Speeding up random walks with neighborhood exploration

Petra Berenbrink
Simon Fraser University*
petra@cs.sfu.ca

Colin Cooper
King's College London
colin.cooper@kcl.ac.uk

Robert Elsässer†
University of Paderborn
elsa@mail.uni-paderborn.de

Tomasz Radzik
King's College London
tomasz.radzik@kcl.ac.uk

Thomas Sauerwald‡
Simon Fraser University
thomas_sauerwald@sfu.ca

Abstract

We consider the following marking process (RW-RAND) made by a random walk on an undirected graph G. Upon arrival at a vertex v, it marks v if unmarked and otherwise it marks a randomly chosen unmarked neighbor of v. We also consider a variant of this process called RW-R-RANK. Here each vertex is assigned a global random rank first and then in each step, the walk marks the lowest ranked unmarked neighbor of the currently visited vertex.

Depending on the degree and the expansion of the graph, we prove several upper bounds on the time required by these processes to mark all vertices. For instance, if G is a hypercube or random graph, our processes mark all vertices in time $O(n)$, significantly speeding up the $\Theta(n \log n)$-cover time of standard random walks.

1 Introduction

In this paper we consider strategies that can be used to speed-up the cover time of a random walk on undirected connected graphs. Speeding up random walks to reduce cover time is a very important task in the theory of computing (cf. [5, 15]). The price of this speed up is normally some extra work that can be performed locally by the walk or by the vertices of the graph. Typical assumptions about what is allowed are as follows:

(A) **Biased transitions.** The use of weighted transition probabilities derived from a knowledge of the local structure of the graph. Using weights that are a function of the degree of neighbor vertices can favor difficult to reach vertices.

(B) **Previous history.** Modify transitions using previous local history of the walk to avoid repetitions. For example, non-backtracking walks exclude the edge used at the previous step. Another way is to let the walk choose an unvisited neighbor whenever possible.

(C) **Local exploration.** At each step the walk can explore among the neighbours of the currently visited vertex. Examples of this are fixed depth look-ahead, or marking a previously unvisited neighbor.

In Case (A), the random walk is no longer simple. Weighted walks can give a better worst case cover time. For example, in Ikeda et al. [15], the transition probability of edge $e = \{u, v\}$ is given a weight $w_{u,v}$ proportional to $1/\sqrt{d(u)d(v)}$. This gives an $O(n^2 \log n)$ upper bound on cover time for any connected n vertex graph G. The success of this approach comes from the fact that at a high degree vertex the biased walk transition gives preference to any low degree neighbours. This speeds up exploration of the low degree vertices, which are often the ones which are hard to reach.

Case (B) is less straightforward. At one extreme non-backtracking walks are still Markovian; at the other, preferring unvisited neighbors is definitely not. Non-backtracking is optimal for covering a cycle, as the walk always moves to the unvisited neighbor. Non-backtracking walks on expanders were studied by the authors of [5], who were able to show that they are rapidly mixing on d-regular expanders. However, for $d \geq 3$, using a non-backtracking walk only improves the cover time of random d-regular graphs by a multiplicative constant (from $(d-1)/(d-2)n \ln n$ to $n \ln n$). This can be shown by applying the techniques of [7] to [5].

Processes for Case (C) have been studied extensively in the context of the coupon collecting. The problem of speeding up coupon collecting process on graphs by marking neighbors was first considered in [1]. Define the process CC-RAND for modified coupon collection on

*partially supported by NSERC.
†partially supported by the German Research Foundation under contract EL 399/2-1.
‡partially supported by PIMS and NSERC.

a graph as follows: The vertices of the graph G have two states, free (unmarked) and marked. Initially all vertices of G are unmarked. At each step, pick a vertex uniformly at random (u.a.r.). If v is free, mark it. If v is marked but has free neighbors, pick a free neighbor u.a.r. and mark it. Else do nothing at this step. Let $cc(G)$ be the expected number of steps to mark all vertices using modified coupon collecting process CC-RAND described above. Adler, Halperin, Karp, and Vazirani [1] show, for the d-dimensional hypercube H with $n = 2^d$, that $cc(H) = O(n)$. Their study was motivated by a question of load balancing in distributed hash tables. The authors of [1] also show that the process CC-RAND covers the vertices of an arbitrary n vertex d-regular graph in time $O(n \cdot (1 + (\log n \cdot \log d)/d))$.

Alon [4] shows that the process covers logarithmic-degree Ramanujan expander graphs in $O(n)$ time and random d-regular graphs in time $n+(1+o(1))\cdot(n \ln n)/d$. He also shows a $(n - n/d + (n/d) \cdot \ln(n/d))$ lower bound for d regular graphs in an off-line model where the whole sequence of vertex selections is known in advance. Later, Dimitrov and Plaxton [8] introduced a related process, *R-RANK*, that we call CC-R-RANK. Each vertex is assigned a rank chosen u.a.r. from $\{1, ..., R\}$. Every time a vertex is chosen (according to the coupon collecting process), the unmarked neighbor with the lowest rank is marked. Ties are broken by choosing the vertex with the lowest label. For d-regular graphs, and $R \geq n$, they bound the cover time by $O(n \cdot (1 + (\log n)/d))$.

In this paper, we study case (C), and in particular, walks that are not only able to mark the currently visited vertex but also a neighbor of it. More precisely, we consider the following marking random walk process RW-RAND. Let v denote the vertex visited by the random walk at the current step. If v is not marked, then mark it. If v is already marked, mark a randomly chosen unmarked neighbor of v, if any, else do nothing. At the next step the walk moves to a random neighbor, and the marking process repeats itself. We also consider a variant process RW-R-RANK that initially assigns a global random rank to each vertex. Then at each step of the random walk, the process RW-R-RANK marks the lowest ranked unmarked neighbor of the currently visited vertex.

Although random walks in general have a cover time of $\Omega(n \log n)$ on any graph, we show that with these simple modifications one can break the $\Omega(n \log n)$ barrier, and cover all vertices of many important graph classes in optimal $O(n)$ steps. As a result these processes are well-suited to speed-up the cover time of a random walk. Furthermore, RW-RAND and RW-R-RANK are comparable with the corresponding graph-based coupon collector processes CC-RAND [1], [4] and CC-R-RANK [8]. Our processes RW-RAND and RW-R-RANK can be viewed as distributed implementations of CC-RAND and CC-R-RANK that avoid "jumping" in each step to a random vertex chosen globally by a central process. We show that the performance of RW-RAND and RW-R-RANK are comparable with the the performance of the corresponding coupon collecting process.

1.1 Model and Definitions

Random walks. Let $G = (V, E)$ be a connected graph with $|V| = n$ and $|E| = m$. The degree of a vertex v is denoted by $d(v)$ and let $N_i(v)$ be the i-th neighborhood of v constructed in the breadth-first manner. $N_0(v) = \{v\}$ and $N_1(v) = N(v)$ is the set of all neighbors of v. We refer to $N_i(v)$ as the level i neighbourhood (in the BFS tree rooted at v).

A *simple random walk* $\mathcal{W}_u$, $u \in V$ on graph G, is a Markov chain $\mathcal{W}_u(t) \in V$, $t \geq 0$, with $\mathcal{W}_u(0) = u$, modeled by a particle moving from vertex to vertex according to the following rule. The probability of transition from vertex v to vertex w is equal to $1/d(v)$, if w is a neighbor of v, and is equal 0 otherwise. We perform a simple random walk $\mathcal{W}_u$ on G, starting from $X_0 = u$. Let $X_t = \mathcal{W}_u(t)$ be the vertex reached at step t. A random walk is *lazy*, if it moves from v to one of its neighbors w with probability $1/(2d(v))$, and stays where it is (at vertex v) with probability $1/2$.

We assume the random walk $\mathcal{W}_u$ on G is ergodic with stationary distribution π, where $\pi_v = d(v)/(2m)$. If this is not the case, e.g. G is bipartite, then the walk can be made ergodic, by making it lazy.

Let $\Phi = \Phi(G)$ denote the *conductance* of graph G defined as

$$\Phi = \min\left\{\frac{|E(X, V \setminus X)|}{\min\{d(X), d(V \setminus X)\}} : \emptyset \neq X \subset V\right\},$$

where $E(Y, Z)$ is the set of edges (y, z) with $y \in Y$ and $z \in Z$, and $d(Y) = \sum_{v \in Y} d(v)$. Let

$$P_u^{(t)}(v) = \mathbf{Pr}(\mathcal{W}_u(t) = v).$$

The rate of convergence to the stationary distribution π is given by (see [17], Corollary 5.6)

$$|P_u^{(t)}(x) - \pi_x| \leq \sqrt{\frac{\pi_u}{\pi_x}}\lambda_{\max}^t \leq \sqrt{\frac{\pi_u}{\pi_x}}\left(1 - \frac{\Phi^2}{8}\right)^t, \tag{1.1}$$

where $\lambda_{\max} = \max\{\lambda_2, |\lambda_n|\}$ with $\lambda_1, \ldots, \lambda_n$ being the eigenvalues of the matrix of transition probabilities $P = P(G)$. In the case of the lazy random walk $\lambda_{\max} = \lambda_2$. Let $\mathbf{T} = \mathbf{T}(G)$ be the smallest t such that the right-hand side of (1.1) is less than n^{-3} for all u and x. We refer to this value $\mathbf{T}$ as the *mixing time* of graph G.

For $v \in V$, let $X(t, v)$ be the number of visits made by a random walk to v during t steps. For $u \in V$, let C_u be the time needed for the non-lazy random walk to visit every vertex of G. The *cover time* $\mathbf{C} = \mathbf{C}(G)$ of G is defined as $\mathbf{C} = \max_{u \in V} \mathbf{E}[C_u]$. The cover time of connected graphs has been extensively studied. It is a classic result of Aleliunas, Karp, Lipton, Lovász and Rackoff [3] that $\mathbf{C}(G) \le 2m(n-1)$. It was shown by Feige [11], [12], that for any connected graph G

$$(1 - o(1)) \cdot n \log n \le \mathbf{C}(G) \le (1 + o(1)) \cdot \frac{4}{27} \cdot n^3.$$

The lower bound is achieved by e.g. the complete graph K_n, whose cover time is determined by the coupon collector problem. A lazy walk doubles the cover time, but asymptotically, half the steps are loops. For our main Theorems 1, 2, and 3, the first two are shown for lazy random walks, while the last one does not require that the random walk is lazy. For now on, we assume that the walk is lazy unless explicitly mentioned to the contrary.

We define $\mathbf{E}_\pi[H_v]$ as the expected hitting time of v from stationarity. The quantity $\mathbf{E}[R_v]$ is the expected number of returns to v during mixing time many steps made by a walk starting from v, more precisely,

$$\mathbf{E}[R_v] \; = \; \sum_{t=0}^{T-1} P_v^{(t)}(v). \tag{1.2}$$

Note that the summations in (1.2) starts from 0, so $\mathbf{E}[R_v] \ge 1$.

Processes. The process RW-RAND for a random walk is defined by analogy with the modified coupon collection process CC-RAND. The vertices of the graph G have two states, free and marked. Initially all vertices of G are free. We perform a simple random walk $\mathcal{W}_u$ on G, starting from an arbitrary vertex $X_0 = u$. If the vertex $X_t = v$ visited at time step t is free, mark it. If v is marked but has free neighbors, pick a free neighbor u.a.r. and mark it. Else do nothing at this step.

Our second process, the RW-R-RANK process is similar to RW-RAND. The difference is that it never marks the currently visited vertex v itself, and that it uses another priority rule to choose between unmarked neighbors. RW-R-RANK first globally assigns a random rank in $\{1, \dots, n^2\}$ to each vertex. Then in each step of the random walk, RW-R-RANK marks the lowest ranked unmarked neighbor of v (ties are broken by choosing the vertex with the lowest label).

The *visited vertices* are the vertices that have been visited (directly) by the random walk. The terms *marked vertices* and *covered vertices* have the same meaning: the vertices that have been marked so far. For the RW-RAND process, the set of visited vertices is a subset of the set of marked (covered) vertices. The *free vertices* are the vertices that are not marked.

Graph Classes. We first consider two graph classes that obey some strong local expansion properties.

A d-regular graph G, is a *Local Expander* with parameter L, if for all $v \in V$, $w \in N_i(v)$, $1 \le i \le L$, and some constant $c > 1$, we have

1. $|N(w) \cap N_{i+1}(v)| \ge c|N(w) \cap N_{i-1}(v)|$ and,
2. $|N(w) \cap N_{i+1}(v)| = \Omega(d)$.

If in addition, for $i \le 4$ and some constant α, we have $|N(w) \cap (N_i(v) \cup N_{i-1}(v))| \le \alpha$ we say G is a *Strong Local Expander*. The d-dimensional hypercube, for example, is a Local Expander (with parameter $L = d/2$) and a Strong Local Expander. We also consider graphs with global expansion properties. We call a d-regular graph G an *expander* if $\Phi(G) = \Omega(1)$, or equivalently $\lambda_2 = 1 - \Theta(1)$. A d-regular graph is called *almost Ramanujan* if $\lambda_2 = 1/d^{\Omega(1)}$. It is known that, e.g., Erdős-Rényi-Random graphs and random d-regular graphs are both Strong Local Expanders [9] and almost Ramanujan graphs [14]. By (1.1), it follows that the mixing time satisfies $\mathbf{T} = O(\log n)$ for these graphs. For hypercubes, we have $\lambda_2 = 1 - 1/d$ [17] and hence again by (1.1), we obtain that $\mathbf{T} = O(\log^2 n)$.

1.2 New Results.

We state the results using **whp** as probability bounds. These bounds are strong in the sense that the results hold with probability $1 - O(n^{-\gamma})$ for some large constant $\gamma > 0$. Our definitions of expansion properties and our main Theorems 1, 2 and 3 are stated for regular graphs only, but they can be easily extended to *pseudo-regular* graphs, in which the ratio of the maximum to minimum degree is constant. We start with a general bound that corresponds to Theorem 4.1 of [1] of the centralized process CC-RAND.

THEOREM 1. *Let G be an arbitrary d-regular graph. Then* RW-RAND *covers G within*

$$O\left(\frac{1}{1-\lambda_2} \cdot \left(n \cdot \frac{\log n \, \log d}{d} + n + d \, \log n\right)\right)$$

steps **whp**.

To establish this theorem, we adopt the following analysis from [1] for the centralized process CC-RAND. We divide the process RW-RAND into $\log d + 1$ phases. As in [1], we then establish that at the end of phase i, every vertex has at most $d/2^{i+1}$ uncovered neighbors. However, in CC-RAND the chosen vertices are independent, while in the process RW-RAND, the chosen vertices are correlated. To cope with the dependencies, we use a Chernoff-type inequality for Markov chains from [16] (see Lemma 10).

The theorem above immediately implies the following corollary.

COROLLARY 1. *Let G be a d-regular expander with $\log n \log\log n \le d \le n/\log n$. Then* RW-RAND *covers G within $O(n)$ steps* **whp**. *If $d = \Omega(\log n)$, then* RW-RAND *covers G within $O(n \log\log n)$ steps* **whp**.

For sparse graphs with strong local expansion properties the following theorem gives a better bound. For $d = \Omega(\log n)$ it gives the asymptotically optimal bound of $O(n)$

THEOREM 2. *Let G be a Strong Local Expander with degree $d \to \infty$ and parameter $L \ge \kappa \cdot \max\{\ln\ln n, \ln \mathbf{T}(G)\}$ ($\kappa > 0$ is a large constant) and let the mixing time $\mathbf{T}(G) = o(n/(d^3 + \log n))$. Then* RW-RAND *covers G in $O(n + (n \log n)/d)$ steps* **whp**.

Since the d-dimensional hypercube is a Strong Local Expander we obtain the following corollary, which should be compared with the $O(n)$ bound shown in [1] for CC-RAND on a hypercube.

COROLLARY 2. RW-RAND *covers all vertices of the d-dimensional hypercube in $O(n)$ steps* **whp**. *For $(1 + \Omega(1)) \cdot \log n \le d \le n^{1/3}/\log n$,* RW-RAND *covers all vertices of Erdős-Rényi-Random graphs with average degree d and of random d-regular graphs in $O(n)$ steps* **whp**.

As regards to the process RW-R-RANK, we prove the following theorem w.r.t. Local Expander or almost Ramanujan graphs. This result is equivalent to the result for CC-R-RANK when $d = \Omega(\ln n)$ and $R = n^2$. This result can easily be extended to the case where the ranks are assigned to the vertices of the graph according to a random permutation of $V = [n]$.

THEOREM 3. *Assume $d = \Omega(\log n)$ and $d \cdot \mathbf{T}(G) = O(n^{1-\epsilon})$ with $\epsilon > 0$ being a small constant. Let G be either a Local Expander with parameter $L \ge \kappa \cdot \max\{\ln\ln n, \ln \mathbf{T}(G)\}$ (κ is a large constant) or an almost Ramanujan graph. Then process* RW-R-RANK *with $R = n^2$ covers G in $O(n)$ steps* **whp**.

COROLLARY 3. RW-R-RANK *covers all vertices of the d-dimensional hypercube in $O(n)$ steps* **whp**. *For $(1 + \Omega(1)) \cdot \log n \le d \le O(n^{1-\epsilon}/\log n)$,* RW-R-RANK *covers all vertices of Erdős-Rényi random graphs with average degree d and of random d-regular graphs in $O(n)$ steps* **whp**.

This result is equivalent to the result for CC-R-RANK when $d = \Omega(\ln n)$ and $R = n^2$. Theorem 3 can easily be extended to the case where the ranks are assigned to the vertices of the graph according to a random permutation of $V = [n]$.

Let us compare our results with the performance of random walks that use look ahead to mark all neighbors of a visited vertex in one step. According to [7], a random walk with look ahead covers all vertices of a random d-regular graph with $d = \Omega(\log n)$ in time $\sim n \ln n/d$. Using the methods of [7], it can be shown that the same process marks all vertices of a hypercube in time $\sim \ln 2 \cdot n$ and Local Expanders with degree $d = \Omega(\log n)$ in $O(n)$ time.

However, our results imply that even if the random walk is only allowed to mark one single neighbor in each step, the cover time is essentially the same on many important graph classes as in the case in which the random walk is endowed with full look ahead abilities.

As in Theorem 1, the main technical challenge to prove Theorem 2 and Theorem 3 are the correlations between the visited vertices by the random walk. The main idea to cope with these correlations is as follows. Assume that t is a multiple of $2\mathbf{T}$. We divide the time interval $[0, t]$ into consecutive subintervals of length $\mathbf{T}$. Now applying Lemma 3 to every even subintervals we obtain

$$\begin{aligned}
&\mathbf{Pr}(v \text{ is unvisited at } t) \\
&\quad\le \left(1 - \frac{\mathbf{T}}{4\mathbf{E}_\pi[H_v]}(1 - o(1))\right)^{t/\mathbf{T}} \\
&\quad\le \left(1 - \Omega\left(\frac{1}{\mathbf{E}_\pi[H_v]}\right)\right)^{t}. \qquad (1.3)
\end{aligned}$$

Thus, after a mixing time $\mathbf{T}$ the random walk behaves like a coupon collecting process with success probability $\Omega(1/\mathbf{E}_\pi[H_v])$, where $\mathbf{E}_\pi[H_v]$ is the expected hitting time of v from stationarity. The results in this paper are based on the following generalization of (1.3) to sets of vertices. For $S \subseteq V$,

$$\begin{aligned}
&\mathbf{Pr}(S \text{ is unvisited at } t) \\
&\quad\le \left(1 - \Omega\left(\frac{1}{\mathbf{E}_\pi[H_S]}\right)\right)^{t}.
\end{aligned}$$

This bound can be derived by analysing a graph obtained by contracting S to a single vertex.

Provided $\mathbf{T} \cdot \pi_v = o(1)$, the quantity $\mathbf{E}_\pi[H_v]$ can be obtained explicitly as $\mathbf{E}_\pi[H_v] \sim \mathbf{E}[R_v]/\pi_v$. where $\mathbf{E}[R_v] \ge 1$ is the expected number of returns to v during $\mathbf{T}$. For many graphs G with moderate expansion we have $\mathbf{E}[R_v] = \Theta(1)$. We establish this property in the case when G is Local Expander or almost Ramanujan.

Finally, we state the following negative result (proof in Appendix). It demonstrates that a certain expansion is necessary to reduce the runtime of RW-R-RANK and RW-RAND to $O(n)$.

THEOREM 4. *There exists a $\Theta(\log n)$-regular graph G with $\mathbf{C}(G) = \Theta(n \log n)$, but* RW-RAND *and* RW-R-RANK *also require $\Omega(n \log n)$ steps.*

1.3 Organization. The remainder of this paper is organized as follows. In Section 2 we derive some auxiliary lemmas about random walks. These are used lateron in Sections 3 and Sections 4 in which we prove our results for the processes RW-R-RANK and RW-RAND, respectively. Theorem 4 is shown in the appendix.

2 Random Walk Properties

We now establish results on the number of visits to a given vertex, or set of vertices in a given time interval. Let $\mathsf{Bin}(r, p)$ and $\mathsf{NegBin}(r, p)$ denote the binomial distribution (the number of successes in r trials) and the negative binomial distribution (the number of trials to get r successes).

2.1 Visits to Single Vertices.

LEMMA 1. *Let vertex $v \in V$ be such that $\mathbf{T} \cdot \pi_v = o(1)$, and let*

$$Y \sim \mathsf{Bin}(\lfloor t/2\mathbf{T} \rfloor, (1 - o(1))\mathbf{T} \cdot \pi_v / 4\mathbf{E}[R_v]).$$

Then for any k

$$\mathbf{Pr}(X(t, v) \geq k) \geq \mathbf{Pr}(Y \geq k).$$

In particular, provided $\mathbf{T} = o(t)$, then

$$\mathbf{Pr}(X(t, v) = 0) \leq \exp(-(1 - o(1))t\pi_v / 8 \cdot \mathbf{E}[R_v]).$$

The proof of Lemma 1 is based on Lemma 3 given below. We first state the following intuitively true result (proof in Appendix).

LEMMA 2. *For any vertex w, $t \geq \mathbf{T}$ and $t' \geq t + \mathbf{T}$,*

$$\begin{aligned} &\mathbf{Pr}(w \in \{X_{t+i} : i = 0, ..., \mathbf{T} - 1\}) \\ &= \ (1 + o(1)) \ \cdot \mathbf{Pr}(w \in \{X_{t'+i} : i = 0, ..., \mathbf{T} - 1\}). \end{aligned}$$

The quantity $\mathbf{E}_\pi[H_w]$ (expected hitting time of a vertex w from the stationary distribution π) can be expressed as $\mathbf{E}_\pi[H_w] = Z_{ww}/\pi_w$, where

$$Z_{ww} = \sum_{t=0}^{\infty} (P_w^{(t)}(w) - \pi_w), \tag{2.4}$$

see e.g. [2]. Using this and Lemma 2, we prove the following lemma (proof in Appendix). Note that a similar result is also proved in [10].

LEMMA 3. *For a vertex w such that $\mathbf{T} \cdot \pi_w = o(1)$ and $t \geq \mathbf{T}$,*

$$\begin{aligned} &\mathbf{Pr}(w \in \{X_{t+i} : i = 0, ..., \mathbf{T} - 1\}) \\ &\geq \ \frac{\mathbf{T}}{4 \cdot \mathbf{E}_\pi[H_w]}(1 - o(1)) \\ &= \ \frac{\mathbf{T} \cdot \pi_w}{4 \cdot \mathbf{E}[R_w]}(1 - o(1)). \end{aligned} \tag{2.5}$$

2.2 Visits to Vertex Sets. We now extend the results presented in Section 2.1 to a subset of vertices. Let $\emptyset \neq S \subseteq V$. From graph G we obtain a (multi)-graph G_S by contracting S to a single vertex γ (we retain in G_S multiple edges and loops). Let $\hat{P}$ denote the transition matrix for a random walk on G_S, and define

$$R_S \ = \ \sum_{t=0}^{\mathbf{T}-1} \hat{P}_\gamma^{(t)}(\gamma). \tag{2.6}$$

Note that we define R_S as the number of returns to γ in the contracted graph G_S rather than a somewhat different notion of the number of returns to S in G. The former is sufficient for our purposes and easier to analyse. Note also that in the summation in (2.6) we use $\mathbf{T}$, the mixing time in G, rather than the mixing time in G_S. Let $\hat{\pi}$ be the stationary distribution of a random walk on G_S. For any vertices u and x in G_S, we have

$$|\hat{P}_u^{(\mathbf{T})}(x) - \hat{\pi}_x| \leq n^{-2}. \tag{2.7}$$

This holds since we use a conductance based definition of the mixing time $\mathbf{T}$. More precisely, to obtain (2.7), observe that $\hat{\pi}_x = \pi_x$, if $x \neq \gamma$, $\hat{\pi}_\gamma = \pi_S = \sum_{v \in S} \pi_v$, and that conductance can only increase under vertex contraction, and then apply (1.1).

Any walk W from $v \notin S$ to S in G with internal vertices not in S corresponds to an identical walk $\hat{W}$ from v to γ in G_S, and both walks have exactly the same probability. Thus we obtain

$$\mathbf{E}_\pi[H_S] \ = \ \mathbf{E}_{\hat{\pi}}[H_\gamma] \cdot \frac{\hat{Z}\gamma\gamma}{\hat{\pi}_\gamma}. \tag{2.8}$$

The following lemma has an analogous proof to the proof of Lemma 3, keeping in mind (2.7) and (2.8).

LEMMA 4. *For a subset S of V such that $\mathbf{T} \cdot \pi_S = o(1)$ and $t \geq \mathbf{T}$,*

$$\begin{aligned} &\mathbf{Pr}(|S \cap \{X_{t+i} : i = 0, ..., \mathbf{T} - 1\}| \geq 1) \\ &\geq \ \frac{\mathbf{T}}{4\mathbf{E}_\pi[H_S]} \cdot (1 - o(1))) \\ &= \ \frac{\mathbf{T} \cdot \pi_S}{4 \cdot \mathbf{E}[R_S]} \cdot (1 - o(1)). \end{aligned}$$

In particular, Lemma 4 implies that if G is a d-regular graph and a subset of vertices S is such that $|S| \cdot \mathbf{T} = o(n)$ and $\mathbf{E}[R_S] = O(1)$, then

(2.9)
$$\mathbf{Pr}(|S \cap \{X_{t+i} : i = 0, ..., \mathbf{T}-1\}| \geq 1) = \Omega\left(\frac{|S| \cdot \mathbf{T}}{n}\right).$$

On the other hand, for any vertex $w \in V$ and $t \geq \mathbf{T}$, $\mathbf{Pr}(w = X_{t+i}) = n^{-1}(1 \pm o(1/n))$. Applying the union bound over all vertices of S and time interval $\mathbf{T}$, we obtain an $O\left(\frac{|S| \cdot \mathbf{T}}{n}\right)$ bound on the probability on the left-hand side in (2.9).

Using Lemma 4, the following Lemma, that is analogous to Lemma 1, can be proven.

LEMMA 5. *Let a subset of vertices $S \subseteq V$ be such that $\mathbf{T} \cdot \pi_S = o(1)$, and let*

$$Y \sim \mathsf{Bin}(\lfloor t/2\mathbf{T} \rfloor, (1 - o(1))\mathbf{T} \cdot \pi_S/4 \cdot \mathbf{E}[R_S]).$$

Then for any k,

$$\mathbf{Pr}(X(t, v) \geq k) \geq \mathbf{Pr}(Y \geq k),$$

In particular, provided $\mathbf{T} = o(t)$, then

$$\mathbf{Pr}(X(t, S) = 0) \leq \exp(-(1 - o(1)) \cdot t\pi_S/8 \cdot \mathbf{E}[R_S]).$$

2.3 Number of Returns to a Vertex. We note the following result (see e.g. [13]), for a random walk on the line $= \{0, ..., a\}$ with absorbing states $\{0, a\}$, and transition probabilities q, p, s for moves left, right and looping respectively. Starting at a vertex z, the probability of absorption at the origin 0 is

$$\rho(z, a) = \frac{(q/p)^z - (q/p)^a}{1 - (q/p)^a} \leq \left(\frac{q}{p}\right)^z, \tag{2.10}$$

provided $q \leq p$. Similarly, for a walk starting at z on the half line $\{0, 1, ...\}$, with absorbing states $\{0, \infty\}$, the probability of absorption at the origin is $\rho(z) = (q/p)^z$.

In this section we assume that the graphs are (Strong) Local Expanders with parameter $L \geq \kappa \max\{\ln \ln n, \ln \mathbf{T}\}$, where κ is a large constant. The constant c below is the constant from the definition of a Local Expander.

LEMMA 6. *Let G be a Local Expander of degree d. For any vertex v we have $\mathbf{E}[R_v] = 1 + O(1/d)$.*

Proof. Assume the random walk starts at v at time step 0. Let r_i be the probability of a return to v at the i-th step. For a walk starting from v we have $r_0 = 1$. Now, we show that the expected number of returns before reaching distance $\omega = L > A \log_c \mathbf{T} + 1$ is $O(1/d)$, if α is large enough.

In the next step, the walk moves to $N(v)$. From $N(v)$ the walk moves to v before moving to $N_2(v)$ with probability $1/\Omega(d)$. Once the walk has moved beyond $N(v)$ it can be coupled with a random walk on a path with loops, and absorbing barriers at $1, \omega$. Conditional on not looping, the probability the walk moves left is $q \leq 1/(1+c) < 1/2$ and moves right is $p \geq c/(1+c) > 1/2$. The probability of absorption at 1 is at most q/p, so the expected number of returns to $N(v)$ before walking to ω is $p/(p-q)$. The expected number of returns to v is therefore $O(1/d)$. The probability of a return to $N(v)$ from distance ω at any step $i \leq \mathbf{T}$ is at most $(1/c)^\omega \leq \mathbf{T}^{-A}$. Thus, provided $A > 1$ the total expected number of returns is $1 + O(1/d)(1 + O(\mathbf{T}^{-A+1}))$. □

The lemma above also holds for Strong Local Expanders, since any Strong Local Expander is a Local Expander. Now we state bounds on the number of returns to certain sets of vertices in Local Expanders and Strong Local Expanders (proofs in Appendix). Such sets are considered in the proofs of Theorems 2 and 3.

LEMMA 7. *Let G be a Local Expander of degree d. For any vertex v we have $\mathbf{E}[R_{N(v)}] = O(1)$.*

LEMMA 8. *Let G be a strong Local Expander of degree d. For any vertex v and $i \leq 4$ we have $\mathbf{E}[R_S] = O(1)$, where $S \subseteq N_i(v)$.*

Lemmas 3 and 4 and the definition of conditional probabilities imply the following corollary.

COROLLARY 4. *If G is d-regular, S is a subset of vertices such that $|S|\mathbf{T} = o(n)$ and $\mathbf{E}[R_S] = O(1)$, for each vertex $w \in S$, $\mathbf{E}[R_w] = O(1)$, and $t \geq \mathbf{T}$, then for each $w \in S$,*

$$\mathbf{Pr}\Big(w \in \{X_{t+i} : i = 0, \ldots, \mathbf{T}-1\} \;\Big|\; |\{X_{t+i} : i = 1, \ldots, \mathbf{T}_S\} \cap S| \geq 1\Big) = \Theta\left(\frac{1}{|S|}\right).$$

3 Analysis of RW-RAND

The proofs are based on [1] with modifications arising from using a random walk to sample vertices instead of coupon collecting.

3.1 d-Regular Graphs. In this section we prove Theorem 1. The outline of the proof is as follows. We divide the process RW-RAND into $\log d + 1$ phases. We prove that at the end of each phase $i, 0 \leq i \leq \log d$, every vertex v has at most $d/2^{i+1}$ uncovered neighbors. Clearly, this implies that at the end of phase $\log d$, all vertices are covered.

To establish the progress of covering neighbors, we divide each phase i into d sub-phases of a certain length.

Using a Chernoff-type bound for Markov Chains, we then establish that in each sub-phase, a new uncovered neighbor of v is marked with constant probability. Taking a sufficiently large number of sub-phases (at least $\Omega(2^{-i} \cdot d)$), we obtain that we visit in phase i at least $2^{-i} \cdot d$ neighbors of v with probability $1 - n^{-\Omega(1)}$. Taking a union bound over all vertices and phases, we conclude that with high probability, all vertices are covered after the last phase.

For the proof of Theorem 1 we use the following Chernoff-type bound by Lezaud [16].

LEMMA 9. ([16, THEOREM 1.1/REMARK 2]) *Consider an ergodic Markov chain $(X_0, X_1, \dots)$ on a graph $G = (V, E)$. Let $f : V \to \mathbb{R}$ be such that $\sum_{u \in V} \pi(u) f(u) = 0$, $\|f\|_\infty \le 1$ and $\|f\|_2^2 \le b^2$ for some value b. Then, for any initial distribution X_0, any positive integer t and all $0 \le \gamma \le 2b^2/5$,*

$$\begin{aligned}
&\mathbf{Pr}\left(\sum_{i=1}^{t} f(X_i) \ge t\gamma\right) \\
&\le e^{(1-\lambda_2)/5} \cdot \sum_{u \in V} \frac{(\mathbf{Pr}(X_0 = u))^2}{\pi(u)} \cdot \\
&\quad \exp\left(-\frac{t\gamma^2(1-\lambda_2)}{4b^2}\left(1 - \frac{5\gamma}{2b^2}\right)\right).
\end{aligned}$$

The following result is an application of the lemma above.

LEMMA 10. *Consider a step t_0 such that the distribution of X_{t_0} is close to stationary, that is, $\mathbf{Pr}(X_{t_0} = u) = (1 + o(1))n^{-1}$ for each vertex u. Let $q = (q_w : w \in V)$ be a vector with $0 \le q_w \le 1$, for each $w \in V$. Let $\mu = (\sum_{w \in V} q_w)/n > 0$ and let*

$$\tilde{q}_w := -q_w + \mu.$$

Then, for a sufficiently large constant $c > 0$ and for any $t \ge c \cdot \frac{1}{1-\lambda_2} \cdot \frac{1}{\mu}$,

$$\mathbf{Pr}\left(\sum_{i=t_0+1}^{t_0+t} q_{X_i} \ge 1\right) \ge \frac{1}{2}. \tag{3.11}$$

Proof. We have $\|\tilde{q}\|_\infty \le 1$ and

$$\begin{aligned}
\|\tilde{q}\|_2^2 &= \frac{1}{n}\sum_{w \in V} \tilde{q}_w^2 = \frac{1}{n}\sum_{w \in V} (-q_w + \mu)^2 \\
&\le \frac{1}{n}\sum_{w \in V} q_w^2 + \mu^2 \le \mu + \mu^2 \le 2\mu. \qquad (3.12)
\end{aligned}$$

We apply Lemma 9 with $f = \tilde{q}$, $t = c \cdot \frac{1}{1-\lambda_2} \cdot \frac{1}{\mu}$, $\gamma = \frac{\mu}{4}$, $b^2 = 2\mu$ to obtain, for a sufficiently large constant c, that

$$\begin{aligned}
&\mathbf{Pr}\left(\sum_{i=t_0+1}^{t_0+t} \tilde{q}_{X_i} \ge \frac{c}{4(1-\lambda_2)}\right) \\
&\le e^{(1-\lambda_2)/5} \cdot (1 + o(1)) \cdot \\
&\quad \exp\left(-\frac{\frac{c}{1-\lambda_2}\left(\frac{1}{\mu}\right) \cdot \left(\frac{\mu}{4}\right)^2 \cdot (1-\lambda_2)}{8\mu} \cdot \left(1 - \frac{5\left(\frac{\mu}{4}\right)}{4\mu}\right)\right) \\
&\le \frac{1}{2}.
\end{aligned}$$

Substituting

$$\sum_{i=t_0+1}^{t_0+t} \tilde{q}_{X_i} = -\left(\sum_{i=t_0+1}^{t_0+t} q_{X_i}\right) + t\mu,$$

we get

$$\mathbf{Pr}\left(\sum_{i=t_0+1}^{t_0+t} q_{X_i} \le \frac{3c}{4(1-\lambda_2)}\right) \le \frac{1}{2}.$$

This gives the claim. □

We are now ready to start the proof of Theorem 1. We consider $\log d + 1$ different phases numbered from 0 to $\log d$. We shall prove that at the end of each phase i, every vertex has at most $d/2^{i+1}$ free neighbors. Phase i consists of $\nu_i = \Theta(\max(\frac{d}{2^i}, \log n))$ iterations, and each iteration consists of two parts. The first part of an iteration consists of $\Theta((\log n)/(1 - {}_2))$ steps to get close to the stationary distribution, i.e., to reach a distribution X_{t_0} with $\mathbf{Pr}(X_{t_0} = u) = (1 + o(1))n^{-1}$ for each u. The second part of an iteration consists of $\tau = 4cn/(d(1 - {}_2))$ steps, where c is a constant which gives (3.11). Thus the number of steps in the whole computation is at most of the order of

$$\begin{aligned}
&\sum_{i=0}^{\log d}\left(\frac{d}{2^i} + \log n\right)\left(\frac{\log n}{1-\lambda_2} + \frac{n}{d \cdot (1-\lambda_2)}\right) \\
&\le \frac{d \log n + \log^2 n \cdot \log d + n + (\log d \log n \cdot n)/d}{1-\lambda_2}.
\end{aligned}$$

The success of phase i will follow from the following lemma.

LEMMA 11. *Consider any phase $i \in \{0, 1, \dots, \log d\}$. Assume that each vertex has at most $d/2^i$ free neighbors at the beginning of phase i. Consider a vertex v that has more than $d/2^{i+1}$ free neighbors at the beginning of an iteration in phase i. Then the probability that a free neighbor of v is marked during this iteration is at least some positive constant.*

Proof. We consider a timestep t_0 in the second part of an iteration where the random walk is at any vertex with probability $(1 \pm o(1))/n$. The remainder of the proof is divided into two cases depending on whether $i = 0$.

1. $i = 0$: We bound the probability that the random walk visits within the next τ steps a free neighbor of v, (that is, we count only "marking by visiting"). Let $q = q(v) = (q_w : w \in V)$ be a vector such that $q_w = 1$ if w is a free neighbor of v at step t_0, and $q_w = 0$ otherwise. By assumption, v has more than $d/2$ free neighbors at step t_0, so $\mu = \frac{1}{n}\sum_{w \in V} q_w \geq \frac{d}{2n}$. The lemma follows since the probability that a free neighbor of v is marked within the next τ steps is at least the probability that $\sum_{i=t_0+1}^{t_0+\tau} q_{X_i} \geq 1$, and (3.11) applies.

2. $1 \leq i \leq \log d$: We bound the probability that within the next τ steps a free neighbor u of v is marked during a visit to a neighbor of u that was already marked by step t_0. We use the following notation. $C(t)$ denotes the set of vertices that are already marked by step t, and $C(v,t)$ denotes those vertices that are additionally a neighbor of v. Similarly, $F(v,t)$ denotes the set of free neighbors of vertex v at the beginning of step t. Let vector $q : V \to \mathbb{R}$ be defined as

$$q_w = \frac{|F(v, t_0+1) \cap F(w, t_0+1)|}{|F(w, t_0+1)|},$$

for each vertex $w \in C(t_0+1)$, and $q_w = 0$, for $w \notin C(t_0+1)$. For a step $t > t_0$, the probability that a free vertex in $F(v, t_0+1)$ is marked in step t, given that the random walk is at this step at a vertex w, this vertex was already covered by step t, and no vertex in $F(v, t_0+1)$ has been marked in steps $t_0+1, \ldots, t-1$, is equal to

$$\frac{|F(v, t_0+1) \cap F(w, t)|}{|F(w, t)|} \geq q_w.$$

Thus the probability of not marking any vertex in $F(v, t_0+1)$ during the steps $t_0+1, \ldots, t_0+\tau$, conditioned on $X_{t_0+1}, \ldots, X_{t_0+\tau}$, is at most

$$\prod_{i=t_0+1}^{t_0+\tau} (1 - q_{X_i}) \leq \left(1 - \frac{\sum_{i=t_0+1}^{t_0+\tau} q_{X_i}}{\tau}\right)^{\tau}$$

The right-hand side above is at most e^{-1}, if $\sum_{i=t_0+1}^{t_0+\tau} q_{X_i} \geq 1$, so the probability of not marking any vertex in $F(v, t_0+1)$ during the steps $t_0+1, \ldots, t_0+\tau$ is at most

$$e^{-1} + \mathbf{Pr}\left(\sum_{i=t_0+1}^{t_0+\tau} q_{X_i} < 1\right). \tag{3.13}$$

We now bound

$$\mathbf{Pr}\left(\sum_{i=t_0+1}^{t_0+\tau} q_{X_i} < 1\right)$$

by applying (3.11). To do so, it remains to verify that $\tau \geq C \cdot \frac{1}{1-\lambda_2} \cdot \frac{1}{\mu}$, that is, that $\mu \geq 4d/n$. Let $\bar{d} := d/2^i$ and recall that every vertex in phase i has at most $\bar{d}$ free neighbors. Since q_w can be positive only for $w \in W = N(F(v,t)) \cap C(t)$, we have

$$\begin{aligned}
\mu &= \frac{1}{n}\sum_{w \in V} q_w \\
&= \frac{1}{n}\sum_{w \in W} \frac{|F(v, t_0+1) \cap F(w, t_0+1)|}{|F(w, t_0+1)|} \\
&\geq \frac{1}{n}\sum_{w \in W} \frac{|F(v, t_0+1) \cap F(w, t_0+1)|}{\bar{d}} \\
&= \frac{1}{n\bar{d}} \sum_{u \in F(v, t_0+1)} |C(u, t_0+1)| \\
&\geq \frac{1}{n\bar{d}} \sum_{u \in F(v, t_0+1)} \frac{|N(v)|}{2} \\
&\geq \frac{1}{n\bar{d}} \cdot \frac{\bar{d}}{2} \cdot \frac{d}{2} = \frac{d}{4n}.
\end{aligned}$$

Thus (3.11) holds, and it implies that (3.13) is at most $e^{-1} + 1/2 < 1$. This means that the probability of marking a vertex in $F(v, t_0+1)$ during the steps $t_0+1, \ldots, t_0+\tau$ is at least some positive constant.

□

LEMMA 12. *For each $i = 0, 1, \ldots, \log d$, if at the beginning of phase i each vertex has at most $d/2^i$ free neighbors, then at the end of this phase each vertex has at most $d/2^{i+1}$ free neighbors* **whp**.

Proof. Consider phase $i \geq 0$ and assume that at the beginning of this phase each vertex has at most $d/2^i$ free neighbors. Let v be a vertex that has more than $d/2^{i+1}$ free neighbors at the beginning of this phase. We say that an iteration of this phase is successful (w.r.t. vertex v), if a free neighbor of v is marked during this iteration, or the number of free neighbors of v is already at most $d/2^{i+1}$ by the beginning of this iteration. Recall that ν_i is the total number of iterations of phase i. Therefore, Lemma 11 implies that the number of successful iterations is stochastically larger than $X \sim \mathsf{Bin}(\nu_i, c)$, for some constant $c > 0$. The probability that at the end of the phase vertex v has more than $d/2^{i+1}$ free neighbors is not greater than the probability that there are less than $d/2^{i+1}$ successful iterations, which is not greater than $\mathbf{Pr}(X < d/2^{i+1})$. Recalling that $\nu_i = \Theta(\max(d/2^i, \log n))$ and using the following Chernoff bound for X

$$\mathbf{Pr}\left(X < \mathbf{E}(X) - \xi\right) \leq \exp\left(-\xi^2/(2\mathbf{E}(X))\right),$$

we get

$$\begin{aligned}\mathbf{Pr}\left(X < d/2^{i+1}\right) &\leq \exp\left(-(\mathbf{E}(X) - d/2^{i+1})^2/(2\mathbf{E}(X))\right) \\ &\leq \exp\left(-\Theta(d/2^i + \log n)\right) = O(n^{-4}).\end{aligned}$$

Therefore, we can take a union bound over all vertices $v \in V$ to conclude that at the end of phase i, each vertex has at most $d/2^{i+1}$ free neighbors **whp**. □

3.2 Strong Local Expander. In this section we prove Theorem 2. Let $F_r(v)$ and $C_r(v)$ denote the set of free vertices and the set of covered vertices, respectively, in the r-th neighborhood $N_r(v)$ of vertex v. For brevity, $F_1(v) = F(v)$. Let $\beta = \max\{1, \lceil(\log n)/d\rceil\}$. We view the process as consisting of the *initial phase* 0 of $O(n\beta)$ steps of the walk, followed by $\log d$ phases numbered $i = 1, 2, \ldots, \log d$. Each phase $i \geq 1$ consists of two parts. The first part has $t_i = 2 + \min\{i, \log d - i\}$ iterations each consisting of $\beta cn/2^i$ steps of the walk, where c is a suitably large constant. The second part consists of one further iteration of $\beta cn/2^{t_i}$ steps of the walk. The iterations in phase $i \geq 1$ are numbered $j = 0, 1, \ldots, t_i$, where iteration t_i is that single iteration in the second part of the phase.

The number of steps of the random walk used in phase i is (observing that $t_i \leq i + 2$)

$$s_i \;=\; \beta cn\left(\frac{t_i}{2^i} + \frac{1}{2^{t_i}}\right) \;\leq\; \beta cn\left(\frac{i+2}{2^i} + \frac{1}{2^{t_i}}\right),$$

so the total number of steps is $O(n\beta) + \sum_{i=1}^{\log d} s_i = O(n\beta) = O(n + (n\log n)/d)$.

We prove by induction that **whp** for each $i = 1, 2, \ldots, \log d + 1$, at the end of phase $i - 1$,

$$|F(v)| \leq \frac{d}{2^i}. \tag{3.14}$$

Thus at the end of the last, $\log d$, phase, **whp** there are no free vertices remaining. To establish that (3.14) holds at the end of phase $i - 1$, we actually prove the following more detailed property.

LEMMA 13. *For each $i = 1, 2, \ldots, \log d + 1$, at the end of phase $i - 1$ (the beginnig of phase i),*

$$|F_1(v)| \leq \frac{d}{2^i}, \quad \textit{and} \quad |F_3(v)| \leq \frac{d^3}{2^{2i}}.$$

The following lemma establishes that after the initial phase 0, **whp** $F(v) \leq d/4$ for all vertices v, that implies the base case in Lemma 13 (observe that $|F_3(v)| \leq |N_2(v)| \cdot \sum_{u\in N_2(v)} |F_1(v)| \leq d^2 \cdot \max\{F_1(u)\}$).

LEMMA 14. *After $O(n\beta)$ steps of the walk,* **whp** *$|F(v)| \leq d/4$ for all vertices v.*

Proof. If quarter of the vertices in $N(v)$ are unvisited, then there exists a subset S of $N(v)$ of size $d/4$ in which no vertex is visited. Using Lemma 1 and Lemma 8, the probability that such a set S exists is at most

$$n \cdot \binom{d}{d/4} \cdot e^{-(1-o(1))\frac{t\pi_S}{8\cdot\mathbf{E}[R_S]}} \leq n2^d e^{-(cd\beta)/2} \leq n^{-B}$$

for large constants c and B. □

Lemma 13 follows by induction on i using the following two lemmas (proven in Appendix).

LEMMA 15. *For a phase $i \geq 1$, assume that $|F(v)| \leq d/2^i$ and $|F_3(v)| \leq d^3/2^{2i}$ for all vertices $u \in V$ at the beginning of phase i. Then* **whp** *at the beginning of each iteration $j = 0, 1, \ldots, t_i$ in this phase, for all vertices*

$$|F_3(v)| \leq \frac{d^3}{2^{2i+j}}. \tag{3.15}$$

LEMMA 16. *For a phase $i \geq 1$, assume that at the beginning of the last iteration $j = t_i$ in this phase (the single iteration in the second part of the phase), $|F(u)| \leq d/2^i$ and $|F_3(u)| \leq d^3/2^{2i+j}$, for all vertices $u \in V$. Then at the end of this iteration (after $Lcn/2^j$ steps of the walk),* **whp** *$|F(u)| \leq \frac{d}{2^{i+1}}$ for all vertices $u \in V$.*

4 Analysis of RW-R-RANK

In this section we prove Theorem 3. In the following we assume that $G = (V, E)$ is either an almost Ramanujan graph or a local expander with parameter $L \geq \kappa \cdot \max\{\ln\ln n, \ln \mathbf{T}\}$, where κ is a large constant. To show Theorem 3, we adapt the techniques of [8] to our case. More precisely, we incorporate the results of Section 2 into the analysis of the corresponding coupon collecting process from [8]. First, we need some definitions. For any sequence of edges $\sigma = (u_1, v_1), \ldots, (u_r, v_r)$, define two sequences $src(\sigma) = u_1, \ldots, u_r$ and $dst(\sigma) = v_1, \ldots, v_r$ called *source* and *destination* nodes. A sequence of vertices is called *rank-sorted* if the associated sequence of vertex ranks is nondecreasing (cf. [8]).

Let $\mathcal{A}$ be the set of time steps in $\bigcup_{i=1}^{\infty}[2i\mathbf{T}, 2i\mathbf{T}+\mathbf{T}]$. An interval $[2i\mathbf{T}, 2i\mathbf{T}+\mathbf{T}]$ is called the ith phase of $\mathcal{A}$. Let $duration(\sigma)$ be defined for a sequence of vertices σ recursively as follows. If σ is empty, then $duration(\sigma) =$

0. Otherwise, let σ be defined by $\sigma = \tau : v$ for some shorter sequence τ and vertex v, i.e., σ consists of the sequence τ followed by vertex v. Let i be the earliest phase in set $\mathcal{A}$ such that $duration(\tau) < i$, and at least one vertex of $N(v)$ is visited by the random walk during that phase. Then, $duration(\sigma) = i$. If some vertex $u \in N(v)$ is visited in phase i, then we call u *corresponding* to v in σ. Furthermore, a sequence $u_1, \ldots, u_r$ is corresponding to $\sigma = v_1, \ldots, v_r$ if u_j is corresponding to v_j in σ for all $j \in \{1, \ldots, r\}$. According to [8] it holds that for any r-sequence of distinct vertices σ,

$$\mathbf{Pr}(\sigma \text{ is rank-sorted}) = \binom{r + n^2 - 1}{r} n^{-2r}.$$

The proof of the first statement of the following lemma is similar to the proof of Lemma 7 from [8]. The only difference is that $N(v)$ is visited by the random walk according to a distribution that can be upper bounded in our analysis by $\text{Geo}(\frac{\mathbf{T}d}{qn})$ (instead of $\text{Geo}(\frac{d}{n})$). Part 2 of the lemma is similar to Lemma 8 of [8]. The proof follows from Lemma 8 of [8] and the second statement of Corollary 5 in the Appendix.

Lemma 17. *For any r-sequence of vertices σ we have for any integer i*

- $\mathbf{Pr}(duration(\sigma) = i) \leq \mathbf{Pr}(X \geq i)$, *where* $X \sim \mathsf{NegBin}\left(r, \frac{\mathbf{T}d}{qn}\right)$.
- $\mathbf{Pr}(src(\sigma)$ *is corresponding to* $dst(\sigma)) = O(d)^{-r}$.

In the rest of this section we assume for simplicity that at the beginning of each phase i of $\mathcal{A}$, the random walk lies on any vertex with the same probability $1/n$ (instead of $1/n \cdot (1 \pm O(1/n^2))$). According to Lemma 9 of [8], for any integer i and r-sequence of edges σ, the events $\mathcal{E}_1 =$ "$dst((\sigma)$ is rank-sorted", $\mathcal{E}_2 =$ "$duration(dst(\sigma)) = i$", and $\mathcal{E}_3 =$ "$src(\sigma)$ is corresponding to $dst(\sigma)$" are mutually independent. Thus, we obtain the following lemma.

Lemma 18. *Let σ be an r-sequence of edges, where the vertices of $dst(\sigma)$ are distinct. Furthermore, let $X \sim \mathsf{NegBin}(r, \frac{\mathbf{T}d}{qn})$, let i be an integer, and let the events $\mathcal{E}_1$, $\mathcal{E}_2$, and $\mathcal{E}_3$ be defined as above. Then,*

$$\mathbf{Pr}(\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3) \leq \binom{r + n^2 - 1}{r} \cdot \mathbf{Pr}(X \geq i) \cdot O(n^2 d)^{-r}.$$

To proceed, we need some definitions introduced in [8]. An r-sequence of edges is called linked, if r equals 0 or 1, or the sequence has the form $\tau : (u, v) : (u', v')$, where $\tau : (u, v)$ is a linked sequence of length $r - 1$, and $(u, v') \in E$. The pointer $parent(v)$ is defined inductively, by setting $parent(v) = Nil$ at the beginning. In phase i, in which v is covered, $parent(v)$ is set to a vertex w that has been covered in a step t of phase $j < i$ with the following property. A vertex of $N(v)$ has been visited by the random walk in step t of phase j, and no vertex of $N(v)$ is visited in any phase between j and i. An r-sequence σ is called *active* if $r = 0$ or $\sigma = v$ with $parent(v) = Nil$, or $\sigma = \tau : v : v'$ with $\tau : v$ active and $parent(v') = v$. An r-sequence of edges σ is active if $dst(\sigma)$ is active and $src(\sigma)$ corresponds to $dst(\sigma)$. An r-sequence of edges σ is i-active if it is active, and either $r = i = 0$ or $\sigma = \tau : (u, v)$, where v is the vertex covered in phase i.

According to Lemma 19 of [8], it holds that some r-sequence of edges is i-active with probability at most

$$n d^{2r-1} \binom{r + n^2 - 1}{r} \mathbf{Pr}(X \geq i) \cdot O(n^2 d)^{-r},$$

where $X \sim \mathsf{NegBin}(r, \frac{\mathbf{T}d}{qn})$. Then, the following lemma holds.

Lemma 19. *If $i \geq q'' n / \mathbf{T}$, where q'' is a large constant, then for any integer r we have*

$$d^{2r-1} \binom{r + n^2 - 1}{r} \mathbf{Pr}(X \geq i) \cdot O(n^2 d)^{-r} \leq \exp(-id\mathbf{T}/(32qn)),$$

where $X \sim \mathsf{NegBin}(r, \frac{\mathbf{T}d}{qn})$.

By setting $i = cqn/\mathbf{T}$, where c is a large constant, we obtain Theorem 3.

References

[1] M. Adler, E. Halperin, R.M. Karp, and V.V. Vazirani. A Stochastic Process on the Hypercube with Applications to Peer-to-Peer Networks. In *35th Annual ACM Symposium on Theory of Computing (STOC '03)*, pages 575–584, 2003.

[2] D.J. Aldous and J.A. Fill. *Reversible Markov Chains and Random Walks on Graphs.* (draft availabe at www.stat.berkeley.edu/aldous/RWG/book.html), 2002.

[3] R. Aleliunas, R.M. Karp, R.J. Lipton, L. Lovász, and C. Rackoff. Random walks, universal traversal sequences, and the complexity of maze problems. In *20th Annual IEEE Symposium on Foundations of Computer Science (FOCS '79)*, pages 218–223, 1979.

[4] N. Alon. Problems and results in extremal combinatorics - II. *Discrete Mathematics*, 308(19):4460–4472, 2008.

[5] N. Alon, I. Benjamini, E. Lubetzky, and S. Sodin. Non-backtracking random walks mix faster. *Communications in Contemporary Mathematics*, 9:585–603, 2007.

[6] A.K. Chandra, P. Raghavan, W.L. Ruzzo, R. Smolensky, and P. Tiwari. The electrical resistance of a graph captures its commute and cover times. *Computational Complexity*, 6(4):312–340, 1997.

[7] C. Cooper and A. Frieze. The cover time of random regular graphs. *SIAM Journal of Discrete Mathematics*, 18(4):728–740, 2005.

[8] N. B. Dimitrov and C. G. Plaxton. Optimal Cover Time for a Graph-Based Coupon Collector Process. In *32nd International Colloquium on Automata, Languages, and Programming (ICALP '05)*, pages 702–716, 2005.

[9] R. Elsässer, L. Gąsieniec. Radio communication in random graphs. *Journal of Computer and System Sciences*, 72(3):490–506, 2006.

[10] R. Elsässer and T. Sauerwald. Tight Bounds for the Cover Time of Multiple Random Walks. In *36nd International Colloquium on Automata, Languages, and Programming (ICALP '09)*, pages 415–426, 2009.

[11] U. Feige. A tight lower bound for the cover time of random walks on graphs. *Random Structures and Algorithms*, 6(4):433–438, 1995.

[12] U. Feige. A tight upper bound for the cover time of random walks on graphs. *Random Structures and Algorithms*, 6(1):51–54, 1995.

[13] W. Feller. *An Introduction to Probability Theory*, volume I. Wiley, 2nd edition, 1960.

[14] J. Friedman. On the second eigenvalue and random walks in random d-regular graphs. *Combinatorica*, 11:331–362, 1991.

[15] S. Ikeda, I. Kubo, N. Okumoto, and M. Yamashita. Impact of Local Topological Information on Random Walks on Finite Graphs. In *30st International Colloquium on Automata, Languages, and Programming (ICALP '03)*, pages 1054–1067, 2003.

[16] P. Lezaud. Chernoff-type bound for finite markov chains. *The Annals of Applied Probability*, 8(3):849–867, 1998.

[17] L. Lovász. Random walks on graphs: A survey. *Combinatorics, Paul Erdös is Eighty (Volume 2)*, pages 1–46, 1993.

A Omitted Proofs from Section 2

A.1 Proof of Lemma 2. Let $A_i = \{X_{t+i} = w\}$ be the event that the walk visits w at step $t+i$, and let

$$C = A_0 \cup A_1 \cup \cdots \cup A_{T-1},$$

so $\mathbf{Pr}(C) = \mathbf{Pr}(w \in \{X_{t+i} : i = 0, ..., \mathbf{T}-1\})$. Define similarly

$$C' = B_0 \cup B_1 \cup \cdots \cup B_{\mathbf{T}-1},$$

so that $\mathbf{Pr}(C') = \mathbf{Pr}(w \in \{X_{t'+i} : i = 0, ..., \mathbf{T}-1\})$. Decompose C according to the last visit to w. Thus

$$\begin{aligned}\mathbf{Pr}(C) = \mathbf{Pr}(&A_0 \cap (\overline{A_1} \cap \cdots \cap \overline{A_{\mathbf{T}-1}})) \\ &+ \mathbf{Pr}(A_1 \cap (\overline{A_2} \cap \cdots \cap \overline{A_{\mathbf{T}-1}})) + \cdots \\ &+ \mathbf{Pr}(A_{\mathbf{T}-2} \cap (\overline{A_{\mathbf{T}-1}})) + \mathbf{Pr}(A_{\mathbf{T}-1}).\end{aligned}$$

By Markov renewal, $\mathbf{Pr}(\overline{A_{i+1}} \cap \cdots \cap \overline{A_{\mathbf{T}-1}} \mid A_i)$ is the probability that a walk starting from w does not return to w during $\mathbf{T}-i$ steps. Thus

$$\begin{aligned}&\mathbf{Pr}(\overline{A_{i+1}} \cap \cdots \cap \overline{A_{\mathbf{T}-1}} \mid A_i) \\ &= \mathbf{Pr}(\overline{B_{i+1}} \cap \cdots \cap \overline{B_{\mathbf{T}-1}} \mid B_i).\end{aligned}$$

Also $\mathbf{Pr}(B_i) = (1+o(1))\mathbf{Pr}(A_i)$ as $\mathbf{T}$ is a mixing time. Thus

$$\begin{aligned}&\mathbf{Pr}(A_i \cap (\overline{A_{i+1}} \cap \cdots \cap \overline{A_{\mathbf{T}-1}})) \\ &= (1+o(1))\mathbf{Pr}(B_i \cap (\overline{B_{i+1}} \cap \cdots \cap \overline{B_{\mathbf{T}-1}})),\end{aligned}$$

completing the proof of the lemma. □

A.2 Proof of Lemma 3. Let ψ denote the distribution of the random walk at step $t \geq \mathbf{T}$. We set $a = 2\lceil \mathbf{E}_\psi[H_w]/\mathbf{T} \rceil$ and have

$$\mathbf{Pr}(H_w > a\mathbf{T}) \leq 1/2. \tag{A.1}$$

Denoting by $\mathcal{E}(r)$ the event that $w \in \{X_{t+r\mathbf{T}+i} : i = 0, ..., \mathbf{T}-1\}$, for an integer $r \geq 0$, and using Lemma 2, we have

$$\begin{aligned}\frac{1}{2} &\leq \mathbf{Pr}(H_w \leq a\mathbf{T}) \\ &\leq \sum_{r=0}^{a-1} \mathbf{Pr}(\mathcal{E}(r)) \leq (1+o(1))a\mathbf{Pr}(\mathcal{E}(0)).\end{aligned} \tag{A.2}$$

Thus

$$\begin{aligned}&\mathbf{Pr}(w \in \{X_{t+i} : i = 0, ..., \mathbf{T}-1\}) \\ &= \mathbf{Pr}(\mathcal{E}(0)) \\ &\geq \frac{1}{2a(1+o(1))} \\ &\geq \frac{\mathbf{T}(1 - O(\mathbf{T}/\mathbf{E}_\psi[H_w]))}{4\mathbf{E}_\psi[H_w](1+o(1))}.\end{aligned} \tag{A.3}$$

Since $t \geq \mathbf{T}$, we know that $\mathbf{E}_\psi[H_w] = (1 \pm o(1))\mathbf{E}_\pi[H_w]$, and $\mathbf{T}/\mathbf{E}_\pi[H_w] = o(1)$ as shown below. Thus (A.3) implies the first bound in (2.5).

Let $d(t) = \max_{u,x} |P_u^{(t)}(x) - \pi_x|$. It follows from e.g. [2] that $d(s+t) \leq 2d(s)d(t)$. Hence, since $\max_{u,x} |P_u^{(\mathbf{T})}(x) - \pi_x| \leq \pi_w$, then for each $k \geq 1$, $\max_{u,x} |P_u^{(k\mathbf{T})}(x) - \pi_x| \leq (2\pi_w)^k$. Thus

$$\begin{aligned} Z_{ww} &= \sum_{t=0}^{\infty} (P_w^{(t)}(w) - \pi_w) \\ &= \sum_{t<\mathbf{T}} (P_w^{(t)}(w) - \pi_w) + \mathbf{T} \sum_{k\geq 1} (2\pi_w)^k \\ &= \mathbf{E}[R_w] - \mathbf{T}\cdot\pi_w + O(\mathbf{T}\cdot\pi_w). \end{aligned} \tag{A.4}$$

Using (A.4), $O(\mathbf{T}\cdot\pi_w) = o(1)$, and $\mathbf{E}[R_w] \geq 1$, we get

$$\frac{\mathbf{T}}{\mathbf{E}_\pi[H_w]} = \frac{\mathbf{T}\cdot\pi_w}{Z_{ww}} = (1+o(1))\frac{\mathbf{T}\cdot\pi_w}{\mathbf{E}[R_w]} = o(1),$$

and the second bound in (2.5) follows. □

A.3 Proof of Lemma 7. We assume that the random walk is at a vertex of $N(v)$ at step 0. Using the same arguments as in the proof of Lemma 6, the random walk moves to $N_2(v)$ with probability $p \geq c/(1+c) > 1/2$. Once the walk has moved to $N_2(v)$ it can be coupled with a random walk on a path with loops, and absorbing barriers at $1, \omega$, where 1 represents $N(v)$ and ω is $N_L(v)$. The probability of absorption at 1 is at most q/p, where $p \geq c/(1+c) > 1/2$ and $q \leq 1/(1+c) < 1/2$. Thus, the walk returns to 1 before hitting ω with probability at most $1 - p + q$.

The probability of a return to $N(v)$ from distance ω at any step $i \leq \mathbf{T}$ is at most $(1/c)^\omega \leq \mathbf{T}^{-A}$. Thus provided $A > 1$ the total expected number of returns is $O(1)$. □

A.4 Proof of Lemma 8. If the random walk is at some vertex of $\cup_{i=0}^4 N_i(v)$ at time 0, where $N_0(v) = \{v\}$, then the walk moves to $N_5(v)$ within the next $5 - i$ steps with probability $(1 - O(1/d))^{5-i} = 1 - O(1/d)$. Once the random walk moves beyond $N_4(v)$, it can be coupled with a random walk on a path with loops, and absorbing barriers at $4, \omega$, where 4 represents $N_4(v)$ and ω is $N_L(v)$. Thus, we may use the same arguments as in the proof of Lemma 7 to show that the expected number returns to $\cup_{i=1}^4 N_i(v)$ (and to S) is $O(1)$. □

B Omitted Proofs from Section 3

B.1 Proof of Lemma 15. The base case, for $j = 0$, follows from the assumptions of the lemma. Let now $j \geq 0$ and assume that at the beginning of iteration j, (3.15) holds for all vertices, and consider a vertex v such that $|F_3(v)| > d^3/2^{2i+j+1}$. Let $S = N(F_3(v)) \cap C_4(v)$ be the covered vertices of $N_4(v)$ adjacent to $F_3(v)$. We count the number x of the edges between sets $F_3(v)$ and S. Since vertices of S have degree at most α in $F_3(v)$ (see the definition of Strong Local Expander), then $x \leq \alpha|S|$. Since each vertex in $F_3(v)$ has at least $d/2$ covered neighbors and at most α of them are in $N_2(v) \cup N_3(v)$, then $x \geq (d/2 - \alpha)|F_3(v)|$. Thus

$$\alpha|S| \geq \left(\frac{d}{2} - \alpha\right)|F_3(v)|. \tag{B.5}$$

Let $\phi = d^3/2^{2i+j}$. Let $l_i = \beta cn/2^i$ be the number of steps in iteration j. Lemma 5 and Lemma 8 imply that for a sequence of $\Theta(n/|S|)$ steps, the probability that at least one vertex in $|S|$ is visited during these steps is at least some positive constant, provided that $\mathbf{T}|S| = o(n)$. When a vertex u of S is visited in the current step, then there is at least $1/|F(u)| \geq 2^i/d$ probability of marking a vertex of $F_3(v)$. Thus with some positive constant probability, a sequence of $\Theta(dn/(|S|2^i))$ steps is successful in marking a vertex of $F_3(v)$. In one iteration we have $\Omega(\beta|S|/d)$ non-overlapping sequences of $\Theta(dn/(|S|2^i))$ steps. If at the end of the iteration we still have $|F_3(v)| > d^3/2^{2i+j+1}$, then this means that fewer than $\phi = d^3/2^{2i+j+1}$ of these sequences have been successful, while throughout this iteration $|S| \geq \beta_1 d\phi$ for some positive constant β_1 (see (B.5)). We have $\phi \geq \beta_2 d$ for a positive constant β_2 (recall that $i + j \leq i + t_i \leq \log d + 2$). Thus the probability that fewer than $\phi = d^3/2^{2i+j+1}$ of these sequences have been successful is at most

$$e^{-\Omega(\beta|S|/d)} = n^{-B}$$

for a large constant B. The union bound over all vertices implies that **whp** at the end of iteration j, $|F_3(v)| \leq d^3/2^{2i+j+1}$ for all vertices. □

B.2 Proof of Lemma 16. Consider a node v that has more than $d/2^{i+1}$ free neighbors at the beginning of this iteration. Let $C_2'(v) = C_2(v) \cap N(F_1(v))$. Note that $|C_2'(v)| \geq d(d/2-\alpha)/(\alpha 2^{i+1})$ as (similarly to (B.5))

$$\alpha|C_2'(v)| \geq \left(\frac{d}{2} - \alpha\right)|F(v)| \geq \left(\frac{d}{2} - \alpha\right)\frac{d}{2^{i+1}}. \tag{B.6}$$

Let

$$S = \left\{u \in C_2'(v) : |N(u) \cap F_3(v)| \leq \frac{9\alpha^2 d}{2^{i+j}}\right\}.$$

We claim that $|S| > |C_2'(v)|/2$. Indeed, supposing that $|S| \leq |C_2'(v)|/2$ and using B.6 and and the assumption that $|F_3(v)| \le d^3/2^{2i+j}$, we get

$$\begin{aligned}\sum_{u\in C_2'(v)} |N(u)\cap F_3(v)| &\ge \frac{9\alpha^2 d}{2^{i+j}}(|C_2'(v)| - |S|)\\ &\ge \frac{9\alpha^2 d}{2^{i+j}}\cdot\frac{d(d/2-\alpha)}{2\alpha\, 2^{i+1}}\\ &\ge \frac{9\alpha(d/2-\alpha)}{4d}|F_3(v)|\\ &> \alpha|F_3(v)|,\end{aligned}$$

which is a contradiction. Thus $|S| \ge d^2/(\alpha 2^{i+3})$.

The rest of the proof follows the same lines as in the proof of Lemma 15. □

C Omitted Proofs from Section 4

In this section we complete the proof of Theorem 3. In order to prove the theorem, we first derive some properties of almost Ramanujan graphs.

Lemma 20. *Let $G = (V, E)$ be an almost Ramanujan graph of degree $d = \Omega(\log n)$. Then, there exists a constant $\epsilon < 1$ such that for any vertex $v \in V$, $j = 1, \dots,$ and constant ρ, we have*

$$\mathbf{Pr}(|\{X_t \in N_1(v) \mid 2j\mathbf{T} \le t \le 2j\mathbf{T} + \mathbf{T}\}| \ge \rho\phi) \le \epsilon^\rho,$$

where ϕ is a constant.

Proof. We know that in any regular graph $P_u^{(t)}(w) \le 1/n + \lambda_2^t$. This implies that for a certain constant ϕ, we get $P_u^{(\phi+\tau)}(w) \le 1/d^4$, where u and w are two vertices of $N_1(v)$ and $\tau \ge 0$. Applying the union bound over all pairs of vertices in $N_1(v)$ and over $\mathbf{T} = O(\log n)$ steps, we obtain that

$$\mathbf{Pr}\left(\left|\bigcup_{\tau=0}^{\mathbf{T}-\phi} X_{t+\phi+\tau} \cap N_1(v)\right| \ge 1 \mid X_t \in N_1(v)\right) \le \epsilon$$

for some $\epsilon < 1$. This implies that if the random walk visits $N_1(v)$ at some time t, it returns to $N_1(v)$ in time period $[t + \phi, t + \mathbf{T}]$ with success probability ϵ, independently of previous returns. Thus,

$$\mathbf{Pr}(|\{X_t \in N_1(v) \mid 2j\mathbf{T} \le t \le 2j\mathbf{T} + \mathbf{T}\}| \ge \rho\phi) \le \epsilon^\rho.$$

□

Lemmas 6-7, 20, and Corollary 4 imply the following result.

Corollary 5. *Let κ be a large constant and let $G = (V, E)$ be an almost Ramanujan graph or a local expander with $L \ge \kappa \cdot \max\{\ln\ln n, \ln \mathbf{T}\}$. Then*

- *there exists a constant q such that for any integer i and vertex v it holds that*

$$\mathbf{Pr}\left(\left|\bigcup_{j=2i\mathbf{T}}^{2i\mathbf{T}+\mathbf{T}} X_j \cap N_1(v)\right| \ge 1\right) \ge \frac{1}{q}\cdot\frac{\mathbf{T}d}{n}.$$

- *for any integer i and neighbor w of a vertex v it holds that*

$$\mathbf{Pr}\left(w \in \bigcup_{j=2i\mathbf{T}}^{2i\mathbf{T}+\mathbf{T}} X_j \;\middle|\; \left|\bigcup_{j=2i\mathbf{T}}^{2i\mathbf{T}+\mathbf{T}} X_j \cap N_1(v)\right| \ge 1\right) = O\left(\frac{1}{d}\right).$$

C.1 Proof of Lemma 19.

According to Lemma 21 of [8], we have

$$\lambda := d^{2r-1}\binom{r+n^2-1}{r}\cdot O(n^2 d)^{-r} \le \exp(\sqrt[4]{q''}d),$$

if q'' is large enough. Consider now the cases $r > id\mathbf{T}/(2qn)$ and $r \le id\mathbf{T}/(2qn)$. If $r > id\mathbf{T}/(2qn)$ and $q'' > q^4 > 64^2$, then

$$\begin{aligned}d^{2r-1}\binom{r+n^2-1}{r}\cdot O(n^2 d)^{-r} &\le e^{-r}\\ &\le \exp(-id\mathbf{T}/(64qn)),\end{aligned}$$

and the lemma follows.

If $r \le id\mathbf{T}/(2qn)$, then define $Y \sim \mathsf{NegBin}(\frac{i\mathbf{T}d}{2qn}, \frac{\mathbf{T}d}{qn})$. We know that

$$\mathbf{Pr}(X \ge i) \le \mathbf{Pr}(Y \ge i),$$

and

$$\mathbf{Pr}(Y \ge i) \ge \exp\left(-\frac{i\mathbf{T}d}{16qn} + \frac{1}{8}\right).$$

If $q'' > q^4 > 64^2$, then

$$\begin{aligned}\lambda\cdot\mathbf{Pr}(X \ge i) &\le \exp(\sqrt[4]{q''}d)\cdot\mathbf{Pr}(Y \ge i)\\ &\le \exp\left(-\frac{i\mathbf{T}d}{16qn} + \frac{1}{8} + \sqrt[4]{q''}d\right)\\ &\le \exp\left(-\frac{id\mathbf{T}}{32qn}\right),\end{aligned}$$

and the lemma follows. □

D Proof of Theorem 4

We first describe the construction of the graph $G = (V(G), E(G))$, which is similar to a Cartesian product of a complete graph with $\log n$ vertices and a three-dimensional Torus graph with $n/\log n$ vertices. More formally, for all $0 \le x, y, z \le \sqrt[3]{n/\log n} - 1$, $0 \le d \le \log n$ we have the following nodes.

$$V(G) := \Big\{(x, y, z, d)\Big\}.$$

For all $0 \le x, y, z \le \sqrt[3]{n/\log n} - 1, 0 \le d, d' < \log n$, we have the following edges.

$$\begin{aligned} E(G) := \Big\{ &\{(x, y, z, d), (x \pm 1, y, z, d)\} \\ &\cup \{(x, y, z, d), (x, y \pm 1, z, d)\} \\ &\cup \{(x, y, z, d), (x, y, z \pm 1, d)\} \\ &\cup \{(x, y, z, d), (x, y, z, d')\} \Big\}. \end{aligned}$$

Observe that $|V(G)| = n \cdot (1 + (1/\log n))$. Hence G is the Cartesian product of a complete graph with $\log n$ vertices and a three-dimensional Torus graph with $n/\log n$ vertices, except that all vertices of a complete graph (with respective coordinates $(x, y, z, .)$) are connected to a distinguished vertex $(x, y, z, \log n)$. Note that each such vertex $(x, y, z, \log n)$ is *not* connected to any other complete graph.

Our first claim is that the cover time of G is $\Theta(n \log n)$. To prove this, we shall distinguish between Torus-steps and complete-graph-steps. A Torus-step of a random walk in G is a step where one of the first three coordinates is changed, the other steps are complete-graph-steps. We observe that if the random walk does a Torus-step, it chooses one of the first coordinates uniformly at random and adds or subtracts one with equal probability. Hence we can project the random walk on $V(G)$ onto a random walk on a three-dimensional Torus graph with $n/(\log n)$ vertices by ignoring the complete-graph-steps.

We know that the expected time to hit any fixed vertex in a three-dimensional Torus graph is $\Theta(n)$ (cf. [6]). To prove that this also holds in G, fix an arbitrary vertex (x_0, y_0, z_0, d_0) of $V(G)$. It follows from the result for the three-dimensional Torus that any vertex of the set $S := \{(x_0, y_0, z_0, d) \mid 0 \le d \le \log n\}$ in G is visited after $O(n/\log n)$ Torus-steps. Using a Chernoff bound, one can easily prove that after $O(n)$ steps in G, the random walk does at least $\Omega(n/\log n)$ Torus-steps. Finally, we observe that once the random walk is in the set S, it will visit the vertex (x_0, y_0, z_0, d_0) with constant probability within the next $O(\log n)$ steps. Combining this, we conclude that the vertex (x_0, y_0, z_0, d_0) is visited in expected $O(n)$ steps. This also implies that the random walk visits all vertices in $\Theta(n \log n)$ steps.

Our next claim is that after $c_1 \cdot (n \log n)$ steps of a random walk (with a sufficiently small constant c_1), there exists a set $\{(x_0, y_0, z_0, d) \mid 0 \le d \le \log n\}$ containing only uncovered vertices. This directly implies that RW-R-RANK and RW-RAND require $\Omega(n \log n)$ steps since the vertex $(x_0, y_0, z_0, \log n)$ is unmarked. Again we use the distinction of Torus-steps and complete-graph-steps. Consider first only Torus-steps. We know that it takes in expectation $c_2 n$ steps ($c_2 > 0$ a constant) to visit all vertices of a three-dimensional Torus graph with $n/\log n$ vertices. Consequently, after $(c_2/2) \cdot n$ steps, at least one uncovered vertex, say (x_0, y_0, z_0) remains with probability at least $1/2$. However, to perform $(c_2/2) \cdot n$ Torus steps, a random walk on G has to take $c_1 \cdot (n \log n)$ steps in total. Hence after $c_1 \cdot (n \log n)$ steps of RW-R-RANK and RW-RAND, the vertex $(x_0, y_0, z_0, \log n)$ will be still unvisited. □