

Vol. 9 (2004), Paper no. 26, pages 825-845.

Journal URL

<http://www.math.washington.edu/~ejpecp/>

MIXING TIMES FOR RANDOM WALKS ON FINITE LAMPLIGHTER GROUPS

YUVAL PERES* AND DAVID REVELLE†

DEPARTMENT OF STATISTICS, UNIVERSITY OF CALIFORNIA
367 EVANS HALL # 3860, BERKELEY, CA, 94720-3860
PERES@STAT.BERKELEY.EDU AND REVELLE@STAT.BERKELEY.EDU

ABSTRACT. Given a finite graph G , a vertex of the lamplighter graph $G^\diamond = \mathbb{Z}_2 \wr G$ consists of a zero-one labeling of the vertices of G , and a marked vertex of G . For transitive G we show that, up to constants, the relaxation time for simple random walk in G^\diamond is the maximal hitting time for simple random walk in G , while the mixing time in total variation on G^\diamond is the expected cover time on G . The mixing time in the uniform metric on G^\diamond admits a sharp threshold, and equals $|G|$ multiplied by the relaxation time on G , up to a factor of $\log |G|$. For $\mathbb{Z}_2 \wr \mathbb{Z}_n^2$, the lamplighter group over the discrete two dimensional torus, the relaxation time is of order $n^2 \log n$, the total variation mixing time is of order $n^2 \log^2 n$, and the uniform mixing time is of order n^4 . For $\mathbb{Z}_2 \wr \mathbb{Z}_n^d$ when $d \geq 3$, the relaxation time is of order n^d , the total variation mixing time is of order $n^d \log n$, and the uniform mixing time is of order n^{d+2} . In particular, these three quantities are of different orders of magnitude.

AMS 2000 subject classifications: Primary: 60J10; Secondary: 60B15.

Key words and phrases: Random walk, mixing time, lamplighter group, cover time

Submitted to EJP on April 6, 2004. Final version accepted on August 23, 2004.

Date: November 17, 2004

*Research supported by NSF Grants DMS-0104073 and DMS-0244479.

†Research partially supported by an NSF postdoctoral fellowship.




FIGURE 1. Lamplighter group over a cycle

1. INTRODUCTION

Given a finite graph $G = (V_G, E_G)$, the wreath product $G^\diamond = \mathbb{Z}_2 \wr G$ is the graph whose vertices are ordered pairs (f, x) , where $x \in V_G$ and $f \in \{0, 1\}^{V_G}$. There is an edge between (f, x) and (h, y) in the graph G^\diamond if x, y are adjacent in G and $f(z) = h(z)$ for $z \notin \{x, y\}$. These wreath products are called *lamplighter graphs* because of the following interpretation: place a lamp at each vertex of G ; then a vertex of G^\diamond consists of a configuration f indicating which lamps are on, and a lamplighter located at a vertex $x \in V_G$.

In this paper we estimate mixing time parameters for random walk on a lamplighter graph G^\diamond by relating them to hitting and covering times in G . When G is the two dimensional discrete torus \mathbb{Z}_n^2 , we prove in Theorem 1.1 that:

- the relaxation time $T_{\text{rel}}((\mathbb{Z}_n^2)^\diamond)$ is of order $n^2 \log n$;
- the mixing time in total variation, $\mathcal{T}_v(\epsilon, (\mathbb{Z}_n^2)^\diamond)$ is asymptotic to $cn^2 \log^2 n$;
- the uniform mixing time $\tau(\epsilon, (\mathbb{Z}_n^2)^\diamond)$ is asymptotic to Cn^4 .

(We recall the definitions of these mixing time parameters in (1)-(3).) The general correspondence between notions of mixing on G^\diamond and properties of random walk on G is indicated in the following table:

underlying graph G	lamplighter graph G^\diamond
maximal hitting time t^*	relaxation time T_{rel}
expected cover time $\mathbb{E}\mathcal{C}$	total variation mixing time \mathcal{T}_v
$\inf\{t : \mathbb{E}2^{ S_t } < 1 + \epsilon\}$	uniform mixing time τ

where S_t is the set of unvisited sites in G at time t . The connections indicated in this table are made precise in Theorems 1.2-1.4 below.

When G is a Cayley graph, the wreath product $G^\diamond = \mathbb{Z}_2 \wr G$ can also be thought of as the Cayley graph of the semi-direct product $G \ltimes \sum_G \mathbb{Z}_2$, where the action on $\sum_G \mathbb{Z}_2$ is by coordinate shift. This means that multiplication in G^\diamond is given by $(f, x)(h, y) = (\psi, xy)$, where $\psi(i) = f(i) + h(x^{-1}i)$. In

the case when G is a cycle or a complete graph, random walks on G^\diamond were analyzed by Häggström and Jonasson [9].

Definitions. Let $\{X_t\}$ be an irreducible Markov chain on a finite graph G with transition probabilities given by $p(x, y)$. Let $p^t(x, y)$ denote the t -fold transition probabilities and μ the stationary distribution. The *relaxation time* is given by

$$(1) \quad T_{\text{rel}} = \max_{i: |\lambda_i| < 1} \frac{1}{1 - |\lambda_i|}$$

where the λ_i are the eigenvalues of the transition matrix $p(x, y)$. The ϵ -mixing time in total variation $\mathcal{T}_v(\epsilon, G)$ and the ϵ -uniform mixing time $\tau(\epsilon, G)$ are defined by:

$$(2) \quad \mathcal{T}_v(\epsilon, G) = \min \left\{ t : \frac{1}{2} \sum_y |p^t(x, y) - \mu(y)| \leq \epsilon \ \forall x \in G \right\}$$

and

$$(3) \quad \tau(\epsilon, G) = \min \left\{ t : \left| \frac{p^t(x, y) - \mu(y)}{\mu(y)} \right| \leq \epsilon \ \forall x, y \in G \right\}.$$

When the graph G is clear, we will often abbreviate

$$(4) \quad \mathcal{T}_v = \mathcal{T}_v(G) = \mathcal{T}_v(1/(2e), G).$$

Another key parameter for us will be the *maximal hitting time*

$$(5) \quad t^* = t^*(G) = \max_{x, y} \mathbb{E}_x T_y,$$

where T_y is the hitting time of y .

The random walk we analyze on $\mathbb{Z}_2 \wr G$ is constructed from a random walk on G as follows. Let p denote the transition probabilities in the wreath product and q the transition probabilities in G . For $a \neq b$, $p[(f, a), (h, b)] = q(a, b)/4$ if f and h agree outside of $\{a, b\}$, and when $a = b$, $p[(f, a), h(a)] = q(a, a)/2$ if f and h agree off of $\{a\}$. A more intuitive description of this is to say that at each time step, the current lamp is randomized, the lamplighter moves, and then the new lamp is also randomized. The second lamp at b is randomized in order to make the chain reversible. To avoid periodicity problems, we will assume that the underlying random walk on G is already aperiodic.

Our first theorem describes the mixing time of the random walk on the wreath product when the lamplighter moves in the d -dimensional discrete torus \mathbb{Z}_n^d .

Theorem 1.1. *For the random walk $\{X_t\}$ on $(\mathbb{Z}_n^2)^\diamond = \mathbb{Z}_2 \wr \mathbb{Z}_n^2$ in which the lamplighter performs simple random walk with holding probability $1/2$ on \mathbb{Z}_n^2 , the relaxation time satisfies*

$$(6) \quad \frac{1}{\pi \log 2} \leq \frac{T_{\text{rel}}((\mathbb{Z}_n^2)^\diamond)}{n^2 \log n} \leq \frac{16}{\pi \log 2} + o(1).$$

For any $\epsilon > 0$, the total variation mixing time satisfies

$$(7) \quad \lim_{n \rightarrow \infty} \frac{\mathcal{T}_v(\epsilon, (\mathbb{Z}_n^2)^\diamond)}{n^2 \log^2 n} = \frac{8}{\pi},$$

and the uniform mixing time satisfies

$$(8) \quad C_2 \leq \frac{\tau(\epsilon, (\mathbb{Z}_n^2)^\diamond)}{n^4} \leq C'_2$$

for some constants C_2 and C'_2 . The uniform mixing time also has a sharp threshold and

$$(9) \quad \tau(1/2, (\mathbb{Z}_n^2)^\diamond) - \tau(\epsilon, (\mathbb{Z}_n^2)^\diamond) = O(n^2 \log n).$$

More generally, for any dimension $d \geq 3$, there are constants C_d and R_d independent of ϵ such that on $\mathbb{Z}_2 \wr \mathbb{Z}_n^d = (\mathbb{Z}_n^d)^\diamond$, the relaxation time satisfies

$$(10) \quad \frac{R_d}{4 \log 2} \leq \frac{T_{\text{rel}}((\mathbb{Z}_n^d)^\diamond)}{n^d} \leq \frac{8R_d}{\log 2} + o(1),$$

the total variation mixing time satisfies

$$(11) \quad \frac{R_d}{2} + o(1) \leq \frac{\mathcal{T}_v(\epsilon, (\mathbb{Z}_n^d)^\diamond)}{n^d \log n} \leq R_d + o(1),$$

and the uniform mixing time satisfies

$$(12) \quad C_d \leq \frac{\tau(\epsilon, (\mathbb{Z}_n^d)^\diamond)}{n^{d+2}} \leq C'_d,$$

with a sharp threshold in that

$$(13) \quad \tau(1/2, (\mathbb{Z}_n^d)^\diamond) - \tau(\epsilon, (\mathbb{Z}_n^d)^\diamond) = O(n^d).$$

The parameter R_d is the expected number of returns to 0 by a simple random walk in \mathbb{Z}^d , and is given by equation (75) in Chapter 5 of [2]. The reason why the exact limit in (7) can be computed is related to the fractal structure of the unvisited set at times up to the covering time. The geometry of this set is not sufficiently well understood in higher dimensions to make it possible to eliminate the factor of two difference between the upper and lower bounds of (11).

Remark. In one dimension, the total variation mixing time on $\mathbb{Z}_2 \wr \mathbb{Z}_n$ was studied in [9], and shown to be on the order of n^2 . More generally, the case of $G \wr \mathbb{Z}_n$ for any finite G has similar behavior [13]. For the walks we consider, it is easy to show that the relaxation time in one dimension is on the order of n^2 , while the uniform mixing time is on the order of n^3 . We compute the asymptotic constant for the uniform mixing time in Section 6.

The d dimensional result for uniform mixing is suggested by analogy with the infinite case of $\mathbb{Z}_2 \wr \mathbb{Z}^d$, where it takes n^{d+2} steps for the probability of being at the identity to become $\exp(-n^d)$, see [11].

The proofs of Theorem 1.1 differ in dimensions 1 and 2 from higher dimensions, partly because a random walk on \mathbb{Z}^d is transient in dimensions 3

and above. Many of the ideas for the proof of the higher dimensional case are actually much more general than what is necessary for the torus.

Theorem 1.2. *Let $\{G_n\}$ be a sequence of vertex transitive graphs and let G_n^\diamond denote $Z_2 \wr G_n$. As $|G_n|$ goes to infinity,*

$$(14) \quad \frac{1}{8 \log 2} \leq \frac{T_{\text{rel}}(G_n^\diamond)}{t^*(G_n)} \leq \frac{2}{\log 2} + o(1).$$

The lower bound in (14) is proved using the variational formula for relaxation time, and the upper bound uses a coupling argument that was introduced in [6] (see also [4]). The geometry of lamplighter graphs allows us to refine this coupling argument and restrict attention to pairs of states such that the position of the lamplighter is the same in both states.

Theorem 1.3. *Let $\{G_n\}$ be a sequence of vertex transitive graphs with $|G_n| \rightarrow \infty$, and \mathcal{C}_n denote the cover time for simple random walk on G_n . For any $\epsilon > 0$, there exist constants c_1, c_2 depending on ϵ such that the total variation mixing time satisfies*

$$(15) \quad [c_1 + o(1)]\mathbb{E}\mathcal{C}_n \leq \mathcal{T}_v(\epsilon, G_n^\diamond) \leq [c_2 + o(1)]\mathbb{E}\mathcal{C}_n.$$

Moreover, if the maximal hitting time satisfies $t^* = o(\mathbb{E}\mathcal{C}_n)$, then for all $\epsilon > 0$,

$$(16) \quad \left[\frac{1}{2} + o(1) \right] \mathbb{E}\mathcal{C}_n \leq \mathcal{T}_v(\epsilon, G_n^\diamond) \leq [1 + o(1)]\mathbb{E}\mathcal{C}_n.$$

Aldous [3] showed that the condition $t^* = o(\mathbb{E}\mathcal{C}_n)$ implies that the cover time has a sharp threshold, that is $\mathcal{C}_n/\mathbb{E}\mathcal{C}_n$ tends to 1 in probability. Theorem 1.3 thus says that in situations that give a sharp threshold for the cover time of G_n , there is also a threshold for the total variation mixing time on G_n^\diamond , although the factor of 2 difference between the bounds means that we have not proved a sharp threshold. When $\mathcal{T}_v(1/(2e), G_n) = o(|G_n|)$, (15) is Corollary 2.13 of [13].

The different upper and lower bounds in (16) cannot be improved without further hypotheses, as the limit exists and is equal to the lower bound when G_n is the complete graph K_n (see [9]) and the upper bound when $G_n = \mathbb{Z}_n^2$ (Theorem 1.1). The reason for this difference has to do with the geometry of the last points that are visited, which are uniformly distributed on K_n but very far from uniformly distributed on \mathbb{Z}_n^2 . This difference will be emphasized again later on. We will see later that the condition that the G_n are vertex transitive can be replaced by a condition on the cover times (Theorem 4.1) or on t^* (Theorem 4.2).

Theorem 1.4. *Let $\{G_n\}$ be a sequence of regular graphs for which $|G_n| \rightarrow \infty$ and the maximal hitting time satisfies $t^* \leq K|G_n|$ for some constant K . Then there are constants c_1, c_2 depending on ϵ and K such that*

$$(17) \quad c_1|G_n|(T_{\text{rel}}(G_n) + \log |G_n|) \leq \tau(\epsilon, G_n^\diamond) \leq c_2|G_n|(\mathcal{T}_v(G_n) + \log |G_n|).$$

Theorem 1.5. *Let $\{G_n\}$ be a sequence of vertex transitive graph such that $|G_n| \rightarrow \infty$. Then the uniform mixing time $\tau(\epsilon, G_n^\diamond)$ has a sharp threshold in the sense that for all $\epsilon > 0$,*

$$(18) \quad \lim_{n \rightarrow \infty} \frac{\tau(\epsilon, G_n^\diamond)}{\tau(1/2, G_n^\diamond)} = 1.$$

The intuition behind both Theorems 1.3 and 1.4 can be most easily explained by considering the case when G_n is the complete graph K_n , with a loop added at each vertex of K_n . The position $\pi(X_t)$ of the lamplighter then performs a simple random walk on K_n with holding probability $1/n$, and every lamp that is visited is randomized. Thus the (random) cover time for the walk on K_n , which is sharply concentrated near $n \log n$, is a strong uniform time (see [1]) for the configuration of the lamps. A strong uniform time (which bounds the mixing time in total variation) for the walk on the wreath product is obtained by adding one more step to randomize the location of the lamplighter.

A special property of the complete graph K_n is that the unvisited set is uniformly distributed. As shown in [9], what is actually needed for mixing in total variation is for the size of the uncovered set to be $O(\sqrt{n})$. At that time, the central limit theorem fluctuations in the number of lamps that are on are of at least the same order as the number of unvisited sites. The amount of time needed for this to happen is $(n \log n)/2$ steps, resulting in a sharp phase transition after $(n \log n)/2$ steps. The factor of 2 difference between the upper and lower bounds in (16) comes from the question of whether or not it suffices to cover all but the last \sqrt{n} sites of the graph. For many graphs, the amount of time to cover all but the last $\sqrt{|G_n|}$ sites is $\mathbb{E}C_n/2$, which will be the lower bound of (16). When the unvisited sites are clumped together instead of being uniformly distributed, it will turn out to be necessary to visit all the sites, and the upper bound of (16) will be sharp.

Convergence in the uniform metric depends upon the moment generating function of the cover time rather than the mean. Let R_t denote the set of lamps that the lamplighter has visited by time t , and $S_t = K_n \setminus R_t$ the set of unvisited lamps. Because any lamp that has been visited is randomized, the probability that the lamps are in any given configuration at time t is at most $\mathbb{E}2^{-|R_t|}$. To get convergence in the uniform metric, we need $\mathbb{E}2^{-|R_t|}$ to be within a factor of $(1 + \epsilon)$ of 2^{-n} , or, equivalently, to have $\mathbb{E}2^{|S_t|} \leq 1 + \epsilon$.

Considering the Markov process in continuous time, the visits to any given site in K_n are independent Poisson processes with rate n^{-1} , so the probability that a given site has not been visited by time t is $\exp(-t/n)$. Therefore $\mathbb{E}2^{|S_t|} = \mathbb{E} \prod 2^{I_i(t)}$, where $I_i(t)$ are indicator functions for the event that site i has not been visited. Since $\mathbb{E}2^{I_i(t)} = [1 + \exp(-t/n)]$, we have

$$\mathbb{E}2^{|S_t|} = [1 + \exp(-t/n)]^{n-1}.$$

For $t = n \log n + cn$, this gives

$$\mathbb{E}2^{|S_t|} = \exp[o(1) + e^{-c}],$$

so there is a sharp threshold at time $n \log n$.

Since the position of the lamplighter is randomized after each step, this indicates that $\tau(\epsilon, K_n^\diamond)$ has a sharp threshold at $n \log n$ as well. In this example, the upper and lower bounds of (17) are of the same order of magnitude and thus are accurate up to constants.

2. MORE EXAMPLES

Example (Hypercube) On the hypercube \mathbb{Z}_2^n , the maximal hitting time is on the order of 2^n . The cover time is on the order of $n2^n$, but $T_{\text{rel}}(\mathbb{Z}_2^n) = n$ and $\mathcal{T}_v(\epsilon, \mathbb{Z}_2^n) \sim (n \log n)/2$ ([2] Chapter 5 Example 15). By Theorem 1.2, $T_{\text{rel}}(\mathbb{Z}_2 \wr \mathbb{Z}_2^n)$ is on the order of 2^n , and Theorem 1.3 shows that the convergence time in total variation on $\mathbb{Z}_2 \wr \mathbb{Z}_2^n$ is on the order of $n2^n$. For uniform convergence, the upper and lower bounds of Theorem 1.4 differ, giving, up to constants, $n2^n$ and $(n \log n)2^n$. While Theorem 1.5 shows that there is a sharp threshold, the location of that threshold is unknown for this example. Thus, while the relaxation time is less than the mixing times, we do not know whether or not the total variation and uniform mixing times are comparable.

Example (Expander graphs) For a graph G_n , the conductance Φ_n is the minimum ratio of edge boundary to volume of all sets $S \subset G_n$ such that $|S| \leq |G_n|/2$. A sequence of graphs $\{G_n\}$ is called a sequence of *expander graphs* if there exists a $\delta > 0$ such that $\Phi_n \geq \delta$ for all n .

For expander graphs, the maximal hitting time on G_n is on the order of $|G_n|$, while the cover time of G_n is of order $|G_n| \log |G_n|$ because T_{rel} is bounded (see [5]). Theorem 4.2 shows that $\mathcal{T}_v(\epsilon, G_n^\diamond)$ is of order $|G_n| \log |G_n|$. The mixing time on such graphs is of order $\log |G_n|$, and so Theorem 1.4 shows that $\tau(\epsilon, G_n^\diamond) \leq C(\epsilon)|G_n| \log |G_n|$. Since $\tau(\epsilon, G) \geq \mathcal{T}_v(\epsilon, G)$ for any G , the matching lower bound holds as well. As a result, $\{G_n\}$ is a sequence of graphs of bounded degree such that the total variation and uniform mixing times for the lamplighter graphs are of the same order.

3. RELAXATION TIME BOUNDS

Proof of Theorem 1.2. For the lower bound, we will use the variational formula for the second eigenvalue. Let $|\lambda_2|$ denote the magnitude of the second largest eigenvalue in absolute value of a transition matrix p , and let π be the stationary distribution of p . The variational characterization for the eigenvalue says that

$$(19) \quad 1 - |\lambda_2| = \min_{\varphi: \text{Var} \varphi > 0} \frac{\mathcal{E}(\varphi \mid \varphi)}{\text{Var} \varphi},$$

where the Dirichlet form $\mathcal{E}(\varphi \mid \varphi)$ is given by

$$(20) \quad \mathcal{E}(\varphi \mid \varphi) = \frac{1}{2} \sum_{x,y} [\varphi(x) - \varphi(y)]^2 \pi(x) p(x, y)$$

$$(21) \quad = \frac{1}{2} \mathbb{E} [\varphi(\xi_1) - \varphi(\xi_0)]^2$$

and $\{\xi_t\}$ is a Markov chain started in the stationary distribution. For the lower bound of (14), we use (19) to show that the spectral gap for the transition kernel p^t is bounded away from 1 when $t = t^*/4$. Fix a vertex $o \in G$, and let $\varphi(f, x) = f(o)$. Then $\text{Var} \varphi = 1/4$ and by running for t steps,

$$\mathcal{E}_t(\varphi \mid \varphi) = \frac{1}{2} \mathbb{E} [\varphi(\xi_t) - \varphi(\xi_0)]^2 = \frac{1}{2} \sum_{x \in G} \nu(x) \frac{1}{2} \mathbb{P}_x[T_o < t],$$

where ν is the stationary measure on G , and $\{\xi_t\}$ is the stationary Markov chain on G^\diamond . For any t ,

$$\mathbb{E}_x T_o \leq t + t^*(1 - \mathbb{P}_x[T_o < t]).$$

For a vertex transitive graph, we have by Lemma 15 in Chapter 3 of [2], that

$$t^* \leq 2 \sum_{x \in G} \nu(x) \mathbb{E}_x T_o.$$

Let $\mathbb{E}_\nu = \sum_x \nu(x) \mathbb{E}_x$ and $\mathbb{P}_\nu = \sum_x \nu(x) \mathbb{P}(x)$. Then

$$t^* \leq 2 \mathbb{E}_\nu T_o \leq 2t + 2t^*[1 - \mathbb{P}_\nu(T_o < t)].$$

Substituting $t = t^*/4$ yields

$$\mathbb{P}_\nu[T_o < t^*/4] \leq \frac{3}{4}.$$

We thus have

$$1 - |\lambda_2|^{t^*/4} \leq \frac{3}{4},$$

and so

$$\log 4 \geq \frac{t^*}{4} (1 - |\lambda_2|),$$

which gives the claimed lower bound on $T_{\text{rel}}(G^\diamond)$.

For the upper bound, we use a coupling argument from [6]. Suppose that φ is an eigenfunction for p with eigenvalue λ_2 . To conclude that $T_{\text{rel}}(G^\diamond) \leq \frac{(2+o(1))t^*}{\log 2}$, it suffices to show that $\lambda_2^{2t^*} \leq 1/2$. For a configuration h on G , let $|h|$ denote the Hamming length of h . Let

$$M = \sup_{f,g,x} \frac{|\varphi(f, x) - \varphi(g, x)|}{|f - g|}$$

be the maximal amount that φ can vary over two elements with the same lamplighter position. If $M = 0$, then $\varphi(f, x)$ depends only on x , and so $\psi(x) = \varphi(f, x)$ is an eigenfunction for the transition operator on G . Since

$T_{\text{rel}}(G) \leq t^*$ (see [2], Chapter 4), this would imply that $|\lambda_2^{2t^*}| \leq e^{-4}$. We may thus assume that $M > 0$.

Consider two walks, one started at (f, x) and one at (g, x) . Couple the lamplighter component of each walk and adjust the configurations to agree at each site visited by the lamplighter. Let (f', x') and (g', x') denote the position of the coupled walks after $2t^*$ steps. Let K denote the transition operator of this coupling. Because φ is an eigenfunction,

$$\begin{aligned} \lambda_2^{2t^*} M &= \sup_{f,g,x} \frac{|p^{2t^*} \varphi(f, x) - p^{2t^*} \varphi(g, x)|}{|f - g|} \\ &\leq \sup_{f,g,x} \sum_{f',g',x'} K^{2t^*}[(f, g, x) \rightarrow (f', g', x')] \frac{|\varphi(f', x') - \varphi(g', x')|}{|f' - g'|} \frac{|f' - g'|}{|f - g|} \\ &\leq M \sup_{f,g,x} \frac{\mathbb{E}|f' - g'|}{|f - g|}. \end{aligned}$$

But at time $2t^*$, each lamp that contributes to $|f - g|$ has probability of at least $1/2$ of having been visited, and so $\mathbb{E}|f' - g'| \leq |f - g|/2$. Dividing by M gives the required bound of $\lambda_2^{2t^*} \leq 1/2$. \square

4. TOTAL VARIATION MIXING

In this section, we prove Theorem 1.3 and then some generalizations. The total variation claims of Theorem 1.1 except for the sharp constant in (7) are applications of Theorem 1.3, and we defer the computation of the sharp constant in two dimensions until later.

As was mentioned in the discussion of $\mathbb{Z}_2 \wr K_n$, the key to the lower bounds comes from running the random walk until the lamplighter visits all but $\sqrt{|G_n|}$ lamps, which takes time $\mathbb{E}\mathcal{C}_n/2$ when $t^* = o(\mathbb{E}\mathcal{C}_n)$.

Proof of Theorem 1.3. We will first prove the lower bound of (16) since the proof is simpler when $t^* = o(\mathbb{E}\mathcal{C}_n)$. Let $S \subset \mathbb{Z}_2 \wr G_n$ be the set of elements (f, x) such that the configuration f contains more than $|G_n|/2 + K|G_n|^{1/2}$ zeroes. Fix $\epsilon \in (0, 1)$. For sufficiently large K and n , we have $\mu(S) \leq (1 - \epsilon)/4$. Fix a basepoint $o \in G_n$, and let $id = (\mathbf{0}, o)$ denote the element of G^\diamond corresponding to all the lamps being off and the lamplighter being at o . The claim is that for large n , at time $t_n = \mathbb{E}\mathcal{C}_n/2$,

$$(22) \quad p^{t_n}(id, S) \geq \frac{1 + \epsilon}{4}$$

and thus $\mathcal{T}_v(\epsilon/2, G_n^\diamond) \geq \mathbb{E}\mathcal{C}_n/2$. To see this, let $\mathcal{C}_n(k)$ be the first time that the lamplighter has visited all but k lamps, and let m_n be given by

$$m_n = \min\{t : \mathbb{P}(\mathcal{C}_n(K|G_n|^{1/2}) > t) \leq (1 + \epsilon)/2\}.$$

Let $0 < s_1 < s_2 < s_3 < s_4$ be stopping times such that s_1 and $s_3 - s_2$ are minimal strong uniform times, and $s_2 - s_1$ and $s_4 - s_3$ are times to hit all but $K|G_n|^{1/2}$ sites. As a result, at time s_1 the walk X_n is uniformly distributed,

the walk covers all but $K|G_n|^{1/2}$ points of G_n on both intervals $[s_1, s_2)$ and $[s_3, s_4)$, and conditioned on X_{s_2} , the walk is uniformly distributed at time s_3 .

For any finite graph G , denote $\mathcal{T}_v = \mathcal{T}_v(1/(2e), G)$. Let $q(\cdot, \cdot)$ be the transition kernel for a random walk on G with stationary distribution ν . Then the separation distance

$$(23) \quad 1 - \min_y \{q^{4\mathcal{T}_v}(x, y)/\nu(y)\}$$

at time $4\mathcal{T}_v$, is at most $1/e$ (see [1] or [2], Chapter 4, Lemma 7).

Returning to G_n , Aldous and Diaconis [1] showed that $\mathbb{P}[s_1 > t]$ is the separation distance at time t . We thus have

$$\mathbb{P}[s_1 > 4\mathcal{T}_v(1/(2e), G_n)] \leq 1/e.$$

Therefore, with probability bounded away from 0,

$$s_4 \leq 2(m_n + 4\mathcal{T}_v(1/(2e), G_n)).$$

But at time s_4 , the probability of a given point being uncovered is at most

$$(K|G_n|^{-1/2})^2,$$

and so the expected number of uncovered points at time s_4 is at most K^2 . In particular, with probability at least $1/2$, there are fewer than $2K^2$ uncovered sites. Each additional run of $2t^*$ steps has probability at least $1/2$ of hitting one of these sites, so with probability $1/4$, all sites are covered after $8K^2t^*$ more steps. Therefore the probability of having covered the entire space by time $2(m_n + \mathcal{T}_v(1/(2e), G_n)) + 8K^2t^*$ is bounded away from 0. When $t^* = o(\mathbb{E}\mathcal{C}_n)$, Aldous [3] showed that $\mathcal{C}_n/\mathbb{E}\mathcal{C}_n \rightarrow 1$ in probability, so

$$\mathbb{E}\mathcal{C}_n \leq (1 + o(1)) [2\mathcal{T}_v(1/(2e), G_n) + 2m_n + 8K^2t^*].$$

Since t^* and thus $\mathcal{T}_v(1/(2e), G_n)$ are much smaller than $\mathbb{E}\mathcal{C}_n$, this means that $m_n \geq \mathbb{E}\mathcal{C}_n(1 + o(1))/2$. But at time $m_n - 1$, the probability of having $K|G_n|^{1/2}$ uncovered points is greater than $(1 + \epsilon)/2$, and so the probability of having at least $(|G_n| + K|G_n|^{1/2})/2$ zeroes in the configuration at time m_n tends to $(1 + \epsilon)/4$. This proves the lower bound of (16).

For the lower bound of (15), we iterate the above process one more time to get better control of the probabilities. More precisely, let

$$r_n = \min \left\{ t : \mathbb{P} \left(\mathcal{C}_n(|G_n|^{7/12}) > t \right) \leq (1 + \epsilon)/2 \right\}.$$

Let \mathbb{P}_w denote the probability measure for the random walk starting at w . By running for $4\mathcal{T}_v(1/(2e), G_n) + r_n$ steps 3 times, a similar argument as before shows that the expected number of unvisited sites is small, meaning that

$$\mathbb{P}_w [\mathcal{C}_n > 3(4\mathcal{T}_v(1/(2e), G_n) + r_n)] \leq c \frac{(1 + \epsilon)^3}{8} \left(1 - n^{-5/2} \right)^3.$$

But if $\mathbb{P}_w(C_n > t) \leq x$ for all w then $\mathbb{E}C_n \leq t(1-x)^{-1}$ simply by running the chain for intervals of length t until all of G_n is covered in one interval. We therefore have

$$\inf \left\{ t : \mathbb{P}(C_n > t) \leq c \frac{(1+\epsilon)^3}{8} \right\} \geq c_1(\epsilon) \mathbb{E}C_n$$

for some constant $c_1(\epsilon)$. Since t^* (and thus $\mathcal{T}_v(1/(2e), G_n)$ by [2], Chapter 4 Theorem 6) is $o(\mathbb{E}C_n)$, this shows that $r_n \geq c_1(\epsilon) \mathbb{E}C_n(1+o(1))$. For large n , however, the total variation distance on G_n^\diamond is still at least ϵ at time r_n , which proves the lower bound of (15).

For the upper bound, let π denote the projection from G^\diamond to G given by the position of the lamplighter. Let μ denote the stationary measure on G^\diamond and ν the stationary measure on G . Note that $\mu(f, x) = 2^{-|G|} \nu(x)$. By the strong Markov property, at time $t = C + 1 + k$ we have $p^t[id, (f, x)] = 2^{-|G|} \mathbb{P}_{X_C}[\pi(X_{k+1}) = x]$. For any $k > 0$,

$$(24) \quad \sum_{(f,x) \in \mathbb{Z}_2 \wr G} \left| p^t[id, (f, x)] - \mu(f, x) \right| \leq \mathbb{P}(t \leq C + k) + \sum_{x \in G} \left| \mathbb{P}_{X_C}[\pi(X_{k+1}) = x] - \nu(x) \right|.$$

Letting $t_1(\epsilon) = \min\{t : \mathbb{P}(C > t) \leq \epsilon/2\}$, we see that $\mathcal{T}_v(\epsilon, G_n^\diamond)$ is at most $t_1(\epsilon) + \mathcal{T}_v(\epsilon/2, G_n)$. The desired upper bound comes from the fact that $t_1 \leq [c_2(\epsilon) + o(1)] \mathbb{E}C_n$ and $\mathcal{T}_v(\epsilon/2, G_n)$ is of a lower order than $\mathbb{E}C_n$. When $t^* = o(\mathbb{E}C_n)$, the fact that $C_n/\mathbb{E}C_n$ tends to 1 in probability implies that $c_2(\epsilon) = 1$ works for all ϵ . \square

In many situations, the assumption of vertex transitivity in Theorem 1.3 can be relaxed.

Theorem 4.1. *Suppose that G_n is a sequence of graphs for which simple random walk satisfies*

$$\lim_{n \rightarrow \infty} \frac{t^* \log |G_n|}{\mathbb{E}C_n} = 1,$$

and $|G_n| \rightarrow \infty$. Then for any $\epsilon > 0$,

$$(25) \quad \left[\frac{1}{2} + o(1) \right] \mathbb{E}C_n \leq \mathcal{T}_v(\epsilon, G_n^\diamond) \leq [1 + o(1)] \mathbb{E}C_n.$$

Proof. Since the proof of the upper bound for Theorem 1.3 did not rely on vertex transitivity, we only need to prove the lower bound. To do so, it suffices to show that for any $\delta > 0$,

$$(26) \quad \lim_{n \rightarrow \infty} \mathbb{P}[C_n(|G_n|^\alpha) < (1 - \alpha - \delta) \mathbb{E}C_n] = 0,$$

and then follow the proof of Theorem 1.3.

The way that we will prove this is by contradiction. If

$$\limsup_{n \rightarrow \infty} \mathbb{P}[\mathcal{C}_n(|G_n|^\alpha) < (1 - \alpha - \delta)\mathbb{E}\mathcal{C}_n] > 0,$$

we will show that we can find an $\epsilon > 0$ such that

$$\limsup_{n \rightarrow \infty} \mathbb{P}[\mathcal{C}_n < (1 - \epsilon)\mathbb{E}\mathcal{C}_n] > 0.$$

But $\mathcal{C}_n/\mathbb{E}\mathcal{C}_n \rightarrow 1$ in probability, so this will be a contradiction. To do this, let S_n denote the final $|G_n|^\alpha$ points in G_n that are hit by $\{X_t\}$. We will first follow Matthews' proof [10] to show that

$$\mathbb{E}[\mathcal{C}_n - \mathcal{C}_n(|G_n|^\alpha) \mid S_n] \leq \alpha\mathbb{E}\mathcal{C}_n.$$

Label the points of S_n from 1 to $|G_n|^\alpha$ with a labelling chosen uniformly from the set of all such labellings, and let $\tilde{\mathcal{C}}_n(k)$ denote the first time after $\mathcal{C}_n(|G_n|^\alpha)$ that the first k elements of S_n (according to our random labelling) are covered by $\{X_t\}$. Note that $\mathbb{P}[\tilde{\mathcal{C}}_n(k+1) > \tilde{\mathcal{C}}_n(k)] = 1/(k+1)$ by symmetry, and thus $\mathbb{E}[\tilde{\mathcal{C}}_n(k+1) - \tilde{\mathcal{C}}_n(k)] \leq t^*/(k+1)$. Summing this telescoping series, we see that $\mathbb{E}\mathcal{C}_n - \mathbb{E}\mathcal{C}_n(|G_n|^\alpha) \leq t^*[\log(|G|^\alpha) + 1] \sim \alpha\mathbb{E}\mathcal{C}_n$.

To complete the proof, note that

$$\begin{aligned} \mathbb{P}[\mathcal{C}_n < (1 - \epsilon)\mathbb{E}\mathcal{C}_n] &\geq \mathbb{P}[\mathcal{C}_n(|G_n|^\alpha) < (1 - \alpha - \delta)\mathbb{E}\mathcal{C}_n] \\ &\quad \times \mathbb{P}[\mathcal{C}_n - \mathcal{C}_n(|G_n|^\alpha) < (\alpha + \delta - \epsilon)\mathbb{E}\mathcal{C}_n \mid S_n]. \end{aligned}$$

For ϵ such that $\alpha(1+\epsilon)^2 < \alpha + \delta - \epsilon$, Markov's inequality shows that the final probability is at least $\epsilon/(1+\epsilon)$ for large n , which completes the proof. \square

Theorem 4.2. *Let G_n be a sequence of regular graphs such that $T_{\text{rel}}(G_n) \leq |G_n|^{1-\delta}$ for some $\delta > 0$. Then there are constants $c_1(\epsilon)$ and $c_2(\epsilon)$ such that*

$$(27) \quad [c_1 + o(1)] |G_n| \log |G_n| \leq \mathcal{T}_v(\epsilon, G_n^\diamond) \leq [c_2 + o(1)] |G_n| \log |G_n|.$$

Proof. This is actually a corollary of a result of Broder and Karlin [5]. They showed that $\mathbb{E}\mathcal{C}_n$ is of the order of $|G_n| \log |G_n|$. Examining the proof of their lower bound (Theorem 13) shows that they actually proved the stronger statement that the time to cover all but $|G_n|^{\alpha+1/2}$ sites is of order $|G_n| \log |G_n|$ for small enough α , from which our result follows as before. \square

5. UNIFORM MIXING

As discussed in the example of a lamplighter graph on the complete graph K_n , the key to the mixing time on the wreath product is the number of sites on the underlying graph that are left uncovered by the projection of the lamplighter's position. We thus first develop some necessary facts about the distribution of the number of uncovered sites, and then prove our main results.

Recall from (23) that the transition kernel q on G and the stationary measure ν satisfy $q^{4\mathcal{T}_v}(x, y) \geq (1 - e^{-1})\nu(y)$, where $\mathcal{T}_v = \mathcal{T}_v(1/(2e), G)$.

Lemma 5.1. *Let $\{X_t\}$ be an irreducible Markov chain on a state space G and let T_x^+ denote the first return time to x . Suppose that there are $\epsilon, \delta > 0$ such that*

$$\mathbb{P}_x(T_x^+ > \epsilon|G|) \geq \delta > 0$$

for all $x \in G$. Let $S \subset G$ be a set of k elements. Then for $k \geq \mathcal{T}_v$, the probability of hitting at least $\delta\epsilon\mathcal{T}_v/4$ elements of S by time

$$4\mathcal{T}_v + \epsilon|G|\mathcal{T}_v k^{-1}$$

is bounded away from 0.

Proof. Let $r = \epsilon|G|\mathcal{T}_v k^{-1}$. For $1 \leq i \leq r$, let I_i be an indicator random variable for the event $\{X_{4\mathcal{T}_v+i} \in S\}$, and let J_i be an indicator for the event that $X_{4\mathcal{T}_v+i}$ is in S but that the walk does not return to $X_{4\mathcal{T}_v+i}$ by time r . Note that $I_i \geq J_i$, and $\sum_1^r J_i$ is the number of distinct elements of S visited between time $4\mathcal{T}_v + 1$ and time $4\mathcal{T}_v + r$, inclusive. By running for an initial amount of time $4\mathcal{T}_v$, the probability of the lamplighter being at any given lamp is at least $(1 - e^{-1})|G|^{-1} \geq (2|G|)^{-1}$. As a result, $\mathbb{E}J_i \geq \delta k(2|G|)^{-1}$ and thus $\mathbb{E}\sum_1^r J_i \geq \epsilon\delta\mathcal{T}_v/2$. To conclude that $\mathbb{P}(\sum J_i > \epsilon\delta\mathcal{T}_v/4)$ is bounded away from 0, it suffices to show that there is a constant C such that $\mathbb{E}(\sum J_i)^2 \leq C\mathcal{T}_v^2$. But

$$\text{Cov}(I_i, I_j) \leq \|I_i\|_2 \|I_j\|_2 \exp[-T_{\text{rel}}^{-1}|i - j|],$$

and summing the geometric series,

$$\sum_{j=1}^r \text{Cov}(I_i, I_j) \leq \frac{kT_{\text{rel}}}{2|G|(1 - \exp[-T_{\text{rel}}^{-1}])}.$$

Therefore,

$$\sum_{i=1}^r \sum_{j=1}^r \text{Cov}(I_i, I_j) \leq \frac{\delta\epsilon\mathcal{T}_v}{2(1 - \exp[-T_{\text{rel}}^{-1}])}.$$

Because $1/2 \leq T_{\text{rel}} \leq \mathcal{T}_v$, there is a constant C such that

$$\mathbb{E}\left(\sum_{i=1}^r I_i\right)^2 \leq C\mathcal{T}_v(G)^2.$$

Since $\mathbb{E}(\sum_{i=1}^r J_i)^2 \leq \mathbb{E}(\sum_{i=1}^r I_i)^2$, this completes the proof. \square

Lemma 5.1 does not apply once we have reduced the number of unvisited sites to less than \mathcal{T}_v , but that situation is easier to control.

Lemma 5.2. *Let $\{X_t\}$ be a Markov chain on G such that $t^* \leq K|G|$, and let $S \subset G$ be a subset of s elements. Then the probability of hitting at least $s/2$ elements of S by time $2K|G|$ is at least $1/2$.*

Proof. Suppose $x \in S$, and let T_x be the first hitting time of x . Since $\mathbb{E}T_x \leq t^* \leq K|G|$, we have $\mathbb{P}[T_x \leq 2K|G|] \geq 1/2$. The expected number of elements of S that are hit by time $2K|G|$ is thus at least $s/2$, and the

number that are hit is at most s , so the probability of hitting at least $s/2$ is bounded below by $1/2$ as claimed. \square

As the next lemma shows, the hypotheses of Lemma 5.1 are easily fulfilled.

Lemma 5.3. *Suppose that G is a regular graph with maximal hitting time t^* . The distribution of T_x^+ , the first return time to x , satisfies*

$$\min_{x \in G} \mathbb{P}_x \left(T_x^+ > \frac{|G|}{2} \right) \geq \frac{|G|}{2t^*} > 0.$$

Proof. For random walk on a regular graph, $\mathbb{E}_x T_x^+ = |G|$. But $\mathbb{E}_x T_x^+ \leq |G|/2 + \mathbb{P}_x(T_x^+ > |G|/2)t^*$, and rearranging terms proves the claim. \square

Lemma 5.4. *Let $\{X_t\}$ be a Markov chain on G such that $t^* \leq K|G|$. Let S_t denote the set of points not covered by time t . Then there are constants c_1, c_2 , and $c_3 > 0$ such that if $t = (1 + a)c_1|G|(\mathcal{T}_v + \log |G|)$, then $\mathbb{E}2^{|S_t|} \leq 1 + c_2 \exp[-ac_3(\mathcal{T}_v + \log |G|)]$.*

Proof. Let $r = \lfloor |G|\mathcal{T}_v^{-1} \rfloor$. For $0 \leq i \leq r - 1$, let $k_i = |G| - i\mathcal{T}_v$, and for $i = r$ let $k_i = 0$. Let $t_i = \min\{t : |S_t| \leq k_i\}$. For $i < r$, Lemmas 5.1 and 5.3 imply that there is a constant $c = c(K)$ such that $t_i - t_{i-1}$ is stochastically dominated by $x_i = (c|G|\mathcal{T}_v k_i^{-1})Z_i$, where Z_i is a geometrically distributed random variable with mean 2. We will begin by bounding the probability that t_i is much larger than $|G|(\mathcal{T}_v + \log |G|)$ when $i < r$. By expressing t_i as a telescoping series and using Markov's inequality, we see that for any $\theta \geq 0$,

$$\begin{aligned} \mathbb{P}[t_i > t] &\leq \mathbb{P} \left[\sum_{j=1}^i x_j > t \right] \\ &= \mathbb{P} \left[\sum_{j=1}^i \frac{c|G|\mathcal{T}_v}{k_j} Z_j > t \right] \\ &\leq \exp[-t\theta] \mathbb{E} \exp \left[\sum_{j=1}^i \frac{c|G|\mathcal{T}_v\theta}{k_j} Z_j \right] \\ &= \exp[-t\theta] \prod_{j=1}^i \mathbb{E} \exp \left[\frac{c|G|\mathcal{T}_v\theta}{k_j} Z_j \right]. \end{aligned}$$

But for $\alpha \leq 1/3 \leq \log(3/2)$,

$$\mathbb{E} \exp[\alpha Z_j] = \frac{\exp(\alpha)}{2 - \exp(\alpha)} \leq \exp(3\alpha).$$

Let $\theta = k_i(3c|G|\mathcal{T}_v)^{-1}$. For $j \leq i$ we have $k_j < k_i$, whence

$$\frac{c|G|\mathcal{T}_v\theta}{k_j} \leq \frac{1}{3}.$$

Consequently,

$$\mathbb{E} \exp \left[\frac{c|G|\mathcal{T}_v\theta}{k_j} Z_j \right] \leq \exp \left[\frac{k_i}{k_j} \right].$$

Since $\mathbb{P}[|S_t| > k_i] = P[t_i > t]$, and $i \leq |G|T_v^{-1} + 1$, we have

$$(28) \quad \mathbb{P}[|S_t| > k_i] \leq \exp \left[-t \left(\frac{k_i}{3c\mathcal{T}_v|G|} \right) + \frac{k_i}{\mathcal{T}_v} \log |G| \right].$$

For $i = r$, Lemma 5.2 shows that $t_r - t_{r-1}$ is stochastically dominated by a sum of at most $\log_2(2\mathcal{T}_v)$ geometric random variables with mean $2K|G|$. As a result,

$$(29) \quad \mathbb{P}[t_r - t_{r-1} > t] \leq \log_2(2\mathcal{T}_v) \exp \left[-\frac{t}{2K|G|} \right].$$

Breaking the possible values of $|S_t|$ into intervals of length \mathcal{T}_v , we get

$$\mathbb{E} 2^{|S_t|} \leq 1 + \sum_{i=0}^r 2^{k_i + \mathcal{T}_v} \mathbb{P}(|S_t| > k_i).$$

Let $c_1 = \max\{7c \log 2, 5K \log 2\}$. For $i < r$, $k_i \geq \mathcal{T}_v$ and

$$2^{k_i + \mathcal{T}_v} \mathbb{P}[|S_t| > k_i] \leq \exp \left[-\frac{7}{3}(\log_2)k_i a - \frac{1}{3}(\log 2)k_i \right]$$

where the second inequality used the fact that $k_i \geq \mathcal{T}_v$ for $i < r$. When $i = r$,

$$\begin{aligned} 2^{k_r + \mathcal{T}_v} \mathbb{P}[|S_t| > k_r] &\leq 2^{\mathcal{T}_v} [\mathbb{P}(t_{r-1} > t/2) + \mathbb{P}(t_r - t_{r-1} > t/2)] \\ &\leq 2^{\mathcal{T}_v} \exp \left[-\frac{7}{3}(1+a)k_{r-1} \log 2 \right] \\ &\quad + 2^{\mathcal{T}_v} \log_2(2\mathcal{T}_v) \exp \left[-\frac{5 \log 2(1+a)(\mathcal{T}_v + \log |G|)}{4} \right] \end{aligned}$$

Summing these terms yields the desired bound. \square

We will defer the proof of Theorem 1.1 to Section 6, but will prove here the somewhat more general Theorem 1.4.

Proof of Theorem 1.4. First, we will prove the lower bound on $\tau(\epsilon, G_n^\diamond)$. The bound $\tau(\epsilon, G_n^\diamond) \geq c|G_n| \log |G_n|$ is only relevant when $T_{\text{rel}}(G_n) \leq |G_n|^{1/2}$, in which case it is a consequence of Theorem 4.2 and the general inequality $\tau(\epsilon, G) \geq \mathcal{T}_v(\epsilon, G)$, applied to $G = G_n^\diamond$. To establish the lower bound, it remains to prove that $\tau(\epsilon, G_n^\diamond) \geq c|G_n| \cdot T_{\text{rel}}(G_n)$. Note that for any $\epsilon > 0$, we have $\tau(\epsilon, G_n^\diamond) \geq |G_n|/2$ when $|G_n|$ is large enough. We may thus assume that $T_{\text{rel}}(G_n)$ is bounded away from 1, and for concreteness will assume $T_{\text{rel}}(G_n) \geq 2$. Let η be a right eigenfunction of the walk $\{Y_t\}$ on G_n with eigenvalue $\lambda = 1 - T_{\text{rel}}^{-1}(G_n)$ and $\sum_g \eta(g) = 0$. Then $\lambda^{-t} \eta(Y_t)$ is a martingale. Let S be the smallest subset consisting of at least half of the vertices of G_n of the form $S = \{x \in G_n : \eta(x) \leq a\}$. Possibly replacing η by $-\eta$, we may

assume that $a \leq 0$. Take M to be $M = \max \eta(x)$, and let x_0 be a vertex of G such that $\eta(x_0) = M$.

Let T_S denote the hitting time of S . By the optional stopping theorem,

$$(30) \quad \mathbb{E}_{x_0} [\lambda^{-T_S \wedge t} \eta(Y_{T_S \wedge t})] = \mathbb{E}_{x_0} \eta(Y_0).$$

Since $a \leq 0$, this yields

$$(31) \quad M \mathbb{P}(T_S > t) \lambda^{-t} \geq M.$$

Since $T_{\text{rel}}(G_n) \geq 2$, we have $\log \lambda \geq -(2 \log 2) T_{\text{rel}}^{-1}$, whence

$$P(T_S > t) \geq \exp[-(2 \log 2) t T_{\text{rel}}^{-1}].$$

At time t , the probability that all lamps are off is at least $P(T_S > t) 2^{-|G_n|/2}$, and for

$$t < T_{\text{rel}}(G_n) \cdot \left(\frac{|G_n|}{4} - \frac{\log(1 + \epsilon)}{2 \log 2} \right)$$

the probability that all lamps are off is at least $(1 + \epsilon) 2^{-|G_n|}$, so $\tau(\epsilon, G^\diamond) \geq |G_n| T_{\text{rel}}(G_n) \cdot (1/4 + o(1))$.

For the upper bound $\tau(\epsilon, G_n^\diamond) \leq t = C(\epsilon) |G_n| (\log |G_n| + \mathcal{T}_v(\epsilon, G_n))$, we need to prove two things for any $(f, x) \in G_n^\diamond$:

First, that $\mathbb{P}(X_t = (f, x)) \geq 2^{-|G_n|} |G_n|^{-1} (1 - \epsilon)$, and second, that $\mathbb{P}(X_t = (f, x)) \leq 2^{-|G_n|} |G_n|^{-1} (1 + \epsilon)$.

For the first bound,

$$\mathbb{P}(X_t = (f, x)) \geq |G_n|^{-1} 2^{-|G_n|} \mathbb{P}(\mathcal{C}_n < t - \tau(\epsilon/2, G_n)) (1 - \epsilon/2),$$

and the hypotheses of the theorem imply that $\mathbb{P}(\mathcal{C}_n < t - \tau(\epsilon/2, G_n))$ tends to 1 (see [3]).

For the second bound, we first run the walk for $r = t - \tau(\epsilon/3, G_n)$ steps. Let π denote the projection from G_n^\diamond to G_n given by the position of the lamplighter, and write $X_t = (f_t, \pi(X_t))$. Since the lamps visited by time r are all randomized,

$$\mathbb{P}(f_t = f) \leq \mathbb{E} 2^{|S_r| - |G_n|},$$

where S_r is the uncovered set of G_n at time r . Considering the next $t - r$ steps, we get

$$\mathbb{P}(\pi(X_t) = x \mid X_j, j \leq r) \leq \max_y \mathbb{P}(\pi(X_{t-r}) = y),$$

which is at most $(1 + \epsilon/3) |G_n|^{-1}$.

By Lemma 5.4,

$$\mathbb{E} 2^{|S_r| - |G_n|} \leq (1 + \epsilon/2) 2^{-|G_n|}.$$

Combining these, we have

$$\mathbb{P}(X_t = (f, x)) \leq (1 + \epsilon) |G_n|^{-1} 2^{-|G_n|},$$

as required. \square

Proof of Theorem 1.5. What we will actually prove is that

$$(32) \quad \tau(7\epsilon/8, G_n^\diamond) \leq \tau(\epsilon, G_n^\diamond) + 2t^* + \tau(\delta, G_n).$$

for a suitable $\delta = \delta(\epsilon)$. Write $T_{\text{rel}} = T_{\text{rel}}(G_n)$. Given (32), proving the theorem only requires that $t^* + \tau(\epsilon/2, G_n) = o(\tau(\epsilon, G_n^\diamond))$. But

$$\tau(\epsilon/2, G_n) \leq \log |G_n| \cdot T_{\text{rel}} = o(\tau(\epsilon, G_n^\diamond)),$$

so we only need the condition that $t^* = o(\tau(\epsilon, G_n^\diamond))$.

For graphs G_n such that $t^*/T_{\text{rel}} \leq |G_n|^{1/2}$, the lower bound of (17) proves the claim. When $t^*/T_{\text{rel}} \geq |G_n|^{1/2}$, Theorem 4 and Corollary 5 in Chapter 7 of [2] show that

$$\mathbb{E}\mathcal{C}_n \geq (1/2 + o(1))t^* \log |G_n|.$$

We are thus in a case covered by Theorem 1.3, and so

$$\tau(\epsilon, G_n^\diamond) \geq \mathcal{T}_v(\epsilon, G_n^\diamond) \geq (1/4 + o(1))t^* \log |G_n|,$$

from which the claim follows.

To show that (32) holds, consider the uncovered set at the times $r = \tau(\epsilon, G_n^\diamond)$ and $t = r + 2t^*$. Conditioned on the $|S_r| > 0$, at time t at least one more point has been covered with probability at least $1/2$. This means that

$$\mathbb{E}2^{|S_t|} - 1 \leq \frac{3}{4} \left(\mathbb{E}2^{|S_r|} - 1 \right).$$

Because $\mathbb{E}2^{|S_r|} \leq 1 + \epsilon$, this implies that $\mathbb{E}2^{|S_t|} \leq 1 + 3\epsilon/4$. Let $u = t + \tau(\delta, G_n)$. Then for any $(f, x) \in G_n^\diamond$,

$$\begin{aligned} \mathbb{P}[(f_u, \pi(X_u)) = (f, x)] &\leq \mathbb{E} \left[2^{|S_t| - |G_n|} \right] \max_{y \in G_n} \mathbb{P}_y[\pi(X_{u-t}) = x] \\ &\leq \left(1 + \frac{3\epsilon}{4} \right) (1 + \delta) |G_n|^{-1} 2^{-|G_n|}. \end{aligned}$$

Taking δ small enough that $(1 + 3\epsilon/4)(1 + \delta) < (1 + 7\epsilon/8)$ proves that (32) holds. \square

6. PROOF OF THEOREM 1.1

We turn now to completing the proof of Theorem 1.1. The statements about relaxation time follow from Theorem 1.2 and standard facts about t^* (see e.g. [2] Chapter 5). The mixing results in dimension $d \geq 3$ follow from Theorems 1.3 and 1.4, so only the dimension $d = 2$ results remain to be shown.

Proof of Theorem 1.1. As usual, we let S_t denote the set of lamps that are unvisited at time t .

For completeness, we will also show that for $d = 1$,

$$\tau(\epsilon, G_n^\diamond) \sim \frac{64 \log 2}{27\pi^2} n^3.$$

For simple random walk with holding probability $1/2$ in one dimension, at time $t = \alpha n^3$,

$$\mathbb{P}(|S_t| = \lambda n) = \exp \left[- \left(\frac{\pi^2}{4} + o(1) \right) \frac{\alpha n}{(1 - \lambda)^2} \right]$$

(see [15], Section 21), and so

$$\mathbb{E}2^{|S_t|} - 1 = n \int_0^1 \exp \left[- \left(\frac{\pi^2}{4} + o(1) \right) \frac{\alpha n}{(1 - \lambda)^2} \right] d\lambda.$$

Taking a Taylor expansion about $\rho = 1 - (\pi^2 \alpha / 2 \log 2)^{1/3}$ to estimate the integral, we see that there is a sharp threshold for the time at which $\mathbb{E}2^{|S_t|} = 1 + \epsilon$ that occurs when

$$t = \frac{32 \log 2}{27 \pi^2}.$$

The argument in the proof of Theorem 1.4 shows that this is also a sharp threshold for $\tau(\epsilon, G_n^\diamond)$.

For $d = 2$, we will give a proof that at first seems wasteful but is surprisingly accurate for the extreme events that are the dominant contribution to $\mathbb{E}2^{|S_t|}$. After n^2 steps, there exists an $\alpha > 0$ such that a simple random walk on \mathbb{Z}_n^2 will cover some ball of radius $n^\alpha \sqrt{2}$ with probability at least $\delta > 0$ [12]. Divide the torus \mathbb{Z}_n^2 into $n^{2-2\alpha}$ boxes of side length n^α . On time intervals of the form $[2kn^2, (2k+1)n^2]$, there is thus a probability of at least δ of the random walk covering at least one of the $n^{2-2\alpha}$ boxes. Moreover, because the mixing time on the torus is of the order n^2 , we see that even conditioned on which boxes are currently covered, if there are i uncovered boxes then the walk will cover one of the uncovered boxes with probability at least $\epsilon i n^{2\alpha-2}$ for some fixed $\epsilon > 0$. Since the probability of covering a new box in any given interval of length $2k^2$ is uniformly bounded away from 0, the number of such intervals needed to cover a new box is stochastically dominated by a geometric random variable. In particular, rescaling implies that if $t_i = \min\{t : \text{there are } i \text{ covered boxes}\}$, then for large enough c

$$\frac{n^{2-2\alpha} - j}{n^{4-2\alpha} c} (t_j - t_{j-1}),$$

is stochastically dominated by a geometric random variable with mean 2. We will use this fact to crudely bound the exponential moment of the number of uncovered sites on the torus.

As before,

$$(33) \quad \mathbb{E}[2^{|S_t|}] \leq 1 + \sum_{i=1}^{n^{2-2\alpha}} \mathbb{P}(t_i > t) 2^{n^2 - i n^{2\alpha}}.$$

Arguing as in the proof of Lemma 5.4,

$$(34) \quad \mathbb{P}(t_i > t) \leq e^{-t\theta} \mathbb{E} \exp \left[\sum_{j=1}^i \frac{\theta n^{4-2\alpha} c}{n^{2-2\alpha} - j} Z_j \right]$$

and so for $\theta = (n^{2-2\alpha} - i)(3n^{4-2\alpha}c)^{-1}$,

$$(35) \quad \mathbb{P}(t_i > t) \leq \exp \left[-t \frac{n^{2-2\alpha} - i}{3n^{4-2\alpha}c} + (n^{2-2\alpha} - i) \log n \right].$$

For $\delta > 0$ and $t = (1 + \delta)n^4 4c \log 2$, substituting into (33) shows that the order of n^4 steps suffice to reduce $\mathbb{E}2^{|S_t|}$ to $(1 + \epsilon)$ and are thus an upper bound for $\tau(\epsilon, G_n^\diamond)$.

For a lower bound on $\tau(\epsilon, G_n^\diamond)$, let S be the set of elements in G_n^\diamond for which the absolute value of the horizontal component of the lamplighter's position is at most $n/4$. The size of S is $n^2/2$, and the probability of remaining in S for the first cn^4 steps decays like $\exp[-cn^2]$, and so for small enough c , we have $\mathbb{E}2^{|S_t|} \geq 2^{n^2/2} \exp[-cn^2] > 1 + \epsilon$. The sharp threshold in (9) is due to Theorem 1.5.

To show that the sharp threshold claimed in (7) actually exists, we need to prove that a lower bound for total variation convergence is $\mathbb{E}\mathcal{C}_n$. Let $\alpha \in (0, 1)$. By [7], Sec. 4, there exists $\beta = \beta(\alpha) > 0$ such that at time $(1 - \alpha)\mathbb{E}\mathcal{C}_n$, the uncovered set contains a ball of radius n^β with high probability. The configuration f at time $(1 - \alpha)\mathbb{E}\mathcal{C}_n$ is thus, with high probability, identically zero on a ball of radius $(1/3)n^{\alpha/2}$. But in the stationary measure μ on G_n^\diamond , the expected number of balls of radius r for which the configuration is identically 0 on the ball is less than $n^2 2^{-r^2}$. Therefore the probability of having a ball of 0's with radius of a greater order than $(\log n)^{1/2}$ tends to 0 as n grows. \square

7. COMMENTS AND QUESTIONS

Random walks on another type of wreath product were analyzed by Schoolfield [14] and by Fill and Schoolfield [8]. When G is the symmetric group S_n , instead of letting S_n act on itself by left-multiplication, it is perhaps more natural to let S_n act on \mathbb{Z}_n by permutation. As such, $G \wr S_n$ is often used to describe $S_n \ltimes \sum_{\mathbb{Z}_n} \mathbb{Z}_2$, while our description yields the Cayley graph of $S_n \ltimes \sum_{S_n} \mathbb{Z}_2$. Mixing times in both total variation and L^2 norm are carefully analyzed for this alternative description of $G \wr S_n$ in [14] and [8].

Although the results of this paper give a good description of the mixing on finite lamplighter groups in both total variation norm and the uniform norm, the discrepancies in the upper and lower bounds raise a number of natural questions.

- (1) For the torus in dimension $d \geq 3$, is $\mathcal{T}_v(\epsilon, G_n^\diamond)$ asymptotic to $\mathbb{E}\mathcal{C}_n$ or $\mathbb{E}\mathcal{C}_n/2$? This question can be stated without reference to the lamplighter: let R_t be the set of points that simple random walk on the torus G_n visits by time t . Suppose that μ is uniformly distributed on $\{0, 1\}^{G_n}$ and that ν_t is uniformly distributed on $\{0, 1\}^{R_t} \times \{0\}^{G_n \setminus R_t}$.

For $1/2 < \alpha < 1$, what is

$$\lim_{n \rightarrow \infty} \|\mu - \nu_{\alpha \mathbb{EC}_n}\|_{TV}?$$

Although the last $n^{d/2}$ points to be covered on \mathbb{Z}_n^d are not exactly uniformly distributed, it is possible that the difference from uniform is dominated by the noise in the uniform measure on $\{0, 1\}^{G_n}$. Proving that this total variation distance tends to 0 for $\alpha \in (1/2, 1)$ would be a way of quantifying that the last points to be visited are nearly uniformly distributed.

- (2) More generally, for which sequences of graphs G_n is $\mathcal{T}_v(\epsilon, G_n^\diamond)$ asymptotic to \mathbb{EC}_n , and for which is it $\mathbb{EC}_n/2$? Can it be asymptotic to $\alpha \mathbb{EC}_n$ for some $\alpha \in (1/2, 1)$?
- (3) Is \mathbb{EC}_n the right order of magnitude for the total variation mixing even when the graphs G_n are not vertex transitive? That is to say, if $t_n = o(\mathbb{EC}_n)$, is $\mathcal{T}_v(\epsilon, G_n^\diamond) > t_n$ for large enough n ?
- (4) For $G_n = \mathbb{Z}_2^n$, is the correct order of magnitude for $\tau(\epsilon, G_n^\diamond)$ equal to $n2^n$ or $n(\log n)2^n$?
- (5) Can the upper bound on the uniform mixing time in Theorem 1.4 be replaced by $C|G_n|(T_{\text{rel}} + \log |G_n|)$?
- (6) To what extent are these results generator dependent? Another type of walk that is often considered on lamplighter groups is one in which at each step either the current lamp is adjusted while the lamplighter holds, or no lamps are adjusted while the lamplighter moves to a neighbor. By comparing Dirichlet forms, this only changes the relaxation time up to constants, but it is not clear what happens to the mixing times. Using log-Sobolev inequalities shows that the uniform mixing time can change by at most a factor of $\log \log |G^\diamond| \simeq \log |G|$.
- (7) What happens if we make the lamps more complicated, meaning that we consider the wreath product $H_n \wr G_n$ and let $|H_n| \rightarrow \infty$?

Acknowledgement We thank Gábor Pete for useful comments.

REFERENCES

- [1] David Aldous and Persi Diaconis, *Strong uniform times and finite random walks*, Adv. in Appl. Math. **8** (1987), no. 1, 69–97. MR 88d:60175
- [2] David Aldous and James Allen Fill, *Reversible markov chains and random walks on graphs*, Available at <http://www.stat.berkeley.edu/~aldous/RWG/book.html>.
- [3] David J. Aldous, *Threshold limits for cover times*, J. Theoret. Probab. **4** (1991), no. 1, 197–211. MR 91m:60123
- [4] Noam Berger, Claire Kenyon, Elchanan Mossel, and Yuval Peres, *Glauber Dynamics on Trees and Hyperbolic Graphs*, Extended abstract by Kenyon, Mossel, and Peres appeared in 42nd IEEE Symposium on Foundations of Computer Science (Las Vegas, NV, 2001) [arXiv:math.PR/0308284](https://arxiv.org/abs/math.PR/0308284).
- [5] Andrei Z. Broder and Anna R. Karlin, *Bounds on the cover time*, J. Theoret. Probab. **2** (1989), no. 1, 101–120. MR 90b:60079

- [6] Mu-Fa Chen, *Trilogy of couplings and general formulas for lower bound of spectral gap*, Probability towards 2000 (New York, 1995), Lecture Notes in Statist., vol. 128, Springer, New York, 1998, pp. 123–136. MR 99h:60130
- [7] Amir Dembo, Yuval Peres, Jay Rosen, and Ofer Zeitouni, *Cover times for brownian motion and random walks in two dimensions*, Ann. Math. (to appear), [arXiv:math.PR/0107191](#).
- [8] James Allen Fill and Clyde H. Schoolfield, Jr., *Mixing times for Markov chains on wreath products and related homogeneous spaces*, Electron. J. Probab. **6** (2001), no. 11, 22 pp. (electronic). MR 2002b:60121
- [9] Olle Häggström and Johan Jonasson, *Rates of convergence for lamplighter processes*, Stochastic Process. Appl. **67** (1997), no. 2, 227–249. MR 98j:60097
- [10] Peter Matthews, *Covering problems for Brownian motion on spheres*, Ann. Probab. **16** (1988), no. 1, 189–199. MR 89a:60190
- [11] Christophe Pittet and Laurent Saloff-Coste, *Amenable groups, isoperimetric profiles and random walks*, Geometric group theory down under (Canberra, 1996), de Gruyter, Berlin, 1999, pp. 293–316. MR 2001d:20041
- [12] P. Révész, *Clusters of a random walk on the plane*, Ann. Probab. **21** (1993), no. 1, 318–328. MR 94f:60047
- [13] J.C. Reyes, *Random walk, semi-direct products, and card shuffling*, Ph.D. thesis, Stanford University, 2002, pp. x+163.
- [14] Clyde H. Schoolfield, Jr., *Random walks on wreath products of groups*, J. Theoret. Probab. **15** (2002), no. 3, 667–693. MR 2003f:60018
- [15] Frank Spitzer, *Principles of random walks*, second ed., Springer-Verlag, New York, 1976, Graduate Texts in Mathematics, Vol. 34. MR 52 #9383