

On sensitivity of uniform mixing times¹

Jonathan Hermon

University of Cambridge, Cambridge, UK. E-mail: jonathan.hermon@statslab.cam.ac.uk

Received 6 July 2016; revised 26 September 2016; accepted 17 October 2016

Abstract. We show that the order of the L_∞ -mixing time of simple random walks on a sequence of uniformly bounded degree graphs of size n may increase by an optimal factor of $\Theta(\log \log n)$ as a result of a bounded perturbation of the edge weights. This answers a question and a conjecture of Kozma.

Résumé. Nous montrons que le temps de mélange pour la distance L_∞ d'une marche aléatoire sur une suite de graphe de taille n et de degré uniformément borné peut être multiplié par un facteur d'ordre $\log \log n$ (optimal) en perturbant le poids des arrêtes du graphe de manière uniformément bornée. Ceci résout une question et une conjecture de Kozma.

MSC: 60J10

Keywords: Sensitivity; Mxing-time; Sensitivity of mixing times; Hitting times

1. Introduction

An important question is whether mixing times are robust under small changes to the geometry of the Markov chain. For instance, can bounded perturbations of the edge weights change the mixing time by more than a constant factor? Similarly, how far apart can the mixing times of lazy simple random walks on two roughly-isometric graphs of bounded degree be? A related question is whether mixing times can be characterized up to universal constants (perhaps only under reversibility) using geometric quantities or extremal characterizations which are robust.² Different variants of this question were asked by various authors such as Pittet and Saloff-Coste [16], Kozma [11, p. 4], Diaconis and Saloff-Coste [6, p. 720] and Aldous and Fill [2, [Open Problem 8.23](#)] (the last two references ask for an extremal characterization of the L_∞ -mixing time in terms of the Dirichlet form).

Denote the L_p mixing time of lazy simple random walk on a finite connected simple graph G by $\tau_p(G)$ (see (1.3)). Kozma [11] made the following conjecture:

Conjecture 1.1 ([11]). *Let G and H be two finite K -roughly isometric graphs (see Definition 1.6) of maximal degree $\leq d$. Then for some $C(K, d)$ depending only on (K, d) ,*

$$\tau_\infty(G) \leq C(K, d)\tau_\infty(H).$$

Our main result, Theorem 1, asserts that this conjecture is false. We shall only consider the following particularly simple type of rough isometries. Let $G_1 := (V_1, E_1)$ be some graph. Let $G_2 = (V_2, E_2)$ be a graph obtained from G_1

¹Financial support by the EPSRC grant EP/L018896/1.

²That is, using quantities which can change by at most some bounded factor under a bounded perturbation of the edge weights, or for lazy simple random walk on a bounded degree graph, under rough-isometries.

by “stretching” some of the edges of G_1 by a factor of at most K (we say that G_2 is a K -stretch of G_1). That is, for some $E \subset E_1$ we replace each edge $\{u, v\} \in E$ by a path of length at most K (whose end-points are still denoted by u and v). Note that $V_1 \subset V_2$. The identity map is a K -rough isometry of G_1 and G_2 .

Theorem 1. *There exist two families of uniformly bounded degree simple connected graphs $(G_n = (V_n, E_n))_{n \geq 0}$ and $(G'_n)_{n \geq 0}$ and some $c > 0$ such that $|V_n| \rightarrow \infty$ and for each n , G'_n is a 2-stretch of G_n and*

$$\tau_1(G'_n) \geq c\tau_\infty(G_n) \log \log |V_n|. \tag{1.1}$$

Remark 1.2. The $\log \log |V_n|$ term in (1.1) is optimal. This follows from the fact that the Log-Sobolev constant, and as shown in [11], also the spectral-profile, provide bounds on $\tau_\infty(G)$ for a graph $G = (V, E)$ of maximal degree d which are sharp up to a $C_d \log \log |V|$ factor. These bounds are *robust* (i.e. invariant, up to a constant factor) under rough isometries in the bounded degree setup ([6, Lemmas 3.3 and 3.4]).

Remark 1.3. Denote the upper bound on $\tau_\infty(G)$ given by the spectral profile [8, Theorem 1.1] by $\rho(G)$ (see (2.6)). Kozma [11] asked whether it is possible for a sequence of bounded degree graphs $H_n = (V(H_n), E(H_n))$ with $|V_n| \rightarrow \infty$ to satisfy $\rho(H_n)/\tau_\infty(H_n) \geq c \log \log |V(H_n)|$ for some absolute constant $c > 0$. Until now, there was no known example exhibiting this behavior (or even one in which $\rho(H_n)/\tau_\infty(H_n)$ diverges). Since in the bounded degree setup the spectral profile is robust (i.e. invariant, up to a constant factor) under rough isometries, for the graphs from Theorem 1 it must be the case that for all n

$$\rho(G_n) \geq c_1 \rho(G'_n) \geq c_1 \tau_\infty(G'_n) \geq c_2 \tau_\infty(G_n) \log \log |V_n|.$$

In other words, to prove Theorem 1 one must first construct a sequence of bounded degree graphs for which the spectral profile bound overshoots the order of the L_∞ mixing time by an optimal factor. It follows from the analysis in [11] that such example must have $\Theta(\log \log |V_n|)$ disjoint sets, whose stationary probabilities are of different logarithmic scales, so that each of which makes roughly the same contribution to $\rho(G_n)$.

We believe that in general, for such example, the aforementioned sets can be chosen so that: (1) the walk can visit only one (or at most some constant number) of them before it is mixed in L_∞ , and (2) after stretching some of the edges by a bounded factor, the walk has to pass through at least some fraction of these sets in order to mix in total-variation.

We say that a family of graphs is *robust* if for every $C > 0$ there exists some $K > 0$ such that if we multiply the edge weights of some of the edges by a factor of at most C on these graphs, the corresponding L_∞ mixing times are preserved up to a factor of K . In [7] Ding and Peres studied robustness of L_1 mixing (see Section 1.3). We note that the L_∞ case is much harder (see the discussion in Section 1.3).

Remark 1.4. Loosely speaking, stretching an edge by a factor of K has the same effect as decreasing its weight to $1/K$. Indeed, in Theorem 1 instead of considering a 2-stretch of G_n we could have considered a bounded perturbation of the edge weights (i.e. the same example works for both setups).

There are numerous works aiming at sharp geometric bounds on the L_∞ -mixing time, τ_∞ , such as Morris and Peres’ evolving sets bound [14], expressed in terms of the expansion profile. The sharpest geometric bounds on τ_∞ are given in terms of the Log-Sobolev constant (see [6] for a survey on the topic) and the spectral profile bound, due to Goel et al. [8]. Both of which determine τ_∞ up to a multiplicative factor of order $\log \log[\max(e^\epsilon, 1/\min \pi(x))]$, where throughout π shall denote the stationary distribution.

These type of geometric bounds on mixing-times are robust under bounded perturbations (and in the bounded degree setup, also under rough isometries). That is, changing some of the edge weights by at most some multiplicative constant factor can change these geometric bounds only by a constant factor. Theorem 1 serves as a cautionary note on the possibility of developing sharp geometric bounds on mixing times.

In contrast with Theorem 1, many well-known families of graphs are robust. Robustness of the L_1 and L_∞ mixing times for all reversible Markov chains under changes to the holding probabilities (i.e. under addition of weighted loops) was established in [15] (by Peres and Sousi) and in [9] (by Peres and the author), respectively. Boczkowski,

Peres and Sousi [5] constructed an example demonstrating that this may fail without reversibility. Robustness of the L_1 and L_∞ mixing times for general (weighted) trees under bounded perturbations of the edge weights was established in [15] (by Peres and Sousi) and in [9] (by Peres and the author), respectively. Robustness of L_1 mixing times for general trees under rough-isometries was recently established in [1] by Addario-Berry and Roberts. Some other examples, which we borrow from [7] (where robustness of the L_1 mixing time is considered, however apart from example (2) below, the analysis of the L_∞ mixing time is identical) are collected in the next proposition.

Proposition 1.5. *The following families of graphs are robust: (1) Tori $\{\mathbb{Z}_n^d : n \geq 1\}$ for every fixed d ; (2) The giant component of a supercritical Erdős–Rényi random graph $\mathcal{G}(n, c/n)$ for a fixed $c > 1$; (3) Maximal connected component of a critical Erdős–Rényi random graph $\mathcal{G}(n, c/n)$; (4) The hypercube $\{0, 1\}^n$.*

1.1. Definitions

Given a (weighted) network $(V, E, (c_e)_{e \in E})$, where each edge $\{u, v\} \in E$ is endowed with a conductance (weight) $c_{u,v} = c_{v,u} > 0$ (with the convention that $c_{u,v} = 0$ if $\{u, v\} \notin E$), a lazy random walk on $G = (V, E)$, (X_t) , repeatedly does the following: when the current state is $v \in V$, the random walk will stay at v with probability $1/2$ and move to vertex u (such that $\{u, v\} \in E$) with probability $c_{u,v}/(2c_v)$, where $c_v := \sum_w c_{v,w}$. The default choice for $c_{u,v}$ is 1 (in which case, we say that the random walk is *unweighted*), which corresponds to lazy simple random walk on G (in which at each step the walk with equal probability either stays put or moves to a new vertex, chosen from the uniform distribution over the neighbors of its current position). Its stationary distribution is given by $\pi(x) := c_x/c_V$, where $c_V := \sum_{v \in V} c_v = 2 \sum_{e \in E} c_e$. This is a *reversible* Markov chain, i.e. $\pi(u)P(u, v) = \pi(v)P(v, u)$ for all $u, v \in V$, where throughout P denotes the transition matrix of the walk.

We denote by \mathbb{P}_x^t (resp. \mathbb{P}_x) the distribution of X_t (resp. $(X_t)_{t \geq 0}$), given that the initial state is x . Let μ, ν be two distributions on the state space Ω . Denote $a_{\mu, \nu, \pi} := \frac{|\mu(x) - \nu(x)|}{\pi(x)}$. The family of L_p distances is defined as follows.

$$\|\mu - \nu\|_{p, \pi} := \begin{cases} (\sum_x \pi(x) a_{\mu, \nu, \pi}^p)^{1/p}, & 1 \leq p < \infty, \\ \max_{x \in \Omega} a_{\mu, \nu, \pi}, & p = \infty \end{cases} \quad (1.2)$$

($p = 1$ gives twice the total-variation distance, i.e. $\|\mu - \nu\|_{1, \pi} = 2\|\mu - \nu\|_{\text{TV}}$). The ε - L_p -mixing-time of the chain is defined as

$$\tau_p(\varepsilon) := \min \left\{ t : \max_x \|\mathbb{P}_x^t - \pi\|_{p, \pi} \leq \varepsilon \right\}. \quad (1.3)$$

When $\varepsilon = 1/2$ we omit it from the notation and terminology (i.e. the L_p mixing time is defined as the $1/2$ - L_p mixing time).

Definition 1.6. Let $G_i := (V_i, E_i)$ ($i = 1, 2$) be two finite graphs. For $u, v \in V_i$, let $d_i(u, v)$ be graph distance (w.r.t. G_i) between u and v (i.e. the number of edges along the shortest path in G_i between u and v). We say that $f : V_1 \rightarrow V_2$ is a K -rough isometry of G_1 and G_2 if

(1)

$$\forall u, v \in V_1, \quad (d_1(u, v) - 1)/K \leq d_2(f(u), f(v)) \leq K(d_1(u, v) + 1).$$

(2) For every $w \in V_2$, there exists some $v \in V_1$ such that $d_2(f(v), w) \leq K$.

We say that G_1 and G_2 are K -roughly isometric if there exists such f as above.

We use the convention that C, C', C_1, \dots (resp. c, c', c_1, \dots) denote positive absolute constants which are sufficiently large (resp. small). Different appearances of the same constant at different places may refer to different numeric values.

1.2. The transitive case

The following question is re-iterated from [7] (over there it is asked for the L_1 mixing time) and essentially also from [16].

Question 1.7. Are transitive graphs robust?

Let $G = (V, E)$ be a finite connected transitive graph. Denote the eigenvalues of $I - P$ (where P is the transition matrix of lazy simple random walk on G) by $0 = \lambda_1 < \lambda_2 \leq \dots \leq \lambda_{|V|}$. Since for transitive graphs the mixing time is the same from all initial states, averaging over the starting position yields that (cf. [2, p. 284]) for all x and $t \geq 0$

$$\|P_x^{2t} - \pi\|_{\infty, \pi} = \|P_x^t - \pi\|_{2, \pi}^2 = \sum_y \pi(y) \|P_y^t - \pi\|_{2, \pi}^2 = \sum_{i=2}^{|V|} \lambda_i^{2t}$$

(transitivity is used only in the middle equality). Since the eigenvalues of $I - P$ are robust (e.g. [2, Corollary 8.4]), it follows that transitive graphs are robust under bounded perturbations which preserve transitivity. A positive answer to Question 1.7 will be obtained by a positive answer to the following question. Is it the case that for transitive graphs, also after a bounded perturbation, the ratio of the L_2 mixing time starting from the worst initial point with that starting from the best initial point is bounded?

1.3. Related work

It is classical that under reversibility the L_1 mixing time can be characterized using hitting times of sets which are “worst” in some sense (e.g. [12, Ch. 24]). Thus in order to show that it is not robust, it suffices to construct an example in which hitting times are not robust. As we explain below, this is somewhat easier.

Recently, the author and Peres obtained a characterization of the L_2 mixing time in terms of hitting time distributions. Namely, Theorem 1.1 in [13] asserts that under reversibility the L_2 mixing time is within some universal constant from the minimal time t such that for every set A of stationary probability at most $1/2$, the probability that A is not escaped from by time t is at most $\pi(A) + \frac{1}{2} \sqrt{\pi(A)(1 - \pi(A))}$. While we do not use this characterization as part of our analysis of the example from the proof of Theorem 1, it guided us in its construction and in the choices of certain parameters.

Ding and Peres [7] constructed a sequence of bounded degree graphs satisfying that the order of the total-variation mixing times strictly increases as a result of a certain sequence of bounded perturbations of the edge weights. Their construction was refined by the author and Peres in [10] (Theorem 3), which contains various additional results concerning sensitivity of mixing times and of the cutoff phenomenon under small changes to the geometry of the chain.

Our construction from the proof of Theorem 1 uses a key observation from [7]. Namely, that the harmonic measure of the walk on a tree may change drastically as a result of a bounded perturbation, and that this can be used to create examples in which hitting times are not robust in the following sense. The order of the expected hitting time of some large set A which is “worst” in some sense (both before and after the perturbation), starting from the worst initial state, may change as a result of the perturbation. The idea of exploiting the non-robustness of the harmonic measure was originally used by Benjamini [4] to study instability of the Liouville property.

Both in Ding and Peres’ construction and in our construction the chain mixes rapidly (there in total-variation and here in L_2) once it reaches a certain “huge” expander, H (and because H carries most of the stationary probability of the walk, the walk cannot mix before reaching it). We use the fact that the harmonic measure is sensitive (in the sense mentioned in the previous paragraph) in order to create “shortcuts” to H for the walk on the original graph, G_n , which are essentially “invisible” for the walk on the 2-stretched graph, G'_n .

In order to change the mixing time in total-variation it suffices to change the order of the expected hitting time of H , starting from the worst initial state. However, this does not suffice in order to change the order of the L_∞ mixing time. Thus we will work much harder and show that for every initial vertex x the following holds. The probability that the walk on the original graph does not reach H in some $t \ll \tau_1(G'_n)$ steps (through one of the aforementioned shortcuts) is much smaller than the L_∞ distance from stationarity of the distribution of the walk at time t , conditioned to not

reach H by that time. In order to achieve this, we will plant more shortcuts to H in regions in which the chain mixes slower.

In [11] Kozma constructed a sequence of reversible Markov chains satisfying that $\rho^{(n)} \geq c\tau_\infty^{(n)} \log \log(1/\min_x \pi_n(x))$, for all n , where $\rho^{(n)}$ is the spectral profile bound (see (2.6)) on the L_∞ mixing time of the n th chain in the sequence, $\tau_\infty^{(n)}$. In Kozma’s construction there are $n - \lceil \log n \rceil + 1$ “islands” $H_{\lceil \log n \rceil}, \dots, H_n$, each of size 2^{2^n} . The weights are chosen so that the stationary distribution is the uniform distribution. In his example the only way for the chain to “escape” from some H_i is by moving to a random state picked according to the uniform distribution on the state space, which occurs in each step w.p. $n2^{-n}$.

As in Kozma’s construction there is a weighted edge between all pairs of vertices, it is not clear that such a construction is possible in the bounded degree setup. Nevertheless, our example uses several ideas from Kozma’s construction.

- In Kozma’s construction, each H_i has vertex set $A_i \times B_i$. The network on H_i can be described as a “Cartesian product” of the complete graphs on A_i and B_i , in which the walk updates its B_i co-ordinate at a rate $\alpha_i(n) = o(1)$.

In our construction we use a bounded degree analog of the aforementioned network, which we denote by U_i . Namely, we replace the complete graphs by expanders. In order to delay the rate of transitions along the expander on the co-ordinate corresponding to B_i in Kozma’s construction, while keeping the graph unweighted, we stretch each of its edges by a factor $\ell_i(n) \rightarrow \infty$.

- In Kozma’s example, for all i , $|A_i| = 2^{2^n - 2^i}$ and so $\pi(A_{i+1}) \approx \pi(A_i)^2$. Moreover, each A_i has roughly the same contribution to the spectral profile bound (see (2.6)). This is achieved by tuning the rates $\alpha_i(n)$ in an appropriate manner. Namely, by setting $\alpha_{i+1}(n) = 2\alpha_i(n)$, for all i . As noted in Remark 1.3, in some sense, such behavior of the spectral profile is necessary in order for it to overshoot τ_∞ by an optimal factor (of order $\log \log |V|$). In our construction we will take $\ell_i^2(n) = 2\ell_{i+1}^2(n)$ in order to obtain the same effect.
- Recall that in Kozma’s construction each pair of vertices are connected by a weighted edge, which has the effect of bringing the chain to stationary “at once” at a fixed rate, and this is the only way the chain can escape from the “island,” H_i , it started at. In our construction we need to somehow imitate this behavior (in a bounded degree, unweighted fashion). At the same time, as noted in Remark 1.3, we need that after stretching some of the edges by a factor of two, the walk must sequentially move through all of the islands in order to mix in total variation.

In order to achieve this behavior, we “stitch” the “islands” together so that the walk can escape from each of them either to a huge expander H (once it is reached, the walk mixes rapidly) or to an island with adjacent index. The islands are glued to each other and to the expander H in a way that allows us to manipulate (by stretching some edges) the probability of escaping an “island” by reaching the expander H .

2. Preliminaries

Recall that the *spectral gap* of a reversible Markov chain with transition matrix P on a finite state space is defined as the smallest non-zero eigenvalue of $I - P$. We say that a graph G is a λ -expander if the spectral gap of lazy simple random walk on G , denoted by $\lambda(G)$, is at least λ . We say that a sequence of graphs $(G_n)_{n \in I}$ is an *expander family* if $\inf_{n \in I} \lambda(G_n) > 0$. As mentioned in Section 1.3, we shall use expanders as building blocks in our construction. By abuse of terminology, below we often refer to a single graph as an “expander.” What we actually mean by that is that all of the graphs we refer to as expanders in the construction of the family of graphs we construct form together an expander family.

Definition 2.1. Consider a Markov chain chain on a finite state space Ω with transition matrix P and stationary distribution π . We define the *Cheeger constant* of the chain as

$$\Phi := \min_{A:0 < \pi(A) \leq 1/2} Q(A)/\pi(A), \quad \text{where } Q(A) := \sum_{x \in A, y \notin A} \pi(x)P(x, y).$$

The following is the well-known discrete analog of Cheeger inequality.

Theorem 2.2 (e.g. [12, Theorem 13.14]). *If P is reversible then*

$$\Phi^2/2 \leq \lambda \leq 2\Phi. \tag{2.1}$$

By (2.1) a sequence of graphs $(G_n)_{n \in I}$ is an expander family iff $\inf_{n \in I} \Phi(G_n) > 0$. The following proposition will be useful in what comes.

Proposition 2.3. *There exists a constant $c_d > 0$ (depending only on d) such that if H is a simple graph of maximal degree d and G is a K -stretch of H , then*

$$\Phi(G) \geq c_d \Phi(H)/K \quad \text{and so} \quad \lambda(G) \geq c_d^2 \Phi^2(H)/(2K^2). \tag{2.2}$$

Proof. The argument in [13, Claim 2.2] covers the case in which $|V(G)| \leq \frac{3}{2}|V(H)|$. We shall reduce the general case to this case. It is easy to see that the minimum in the definition of Φ is always attained by a connected set. For $e = \{u, v\} \in E(H)$ let γ_e be the collection of internal vertices along the segment from u to v in G which replaced the edge e (if e was not stretched, then γ_e is empty). Let π_G be the stationary distribution of the walk on G . Similar reasoning as in [13, Claim 2.2] shows that for some connected set B

- (i) $Q_G(B)/\pi_G(B) \leq C_d \Phi(G)$ (where $Q_G(B)$ denotes Q w.r.t. the graph G).
- (ii) There is at most one $e \in E(H)$ so that $\gamma_e \setminus B$ and $\gamma_e \cap B$ are both non-empty. However, even if such e exists, $\gamma_e \cap B$ is connected.

The operation of contracting a set of vertices D is defined as follows. Replace all of D by a single vertex x and for every edge $\{u, v\}$ with u in D and $v \notin D$, replace it with an edge $\{x, v\}$ and if also $v \in D$ replace it by a loop at x of weight 2. Let G_2 be the network obtained by contracting each $\gamma_{e'}$, for all $e' \in E(H)$, apart from γ_e for the e so that $\gamma_e \setminus B$ and $\gamma_e \cap B$ are both non-empty (if such e exists). If such e exists, we contract $\gamma_e \cap B$ and also $(\gamma_e \setminus B) \cup \{v\}$, where v is the endpoint of e incident to $(\gamma_e \setminus B)$.

Let G_3 be the graph obtained by deleting all loops from G_2 . It is straightforward to check that $\Phi(G_2) \geq \Phi(G_3)/(2K) \geq c'_d \Phi(H)/K$ (where the second inequality follows from [13, Claim 2.2]). Conversely, by (i)–(ii), $\Phi(G_2) \leq Q_G(B)/\pi_G(B) \leq C_d \Phi(G)$ (where the first inequality is obtained by considering the set B' in G_2 obtained from B by replacing each $\gamma_{e'} \subset B$ by the corresponding vertex in G_2) and so indeed $\Phi(G) \geq c_d \Phi(H)/K$. \square

The Poincaré (spectral gap) inequality asserts that when time is scaled according to the inverse of the spectral gap, the L_2 distance from stationarity of every distribution decays exponentially in the number of (scaled) time units.

Lemma 2.4. *Let (Ω, P, π) be a finite lazy irreducible reversible Markov chain with spectral gap λ . Let μ be a distribution on Ω . Then*

$$\|P^t_\mu - \pi\|_{2,\pi} \leq e^{-\lambda t} \|\mu - \pi\|_{2,\pi}, \quad \text{for all } t \geq 0. \tag{2.3}$$

Definition 2.5. Let (Ω, P, π) be reversible. Let $A \subsetneq \Omega$. We define $\lambda(A)$ (“the spectral gap of the set A ”) to be the smallest eigenvalue of the substochastic matrix obtained by restricting $I - P$ to A . Similarly, define $\Phi(A) = \min_{B \subset A} Q(B)/\pi(B)$.

The following extension of (2.1) is due to Goel et al. [8, (1.4) and Lemma 2.4]. For every irreducible reversible chain, and every set A with $\pi(A) \leq 1/2$ we have that

$$\Phi^2(A)/4 \leq \lambda(A) \leq \Phi(A). \tag{2.4}$$

The *hitting time* of a set D is defined as $T_D := \inf\{t : X_t \in D\}$. Using the spectral decomposition of P_A , the restriction of P to the set A (e.g. [3, Lemma 3.8]), it is easy to show that for every set A and $a, a' \in A$ we have that $P^t_A(a, a') \leq \sqrt{\frac{\pi(a')}{\pi(a)}} e^{-\lambda(A)t}$ and so

$$\forall a \in A, \quad P_a[T_{\Omega \setminus A} > t] = \sum_{b \in A} P^t_A(a, b) \leq |A| \max_{b \in A} \sqrt{\frac{\pi(b)}{\pi(a)}} e^{-\lambda(A)t}. \tag{2.5}$$

Finally, we recall that the spectral profile upper bound on τ_∞ is [8]

$$\rho := 8\lambda^{-1} \log 2 + \int_{\min_x \pi(x)}^{1/2} \frac{4dv}{v\Lambda(v)}, \quad \text{where } \Lambda(v) := \inf_{A \subset \Omega: \pi(A) \leq v} \lambda(A). \quad (2.6)$$

3. Proof of Theorem 1

3.1. The construction

For notational convenience we often omit ceiling signs. As described in Section 1.3 we shall construct graphs U_n, \dots, U_1 . For all i we will have

$$2^{2^{3n}+5n} \leq |V(U_i)| \leq C2^{2^{3n}+6n}.$$

We will then “stitch” them together to obtain the ultimate graph $G_n = (V_n, E_n)$. Hence

$$n2^{2^{3n}+5n} \leq |V_n| \leq Cn2^{2^{3n}+6n}. \quad (3.1)$$

For all i , the graph U_i will be a Cartesian product of an expander $H_i = (V(H_i), E(H_i))$ with a graph W_i (obtained by making a small modification to a certain tree \mathcal{T}_i), of sizes

$$2^{2^{3n}-2^{2n+i}} \leq |V(H_i)| \leq 2^{2^{3n}-2^{2n+i}+n-i}, \quad 2^{2^{2n+i}+5n} \leq |V(W_i)| \leq C2^{2^{2n+i}+5n}. \quad (3.2)$$

Step 1.1 Let $1 \leq i \leq n$. We now construct the tree \mathcal{T}_i which shall have roughly $2^{2^{2n+i}}$ *good leaves* GL_i and roughly $2^{2^{2n+i-1}}$ *bad leaves* BL_i (apart from $i = 1$ which only has good leaves). It will be obtained by stretching the edges of a tree $\mathcal{T}_{bs,i}$ (bs stands for “before stretching”) which is a “binary tree,” rooted at o_i , whose good leaves, GL_i , are of depth 2^{2n+i} , while its bad leaves, BL_i , are all of some other depth, j_i . The sets of good leaves of the two trees $\mathcal{T}_{bs,i}$ and \mathcal{T}_i are the same, and likewise for the sets of bad leaves. Note that usually a finite binary tree is defined so that all of its leaves are of the same depth, while here, crucially, the leaf set is partitioned into two sets of different depths. By abuse of terminology, we still refer to such a tree as a finite binary tree.

- We first describe the construction of \mathcal{T}_1 as it is simpler. Take a binary tree, $\mathcal{T}_{bs,1}$, of depth 2^{2n+1} rooted at o_1 . Denote its leaves by GL_1 . Then stretch each of its edges by a factor of 2^{5n} .
- We now construct \mathcal{T}_i for $1 < i \leq n$ in several steps. Before describing $\mathcal{T}_{bs,i}$ we consider a binary tree of depth 2^{2n+i-2} , rooted at o_i , denoted by $\mathcal{T}_{fh,i}$ (fh stands for “first half,” as $\mathcal{T}_{fh,i}$ is the “first half” of $\mathcal{T}_{bs,i}$).

Step 1.2 For every vertex u which is not a leaf of $\mathcal{T}_{fh,i}$ we distinguish its two children by *left* and *right* child. For every vertex $u \in \mathcal{T}_{fh,i}$ let $\text{Left}(u)$ (resp. $\text{Right}(u)$) be the number of left (resp. right) children along the path from o_i to u . Let

$$g(u) = \text{Left}(u) - \text{Right}(u). \quad (3.3)$$

Denote the k th level of a rooted tree \mathcal{T} by $\mathcal{L}_k(\mathcal{T})$. We partition the leaf set of $\mathcal{T}_{fh,i}$ $\mathcal{L}_{2^{2n+i-2}}(\mathcal{T}_{fh,i})$, into two parts: GMP_i , the set of *good middle points* and BMP_i , the set of *bad middle points* (they are “middle points” w.r.t. $\mathcal{T}_{bs,i}$) defined as follows:

$$\text{GMP}_i := \{u \in \mathcal{L}_{2^{2n+i-2}}(\mathcal{T}_{fh,i}) : g(u) \leq 2^{2^{2n+i}-6}\} \quad \text{and} \quad \text{BMP}_i := \mathcal{L}_{2^{2n+i-2}}(\mathcal{T}_{fh,i}) \setminus \text{GMP}_i.$$

We now extend $\mathcal{T}_{fh,i}$ so that the good leaves, GL_i (resp. bad leaves BL_i), of the resulting tree, $\mathcal{T}_{bs,i}$, will be the leaves which are decedents of GMP_i (resp. BMP_i).

- Attach to each vertex in GMP_i a binary tree of depth $2^{2n+i} - 2^{2n+i-2}$. This makes the total number of leaves that have a vertex in GMP_i as an ancestor $2^{2^{2n+i}}(1 - o(1))$.

- Attach to each vertex in BMP_i a binary tree of depth $\lceil \log_2(2^{2^{2n+i-1}}/|\text{BMP}_i|) \rceil$ so that the set of leafs that have a vertex in BMP_i as an ancestor, BL_i , is of size

$$2^{2^{2n+i-1}} \leq |\text{BL}_i| \leq 2^{2^{2n+i-1}+1}. \quad (3.4)$$

Call the resulting tree $\mathcal{T}_{\text{bs},i}$.

Step 1.3 Denote the union of GMP_i with the collection of vertices of $\mathcal{T}_{\text{bs},i}$ which have an ancestor in GMP_i by GS_i (a shorthand for “good side”).

- For each $1 < i \leq n$, stretch each edge of $\mathcal{T}_{\text{bs},i}$ that both of its end-points lie in GS_i by a factor of 2^{5n} . Stretch each of the rest of the edges of $\mathcal{T}_{\text{bs},i}$ by a factor of

$$q_i := \lceil 2^{8n - \frac{i}{2}} \rceil. \quad (3.5)$$

Call the resulting tree \mathcal{T}_i .

Step 2 We now modify the tree \mathcal{T}_i in the region close to its leafs in order to make it an expander. We note that for $i = 1$ the below modification is done only to the region close to GL_1 as \mathcal{T}_1 only has good leafs. Let $\text{GE}_i = (V(\text{GE}_i), E(\text{GE}_i))$ and $\text{BE}_i = (V(\text{BE}_i), E(\text{BE}_i))$ (E stands for expander) be 3-regular expanders of size $|\text{GL}_i|/2$ and $|\text{BL}_i|/2$, resp. Let PGL_i (resp. PBL_i) (a shorthand for “parents of good (resp. bad) leafs” w.r.t. $\mathcal{T}_{\text{bs},i}$), be the collection of the $|\text{GL}_i|/2$ (resp. $|\text{BL}_i|/2$) vertices of distance 2^{5n} (resp. q_i) w.r.t. \mathcal{T}_i from GL_i (resp. BL_i). We naturally identify PGL_i (resp. PBL_i) with the collection of parents w.r.t. $\mathcal{T}_{\text{bs},i}$ of the vertices in GL_i (resp. BL_i). Identify each vertex of PGL_i (resp. PBL_i) with a vertex of GE_i (resp. BE_i) in a bijective manner. Let $\mathcal{T}_v = (V(\mathcal{T}_v), E(\mathcal{T}_v))$ be the induced tree at v w.r.t. \mathcal{T}_i (i.e. the induced tree on the set of vertices that the path from them to o_i goes through v). For each $u, v \in \text{PGL}_i$ (resp. PBL_i) such that $\{u, v\} \in E(\text{GE}_i)$ (resp. $E(\text{BE}_i)$), we connect each vertex $w \in V(\mathcal{T}_u)$ to $\phi_{u,v}(w) \in V(\mathcal{T}_v)$ by an edge, where $\phi_{u,v}$ is the trivial isomorphism of \mathcal{T}_u and \mathcal{T}_v . Call the resulting graph $W_i = (V(W_i), E(W_i))$.

Step 3 Let $1 \leq i \leq n-1$. Denote $s_i := \prod_{j=i+1}^n |\text{BL}_j|$. Note that by (3.4)

$$2^{2^{3n} - 2^{2n+i}} \leq s_i \leq 2^{2^{3n} - 2^{2n+i} + n - i}.$$

Let $H_i = (V(H_i), E(H_i))$ be a 3-regular expander of size s_i . Let $U_i = W_i \times H_i$ (a Cartesian product of W_i and H_i). That is, $V(U_i) = V(W_i) \times V(H_i)$ and $\{(u, h), (u', h')\} \in E(U_i)$ if either $u = u'$ and $\{h, h'\} \in E(H_i)$ or $h = h'$ and $\{u, u'\} \in E(W_i)$.

We now “stitch” together the graphs U_n, \dots, U_1 . We refer to

$$R_i := \{o_i\} \times V(H_i)$$

(for $i = n$, $R_n := \{o_n\}$) as the *roots* of U_i . We refer to the set

$$\text{Bad}_i := \text{BL}_i \times V(H_i) \quad (\text{for } i = n, \text{Bad}_n := \text{BL}_n)$$

as the *bad leafs* of U_i (even though it is not a tree). Similarly, we refer to

$$\text{Good}_i := \text{GL}_i \times V(H_i)$$

(for $i = n$, $\text{Good}_n := \text{GL}_n$) as the *good leafs* of U_i .

Note that $|\text{Bad}_{i+1}| = |R_i| = s_i$, for all i . We shall connect U_{i+1} to U_i for all $1 \leq i < n$, by identifying Bad_{i+1} with R_i (step 4). We shall also connect all of the U_i 's “at once” by connecting $\text{Good} := \bigcup_{i=1}^n \text{Good}_i$ using one “huge” expander H of size $|\text{Good}|$ (step 5).

Step 4 For all $i < n$, we take an arbitrary bijection $\phi : \text{Bad}_{i+1} \rightarrow R_i$. We then replace each $u \in \text{Bad}_{i+1}$ and $\phi(u)$ with a new vertex, u' , that the set of edges which are incident to it is the union of the edges which are incident to u (in U_{i+1}) and to $\phi(u)$ (in U_i). By abuse of notation we shall not distinguish between the set $\{u' : u \in \text{Bad}_{i+1}\}$ and the sets Bad_{i+1} and R_i .

Step 5 We identify the set Good with a 3-regular expander $H = (V(H), E(H))$ of size $|\text{Good}|$ as follows. We label Good by the set $V(H)$ and connect $u, v \in \text{Good}$ if $\{u, v\} \in E(H)$. Call the obtained graph $G_n = (V_n, E_n)$.

Step 6 We now describe $G'_n = (V'_n, E'_n)$. The only difference between it and G_n is that for all $1 < i \leq n$ we perform the following step in between step 1.1 and step 1.2:

- For every non-leaf (w.r.t. $\mathcal{T}_{\text{fh},i}$) vertex $u \in V(\mathcal{T}_{\text{fh},i})$ and its right child v , we stretch the edge $\{u, v\}$ by a factor of 2. Call the obtained graph $\mathcal{T}'_{\text{fh},i}$. The remaining steps are analogous to the ones in the construction of G_n . We spell them out for the sake of concreteness.
- The leaf set of $\mathcal{T}'_{\text{fh},i}$ can be identified with that of $\mathcal{T}_{\text{fh},i}$. Hence we may partition the leaf set of $\mathcal{T}'_{\text{fh},i}$ into the same two sets GMP_i and BMP_i defined in step 1.2 using the function g , taken again w.r.t. the tree $\mathcal{T}_{\text{fh},i}$.
- After this is done, we can proceed with step 1.2 and obtain the tree $\mathcal{T}'_{\text{bs},i}$.
- We denote the tree of index i obtained at the end of step 1.3 (when before step 1.2 one performs the aforementioned intermediate step) by \mathcal{T}'_i . The only difference between \mathcal{T}'_i and \mathcal{T}_i is that the edges that were stretched by a factor 2 in $\mathcal{T}'_{\text{fh},i}$ will be stretched ultimately by a total factor of $2q_i$ rather than just q_i (as in \mathcal{T}_i).
- The construction of G'_n is concluded by completing the remaining steps of the construction of G_n with \mathcal{T}'_i now playing the role of \mathcal{T}_i in construction of G_n .

3.2. Analysis of the construction

We start with the analysis of G'_n . We write \mathbb{P}' and \mathbb{E}' to denote probabilities and expectations w.r.t. the walk on either G'_n , or \mathcal{T}'_i , for some i (where the identity of the graph will be clear from context, and otherwise specified). Denote the stationary distribution of the walk on G'_n by π' . We will show that (for all sufficiently large n)

$$\tau_1(G'_n) \geq (n-1)2^{18n-2} =: \tau. \quad (3.6)$$

Denote $\text{Nice}_i = \text{GMP}_i \times V(H_i)$ for $2 \leq i \leq n$ and $\text{Nice}_1 = R_1$. Denote the connected component of o_n w.r.t. the cut $\text{Nice} := \bigcup_{i=1}^n \text{Nice}_i$ of the graph G'_n by Small . Recall that for every distribution μ on the state space we have that $\frac{1}{2}\|\mu - \pi'\|_{1,\pi'} = \|\mu - \pi'\|_{\text{TV}} = \max_A \mu(A) - \pi'(A)$, where $\|\mu - \pi'\|_{\text{TV}}$ is the total variation distance. It follows that

$$\|\mathbb{P}'_{o_n}(X_\tau \in \cdot) - \pi'(\cdot)\|_{\text{TV}} \geq \mathbb{P}'_{o_n}[X_\tau \in \text{Small}] - \pi'(\text{Small}) \geq \mathbb{P}'_{o_n}[T_{\text{Nice}} > \tau] - C_0 2^{-2^{n/2}},$$

where we have used the following estimate in the last inequality

$$\pi'(\text{Small}) \leq \max_{u,v} \frac{\deg(u)}{\deg(v)} \frac{|\text{Small}|}{|V(G'_n)|} \leq 15|\text{Small}|/|V(G'_n)| \leq C_0 2^{-2^{n/2}}.$$

Hence, in order to prove (3.6) it suffices to show that

$$\mathbb{P}'_{o_n}[T_{\text{Nice}} \geq \tau] = 1 - o(1). \quad (3.7)$$

In order to establish (3.7) we use the following lemma.

Lemma 3.1. *Uniformly in $2 \leq i \leq n$ and $r \in R_i$, we have that*

$$\mathbb{P}'_r[T_{\text{Bad}_i} \leq 2^{18n-2} \mid T_{\text{Bad}_i} < T_{\text{Nice}}] \leq \mathbb{P}'_{o_i}[T_{\text{BL}_i} \leq 2^{18n-2} \mid T_{\text{BL}_i} < T_{\text{GMP}_i}] = o(1/n), \quad (3.8)$$

$$\mathbb{P}'_r[T_{\text{Nice}} < T_{\text{Bad}_i}] \leq \mathbb{P}'_{o_i}[T_{\text{GMP}_i} < T_{\text{BL}_i}] + o(1/n) = o(1/n), \quad (3.9)$$

where in the l.h.s of both (3.8)–(3.9) the probability is taken w.r.t. the walk on G'_n and in the middle terms w.r.t. the walk on \mathcal{T}'_i .

As $R_i = \text{Bad}_{i+1}$ for all $1 \leq i < n$, it follows from Lemma 3.1 that w.p. $1 - o(1)$, started from o_n , before the walk on G'_n reaches Nice it has to make its way from R_{i+1} to R_i , for all $1 \leq i < n$, and each of these $n - 1$ “stages” will take it at least 2^{18n-2} steps (in which case, it must be the case that $\tau_1(G'_n) \geq (n-1)2^{18n-2}$, as desired). We believe that the assertion of Lemma 3.1 is intuitive and that from a high level perspective the proof is not complicated. For the sake of completeness we choose to present a relatively detailed proof of Lemma 3.1.

Proof of Lemma 3.1. The first inequalities in both (3.8)–(3.9) are obtained via a straightforward coupling argument and symmetry (of the sets Nice_i , Bad_i and R_i w.r.t. the H_i co-ordinate of the walk). The additive $o(1/n)$ term in the middle term of (3.9) is there to cover the following two scenarios w.r.t. the walk on G'_n , started from $r \in R_i$:

- $T_{\bigcup_{j=i+1}^n \text{Nice}_j} < T_{\text{Nice}_i \cup \text{Bad}_i}$.
- $T_{\text{PBL}_i} < T_{\text{Nice}} < T_{\text{Bad}_i}$.

Working out the details of the aforementioned couplings is left as an exercise.

We now prove the equality in (3.8). Recall the notation from step 6 of the construction. Let $(X_s^i)_{s \geq 0}$ be lazy simple random walk on \mathcal{T}'_i . We may view the walk $(X_s^i)_{s \geq 0}$ only when it visits distinct vertices of $\mathcal{T}_{\text{bs},i}$ (recall that $V(\mathcal{T}_{\text{bs},i}) \subset V(\mathcal{T}'_{\text{bs},i}) \subset V(\mathcal{T}'_i)$). That is, consider the walk $Y_j = X_{S_j}^i$, for $j \geq 0$, where $S_0 = T_{V(\mathcal{T}_{\text{bs},i})}$ and for $j \geq 1$

$$S_j = \inf\{s > S_{j-1} : X_s^i \in V(\mathcal{T}_{\text{bs},i}) \setminus \{X_{S_{j-1}}^i\}\}.$$

For a set $D \subset V(\mathcal{T}_{\text{bs},i})$ denote its hitting time w.r.t. (Y_s) by

$$\tau_D := \inf\{s : Y_s \in D\}.$$

Recall that $V(\mathcal{T}_{\text{bs},i}) \setminus \text{GS}_i$ (where GS_i is defined at step 1.3 of the construction) is the set of vertices of $\mathcal{T}_{\text{bs},i}$ which are connected in \mathcal{T}'_i to their neighbors w.r.t. $\mathcal{T}_{\text{bs},i}$ by paths of length either q_i or $2q_i$. It is easy to see that

- (a) If $Y_{j-1} \in V(\mathcal{T}_{\text{bs},i}) \setminus \text{GS}_i$, then the law of $S_j - S_{j-1}$ stochastically dominates the law of the time it takes lazy simple random walk on \mathbb{Z} to reach distance q_i from its starting position (call this law $\bar{\xi}$). Moreover, started from o_i , conditioned on $\tau_{\text{GMP}_i} \geq j$, we have that S_j stochastically dominates

$$S'_j := \sum_{k=1}^j \xi_k,$$

where ξ_1, ξ_2, \dots are i.i.d. random variables with law $\bar{\xi}$.

- (b) Started from o_i , conditioned on $\tau_{\text{BL}_i} < \tau_{\text{GMP}_i}$ we have that the law of τ_{BL_i} stochastically dominates the law of the hitting time of the 2^{2n+i-2} th level of a binary tree, by simple random walk started from its root.
- (c) The law of the last hitting time is highly concentrated around $3 \times 2^{2n+i-2}$. Hence

$$\mathbb{P}_{o_i}[\tau_{\text{BL}_i} \leq b_i \mid \tau_{\text{BL}_i} < \tau_{\text{GMP}_i}] \leq 2^{-c_0 2^{2n+i}} = o(1/n), \quad \text{where } b_i := 3 \times 2^{2n+i-3}.$$

Note that (using (3.5))

$$2^{18n-2} \leq \frac{2}{3} b_i q_i^2 = \frac{2}{3} \mathbb{E}[S'_{b_i}].$$

Using (a)–(c) above, it is not hard to verify that, for all $2 < i \leq n$,

$$\begin{aligned} & \mathbb{P}'_{o_i} [T_{\text{BL}_i} \leq 2^{18n-2} \mid T_{\text{BL}_i} < T_{\text{GMP}_i}] \\ & \leq \mathbb{P}'_{o_i} [\tau_{\text{BL}_i} \leq b_i \mid \tau_{\text{BL}_i} < \tau_{\text{GMP}_i}] + \mathbb{P}'_{o_i} [S_{b_i} \leq 2^{18n-2} \mid \tau_{\text{BL}_i} < \tau_{\text{GMP}_i}] \\ & \leq 2^{-c_0 2^{2n}} + \mathbb{P} \left[S'_{b_i} \leq \frac{2}{3} \mathbb{E}[S'_{b_i}] \right] = o(1/n), \end{aligned}$$

where the first probability is taken w.r.t. the walk on \mathcal{T}'_i and $\mathbb{P}'_{o_i} [S_{b_i} \leq 2^{18n-2} \mid \tau_{\text{BL}_i} < \tau_{\text{GMP}_i}]$ w.r.t. $(X_s^i)_{s \geq 0}$ (in the sense that (S_j) and the walk (Y_j) are both determined by $(X_s^i)_{s \geq 0}$). This concludes the proof of (3.8). We now prove the equality in (3.9).

Fix $1 < i \leq n$. Consider an infinite tree \mathcal{T}' obtained by starting with \mathcal{T} , an infinite binary tree, rooted at o , and stretching every edge between each vertex and its left (resp. right) child by a factor of q_i (resp. $2q_i$). Let v_k (resp. u_k) be the last (resp. first) vertex in $\mathcal{L}_k(\mathcal{T})$ to be visited by a simple random walk on \mathcal{T}' started from o .

As before, for every $v \in V(\mathcal{T})$ let $g(v) = \text{Left}(v) - \text{Right}(v)$ (this function is defined w.r.t. \mathcal{T}). Using a network reduction and the well-known connection between random walks and electrical networks (cf. [12, Example 9.9] for a similar but simpler network reduction), we argue that

Lemma 3.2. $(g(v_{k+1}) - g(v_k))_{k \geq 0}$ is a sequence of i.i.d. random variables, each equals to 1 w.p. $\frac{\sqrt{2}}{1+\sqrt{2}}$ and to -1 w.p. $\frac{1}{1+\sqrt{2}}$.

We do not care so much about the exact value of the constant $\frac{\sqrt{2}}{1+\sqrt{2}}$. The important point is that it is larger than $1/2$. Before proving the lemma, we explain how it implies (3.9). It follows from Lemma 3.2 that $g(v_{2^{2n+i-2}})$ is highly concentrated around $2^{2n+i-2}(\frac{\sqrt{2}}{1+\sqrt{2}})$. It is not hard to verify that this implies that $g(u_{2^{2n+i-2}}) \geq 2^{2n+i-5}$ w.p. $1 - o(1/n)$ (uniformly in i).

Using a coupling argument (in which the walks on \mathcal{T}' and on \mathcal{T}'_i , started from o and o_i , resp., are coupled so that they follow the same trajectory until they reach the 2^{2n+i-2} th level of the corresponding non-stretched trees) we obtain (3.9) (we leave the details as an exercise). This concludes the proof of $\tau_1(G'_n) \geq (n-1)2^{18n-2}$.

Proof of Lemma 3.2. The fact that $(g(v_{k+1}) - g(v_k))_{k \geq 0}$ is a sequence of i.i.d. random variables follows by symmetry and the definition of the sequence $(v_i)_{i \geq 0}$. Let w be the effective conductance (e.g. [12, Ch. 9]) from the root of \mathcal{T}' to “infinity” (i.e. the limit of the effective conductance between the root and $\mathcal{L}_\ell(\mathcal{T}')$ as $\ell \rightarrow \infty$). Similarly, let w_R and w_L be the effective conductances between the root and infinity in the right and left subtrees of \mathcal{T}' (where the right subtree is obtained by deleting from \mathcal{T}' the left child of the root and all of its descendants, and the left subtree is similarly defined), resp.. Then $w = w_R + w_L$ and by symmetry $\frac{1}{w_R} = 2q_i + \frac{1}{w}$ and $\frac{1}{w_L} = q_i + \frac{1}{w}$. The proof is concluded by solving these equations (which yields $w = \frac{1}{\sqrt{2}q_i}$ and $w_L = \frac{1}{(1+\sqrt{2})q_i}$) and noting that $\mathbb{P}[g(v_1) - g(v_0) = 1] = \frac{w_L}{w}$, where the probability is taken w.r.t. SRW on \mathcal{T}' , started at its root. \square

We now show that $\tau_\infty(G_n) \leq C_1 2^{18n}$. By standard results it suffices to show that there exist $C > 0$ and $\ell > 1$ such that for all n ,

$$\tau_\ell(G_n) \leq \lceil C 2^{18n} \rceil =: t$$

(C shall and ℓ shall be determine later, and throughout we assume that C is sufficiently large and $\ell - 1$ is sufficiently small so that all of the equations involving them below are satisfied). This follows from the following standard fact (e.g. [17, Lemma 2.4.6]).

Fact 3.3. For every reversible Markov chain we have that for all $a > 0$

$$\begin{aligned} 2\tau_2(\sqrt{a}) &= \tau_\infty(a), \\ \forall 1 < \ell < 2, \quad \tau_2(a^{m_\ell}) &\leq m_\ell \tau_\ell(a), \quad \text{where } m_\ell := 1 + \lceil (2 - \ell)/(2\ell - 2) \rceil. \end{aligned} \tag{3.10}$$

We argue that the spectral gap of G_n , denoted by $\lambda(G_n)$, satisfies

$$\lambda(G_n) \geq c_1 / \left(\max_i q_i \right)^2 = c_1 q_2^{-2} \geq c_2 2^{-16n}. \tag{3.11}$$

This follows from Proposition 2.3, as G_n is a q_2 -stretch of a bounded degree expander.

We start with an elementary observation which shall be used below repeatedly. Let $x \in V_n$ and $t > 0$. Let A be some event (which is determined by (X_0, \dots, X_t)). Denote its complement by A^c . Then for all $1 < \ell \leq 2$

$$\|P'_x - \pi\|_{\ell, \pi} = P_x[A] \|P'_x[\cdot | A] - \pi(\cdot)\|_{\ell, \pi} + P_x[A^c] \|P'_x[\cdot | A^c] - \pi(\cdot)\|_{\ell, \pi}. \tag{3.12}$$

In particular (using the fact that for every distribution μ on V_n and every $\ell > 1$ we have that $\|\mu - \pi\|_{\ell, \pi} \leq C'|V_n|^{(\ell-1)/\ell}$ if $P_x[A^c] \ll |V_n|^{-(\ell-1)/\ell}$, we may neglect the second term in the r.h.s. above, and concentrate on bounding $P_x[A] \|P'_x[\cdot | A] - \pi(\cdot)\|_{\ell, \pi}$ from above.

We now argue that starting from a vertex in the expander H the walk mixes in L_2 in at most 2^{17n} steps (this is a wasteful estimate, but it suffices for our purposes). Hence, by the Markov property and the triangle inequality, it suffices to show that for $t = \lceil C2^{18n} \rceil$,

$$\forall x, \quad \mathbb{P}_x[T_{V(H)} > t] \left\| \mathbb{P}_x^t[\cdot \mid T_{V(H)} > t] - \pi(\cdot) \right\|_{\ell, \pi} \leq 1/4. \tag{3.13}$$

Lemma 3.4.

$$\lim_{n \rightarrow \infty} \max_{x \in V(H)} \left\| \mathbb{P}_x^{2^{17n}} - \pi \right\|_{2, \pi} = 0.$$

Proof. Let

$$J := \text{PGL}_n \cup \left(\bigcup_{i=1}^{n-1} \text{PGL}_i \times V(H_i) \right)$$

be the collection of vertices of distance (w.r.t. G_n) 2^{5n} from $V(H)$. Each vertex in J is connected to two vertices in $V(H)$ by a path of length 2^{5n} . Note that (started from $V(H)$) by time 2^{14n} the walk reaches J w.p. at least $1 - C_2|V_n|^{-2}$. Hence (by the above discussion regarding (3.12)) we can neglect the case this fails. Let i be the index such that $X_{T_J} \in V(U_i)$. Trivially, by the time the walk can cross a path of length 2^{5n} it must make at least 2^{5n} steps. But this means that, w.p. at least $1 - C_2|V_n|^{-2}$, the walk will make at least 2^{5n-6} steps along each of the expanders H_i and GE_i between the last visit to $V(H)$ prior to T_J and T_J . Using (3.1)–(3.2) and (2.3) it is not hard to show that this means that at time 2^{14n} the L_2 distance of the walk (started from H) from π is $O(n2^{6n})$ (as we later perform a similar more subtle calculation, we leave this as an exercise). Finally, the claim is obtained using (3.11) and the Poincaré (spectral gap) inequality (2.3). \square

To conclude the proof we now verify (3.13). The case that $x \in V(U_1)$ can be treated separately in a similar manner to the analysis below. The following lemma asserts that w.l.o.g. we may assume that the initial position of the walk is in R_i for some i .

Lemma 3.5. *Provided that $\ell - 1$ is sufficiently small and C is sufficiently large, for all $1 < i \leq n$ and $x \in V(U_i)$ (uniformly)*

$$\mathbb{P}_x[T_{V(H)} > t/2, T_{R_i \cup R_{i-1}} > t/2] \left\| \mathbb{P}_x^{t/2}[\cdot \mid T_{V(H)} > t/2, T_{R_i \cup R_{i-1}} > t/2] - \pi(\cdot) \right\|_{\ell, \pi} = o(1). \tag{3.14}$$

Proof. We first argue that for all $1 < i \leq n$ and $x \in V(U_i)$ (provided that C is taken to be sufficiently large)

$$\mathbb{P}_x[T_{V(H)} > t/2, T_{R_i \cup R_{i-1}} > t/2] \leq C_3 2^{-2^{2n+i}}. \tag{3.15}$$

Let Z_i (resp. \bar{Z}_i) be the number of steps the walk made by time $t/2$ along the W_i (resp. H_i) co-ordinate of U_i . As $\mathbb{P}_x[Z_i < t/30] < C_2|V(G_n)|^{-2}$, in order to prove (3.15) it suffices to show that (provided that C is taken to be sufficiently large)

$$\max_{u \in V(W_i)} \mathbb{P}_u[T_{\text{GL}_i \cup \text{BL}_i \cup \{o_i\}} > t/30] \leq 2^{-2^{2n+i}}, \tag{3.16}$$

where the probability is taken w.r.t. the walk on W_i . Let \tilde{P}_i and $\tilde{\pi}_i$ be the transition matrix and stationary distribution (resp.) of lazy simple random walk on W_i . Let $\tilde{\lambda}_i$ be the smallest eigenvalue of the substochastic matrix obtained by restricting $I - \tilde{P}_i$ to $V(W_i) \setminus (\text{GL}_i \cup \text{BL}_i \cup \{o_i\})$. Using (2.4) and an obvious extension of Proposition 2.3 it is not hard to verify that

$$\tilde{\lambda}_i \geq \tilde{c} 2^{-(16n-i)}.$$

The proof of (3.16) (and so also of (3.15)) can now be concluded using (2.5).

Our next goal is to show that

$$\left\| \mathbb{P}_x[X_{t/2} \in \cdot \mid T_{V(H)} > t/2, T_{R_i \cup R_{i-1}} > t/2] - \pi(\cdot) \right\|_{\ell, \pi} \leq C_7 2^{(\ell-1)2^{2n+i+2}}. \quad (3.17)$$

Observe that this, in conjunction with (3.15), implies the assertion of the lemma.

For every $(u, h) \in V(W_i) \times V(H_i)$ denote

$$p_{u,h} := \mathbb{P}_x[X_{t/2} \in (u, h) \mid T_{V(H)} > t/2, T_{R_i \cup R_{i-1}} > t/2].$$

As before, since $\mathbb{P}_x[\bar{Z}_i < t/20] < C_2 |V(G_n)|^{-2}$ we can neglect the case that $\bar{Z}_i < t/20$. Hence we can make the following estimate: for every $h \in V(H_i)$ we have that for all $1 < i \leq n$ and $x \in V(U_i)$

$$\sum_{u \in V(W_i)} p_{u,h} \leq C'_4 \sum_{u \in V(W_i)} \mathbb{P}_x[X_{t/2} \in (u, h) \mid T_{V(H)} > t/2, T_{R_i \cup R_{i-1}} > t/2, \bar{Z}_i \geq t/20].$$

Note that the conditioning on $T_{V(H)} > t/2, T_{R_i \cup R_{i-1}} > t/2$, does not affect the projection of the walk onto its H_i co-ordinate, given the number of steps made on that co-ordinate. Also observe that conditioned on $T_{V(H)} > t/2, T_{R_i \cup R_{i-1}} > t/2$, the aforementioned projection (up to time $t/2$, viewed in times in which the H_i co-ordinate changes) is itself a random walk on the expander H_i . This, in conjunction with the fact that the L_∞ mixing time of SRW on H_i is at most $t/20$ (to see this, use (2.3), (3.10) and the fact that $|V(H_i)| = |R_i| \leq 2^{2^{3n-2^{2n+i}+n-i}}$), implies that

$$\sum_{u \in V(W_i)} \mathbb{P}_x[X_{t/2} \in (u, h) \mid T_{V(H)} > t/2, T_{R_i \cup R_{i-1}} > t/2, \bar{Z}_i \geq t/20] \leq \bar{C}_4 / |V(H_i)|.$$

Combining the last two inequalities, we get that for every $h \in V(H_i)$, for all $1 < i \leq n$ and $x \in V(U_i)$

$$\sum_{u \in V(W_i)} p_{u,h} \leq C_4 / |V(H_i)|. \quad (3.18)$$

Since $\sum_i^k x_i^\ell$ subject to the constraints $x_i \in \mathbb{R}_+$, for all i , and $\sum_{i=1}^k x_i = a > 0$ is at most a^ℓ , provided that $\ell \geq 1$, using (3.18) we get that for all $1 < i \leq n$ and $x \in V(U_i)$

$$\begin{aligned} & \left\| \mathbb{P}_x[X_{t/2} \in \cdot \mid T_{V(H)} > t/2, T_{R_i \cup R_{i-1}} > t/2] - \pi(\cdot) \right\|_{\ell, \pi}^\ell \\ & \leq C_5 \sum_{h \in V(H_i)} 1/|V(G_n)| \sum_{u \in V(W_i)} p_{u,h}^\ell / (1/|V(G_n)|)^\ell \\ & \leq C_6 |V(G_n)|^{\ell-1} \sum_{h \in V(H_i)} |V(H_i)|^{-\ell} = C_6 \left(\frac{|V(G_n)|}{|V(H_i)|} \right)^{\ell-1} \leq C_7 2^{(\ell-1)2^{2n+i+2}}, \end{aligned} \quad (3.19)$$

as desired. □

Using (3.14), in order to prove (3.13) it suffices to show that for all i and all $x \in R_i$ we have (provided that $\ell - 1$ is sufficiently small and C is sufficiently large) that

$$\mathbb{P}_x[T_{V(H)} > t/2] \left\| \mathbb{P}_x^{t/2}[\cdot \mid T_{V(H)} > t/2] - \pi(\cdot) \right\|_{\ell, \pi} \leq 1/8. \quad (3.20)$$

The case $i = 1$ can be treated separately in a similar manner to the analysis below. Hence below we assume that $i > 1$ and $x \in R_i$.

We argue that (provided that C is sufficiently large) for all $i > 1$ and $x \in R_i$,

$$c_3 \mathbb{P}_x \left[T_{V(H)} > \frac{t}{2} \right] \leq \mathbb{P}_x [T_{R_{i-1}} < T_{V(H)}] \leq \frac{|BMP_i|}{|GMP_i \cup BMP_i|} = \frac{|BMP_i|}{2^{2^{2n+i-2}}} \leq 2^{-c_2 2^{2n+i}}. \quad (3.21)$$

The proof of the first inequality in (3.21) is similar to the proof of (3.15) and hence omitted. The last inequality in (3.21) is trivial. Finally, using a coupling argument, the second inequality in (3.21) can be proven by noting that started from o_i , the hitting distribution of $\text{GMP}_i \cup \text{BMP}_i = \mathcal{L}_{2^{2n+i-2} \times q_i}(\mathcal{T}_i) = \mathcal{L}_{2^{2n+i-2}}(\mathcal{T}_{\text{bs},i})$ w.r.t. simple random walk on \mathcal{T}_i is the uniform distribution on this set. We omit the details.

It is easy to verify that starting at $x \in R_i$, the walk will make by time $t/2$ at least 2^{5n} consecutive steps on either U_i or U_{i+1} w.p. at least $1 - C_8|V(G_n)|^{-2}$. On this event, as before, the walk will make by time $t/2$ at least 2^{5n-6} steps on either H_i or on both of H_{i+1} and BE_{i+1} (in one of its visits to either U_i or U_{i+1} , resp., in which it stays in it for at least 2^{5n} steps) w.p. at least $1 - C_9|V(G_n)|^{-2}$. Using the same reasoning as in (3.18)–(3.19), we get that for all $i \geq 1$ and $x \in R_i$,

$$\left\| \mathbb{P}_x^{t/2}[\cdot \mid T_{V(H)} > t/2] - \pi(\cdot) \right\|_{\ell, \pi}^{\ell} \leq C_7 2^{(\ell-1)2^{2n+i+2}}. \quad (3.22)$$

We leave the verification of (3.22) as an exercise. Using (3.21)–(3.22) we get that for all $i > 1$ and $x \in R_i$ (provided that $\ell - 1$ is sufficiently small and C is sufficiently large),

$$\mathbb{P}_x[T_{V(H)} > t/2] \left\| \mathbb{P}_x^{t/2}[\cdot \mid T_{V(H)} > t/2] - \pi(\cdot) \right\|_{\ell, \pi} \leq 1/8. \quad (3.23)$$

This concludes the proof of $\tau_2(G_n) \leq \lceil C2^{18n} \rceil$. □

4. Proof of Proposition 1.5

(1) Torus \mathbb{Z}_n^d (the argument is almost identical to that in [7]): The lower bound follows from the fact that the inverse of the spectral gap (which is robust and up to a $\log 2$ factor, bounds τ_1 from below) is at least $c_d n^2$, while the upper bound can be deduced from the Morris and Peres' Evolving sets bound [14], which is also robust. (2) The giant component of an Erdős–Rényi supercritical random graph $\mathcal{G}(n, c/n)$: the existence of paths of length $\Theta(\log n)$ implies that the inverse of the Log-Sobolev constant, is at least $c \log^3 n$ (this can be derived using Lemma 4.2 in [8]). This provides a robust lower bound on τ_{∞} . The upper bound follows from the fact that in this case the spectral gap satisfies $\lambda := \Theta(1/\log^2 n)$ and so $\tau_{\infty} \leq C\lambda^{-1} \log n$ (since c is fixed, the average degree is uniformly bounded). Finally, for example (3) and (4) the argument is identical to that in [7] and hence omitted.

Acknowledgements

The author is grateful to Gady Kozma and Yuval Peres for useful discussions. The author would like to thank them and the anonymous referee for making suggestions which improved the presentation of the note.

References

- [1] L. Addario-Berry and M. I. Roberts. Mixing time bounds via bottleneck sequences. *J. Stat. Phys.* (2017) 1–27.
- [2] D. Aldous and J. Fill. Reversible Markov chains and random walks on graphs. Unfinished manuscript. Available at <http://www.stat.berkeley.edu/~aldous/RWG/book.html>.
- [3] R. Basu, J. Hermon and Y. Peres. Characterization of cutoff for reversible Markov chains. *Ann. Probab.* **45** (3) (2017) 1448–1487. [MRMR3650406](#)
- [4] I. Benjamini. Instability of the Liouville property for quasi-isometric graphs and manifolds of polynomial volume growth. *J. Theoret. Probab.* **4** (3) (1991) 631–637. [MR1115166](#)
- [5] L. Boczkowski, Y. Peres and P. Sousi. Sensitivity of mixing times in Eulerian digraphs. Preprint, 2016. Available at [arXiv:1603.05639](#).
- [6] P. Diaconis and L. Saloff-Coste. Logarithmic Sobolev inequalities for finite Markov chains. *Ann. Appl. Probab.* **6** (3) (1996) 695–750. [MR1410112](#)
- [7] J. Ding and Y. Peres. Sensitivity of mixing times. *Electron. Commun. Probab.* **18** (2013) 1–6. [MR3141797](#)
- [8] S. Goel, R. Montenegro and P. Tetali. Mixing time bounds via the spectral profile. *Electron. J. Probab.* **11** (1) (2006) 1–26. [MR2199053](#)
- [9] J. Hermon and Y. Peres. A characterization of L_2 mixing and hypercontractivity via hitting times and maximal inequalities. *Probab. Theory Related Fields* (2017) 1–32.
- [10] J. Hermon and Y. Peres. On sensitivity of mixing times and cutoff. Preprint, 2016. Available at [arXiv:1610.04357](#).
- [11] G. Kozma. On the precision of the spectral profile. *ALEA Lat. Am. J. Probab. Math. Stat.* **3** (2007) 321–329. [MR2372888](#)
- [12] D. Levin, Y. Peres and E. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, Providence, RI, 2017. [MR2466937](#)

- [13] E. Lubetzky and A. Sly. Explicit expanders with cutoff phenomena. *Electron. J. Probab.* **16** (15) (2011) 419–435. [MR2774096](#)
- [14] B. Morris and Y. Peres. Evolving sets, mixing and heat kernel bounds. *Probab. Theory Related Fields* **133** (2) (2005) 245–266. [MR2198701](#)
- [15] Y. Peres and P. Sousi. Mixing times are hitting times of large sets. *J. Theoret. Probab.* **28** (2) (2015) 488–519. [MRMR3370663](#)
- [16] C. Pittet and L. Saloff-Coste. On the stability of the behavior of random walks on groups. *J. Geom. Anal.* **10** (4) (2000) 713–737. [MR1817783](#)
- [17] L. Saloff-Coste. Lectures on Finite Markov Chains. In *Lectures on Probability Theory and Statistics* 301–413. Springer, Berlin, 1997. [MR1490046](#)