(d(d-1)^{t-1}) .$$

We deduce that $|\partial B_t^*(u) \cap B_{t+1}(v)| = o(d(d-1)^t)$, and the proof follows. \blacksquare

Lemma 3.5. *Let $K = \lfloor \log_{d-1} \log n \rfloor$ and $T = \lfloor \frac{1}{2} \log_{d-1} n \rfloor$. With high probability, any two K -roots u and v with $\text{dist}(u, v) > 2K$ satisfy*

$$\mathcal{S}_{2T+\ell}(u, v) \geq (1 - o(1)) \frac{1}{n} d(d-1)^{2T+\ell-1}$$

for all $2K \leq \ell \leq \frac{1}{20} \log_{d-1} n$, where $\mathcal{S}_k(u, v)$ denotes the number of simple paths of length k between u and v , and the $o(1)$ -term tends to 0 as $n \rightarrow \infty$.

Proof. Fix ℓ as above and expose the neighborhoods of u and v up to distance

$$t_u = \lceil \frac{1}{2}(2T + \ell - 1) \rceil , \quad t_v = \lfloor \frac{1}{2}(2T + \ell - 1) \rfloor$$

respectively. Notice that this selection gives

$$2T + \ell - 1 = t_u + t_v , \quad 0 \leq t_u - t_v \leq 1 .$$

We further define

$$A_u = \partial B_{t_u}^*(u) \setminus B_{t_u}(v) , \quad A_v = \partial B_{t_v}^*(v) \setminus B_{t_v}(u) .$$

We may now assume that the statement of Lemma 3.4 holds with respect to the neighborhoods of u and v already revealed (and them alone), that is

$$\begin{aligned} |A_u| &= (1 - o(1))d(d-1)^{t_u-1} , \\ |A_v| &= (1 - o(1))d(d-1)^{t_v-1} . \end{aligned}$$

In other words, A_u has $(1 - o(1))d(d-1)^{t_u}$ unmatched points and similarly, A_v has $(1 - o(1))d(d-1)^{t_v}$ unmatched points.

Now, sequentially match each of the points in A_u , and let $M_{u,v}$ denote the number of points of A_u matched with points in A_v . To obtain an upper bound on $M_{u,v}$, we once again repeat the arguments of Lemma 2.1, implying that it is stochastically bounded from above by a binomial variable as follows

$$M_{u,v} \preceq \text{Bin} \left((d-1)|A_u|, \frac{(d-1)|A_v|}{(1 - o(1))dn} \right) .$$

Since

$$\frac{(d-1)^2 |A_u| |A_v|}{dn} \leq O(n^{1/10}) ,$$

Chernoff bounds (cf., e.g., [6]) give that $M_{u,v} \leq n^{1/4}$ except with probability $e^{-\Omega(n^{1/4})}$. We thus assume that indeed $M_{u,v} \leq n^{1/4}$.

In this case, as we sequentially match points, each point in A_u has at least $|A_v| - n^{1/4}$ remaining points in A_v which it could potentially be matched to.

That is, conditional on previous matchings each point has at least $\frac{|A_v| - n^{1/4}}{dn}$ probability of being matched to a point in A_v . It follows that $M_{u,v}$ is stochastically bounded from below by a binomial variable

$$M_{u,v} \succeq \text{Bin} \left((d-1)|A_u|, \frac{(d-1)(|A_v| - n^{1/4})}{dn} \right).$$

Now

$$\frac{(d-1)^2|A_u|(|A_v| - n^{1/4})}{dn} = (1 - o(1)) \frac{1}{n} d(d-1)^{2T+\ell-1} = \Omega(\log_{d-1}^2 n),$$

and again by Chernoff bounds we have that the number of matchings is at least $(1 - o(1)) \frac{1}{n} d(d-1)^{2T+\ell-1}$ except with probability

$$\exp(-\Omega(\log_{d-1}^2 n)) = o(n^{-3}).$$

Each matching between a point in A_u and a point in A_v determines a simple path from u to v of length $2T + \ell$, thus

$$\mathcal{S}_{2T+\ell}(u, v) \geq M_{u,v} \geq (1 - o(1)) \frac{1}{n} d(d-1)^{2T+\ell-1}.$$

Taking a union bound over all u, v and ℓ completes the result. ■

Proof of Theorem 1. Set $K = \lfloor \log_{d-1} \log n \rfloor$ and set $T = \lfloor \frac{1}{2} \log_{d-1} n \rfloor$. By Lemma 3.2, after $4K$ steps with high probability the random walk is at a K -root. Since we are only seeking to establish t_{MIX} up to an accuracy of $o(\sqrt{\log_{d-1} n})$ and since $K = o(\sqrt{\log_{d-1} n})$ it is enough to consider the worst case mixing from a K -root to establish the result.

Let us assume that the statement of Lemma 3.5 holds. Let u and v be K -roots with $\text{dist}(u, v) > 2K$. By Lemma 3.5,

$$\mathcal{S}_{2T+\ell}(u, v) \geq \frac{1 - o(1)}{n} d(d-1)^{2T+\ell-1} \quad \text{for } 2K \leq \ell \leq \frac{1}{20} \log_{d-1} n.$$

Now let \mathcal{T} be the cover tree for G at u with a map φ , as defined in (2.5). Since each simple path in G corresponds to a distinct simple path in \mathcal{T} ,

$$\begin{aligned} \#\{w \in \mathcal{T} : \varphi(w) = v, \text{dist}(\rho, w) = 2T + \ell\} &\geq \mathcal{S}_{2T+\ell}(u, v) \\ &\geq \frac{1 - o(1)}{n} d(d-1)^{2T+\ell-1}, \end{aligned}$$

when $2K \leq \ell \leq \frac{1}{20} \log_{d-1} n$. Let X_t be a SRW on \mathcal{T} started from ρ and let $W_t = \varphi(X_t)$ be the corresponding SRW on \mathcal{G} started from u . Note that, by symmetry, conditioned on $\text{dist}(\rho, X_t) = k$ the random walk is uniform on the $d(d-1)^{k-1}$ points $\{w \in \mathcal{T} : \text{dist}(\rho, w) = k\}$. In addition, a random walk on a d -regular tree with $d \geq 3$ is transient, since the distance from the root

is a biased random walk with positive speed. In particular, the random walk returns to ρ only a finite number of times almost surely. If $X_t \neq \rho$ then

$$(\text{dist}(X_{t+1}, \rho) - \text{dist}(X_t, \rho)) \sim \begin{cases} -1 & 1/d, \\ 1 & (d-1)/d. \end{cases}$$

Therefore, the Central Limit Theorem gives that

$$\frac{\text{dist}(X_t, \rho) - \frac{(d-2)t}{d}}{\frac{2\sqrt{d-1}}{d}\sqrt{t}} \xrightarrow{d} N(0, 1). \quad (3.4)$$

Let A be the set of vertices which are K -roots and whose distance from u is greater than $2K$. Since there are at most $d(d-1)^{2K-1} = o(n)$ vertices within distance $2K$ of u , and since by Lemma 3.2 there are $n - o(n)$ K -roots in total, it follows that $|A| \geq n - o(n)$.

Combining these arguments, we deduce that if $v \in A$ and

$$t = \left\lfloor \frac{d}{d-2} \log_{d-1} n + k \sqrt{\log_{d-1} n} \right\rfloor \quad (3.5)$$

then

$$\begin{aligned} \mathbb{P}(W_t = v) &= \sum_{j=0}^t \mathbb{P}(\text{dist}(\rho, X_t) = j) \frac{\#\{w \in \mathcal{T} : \varphi(w) = v, \text{dist}(\rho, w) = j\}}{d(d-1)^{j-1}} \\ &\geq \sum_{\ell=2K}^{\frac{1}{20} \log_{d-1} n} \mathbb{P}(\text{dist}(\rho, X_t) = 2T + \ell) \frac{\frac{1+o(1)}{n} d(d-1)^{2T+\ell-1}}{d(d-1)^{2T+\ell-1}} \\ &= (1 + o(1)) \frac{1}{n} \mathbb{P} \left(2T + 2K \leq \text{dist}(\rho, X_t) \leq 2T + \frac{1}{20} \log_{d-1} n \right) \\ &= (1 + o(1)) \frac{1}{n} \left(1 - \Phi \left(\frac{-k}{\Lambda} \right) \right), \end{aligned}$$

where the final equality follows from equation (3.4) and where Φ is the distribution function of the standard normal and $\Lambda = \frac{2\sqrt{d-1}}{d-2} \sqrt{\frac{d}{d-2}}$. Then

$$\begin{aligned} \|\mathbb{P}(W_t \in \cdot) - \pi\|_{\text{TV}} &= \sum_{v \in V} \max \left\{ \frac{1}{n} - \mathbb{P}(W_t = v), 0 \right\} \\ &\leq \frac{n - |A|}{n} + \sum_{v \in A} \max \left\{ \frac{1}{n} - \mathbb{P}(W_t = v), 0 \right\} \\ &\leq o(1) + (1 + o(1)) |A| \frac{1}{n} \Phi \left(\frac{-k}{\Lambda} \right) = (1 + o(1)) \Phi \left(\frac{-k}{\Lambda} \right). \quad (3.6) \end{aligned}$$

It remains to provide a matching lower bound for $\|\mathbb{P}(W_t \in \cdot) - \pi\|_{\text{TV}}$. To this end, let $M = \log_{d-1} n - K$ and note that

$$\pi(B_M(u)) \leq \frac{1}{n} d(d-1)^{M-1} = o(1).$$

If $w \in T$ and $\text{dist}(\rho, w) \leq M$ then $\varphi(w) \in B_M$. For the same choice of t as given in (3.5), equation (3.4) gives that

$$\mathbb{P}(\text{dist}(X_t, \rho) \leq M) = (1 + o(1))\Phi\left(\frac{-k}{\Lambda}\right),$$

and so

$$\mathbb{P}(W_t \in B_M) \geq (1 + o(1))\Phi\left(\frac{-k}{\Lambda}\right).$$

It follows that

$$\|\mathbb{P}(W_t \in \cdot) - \pi\|_{\text{TV}} \geq \mathbb{P}(W_t \in B_M) - \pi(B_M) = (1 + o(1))\Phi\left(\frac{-k}{\Lambda}\right). \quad (3.7)$$

Combining equations (3.6) and (3.7) establishes that for any $0 < s < 1$

$$t_{\text{MIX}}(s) = \log_{d-1} n - (\Lambda + o(1))\Phi^{-1}(s)\sqrt{\log_{d-1} n},$$

completing the proof. ■

4. CUTOFF FOR THE NON-BACKTRACKING RANDOM WALK

In this section, we prove Theorem 2 that establishes the cutoff of the NBRW on a typical random d -regular graph for $d \geq 3$ fixed. Throughout this section, let $d \geq 3$ be some fixed integer, and consider some $G \sim \mathcal{G}(n, d)$.

Since the SRW induces a cutoff window of order $\sqrt{\log n}$ merely on account of its backtracking ability, throughout our arguments in Section 3 we could easily afford burn-in periods of order $\log \log n$. On the other hand, our statements for the NBRW establish a constant cutoff window (and moreover, logarithmic in $1/\varepsilon$), and therefore require a far more delicate approach.

Recall that the NBRW is a Markov chain on the set of directed edges; we thus begin by defining a *directed K -root*, analogous to Definition 3.1.

Definition 4.1 (*directed K -root*). A directed edge $\bar{x} \in \bar{E}$ is a directed K -root iff the induced subgraph on $B_K(\bar{x})$ is a tree, i.e., $\text{tx}(B_K(\bar{x})) = 0$.

As before, it is straightforward to show that the directed edges of G have locally-tree-like neighborhoods. This is stated by the next lemma.

Lemma 4.2. *Let $L = \lfloor \frac{1}{5} \log_{d-1} n \rfloor$. Then **whp**, $\text{tx}(B_L(\bar{x})) \leq 1$ for all $\bar{x} \in \bar{E}$. In addition, for any $r = r(n)$ and $h = h(n) \rightarrow \infty$ arbitrarily slowly, **whp** at least $dn - h(d-1)^{2r}$ directed edges satisfy $\text{tx}(B_r) = 0$.*

Proof. Clearly, if $\bar{x} = (u, v) \in \bar{E}$ we have $\text{tx}(B_t(\bar{x})) \leq \text{tx}(B_t(v))$ for any t , thus the first statement of the lemma follows immediately from Lemma 2.1.

To show the second statement, recall the exploration process performed in the proof Lemma 2.1, where $A_{i,k}$ denoted the event that the k -th matching generated in the i -th layer already belongs to our exposed neighborhood. In our setting, we perform a similar exploration process on a random $\bar{x} =$

$(u, v) \in \bar{E}$, only this time the initial vertex v corresponds to $d - 1$ points rather than d (having excluded its edge to u). Thus, (2.4) translates into

$$\sum_{i=0}^{t-1} \sum_{k=1}^{m_i} \mathbf{1}_{A_{i,k}} \leq \text{Bin} \left((d-1)^{t+1}, \frac{(d-1)^t}{n} \right).$$

It follows that the probability that $\mathbf{tx}(B_r(\bar{x})) > 0$ is at most $O((d-1)^{2r}/n)$, and the expected number of such $\bar{x} \in \bar{E}$ is $O((d-1)^{2r})$, as required. \blacksquare

The following lemma, which is the analogue of Lemma 3.2, shows that a small burn-in period typically brings the NBRW to a directed L -root for a certain L (and allows us to restrict our attention to such starting positions).

Lemma 4.3. *Let $\varepsilon > 0$, set $K = \lceil \log_{d-1}(2/\varepsilon) \rceil$ and $L = \lfloor \frac{1}{6} \log_{d-1} n \rfloor$. Let $\bar{x} \in \bar{E}$ be such that $\mathbf{tx}(B_{K+L}(\bar{x})) \leq 1$. Then the non-backtracking walk of length K from \bar{x} ends at a directed L -root with probability at least $1 - \varepsilon$.*

Proof. Let H be the subgraph formed by the elements (directed edges) of $B_{K+L}(\bar{x})$, and notice that the L -radius neighborhoods of all possible endpoints \bar{y} of a non-backtracking walk of length K from \bar{x} are all contained in H . Thus, if $\mathbf{tx}(B_{K+L}(\bar{x})) = 0$ then clearly every such endpoint is a directed L -root.

Otherwise, consider the undirected underlying graph of H . This graph contains a single simple cycle C (by the assumption that $\mathbf{tx}(B_{K+L}(\bar{x})) \leq 1$), therefore the distance of any vertex $u \in H$ from C is well defined. Let (\bar{W}_t) denote the non-backtracking random walk started at $\bar{W}_0 = \bar{x}$. For some $1 \leq t < K$, write $\bar{W}_t = (u, v)$ and $\bar{W}_{t+1} = (v, w)$. Crucially, we claim that if $\text{dist}(v, C) < \text{dist}(w, C)$, then \bar{W}_j is a directed L -root for all $j \in \{t+1, \dots, K\}$. Indeed, our subgraph consists of a cycle C with disjoint trees rooted at some of its vertices. Therefore, as soon as the non-backtracking walk makes a single step away from C , by definition it can only traverse further away from C with each additional step (as long as it is in H).

Furthermore, if $v \notin C$ (that is, v belongs to one of the trees rooted on C), then with probability $\frac{1}{d-1}$ the distance to C decreases by 1 in \bar{W}_{t+1} , otherwise it increases by 1. Similarly,

$$\mathbb{P}(w \in C \mid u, v \in C) = 1/(d-1).$$

The remaining case is the *single* step immediately following the first visit to the cycle C , if such exists, where the probability of remaining on C (traversing along one of the two possible directions on it) is $\frac{2}{d-1}$. Altogether,

$$\mathbb{P}_{\bar{x}}(\bar{W}_K \text{ is not a directed } L\text{-root}) \leq 2(d-1)^{-K} \leq \varepsilon,$$

as required. \blacksquare

The next two lemmas are the analogues of Lemmas 3.3 and 3.4 for directed K -roots, and both follow by essentially repeating the original arguments.

Lemma 4.4. *Set $T = \frac{51}{100} \log_{d-1} n$ and $K = K(n)$. Then with probability $1 - o(n^{-3})$, every directed K -root \bar{x} satisfies*

$$|\partial B_t(\bar{x})| \geq \left(1 - (d-1)^{-K} - O(n^{-1/5})\right) (d-1)^t \text{ for all } t \leq T .$$

Lemma 4.5. *Let $\varepsilon > 0$, $T = \frac{51}{100} \log_{d-1} n$ and $L = \lceil \frac{1}{6} \log_{d-1} n \rceil$. With probability $1 - o(n^{-3})$, any two directed L -roots \bar{x} and \bar{y} with $\text{dist}(\bar{x}, \bar{y}) > 2L$ satisfy*

$$|B_t(\bar{x}) \cap B_t(\bar{y})| < n^{-1/7} (d-1)^t \text{ for all } t \leq T .$$

We now turn to prove the Poissonization argument, on which the entire proof of Theorem 2 hinges. Recall that in Theorem 1 we could afford a relatively large (order $\log \log n$) error, which enabled us to apply standard large deviation arguments for the size of cuts between certain neighborhoods of two vertices u, v (as studied in Lemma 3.5). On the other hand, here we can only afford an $O(1)$ error, so the number of paths of length the mixing time between two random vertices will approximately be a Poisson random variable with constant mean. In order to bypass this obstacle and derive the concentration results needed for proving cutoff, we instead consider the joint distribution of u and vertices v_1, \dots, v_M for some large (poly-logarithmic) M . This approach, incorporated in the next proposition, amplifies the error probabilities as required.

Proposition 4.6. *Let $\varepsilon > 0$, set*

$$K = \lceil 2 \log_{d-1}(1/\varepsilon) \rceil , \quad T = \lceil \log_{d-1}(dn) \rceil , \quad \mu = (d-1)^{T+K}/dn ,$$

and for each $\bar{x} \in \bar{E}$, define the random variable $Z = Z(\bar{x})$ by

$$\mathbb{P}(Z = k) = \frac{1}{dn} \left| \left\{ \bar{y} \in \bar{E} : \mathcal{N}_{T+K-1}(\bar{x}, \bar{y}) = k \right\} \right| ,$$

*where $\mathcal{N}_\ell(\bar{x}, \bar{y})$ is the number of ℓ -long non-backtracking paths from \bar{x} to \bar{y} . Then **whp**, every \bar{x} that is a directed L -root for $L = \lceil \frac{1}{6} \log_{d-1}(dn) \rceil$ satisfies*

$$\mathbb{E} \left[\left| \frac{Z(\bar{x})}{\mu} - 1 \right| \mid \mathcal{F}_G \right] < 2\varepsilon + \frac{5}{\log \log n} ,$$

where \mathcal{F}_G is the σ -field generated by the graph $G \sim \mathcal{G}(n, d)$.

Proof. Condition on the statement of Lemma 4.2 for the choices $r(n) = L$ and $h(n) = \log n$. That is, we assume that there are at least $dn - (\log n)n^{1/3}$ directed L -roots in \bar{E} .

Let \bar{x} be a uniformly chosen directed edge, and expose its L -radius neighborhood according to the configuration model. As the statement of the

proposition only refers to directed L -roots, we may at this point assume that \bar{x} is indeed such an edge (recall that the property of being a directed L -root is solely determined by the structure of the induced subgraph on $B_L(\bar{x})$, and thus this conditioning does not affect the distribution of the future pairings). With this assumption in mind, continue exposing the neighborhood of \bar{x} to obtain $B_{2L}(\bar{x})$.

Our goal is to show that

$$\mathbb{P} \left(\mathbb{E} [|(Z(\bar{x})/\mu) - 1| \mid \mathcal{F}_G] \geq 2\varepsilon + \frac{5}{\log \log n} \right) = o(1/n) ,$$

in which case a first moment argument will immediately complete the proof of the proposition.

We next consider a uniformly chosen set of M directed edges, $\mathcal{B} \subset \bar{E}$, for some $\log^2 n \leq M \leq 2 \log^2 n$ (to be specified later), by selecting its elements one by one. That is, after i steps ($0 \leq i < M$), $|\mathcal{B}| = i$ and we add a directed edge uniformly chosen over the $dn - i$ remaining elements of \bar{E} . With the addition of every new element, we also develop its $2L$ -radius neighborhood.

Notice that, after i steps, there are at most $(\log n)n^{1/3}$ directed edges which are *not* directed L -roots in \bar{E} , and furthermore,

$$|B_{2L}(\bar{x}) \cup (\cup_{\bar{y} \in \mathcal{B}} B_{2L}(\bar{y}))| \leq (i+1)n^{1/3} \leq Mn^{1/3} .$$

Therefore, the probability that the $(i+1)$ -th element of \mathcal{B} either belongs to one of the existing $2L$ -radius neighborhoods, or is *not* a directed L -root, is at most $2Mn^{-2/3}$. Clearly, the probability that 4 such ‘‘bad’’ edges are selected is at most $O(M^4 n^{-8/3}) = o(n^{-2})$.

Altogether, we may assume with probability $1 - o(n^{-2})$, the set \mathcal{B} contains a subset $\mathcal{B}' = \{\bar{y}_1, \dots, \bar{y}_{M'}\}$ of size $M' \geq M - 3$, such that the following holds:

- (i) Every member of $\{\bar{x}\} \cup \mathcal{B}'$ is an L -root.
- (ii) The pairwise distances of $\{\bar{x}\} \cup \mathcal{B}'$ all exceed $2L$.

For any $\bar{y} \in \bar{E}$, let $Z_{\bar{y}} = \mathcal{N}_{T+K-1}(\bar{x}, \bar{y})$, and for any $S \subset \bar{E}$, let Z_S be the random variable that accepts the value $Z_{\bar{y}}$ with probability $1/|S|$ for each $\bar{y} \in S$. We will use an averaging argument to show that Z can be well approximated by $Z_{\mathcal{B}}$, which in turn is well approximated by $Z_{\mathcal{B}'}$.

Setting

$$T_1 = \lfloor (T+K)/2 \rfloor , \quad T_2 = \lceil (T+K)/2 \rceil - 2 ,$$

we wish to develop the T_1 -radius neighborhood of \bar{x} as well as the T_2 -radius neighborhoods of every $\bar{y} \in \mathcal{B}'$. To this end, put

$$\begin{aligned} U &\triangleq \partial B_{T_1}(\bar{x}) , & V_i &\triangleq \partial B_{T_2}(\bar{y}_i) , \\ \tilde{U} &\triangleq U \setminus \cup_i B_{T_2}(\bar{y}_i) , & \tilde{V}_i &\triangleq V_i \setminus (B_{T_1}(\bar{x}) \cup (\cup_{j \neq i} B_{T_2}(\bar{y}_j))) . \end{aligned}$$

Recalling Lemma 4.4 (and the fact that $\{\bar{x}\} \cup \mathcal{B}'$ are all directed L -roots), with probability $1 - o(n^{-3})$ we have

$$\begin{aligned} |U| &\geq \left(1 - O(n^{-\frac{1}{5}})\right) (d-1)^{T_1}, \\ |V_i| &\geq \left(1 - O(n^{-\frac{1}{5}})\right) (d-1)^{T_2} \quad \text{for all } i \in [M']. \end{aligned}$$

Combining this with Lemma 4.5, we deduce that for any sufficiently large n the following holds with probability $1 - o(n^{-3})$:

$$\begin{aligned} (1 - 2n^{-\frac{1}{7}}) (d-1)^{T_1} &\leq |\tilde{U}| \leq (d-1)^{T_1}, \\ (1 - 2n^{-\frac{1}{7}}) (d-1)^{T_2} &\leq |\tilde{V}_i| \leq (d-1)^{T_2} \quad \text{for all } i \in [M']. \end{aligned}$$

We will use a standard Poissonization approach in order to approximate the joint distribution of the variables $\{Z_{\bar{y}} : \bar{y} \in \mathcal{B}'\}$ (that are fully determined by the graph G) using the following set of variables:

$$\tilde{Z}_{\bar{y}_i} \triangleq \left| \{u, v \in E : u \in \tilde{U}, v \in \tilde{V}_i\} \right| \quad (i \in [M']).$$

We claim that $\tilde{Z}_{\bar{y}_i} \leq Z_{\bar{y}_i}$ for all i . To see this, recall that $Z_{\bar{y}_i}$ counts the number of non-backtracking paths of length $T+K-1$ from \bar{x} to \bar{y}_i . Since \tilde{U} and \tilde{V}_i are disjoint subsets of the boundaries of the T_1 -radius neighborhood U and the T_2 -radius neighborhood V_i respectively, every edge between them corresponds to at least one distinct such path of length $T_1+T_2+1 = T+K-1$. Therefore, by the triangle inequality,

$$\begin{aligned} \mathbb{M}\mathbb{E} \left[\left| \frac{Z_{\mathcal{B}}}{\mu} - 1 \right| \middle| \mathcal{F}_G \right] &= \sum_{\bar{y} \in \mathcal{B}} \left| \frac{Z_{\bar{y}}}{\mu} - 1 \right| \leq \sum_{i=1}^{M'} \left| \frac{Z_{\bar{y}_i}}{\mu} - 1 \right| + \sum_{\bar{y} \in \mathcal{B} \setminus \mathcal{B}'} \frac{Z_{\bar{y}}}{\mu} + |\mathcal{B} \setminus \mathcal{B}'| \\ &\leq \sum_{i=1}^{M'} \left(\left| \frac{\tilde{Z}_{\bar{y}_i}}{\mu} - 1 \right| + \frac{Z_{\bar{y}_i} - \tilde{Z}_{\bar{y}_i}}{\mu} \right) + \sum_{\bar{y} \in \mathcal{B} \setminus \mathcal{B}'} \frac{Z_{\bar{y}}}{\mu} + 3 \\ &= \sum_{i=1}^{M'} \left| \frac{\tilde{Z}_{\bar{y}_i}}{\mu} - 1 \right| - \sum_{i=1}^{M'} \frac{\tilde{Z}_{\bar{y}_i}}{\mu} + \sum_{\bar{y} \in \mathcal{B}} \frac{Z_{\bar{y}}}{\mu} + 3. \end{aligned} \tag{4.1}$$

Let $\tilde{\mathcal{Z}}$ denote the first summand in the last expression:

$$\tilde{\mathcal{Z}} \triangleq \sum_{i=1}^{M'} \left| \frac{\tilde{Z}_{\bar{y}_i}}{\mu} - 1 \right|.$$

The following lemma estimates $\tilde{\mathcal{Z}}$, as well as the second summand in (4.1).

Lemma 4.7. *Define $\tilde{\mathcal{Z}}$ and $\tilde{Z}_{\bar{y}_i}$ for $i = 1, \dots, M'$ as above. Then:*

$$\mathbb{P} \left(\tilde{\mathcal{Z}} > \varepsilon + \frac{4}{\log \log n} \right) = o(n^{-2}), \tag{4.2}$$

and

$$\mathbb{P} \left(\frac{1}{M} \sum_{i=1}^{M'} \frac{\tilde{Z}_{\bar{y}_i}}{\mu} \leq 1 - \varepsilon - \frac{1}{\log \log n} \right) = o(n^{-2}) . \quad (4.3)$$

Proof. We claim that, with probability $1 - o(n^{-2})$, each of the variables $\tilde{Z}_{\bar{y}_i}$ is stochastically dominated from below and from above by i.i.d. pairs of binomial variables, $R_i^- \leq R_i^+$ (coupled in the obvious manner), defined as:

$$\begin{aligned} R_i^- &\sim \text{Bin} \left((1 - n^{-\frac{1}{8}})(d-1)^{T_2+1}, p^- \right), & p^- &\triangleq (1 - n^{-\frac{1}{8}}) \frac{(d-1)^{T_1+1}}{dn}, \\ R_i^+ &\sim \text{Bin} \left((d-1)^{T_2+1}, p^+ \right), & p^+ &\triangleq (1 + n^{-\frac{1}{4}}) \frac{(d-1)^{T_1+1}}{dn}, \\ \Delta_i &\triangleq R_i^+ - R_i^- \geq 0. \end{aligned}$$

To see this, consider the configuration model at the starting phase where the vertices in $\tilde{U} \cup (\cup_i \tilde{V}_i)$ all have degree 1 (that is, each of these vertices comprise $(d-1)$ points that still wait to be paired), and expose the pairings of the points in \tilde{V}_i sequentially. Suppose that for all $j < i$ we have already constructed a coupling where $R_j^- \leq \tilde{Z}_{\bar{y}_j} \leq R_j^+$, and next wish to do the same for $\tilde{Z}_{\bar{y}_i}$.

By Lemma 4.5, with probability $1 - o(n^{-3})$ there still remain at least $(1 - n^{-1/8})(d-1)^{T_2}$ vertices of degree 1 in \tilde{V}_i and at least $(1 - n^{-1/8})(d-1)^{T_1}$ such vertices in \tilde{U} (otherwise the intersection of either $B(\bar{y}_i)$ or $B(\bar{x})$ with one of $B(\bar{y}_1), \dots, B(\bar{y}_{i-1})$ would contain at least $n^{-1/7}(d-1)^{T_1}$ vertices). We thus have at least $(1 - n^{-1/8})(d-1)^{T_2+1}$ unmatched points corresponding to \tilde{V}_i , and at least $(1 - n^{-1/8})(d-1)^{T_1+1}$ unmatched points corresponding to \tilde{U} . Associating each such point corresponding to \tilde{V}_i with a Bernoulli variable, which succeeds if and only if it is matched to \tilde{U} , clearly establishes the coupling of $\tilde{Z}_{\bar{y}_i} \geq R_i^-$.

Conversely, $\tilde{V}_i \leq (d-1)^{T_2}$ and $\tilde{U} \leq (d-1)^{T_1}$, hence there are at most $(d-1)^{T_2+1}$ unmatched points corresponding to \tilde{V}_i and at most $(d-1)^{T_1+1}$ unmatched points corresponding to \tilde{U} . Since both the T_1 -radius and the T_2 -radius neighborhoods of any element contains $O(\sqrt{n})$ distinct vertices, the probability of a point corresponding to \tilde{V}_i being matched to \tilde{U} is at most

$$\frac{(d-1)^{T_1+1}}{dn - O(M\sqrt{n})} \leq \frac{(d-1)^{T_1+1}}{(1 - o(n^{-1/4}))dn} .$$

Therefore, we can readily construct the coupling $\tilde{Z}_{\bar{y}_i} \leq R_i^+$.

Since it was possible to construct each of the above couplings with probability $1 - o(n^{-3})$, clearly all M' variables can be coupled as above with probability $1 - o(n^{-2})$.

Finally, consider a set of i.i.d. binomial random variables Q_i with means $\mathbb{E}Q_1 = \mu = (d-1)^{T+K}/dn$, defined by

$$Q_i \sim \text{Bin} \left((d-1)^{T_2+1}, \frac{(d-1)^{T_1+1}}{dn} \right),$$

and coupled in the obvious manner such that $R_i^- \leq Q_i \leq R_i^+$. Clearly, as $|\tilde{Z}_{y_i} - Q_i| \leq R_i^+ - R_i^- = \Delta_i$, it follows that

$$\tilde{Z} = \frac{1}{M'} \sum_{i=1}^{M'} \left| \frac{\tilde{Z}_{v_i}}{\mu} - 1 \right| \leq \frac{1}{M'} \sum_{i=1}^{M'} \left| \frac{Q_i}{\mu} - 1 \right| + \frac{1}{M'} \sum_{i=1}^{M'} \frac{\Delta_i}{\mu}. \quad (4.4)$$

Since $\mu \geq (d-1)^K \geq 1/\varepsilon^2$, for all $i \in [M']$ we have

$$\begin{aligned} \mathbb{E} \left| \frac{Q_i}{\mu} - 1 \right| &\leq \frac{1}{\mu} \sqrt{\text{Var}(Q_i)} = \frac{1 + O(n^{-\frac{1}{4}})}{\sqrt{\mu}} \leq \left(1 + \frac{1}{\log n} \right) \varepsilon, \\ \frac{\mathbb{E}\Delta_i}{\mu} &\leq (1 - n^{-\frac{1}{8}}) \left(n^{-\frac{1}{4}} + n^{-\frac{1}{8}} \right) + n^{-\frac{1}{8}} \left(1 + n^{-\frac{1}{4}} \right) = O\left(n^{-\frac{1}{4}} \right). \end{aligned}$$

where the last inequalities in both estimates hold for any sufficiently large n . Furthermore, since the $\{Q_i\}$ -s are i.i.d. binomial variables, Chernoff's inequality implies that

$$\mathbb{P} \left(\frac{1}{M'} \sum_{i=1}^{M'} \frac{Q_i}{\mu} > 1 + \frac{1}{\log \log n} \right) < e^{-\frac{\mu M'}{4(\log \log n)^2}} = e^{-\Omega\left(\left(\frac{\log n}{\log \log n}\right)^2\right)} = o(n^{-2}), \quad (4.5)$$

and an analogous argument for the $\{\Delta_i\}$ -s (recall that by definition, we have $\Delta_i = \Delta'_i + \Delta''_i$, where the $\{\Delta'_i\}$ -s and $\{\Delta''_i\}$ -s are two sequences of i.i.d. binomial variables, independent of each other), combined with the fact that $\mathbb{E}\Delta_i/\mu = O(n^{-1/4})$, gives

$$\mathbb{P} \left(\frac{1}{M'} \sum_{i=1}^{M'} \frac{\Delta_i}{\mu} > \frac{1}{\log \log n} \right) \leq e^{-\Omega\left(\left(\frac{\log n}{\log \log n}\right)^2\right)} = o(n^{-2}). \quad (4.6)$$

Define

$$X_t \triangleq \sum_{i=1}^t \left| \frac{Q_i}{\mu} - 1 \right| - \left(\frac{Q_i}{\mu} - 1 \right) - \mathbb{E} \left| \frac{Q_i}{\mu} - 1 \right|.$$

Since $\mathbb{E} |(Q_i/\mu) - 1| \leq (1 + \frac{1}{\log n})\varepsilon < 2$ for large n (with room to spare), we deduce that X_t is a martingale with bounded increments:

$$|X_{t+1} - X_t| \leq 2 + \mathbb{E} \left| \frac{Q_i}{\mu} - 1 \right| \leq 4.$$

Therefore, Azuma's inequality (cf., e.g., [6, Chapter 7.2]) implies that

$$\mathbb{P} \left(X_{M'}/M' > \frac{1}{\log \log n} \right) < e^{-\frac{1}{2}M'/(4 \log \log n)^2} = o(n^{-2}). \quad (4.7)$$

Since $\mathbb{E} |(Q_1/\mu) - 1| < (1 + \frac{1}{\log n})\varepsilon$ and

$$\frac{1}{M'} \sum_{i=1}^{M'} \left| \frac{Q_i}{\mu} - 1 \right| = \mathbb{E} \left| \frac{Q_1}{\mu} - 1 \right| + (X_{M'}/M') + \frac{1}{M'} \sum_{i=1}^{M'} \left(\frac{Q_i}{\mu} - 1 \right),$$

the bounds in (4.5) and (4.7) now imply that

$$\mathbb{P} \left(\frac{1}{M'} \sum_{i=1}^{M'} \left| \frac{Q_i}{\mu} - 1 \right| > \varepsilon + \frac{3}{\log \log n} \right) = o(n^{-2}).$$

Together with (4.4) and (4.6), we obtain that (4.2) indeed holds.

Similarly, since $\tilde{Z}_{\tilde{y}_i} \geq R_i^-$ for all i , and the $\{R_i^-\}$ -s are i.i.d. binomial variables with $\mathbb{E} R_i^- \geq (1 - \varepsilon - 3n^{-1/8})\mu$, we can apply Chernoff's inequality to derive a lower bound on $\sum_{i=1}^{M'} (\tilde{Z}_{\tilde{y}_i}/\mu)$. Keeping in mind that

$$\frac{1}{M} \sum_{i=1}^{M'} \frac{\tilde{Z}_{\tilde{y}_i}}{\mu} \geq \left(1 - \frac{3}{M} \right) \sum_{i=1}^{M'} \frac{\tilde{Z}_{\tilde{y}_i}}{\mu},$$

we obtain that (4.3) holds, as

$$\mathbb{P} \left(\frac{1}{M} \sum_{i=1}^{M'} \frac{\tilde{Z}_{\tilde{y}_i}}{\mu} \leq 1 - \varepsilon - \frac{1}{\log \log n} \right) \leq e^{-\Omega\left(\left(\frac{\log n}{\log \log n}\right)^2\right)} = o(n^{-2}).$$

This completes the proof of Lemma 4.7. \blacksquare

We can now combine (4.2) and (4.3) with (4.1), and deduce that the following statement holds with probability $1 - o(n^{-2})$:

$$\mathbb{E} \left[\left| \frac{Z_{\mathcal{B}}}{\mu} - 1 \right| \mid \mathcal{F}_G \right] \leq 2\varepsilon - 1 + \frac{5}{\log \log n} + \frac{1}{M} \sum_{\tilde{y} \in \mathcal{B}} \frac{Z_{\tilde{y}}}{\mu}. \quad (4.8)$$

To transform the above into the required bound on Z , take $M = \lceil \log^2 n \rceil$, and consider a collection of bins, each of size either M or $M + 1$, such that the total of their sizes is dn . Let $\mathcal{B}'_1, \dots, \mathcal{B}'_{\ell_1}$ denote the M -element bins, and let $\mathcal{B}''_1, \dots, \mathcal{B}''_{\ell_2}$ denote the $(M + 1)$ -element bins. Next, randomly partition the elements of \bar{E} into these bins (i.e., each bin \mathcal{B} will contain a uniformly chosen set of $|\mathcal{B}|$ directed edges).

Since there are at most $\lfloor dn/M \rfloor = O(n/M)$ different bins, and for each bin the corresponding $Z_{\mathcal{B}}$ satisfies (4.8) with probability at least $1 - o(n^{-2})$, we deduce that all the variables $Z_{\mathcal{B}'_j}$ and $Z_{\mathcal{B}''_j}$ satisfy this inequality with

probability at least $1 - o(1/n)$. Therefore, with probability at least $1 - o(1/n)$,

$$\begin{aligned} \mathbb{E} \left[\left| \frac{Z}{\mu} - 1 \right| \middle| \mathcal{F}_G \right] &= \frac{1}{dn} \sum_{\bar{y} \in \bar{E}} \left| \frac{Z_{\bar{y}}}{\mu} - 1 \right| \\ &= \frac{M}{dn} \sum_{j=1}^{\ell_1} \mathbb{E} \left[\left| \frac{Z_{\mathcal{B}'_j}}{\mu} - 1 \right| \middle| \mathcal{F}_G \right] + \frac{M+1}{dn} \sum_{j=1}^{\ell_2} \mathbb{E} \left[\left| \frac{Z_{\mathcal{B}''_j}}{\mu} - 1 \right| \middle| \mathcal{F}_G \right] \\ &\leq 2\varepsilon - 1 + \frac{5}{\log \log n} + \frac{1}{dn} \sum_{\bar{y} \in \bar{E}} \frac{Z_{\bar{y}}}{\mu} = 2\varepsilon + \frac{5}{\log \log n} , \end{aligned}$$

where the last equality follows from the fact that

$$\sum_{\bar{y}} Z_{\bar{y}} = \sum_{\bar{y}} \mathcal{N}_{T+K-1}(\bar{x}, \bar{y}) = (d-1)^{T+K} = \mu dn .$$

This completes the proof. ■

Proof of Theorem 2. Let (\bar{W}_t) be the non-backtracking random walk, and let π denote the stationary distribution on \bar{E} .

The lower bound is a consequence of the following simple claim:

Claim 4.8. *Every d -regular graph on n vertices satisfies*

$$t_{\text{MIX}}(1 - \varepsilon) \geq \lceil \log_{d-1}(dn) \rceil - \lceil \log_{d-1}(1/\varepsilon) \rceil \quad \text{for any } 0 < \varepsilon < 1 .$$

Proof of claim. Let $\varepsilon > 0$ and let $\bar{x}_0 \in \bar{E}$ be any starting position. Clearly, at time $T = \lceil \log_{d-1}(\varepsilon dn) \rceil$ we have

$$|\partial B_T(\bar{x}_0)| \leq (d-1)^T \leq \varepsilon dn ,$$

and the set $A \triangleq \bar{E} \setminus \partial B_T(\bar{x}_0)$ has stationary measure at least $1 - \varepsilon$. Thus,

$$\|\mathbb{P}_{\bar{x}_0}(\bar{W}_T \in \cdot) - \pi\|_{\text{TV}} \geq |\mathbb{P}_{\bar{x}_0}(\bar{W}_T \in A) - \pi(A)| \geq 1 - \varepsilon ,$$

implying that $t_{\text{MIX}}(1 - \varepsilon) > T$. The proof now follows from the fact that

$$\begin{aligned} \lceil \log_{d-1}(dn) \rceil - \lceil \log_{d-1}(1/\varepsilon) \rceil &= \lceil \log_{d-1}(dn) \rceil + \lceil \log_{d-1} \varepsilon \rceil \\ &\leq \lceil \log_{d-1}(\varepsilon dn) \rceil \leq T + 1 . \end{aligned} \quad \blacksquare$$

For the upper bound, let \bar{x}_0 be the worst starting position, and let $\bar{x} = \bar{W}_{t_0}$, where $t_0 = \lceil \log_{d-1}(2/\varepsilon) \rceil$. Let LR denote the event that \bar{x} is a directed L -root, where $L = \lceil \frac{1}{6} \log_{d-1}(dn) \rceil$. Conditioning on the statements of Lemma 4.2 and Lemma 4.3 (and recalling that both hold **whp**) we obtain that $\mathbb{P}_{\bar{x}_0}(\text{LR}) \geq 1 - \varepsilon$.

Condition on the statement of Proposition 4.6, and following its notation, let $Z(\bar{x})$ accept the value $\mathcal{N}_{T+K-1}(\bar{x}, \bar{y})$ with probability $1/dn$, where

$$K = \lceil 2 \log_{d-1}(1/\varepsilon) \rceil , \quad T = \lceil \log_{d-1}(dn) \rceil , \quad \mu = (d-1)^{T+K} / dn .$$

The following then holds:

$$\begin{aligned}
& \sum_{\bar{y} \in \bar{E}} \left| \mathbb{P}_{\bar{x}}(\bar{W}_{T+K} = \bar{y} \mid \text{LR}) - \frac{1}{dn} \right| \\
&= \sum_k |\{\bar{y} : \mathcal{N}_{T+K-1}(\bar{x}, \bar{y}) = k\}| \left| \frac{k}{(d-1)^{T+K}} - \frac{1}{dn} \right| \\
&= \sum_k \mathbb{P}(Z = k \mid \mathcal{F}_G) \left| \frac{k}{\mu} - 1 \right| = \mathbb{E} [|(Z/\mu) - 1| \mid \mathcal{F}_G] \leq 2\varepsilon + o(1),
\end{aligned} \tag{4.9}$$

where in the last inequality we applied Proposition 4.6 onto the directed L -root \bar{x} (given the event LR). We deduce that for $t(\varepsilon) = t_0 + T + K$:

$$\begin{aligned}
\left\| \mathbb{P}_{\bar{x}_0}(\bar{W}_t \in \cdot) - \pi \right\|_{\text{TV}} &= \frac{1}{2} \sum_{\bar{y} \in \bar{E}} \left| \mathbb{P}_{\bar{x}_0}(\bar{W}_t = \bar{y}) - \frac{1}{dn} \right| \\
&\leq \frac{1}{2} \mathbb{P}_{\bar{x}_0}(\text{LR}) \sum_{\bar{y} \in \bar{E}} \left| \mathbb{P}_{\bar{x}_0}(\bar{W}_t = \bar{y} \mid \text{LR}) - \frac{1}{dn} \right| + \mathbb{P}_{\bar{x}_0}(\text{LR}^c) \\
&\leq \varepsilon + (1 - \varepsilon) \mathbb{P}_{\bar{x}_0}(\text{LR}^c) + o(1) \leq 2\varepsilon - \varepsilon^2 + o(1) < 2\varepsilon,
\end{aligned} \tag{4.10}$$

where the first inequality in the last line is by (4.9), the second one is due to the fact that $\mathbb{P}(\text{LR}^c) \leq \varepsilon$, and the third inequality holds for sufficiently large values of n . Therefore, for any large n we have

$$\begin{aligned}
t_{\text{MIX}}(\varepsilon) &\leq t(\varepsilon/2) \leq \lceil \log_{d-1}(dn) \rceil + 3 \lceil \log_{d-1}(2/\varepsilon) \rceil + \lceil \log_{d-1} 2 \rceil \\
&\leq \lceil \log_{d-1}(dn) \rceil + 3 \lceil \log_{d-1}(1/\varepsilon) \rceil + 4
\end{aligned}$$

(where in the last inequality we used the fact that $d \geq 3$), as required. \blacksquare

5. CUTOFF FOR RANDOM REGULAR GRAPHS OF LARGE DEGREE

In this section, we prove Theorem 3 and Corollary 4, which extend our cutoff result for the SRW and NBRW on almost every random regular graph of fixed degree $d \geq 3$ to the case of d large. To prove cutoff for the NBRW, we adapt our original arguments (from the case of d fixed) to the new delicate setting where our error probabilities are required to be exponentially small in d . The behavior of the SRW is then obtained as a corollary of this result.

Throughout the section, let $d = d(n) \rightarrow \infty$ with n , and recall that we further assume that $d = n^{o(1)}$, since otherwise the the mixing time is $O(1)$ and cutoff is impossible.

5.1. NBRWs on random regular graphs of large degree. As we will soon show, when d is large we no longer need to deal with K -roots (and the locally-tree-like geometry of the starting point of our walk), as all vertices

will have sufficient expansion **whp**. However, the analysis of the configuration model becomes more delicate, as the probability that it produces a simple graph is $(1 + o(1)) \exp\left(\frac{1-d^2}{4}\right)$ (see (2.1)), which now decays with n . Thus, to prove that the probability of an event goes to 0 on $\mathcal{G}(n, d)$, we must now show that its probability is $o(\exp(-d^2/4))$ in the configuration model.

Lemma 5.1. *With high probability, for all $\bar{x} \in \bar{E}$ and all $t \leq \frac{4}{7} \log_{d-1} n$,*

$$|\partial B_t(\bar{x})| \geq (1 - o(1))(d - 1)^t . \tag{5.1}$$

Proof. The proof is an adaption of Lemma 3.3. Pick a directed edge \bar{x} uniformly at random and expose its first level. Since we are interested in probabilities conditioned on the graph G being simple, we may assume that $|\partial B_1(\bar{x})| = d - 1$, that is, there are no self-loops or multiple edges from \bar{x} .

We will show that (5.1) holds with probability $1 - o(n^{-1} \exp(-d^2/4))$ for the above \bar{x} in the configuration model. Clearly, for any $t < t'$ we have $|\partial B_t(\bar{x})| \geq (d-1)^{t'-t} |\partial B_{t'}(\bar{x})|$, hence we can restrict our attention to $\partial B_T(\bar{x})$ where $T = \lfloor \frac{4}{7} \log_{d-1} n \rfloor$.

Following the notation in the proof of Lemma 2.1, let $A_{i,k}$ be the event that, in the process of sequentially matching points, the newly exposed pair of the k -th unmatched point in ∂B_i belongs to some vertex already in B_{i+1} . Further recall that, by (2.3) and the discussion thereafter, the number of events $\{A_{i,k} : 0 \leq i < T\}$ that occur is stochastically dominated by a binomial variable with parameters $\text{Bin}\left((d-1)^{T+1}, \frac{(d-1)^T}{n}\right)$. By our choice of T , the expectation of this random variable is

$$(d-1)^{2T+1}/n \leq dn^{1/7} \leq n^{1/7+o(1)} ,$$

hence the number of events $A_{i,k}$ with $0 \leq i < T$ that occur is less than $n^{1/6}$ (with room to spare) with probability at least $1 - \exp(-\Omega(n^{1/6}))$. Next, set

$$L = \lfloor \frac{1}{5} \log_{d-1} n \rfloor , \quad \rho = \lceil 4 + 2d^2 / \log n \rceil = o(d^2) .$$

As before, we can stochastically dominate the number of events $A_{i,k}$ that occur in the first L levels, $\{A_{i,k} : 0 \leq i < L\}$, by a binomial variable $X_L \sim \text{Bin}\left((d-1)^{L+1}, \frac{(d-1)^L}{n}\right)$. Since the expected value of X_L is

$$(d-1)^{2L+1}/n = o(n^{-1/2}) ,$$

and since $L \rightarrow \infty$ with n (by our assumption on d), it is easy to verify that

$$\mathbb{P}(X_L \geq \rho) = (1 + o(1))\mathbb{P}(X_L = \rho) = o(n^{-\rho/2}) .$$

Recalling the definition of ρ , it now follows that the number of events $A_{i,k}$ with $0 \leq i < L$ that occur is less than ρ except with probability $o(n^{-2}e^{-d^2})$.

Each event $A_{i,k}$ reduces the number of leaves in level $i + 1$ by at most 2, hence it reduces the number of leaves in level $t > i$ by at most $2(d - 1)^{t-i-1}$ vertices. It then follows that for each $0 \leq t < T$,

$$|\partial B_t(\bar{x})| \geq (d - 1)^t - \sum_{i < t} \sum_k \mathbf{1}_{A_{i,k}} 2(d - 1)^{t-i-1}. \quad (5.2)$$

As $|\partial B_1(\bar{x})| = d - 1$, there are no events of the form $A_{0,k}$. Therefore, by the discussion above, with probability $1 - o(n^{-2}e^{-d^2})$ we have

$$\sum_{i < L} \sum_k \mathbf{1}_{A_{i,k}} 2(d - 1)^{t-i-1} \leq 2(d - 1)^{t-2} \rho = o((d - 1)^t).$$

Furthermore, by the above discussion on the number of events $A_{i,k}$ that occur, we deduce that with probability at least $1 - \exp(-\Omega(n^{1/6}))$

$$\sum_{i=L}^{t-1} \sum_k \mathbf{1}_{A_{i,k}} 2(d - 1)^{t-i-1} \leq 2(d - 1)^{t-L-1} n^{1/6} = o((d - 1)^t).$$

Plugging the above in (5.2), we obtain that with probability $1 - o(n^{-2}e^{-d^2})$

$$|\partial B_t(\bar{x})| \geq (1 - o(1))(d - 1)^t, \quad (5.3)$$

and a union bound implies that (5.3) holds for all directed edges \bar{x} which satisfy $|\partial B_1(\bar{x})| = d - 1$ except with probability $O(\frac{d}{n} \exp(-d^2)) = o(\exp(-d^2))$. By (2.1), it now follows that (5.1) also holds **whp** over $\mathcal{G}(n, d)$. \blacksquare

The following lemma, the analogue of Lemma 3.4, is proved by essentially following the same argument as in the proof of Lemma 3.4, i.e., calculating the size of the common neighborhood of two vertices. The difference is again that here we need to deal with the fact that the probability that the configuration model is a simple graph is exponentially small in d . This is achieved by repeating the approach, demonstrated in Lemma 5.1 above, of treating $B_1(\bar{x})$ separately. Applying this analysis to the neighborhoods of the 2 starting directed edges \bar{x}, \bar{y} gives the required result, with the remaining arguments of Lemma 3.4 left unchanged (we omit the full details).

Lemma 5.2. *Set $T = \frac{51}{100} \log_{d-1} n$ and $L = \frac{1}{6} \log_{d-1} n$. Then **whp**, for every $\bar{x}, \bar{y} \in \bar{E}$ with $\text{dist}(\bar{x}, \bar{y}) > 2L$ and every $t \leq T$,*

$$|B_t(\bar{x}) \cup B_t(\bar{y})| \geq n^{-1/7} (d - 1)^t.$$

The final ingredient needed is the analogue of the Poissonization result of Proposition 4.6, as given by the following proposition.

Proposition 5.3. *Let $\varepsilon > 0$, set*

$$T = \lceil \log_{d-1}(dn) + 2 \log_{d-1}(1/\varepsilon) \rceil, \quad \mu = (d - 1)^T / dn,$$

and for each $\bar{x} \in \bar{E}$, define the random variable $Z = Z(\bar{x})$ by

$$\mathbb{P}(Z = k) = \frac{1}{dn} \left| \{ \bar{y} \in \bar{E} : \mathcal{N}_{T-1}(\bar{x}, \bar{y}) = k \} \right| ,$$

where $\mathcal{N}_\ell(\bar{x}, \bar{y})$ is the number of ℓ -long non-backtracking paths from \bar{x} to \bar{y} . Then **whp**, every \bar{x} satisfies

$$\mathbb{E} \left[\left| (Z(\bar{x})/\mu) - 1 \right| \mid \mathcal{F}_G \right] < 2\varepsilon + \frac{5}{\log \log n} ,$$

where \mathcal{F}_G is the σ -field generated by the graph $G \sim \mathcal{G}(n, d)$.

The proof of the above proposition is essentially the same as the proof of Proposition 4.6, with some minor adjustments to the estimates to ensure that they hold with probability $o(\exp(-d^2/4))$. The main necessary change is to let the bin sizes depend on d , namely to set $M = d^3 \log^2 n$. As only minor adjustments to some of the bounds are required elsewhere, we omit the details.

Proof of Theorem 3. The lower bound of $t_{\text{MIX}}(s) \geq \lceil \log_{d-1}(dn) \rceil$ follows immediately from Claim 4.8, whose proof remains valid without change, even when d is allowed to grow with n .

To obtain the upper bound, let (\bar{W}_t) denote the non-backtracking random walk started at $\bar{W}_0 = \bar{x}$. Set $\varepsilon = 3s$, and

$$T = \lceil \log_{d-1}(dn) + 2 \log_{d-1}(1/\varepsilon) \rceil , \quad \mu = (d-1)^T / dn .$$

By Proposition 5.3 we have that **whp**,

$$\begin{aligned} & \sum_{\bar{y} \in \bar{E}} \left| \mathbb{P}_{\bar{x}}(\bar{W}_T = \bar{y}) - \frac{1}{dn} \right| \\ &= \sum_k |\{ \bar{y} : \mathcal{N}_{T-1}(\bar{x}, \bar{y}) = k \}| \left| \frac{k}{(d-1)^T} - \frac{1}{dn} \right| \\ &= \sum_k \mathbb{P}(Z = k \mid \mathcal{F}_G) \left| \frac{k}{\mu} - 1 \right| \\ &= \mathbb{E} \left[\left| (Z/\mu) - 1 \right| \mid \mathcal{F}_G \right] \leq 2\varepsilon + o(1) \leq s \end{aligned}$$

for large n . We conclude that $t_{\text{MIX}}(s) \leq T \leq \lceil \log_{d-1}(dn) \rceil + 1$, since $\log_{d-1}(1/\varepsilon) = o(1)$ by our assumption on d . \blacksquare

5.2. Duality between non-backtracking and simple random walks.

The following observation is attributed to Yuval Peres:

Observation 5.4. *Conditioning on being in level k of the simple random walk on the tree, we are uniform over k -long non-backtracking random walks.*

More specifically, let \mathcal{T} be the cover tree for G at u with a map φ , as defined in (2.5). Let X_t be a SRW on \mathcal{T} started from ρ and let $W_t = \varphi(X_t)$ be the corresponding SRW on G started from u . Compare this with a NBRW random walk \overline{W}_t started from $\bar{x} = (w, u)$ where w is chosen uniformly from the neighbors of u . For a directed edge (y, z) let $\psi(\cdot)$ denote the projection $\psi((y, z)) = z$, giving the vertex the NBRW is presently situated at.

Note that, by symmetry, conditioned on $\text{dist}(\rho, X_t) = k$ the random walk is uniform on the $d(d-1)^{k-1}$ points $\{w \in \mathcal{T} : \text{dist}(\rho, w) = k\}$. By the obvious one-to-one correspondence between paths of length k from ρ in \mathcal{T} and non-backtracking paths of length k in G from u , the following holds: Conditioned on $\text{dist}(\rho, X_t) = k$ we have that W_t is distributed as $\psi(\overline{W}_k)$. Thus, if \overline{W}_t is mixed at time k then a SRW will be mixed once its lift to the cover tree reaches distance k from the root.

Proof of Corollary 4. In our proof of Theorem 1, it was shown using the Central Limit Theorem (see equation (3.4)) that the distance from the root of the walk in the cover tree is given by

$$\frac{\text{dist}(X_t, \rho) - \frac{(d-2)t}{d}}{\frac{2\sqrt{d-1}}{d}\sqrt{t}} \xrightarrow{d} N(0, 1). \quad (5.4)$$

When d grows with n this Gaussian approximation still holds provided the variance satisfies $\frac{2\sqrt{d-1}}{d}\sqrt{t} \rightarrow \infty$ or equivalently $(t/d) \rightarrow \infty$. When d and t are of the same order, the number of backtracking steps is asymptotically a Poisson random variable with mean (t/d) , therefore $(t - \text{dist}(X_t, \rho))$ is distributed as twice a $\text{Po}(t/d)$ random variable. In both of these cases, whenever t has order $\log_{d-1} n$, the variance of $\text{dist}(X_t, \rho)$ is of order $\frac{\log n}{d \log d}$. Finally, when $t/d \rightarrow 0$, the number of backtracking steps goes to 0 as well. This understanding of $\text{dist}(X_t, \rho)$ will allow us to translate our results on NBRWs into statements on SRWs.

If $w \in \mathcal{T}$ and $\text{dist}(\rho, w) \leq R$ then $\varphi(w) \in B_R$ and hence,

$$\|\mathbb{P}(W_t \in \cdot) - \pi\|_{\text{TV}} \geq \mathbb{P}(W_t \in B_R) - \pi(B_R) \geq \mathbb{P}(\text{dist}(X_t, \rho) \leq R) - \pi(B_R).$$

In particular, as $|B_R| \leq O(\frac{n}{d-1}) = o(1)$ for $R \leq \log_{d-1}(n) - 1$, we have that

$$\|\mathbb{P}(W_t \in \cdot) - \pi\|_{\text{TV}} \geq \mathbb{P}(\text{dist}(X_t, \rho) \leq \log_{d-1}(n) - 1) - o(1). \quad (5.5)$$

Next, let $\varrho_k = d_{\text{TV}}(\overline{W}_k, \pi)$ be the total variation distance between the NBRW and the stationary distribution. According to Observation 5.4 (the correspondence between walks on the cover tree conditioned to be at distance

k and NBRWs of length k), the following holds:

$$\begin{aligned} \|\mathbb{P}(W_t \in \cdot) - \pi\|_{\text{TV}} &\leq \sum_{k=0}^t \|\mathbb{P}(W_t \in \cdot \mid \text{dist}(X_t, \rho) = k) - \pi\|_{\text{TV}} \\ &\quad \cdot \mathbb{P}(\text{dist}(X_t, \rho) = k) \\ &= \sum_{k=0}^t \varrho_k \mathbb{P}(\text{dist}(X_t, \rho) = k) . \end{aligned}$$

Now, by Theorem 3, when $k > \lceil \log_{d-1}(dn) \rceil$ we have $\varrho_k = o(1)$, hence

$$\|\mathbb{P}(W_t \in \cdot) - \pi\|_{\text{TV}} \leq \mathbb{P}(\text{dist}(X_t, \rho) \leq \lceil \log_{d-1}(dn) \rceil) + o(1). \quad (5.6)$$

Equations (5.5) and (5.6) imply that mixing takes place when $\text{dist}(X_t, \rho)$ is $\log_{d-1} n + O(1)$. By the above discussion on the distribution of $\text{dist}(X_t, \rho)$ this occurs when t is around $\frac{d}{d-2} \log_{d-1} n$ with window $\sqrt{\frac{\log n}{d \log d}}$.

It remains to address the case where $\frac{d \log \log n}{\log n} \rightarrow \infty$. Notice that here, as the probability of the SRW on G making a backtracking step is $1/d$, the probability of backtracking anywhere in its first $\lceil \log_{d-1}(dn) \rceil + 1$ steps is $o(1)$. Hence, we can couple the SRW and NBRW in their first $\lceil \log_{d-1}(dn) \rceil + 1$ steps **whp**, implying they have the same mixing time. In particular, we may conclude that for any fixed $0 < s < 1$, the worst case total-variation mixing time of the SRW on G **whp** satisfies

$$t_{\text{MIX}}(s) \in \{ \lceil \log_{d-1}(dn) \rceil, \lceil \log_{d-1}(dn) \rceil + 1 \} ,$$

as required. ■

6. CONCLUDING REMARKS AND OPEN PROBLEMS

- We have established the cutoff phenomenon for SRWs and NBRWs on almost every d -regular graph on n vertices, where $3 \leq d \leq n^{o(1)}$ (beyond which the mixing time is $O(1)$ and we cannot have cutoff). For both walks, we obtained the precise cutoff location and window:
 1. For the SRW, the cutoff point is **whp** at $\frac{d}{d-2} \log_{d-1} n$, and in fact, we obtained the *two* leading order terms of $t_{\text{MIX}}(s)$ for any fixed s .
 2. For the NBRW, cutoff occurs at $\log_{d-1}(dn)$ **whp** ($\frac{d}{d-2}$ times faster than the SRW) with an $O(1)$ window. Moreover, for large d , the entire mixing transition takes place within a 2-step cutoff window.
- Given our discussion in Section 1 on expander graphs (and the product-criterion for cutoff), it would be interesting to extend our results to any arbitrary family of expanders. While one may design such graphs where the SRW has no cutoff, such constructions seem highly asymmetric, and the following conjecture seems plausible (see also [15, Question 5.2]):

Conjecture 6.1. *The SRW on any family of vertex-transitive expander graphs exhibits cutoff.*

- Similarly, recalling the above comparison of t_{MIX} of the SRW and the NBRW on random regular graphs, it would be interesting to extend this result to any family of vertex-transitive expander graphs.

ACKNOWLEDGMENTS

We are grateful to Yuval Peres for sharing his deep insight on cutoff with us over many fruitful discussions, and in particular, for simplifying our proof of Theorem 1 via the duality between NBRWs and SRWs (Observation 5.4). We thank Laurent Saloff-Coste for useful discussions, as well as Jian Ding for comments on an earlier version. We also wish to thank the anonymous referees for helpful comments.

REFERENCES

- [1] M. Ajtai, *Recursive construction for 3-regular expanders*, *Combinatorica* **14** (1994), no. 4, 379–416.
- [2] D. Aldous, *Random walks on finite groups and rapidly mixing Markov chains*, *Seminar on probability*, XVII, 1983, pp. 243–297.
- [3] D. Aldous and P. Diaconis, *Shuffling cards and stopping times*, *Amer. Math. Monthly* **93** (1986), 333–348.
- [4] D. Aldous and J. A. Fill, *Reversible Markov Chains and Random Walks on Graphs*. In preparation, <http://www.stat.berkeley.edu/~aldous/RWG/book.html>.
- [5] N. Alon, I. Benjamini, E. Lubetzky, and S. Sodin, *Non-backtracking random walks mix faster*, *Commun. Contemp. Math.* **9** (2007), no. 4, 585–603.
- [6] N. Alon and J. H. Spencer, *The probabilistic method*, 3rd ed., John Wiley & Sons Inc., Hoboken, NJ, 2008.
- [7] O. Angel, J. Friedman, and S. Hoory, *The non-backtracking spectrum of the universal cover of a graph*, preprint.
- [8] N. Berestycki and R. Durrett, *Limiting behavior for the distance of a random walk*, *Electron. J. Probab.* **13** (2008), 374–395.
- [9] A. Broder and E. Shamir, *On the second eigenvalue of random regular graphs*, *Proc. of the 28th Annual Symposium on the Foundations of Computer Science* (1987), pp. 286–294.
- [10] B. Bollobás, *A probabilistic proof of an asymptotic formula for the number of labelled regular graphs*, *European J. Combin.* **1** (1980), no. 4, 311–316.
- [11] B. Bollobás, *Random graphs*, 2nd ed., *Cambridge Studies in Advanced Mathematics*, vol. 73, Cambridge University Press, Cambridge, 2001.
- [12] G.-Y. Chen and L. Saloff-Coste, *The cutoff phenomenon for ergodic Markov processes*, *Electronic Journal of Probability* **13** (2008), 26–78.
- [13] P. Diaconis, *The cutoff phenomenon in finite Markov chains*, *Proc. Nat. Acad. Sci. U.S.A.* **93** (1996), no. 4, 1659–1664.
- [14] P. Diaconis and M. Shahshahani, *Generating a random permutation with random transpositions*, *Z. Wahrsch. Verw. Gebiete* **57** (1981), no. 2, 159–179.

- [15] J. Ding, E. Lubetzky, and Y. Peres, *Total-variation cutoff in birth-and-death chains*, Probab. Theory Related Fields, to appear.
- [16] C. C. Z. Dou, *Studies of random walks on groups and random graphs*, Ph.D. thesis, Massachusetts Institute of Technology, Dept. of Mathematics, 1992.
- [17] R. Durrett, *Random graph dynamics*, Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, Cambridge, 2007.
- [18] J. Friedman, *A proof of Alon's second eigenvalue conjecture and related problems*, Mem. Amer. Math. Soc. **195** (2008), no. 910, viii+100.
- [19] M. Hildebrand, *Generating random elements in $SL_n(\mathbf{F}_q)$ by random transvections*, J. Algebraic Combin. **1** (1992), no. 2, 133–150.
- [20] M. Hildebrand, *Random walks on random simple graphs*, Random Structures Algorithms **8** (1996), no. 4, 301–318.
- [21] S. Janson, T. Luczak, and A. Rucinski, *Random graphs*, Wiley-Interscience Series in Discrete Mathematics and Optimization, Wiley-Interscience, New York, 2000.
- [22] L. Lovász, *Random walks on graphs: a survey*, Combinatorics, Paul Erdős is eighty, Vol. 2 (Keszthely, 1993), Bolyai Soc. Math. Stud., vol. 2, János Bolyai Math. Soc., Budapest, 1996, pp. 353–397.
- [23] B. D. McKay, *Asymptotics for symmetric 0-1 matrices with prescribed row sums*, Ars Combin. **19** (1985), no. A, 15–25.
- [24] 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>.
- [25] M. S. Pinsker, *On the complexity of a concentrator*, 7th Annual Teletraffic Conference (Stockholm, 1973).
- [26] O. Reingold, S. Vadhan, and A. Wigderson, *Entropy waves, the zig-zag graph product, and new constant-degree expanders*, Ann. of Math. (2) **155** (2002), no. 1, 157–187.
- [27] L. Saloff-Coste, *Random walks on finite groups*, Probability on discrete structures, 2004, pp. 263–346.
- [28] N. C. Wormald, *Models of random regular graphs*, Surveys in combinatorics, 1999 (Canterbury), London Math. Soc. Lecture Note Ser., vol. 267, Cambridge Univ. Press, Cambridge, 1999, pp. 239–298.

EYAL LUBETZKY

MICROSOFT RESEARCH, ONE MICROSOFT WAY, REDMOND, WA 98052-6399, USA.

E-mail address: eyal@microsoft.com

ALLAN SLY

MICROSOFT RESEARCH, ONE MICROSOFT WAY, REDMOND, WA 98052-6399, USA.

E-mail address: allansly@microsoft.com