Electronic Journal of Probability

Vol. 16 (2011), Paper no. 15, pages 419–435.

Journal URL
http://www.math.washington.edu/~ejpecp/

Explicit expanders with cutoff phenomena

Eyal Lubetzky* Allan Sly†

Abstract

The cutoff phenomenon describes a sharp transition in the convergence of an ergodic finite Markov chain to equilibrium. Of particular interest is understanding this convergence for the simple random walk on a bounded-degree expander graph. The first example of a family of bounded-degree graphs where the random walk exhibits cutoff in total-variation was provided only very recently, when the authors showed this for a typical random regular graph. However, no example was known for an explicit (deterministic) family of expanders with this phenomenon. Here we construct a family of cubic expanders where the random walk from a worst case initial position exhibits total-variation cutoff. Variants of this construction give cubic expanders without cutoff, as well as cubic graphs with cutoff at any prescribed time-point.

Key words: Cutoff phenomenon, Random walks, Expander graphs,
Explicit constructions.

AMS 2000 Subject Classification: Primary 60B10, 60J10, 60G50, 05C81.

Submitted to EJP on April 21, 2010, final version accepted December 12, 2010.

*Microsoft Research, One Microsoft Way, Redmond, WA 98052. Email: eyal@microsoft.com
†Microsoft Research, One Microsoft Way, Redmond, WA 98052. Email: allansly@microsoft.com

1 Introduction

A finite ergodic Markov chain is said to exhibit *cutoff* in total-variation if its L^1-distance from the stationary distribution drops abruptly from near its maximum to near 0. In other words, one should run the Markov chain until the cutoff point for it to even slightly mix in L^1 whereas running it any further is essentially redundant.

Let (X_t) be an aperiodic irreducible discrete-time Markov chain on a finite state space Ω with stationary distribution π. The worst-case total-variation distance to stationarity at time t is defined as

$$d(t) \stackrel{\triangle}{=} \max_{x\in\Omega} \|\mathbb{P}_x(X_t \in \cdot) - \pi\|_{\mathrm{TV}}\,,$$

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

$$\|\mu-\nu\|_{\mathrm{TV}} \stackrel{\triangle}{=} \sup_{A\subset\Omega} \big|\mu(A)-\nu(A)\big| = \frac{1}{2}\sum_{x\in\Omega} |\mu(x)-\nu(x)|\,.$$

Define $t_{\mathrm{MIX}}(\varepsilon)$, the total-variation *mixing-time* of (X_t) for $0<\varepsilon<1$, to be

$$t_{\mathrm{MIX}}(\varepsilon) \stackrel{\triangle}{=} \min\{t : d(t) < \varepsilon\}\,.$$

Let $\big(X_t^{(n)}\big)$ be a family of such chains, each with its total-variation distance from stationarity $d_n(t)$, its mixing-time $t_{\mathrm{MIX}}^{(n)}$, etc. This family exhibits *cutoff* iff the following sharp transition in its convergence to equilibrium occurs:

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

The rate of convergence in (1.1) is addressed by the following notion of a *cutoff window*: For two sequences t_n, w_n with $w_n = o(t_n)$ we say that $\big(X_t^{(n)}\big)$ has cutoff at time t_n with window w_n if and only if

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

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

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

The cutoff phenomenon was first identified for random transpositions on the symmetric group in [10] and for random walks on the hypercube in [2]. The term "cutoff" was coined by Aldous and Diaconis in [3], where cutoff was shown for the top-in-at-random card shuffling process. While believed to be widespread, there are relatively few examples where the cutoff phenomenon has been rigorously confirmed. Even for fairly simple chains, determining whether there is cutoff often requires the full understanding of their delicate behavior around the mixing threshold. See [8,9,19] and the references therein for more on the cutoff phenomenon.

A specific Markov chain which found numerous applications in a wide range of areas in mathematics over the last quarter of a century is the simple random walk (SRW) on a bounded-degree *expander* graph. A finite graph is called an expander if every small subset of the vertices has a relatively large

edge boundary. Formally, the Cheeger constant of a d-regular graph G on n vertices (also referred to as the edge isoperimetric constant) is defined as

$$\mathrm{ch}(G) = \min_{\emptyset \neq S \subsetneq V(G)} \frac{|\partial S|}{|S| \wedge |V(G) \setminus S|},$$

where ∂S is the set of edges with exactly one endpoint in S. We say that G is a c-edge-expander for some fixed $c > 0$ if it satisfies $\mathrm{ch}(G) > c$. The well-known discrete analogue of Cheeger's inequality [5, 6, 12, 20] implies that the spectral-gap of the SRW on a family of c-edge-expander graphs on n vertices is uniformly bounded away from 0, hence these chains rapidly converge to equilibrium within $O(\log n)$ steps. See the survey [14] for more on the applications of random walks on expanders.

In 2004, Peres[16] observed that for any family of reversible Markov chains, total-variation cutoff can only occur if the inverse spectral-gap has smaller order than the mixing time. Note that this condition clearly holds for the simple random walk on an n-vertex expander, where the inverse-gap is $O(1)$ whereas $t_{\mathrm{MIX}} \asymp \log n$. It was shown by Chen and Saloff-Coste [8] that when measuring convergence in L^p-distance for $p > 1$ this criterion does ensure cutoff, however the case $p = 1$ (cutoff in total-variation) has proved to be significantly more complicated. There are known examples where the above condition does not imply cutoff (see [8, Section 6]), yet it was conjectured by Peres to be sufficient in many natural families of chains (e.g. [11] confirming this for birth-and-death chains). In particular, this was conjectured for the lazy random walk on bounded-degree transitive graphs.

The first example of a family of bounded-degree graphs where the random walk exhibits cutoff in total-variation was provided only very recently [15], when the authors showed this for a typical random regular graph. It is well known that for any fixed $d \geq 3$, a random d-regular graph is with high probability (w.h.p.) a very good expander, hence the simple random walk on almost every d-regular expander exhibits worst-case total-variation cutoff. However, to this date there were no known examples for an *explicit* (deterministic) family of expanders with this phenomenon.

In Section 2.1 we provide what is, to the best of our knowledge, the first explicit construction of a family of bounded-degree expanders where the simple random walk has worst-case total-variation cutoff.

Theorem 1. *There is an explicit family of* 3*-regular expanders on which the* SRW *from a worst case initial position exhibits total-variation cutoff.*

The construction mimics the behavior of the SRW on random regular graphs, whose mixing was shown in [15] (as conjectured by Durrett [13] and Berestycki [7]) to resemble that of a walk started at a root of a d-regular tree. Two smaller expanders that are embedded into the graph structure allow careful control over the mixing time from all possible initial positions.

A straightforward modification of the above construction yields an explicit family of cubic expanders where the SRW from a worst-case initial position does *not* exhibit cutoff in total-variation (despite Peres' cutoff criterion). Note that Peres and Wilson [17] had already sketched an example for a family of expanders without total-variation cutoff. We describe our simple construction achieving this in Section 2.2 for completeness.

A final variant of the construction, presented in Section 2.3, provides cubic graphs with cutoff occurring at essentially any prescribed order of location. Namely, there is an explicit family of cubic

graphs where the SRW has cutoff at any specified order between $[\log n, n^2)$ whereas for $t_{\text{MIX}} \asymp n^2$ there cannot be cutoff (it is well-known that on any family of bounded-degree graphs on n vertices the SRW has $c \log n \leq t_{\text{MIX}} \leq c' n^2$ for some fixed $c, c' > 0$).

Theorem 2. *Let t_n be a monotone sequence with $t_n \geq \log n$ and $t_n = o(n^2)$. There is an explicit family (G_n) of 3-regular graphs with $|G_n| \asymp n$ vertices where the* SRW *from a worst-case initial position exhibits total-variation cutoff at $t_{\text{MIX}} \asymp t_n$.*

Furthermore, for any family of bounded-degree n-vertex graphs where the SRW *has $t_{\text{MIX}} \asymp n^2$ (largest possible order of mixing) there cannot be cutoff.*

2 Explicit constructions achieving cutoff

2.1 Proof of Theorem 1: explicit expanders with cutoff

To simplify the exposition, we will first construct a family of 5-regular expanders where the SRW from a worst initial position exhibits cutoff. Subsequently, we will describe how to modify the construction to yield a family of cubic expanders with this property (as per the statement of Theorem 1).

The graph we construct will contain a smaller explicit expander on a fixed proportion of its vertices, connected to what is essentially a product of another expander with a "stretched" regular tree (one where the edges in certain levels are replaced by paths).

Let $h \to \infty$ and let H_1, H_2 be two explicit expanders as follows (cf. e.g. [1] for an explicit construction of a 3-regular graph, as well as [18] and the references therein for additional explicit constructions of constant-degree expander graphs):

- H_1 : An explicit 3-regular expander on $20 \cdot 2^{2h}$ vertices.
- H_2 : An explicit 4-regular expander on $20 \cdot 2^{6h}$ vertices.

Let $\lambda(H_i)$ denote the largest absolute-value of any nontrivial eigenvalue of H_i for $i = 1, 2$. Finally, let L be some sufficiently large fixed integer whose value will be specified later.

Our final construction for 5-regular expander will be based on a (modified) regular tree, hence it will be convenient to describe its structure according to the various levels of its vertices. Let the vertex ρ denote the root of the tree, and construct the graph G as follows:

1. Levels $0, 1, 2$: First levels of a 5-regular tree rooted at ρ.
   - Denote by $U = \{u_1, \ldots, u_{20}\}$ the vertices comprising level 2.
2. Levels $3, \ldots, h+2$: Stretched 4-ary trees rooted at each vertex of U, defined as follows:
   - For each $u_i \in U$ place an h-level 4-ary tree $\mathcal{T}_{u_i}$ rooted at u_i and identify the vertices of $\mathcal{T}_{u_i}$ and $\mathcal{T}_{u_j}$ via the trivial isomorphism.
   - Replace every edge of each $\mathcal{T}_{u_i}$ by a (disjoint) path of length L.
     Connect $\mathcal{T}^*_{u_1}$, the new interior vertices in $\mathcal{T}_{u_1}$ (with initial degree 2) to their isomorphic counterparts in $\mathcal{T}^*_{u_2}, \mathcal{T}^*_{u_3}, \mathcal{T}^*_{u_4}$ (add 4-cliques between identified interior vertices) and similarly for $\{\mathcal{T}^*_{u_5}, \ldots, \mathcal{T}^*_{u_8}\}$ etc.

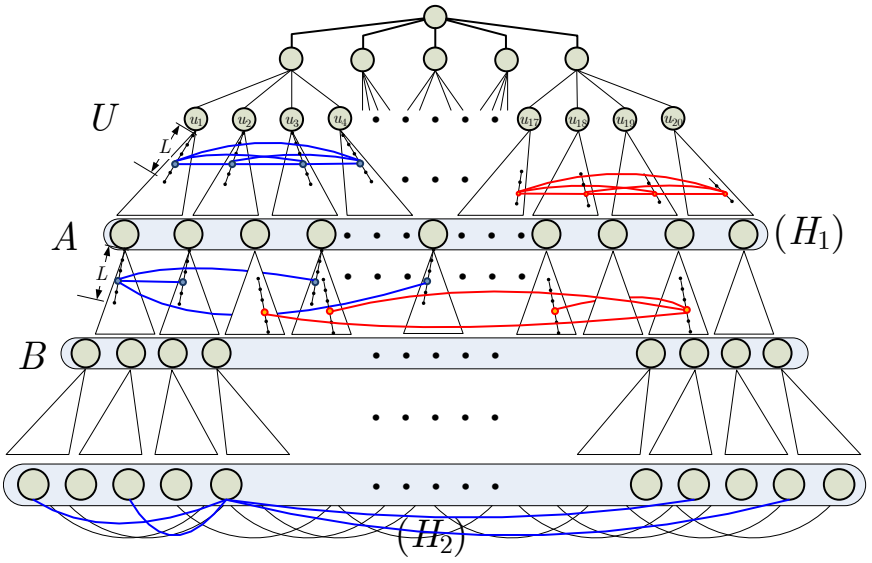

Figure 1: Explicit construction of a 5-regular graph on which the random walk exhibits total-variation cutoff.

- Let A denote the final $20 \cdot 4^h$ vertices comprising level $h+2$, and associate the vertices of A with those of H_1.

3. Levels $h+3, \ldots, 2h+2$: Product of H_1 and a stretched 4-ary tree, defined as follows:
   - For each $a \in A$ place an h-level L-stretched 4-ary tree $\mathcal{T}_a$.
     Connect vertices in $\mathcal{T}_a^*$ with their counterparts in $\mathcal{T}_b^*$ for $ab \in E(H_1)$.
   - Let B denote the final $20 \cdot 4^{2h}$ vertices comprising level $2h+2$.
4. Levels $2h+3, \ldots, 3h+2$: A forest of 4-ary trees rooted at each vertex of B.
5. Last level: Associate leaves with H_2 and interconnect them accordingly.

Finally, the aforementioned parameter L is chosen as follows: Denote by

$$\mathtt{gap}_1 = \inf_{|H_1|} \left(1 - \lambda(H_1)/3\right) \ , \ \mathtt{gap}_2 = \inf_{|H_2|} \left(1 - \lambda(H_2)/4\right)$$

the minimum spectral-gaps in the explicit expanders that were embedded in our construction (recalling from the introduction that $\mathtt{gap}_i > 0$ for both $i = 1, 2$ by the definition of expanders together with the discrete analogue of Cheeger's inequality), and define

$$L = \left\lceil \frac{2}{\sqrt{\mathtt{gap}_1}} \vee \frac{16}{\mathtt{gap}_2} \vee 32 \right\rceil . \tag{2.1}$$

See Fig. 1 for an illustration of the above construction.

For some insight to the various building blocks of the construction, observe that the height of the SRW on a regular tree behaves as a one-dimensional biased random walk with the bias in the direction of the leaves. As the distribution is concentrated on the levels closest to the leaves, the

SRW mixes once it is reaches the leaves and is approximately uniform on its level. When the walk is started from the first $h+2$ levels, the addition of the expander H_1 ensures that it is uniform on its level by the time it reaches the leaves. Walks started from below level $h+2$ are not necessarily mixed upon reaching the leaves yet subsequently mix after that point due to the expander H_2 in a total time which is strictly less than the time it takes the walk from the root to merely reach the leaves. Our construction thus ensures that the root is the worst case starting position for mixing and that the SRW started at it exhibits cutoff.

We begin by establishing the expansion of the above constructed G. Throughout the proof we omit ceilings and floors in order to simplify the exposition.

Lemma 2.1. *Let $\kappa = (\mathrm{ch}(H_2) \wedge 1)/3$ for H_2 our explicit 4-regular expander. For any integer $L > 0$, the Cheeger constant of the above described 5-regular graph G with parameter L satisfies $\mathrm{ch}(G) \geq \kappa/25L$. Moreover, the induced subgraph $\tilde{G}$ on the last h levels (i.e., levels $2h+2, \ldots, 3h+2$) has $\mathrm{ch}(\tilde{G}) \geq \kappa$.*

Proof. First consider the entire graph G. Since we are only interested in a lower bound on $\mathrm{ch}(G)$, clearly it is valid to omit edges from the graph, in particular we may erase the cross edges between any subtrees $\mathscr{T}_u^*, \mathscr{T}_v^*$ described in Items 2,3 of the construction. This converts every stretched edge of the 5-regular tree of G simply into a 2-path (one where all interior vertices have degree 2) of length L.

Next, we contract all the above mentioned 2-paths into single edges and denote the resulting graph by F. The next simple claim shows that this decreases the Cheeger constant by at most $O(L)$.

Claim 2.2. *Let F be a connected graph with maximal degree Δ and let G be a graph on at most $\frac{3}{2}|F|$ vertices obtained via replacing some of the edges of F by 2-paths of length L. Then $\mathrm{ch}(G) \geq \mathrm{ch}(F)/\Delta^2 L$.*

Proof. Let $X \subset V(G)$ be a set of cardinality at most $|G|/2$ achieving the Cheeger constant of G. We may assume that $\Delta \geq 3$ otherwise F is a disjoint union of paths and cycles and the result holds trivially.

Notice that if X contains two endpoints of a 2-path $\mathscr{P} = (x_0, \ldots, x_L)$ while only containing $k < L-1$ interior vertices of $\mathscr{P}$ then we can assume that $X \cap \mathscr{P} = \{x_0, \ldots, x_k, x_L\}$, i.e. all the interior vertices are adjacent (this maintains the same cardinality of X while not increasing ∂X). With this in mind, modify the set X into the set X' by repeating the following operation: As long as there is a 2-path $\mathscr{P}$ as above (with $x_0, \ldots, x_k$ and x_L in X for some $k < L-1$) we replace x_L by x_{k+1}. This maintains the cardinality of the set while increasing its edge-boundary by at most $\Delta - 2$ (as x_L formerly contributed at least 1 edge to this boundary due to $x_{L-1} \notin X$). Altogether, this yields a set X' where no 2-path $\mathscr{P} = (x_0, \ldots, x_L) \nsubseteq X'$ has both $x_0, x_L \in X'$, while X' satisfies

$$|\partial X'|/|X'| \leq (\Delta - 1)\,\mathrm{ch}(G)\,.$$

The obtained subset X' is possibly disconnected, and we will next argue that its connected components satisfy an appropriate isoperimetric inequality. Consider X'', the connected component of X' that minimizes $|\partial X''|/|X''|$. If X'' is completely contained in the interior of one of the new 2-paths then the statement of the claim immediately holds since

$$\mathrm{ch}(G) \geq \frac{|\partial X'|}{(\Delta-1)|X'|} \geq \frac{|\partial X''|}{(\Delta-1)|X''|} \geq \frac{2}{(L-1)(\Delta-1)} \geq \frac{\mathrm{ch}(F)}{\Delta^2 L}\,,$$

with the last inequality due to the fact $\mathrm{ch}(F) \leq \Delta$. Suppose therefore that this is not the case hence we may now assume that X'' contains at least one endpoint of any 2-path it intersects.

Let $Y = X'' \cap V(F)$, i.e. the subset of the vertices of F obtained from X'' by excluding any vertex that was created in G due to subdivision of edges. Observe that our assumption on X' implies that

$$|\partial Y| = |\partial X''|\,,$$

since either a 2-path $\mathscr{P}$ is completely contained in X'' (not contributing to $\partial X''$) or $\mathscr{P} \cap X = \{x_0, \ldots, x_k\}$ for some $k < L$ (contributing the edge x_k, x_{k+1} to $\partial X''$, corresponding to the edge x_0, x_L in ∂Y).

It remains to consider $|Y|$. Clearly, X'' can be obtained from Y by adding at most Δ new 2-paths with $L-1$ new interior points per vertex, hence

$$|Y| \geq |X''|/\Delta L\,.$$

On the other hand, since $|Y| \leq |X''| \leq |G|/2$ and $|G| \leq \frac{3}{2}|F|$ we have

$$|Y|/|F| \leq \tfrac{3}{2}|X''|/|G| \leq \tfrac{3}{4}\,,$$

which together with the fact that $|X''| \leq |G|/2$ implies that

$$|V(F) \setminus Y| \geq \tfrac{1}{4}|F| \geq \tfrac{1}{6}|G| \geq \tfrac{1}{3}|X''| \geq |X''|/\Delta L\,.$$

Altogether,

$$\mathrm{ch}(F) \leq \frac{|\partial Y|}{|Y| \,\wedge\, |V(F) \setminus Y|} \leq \Delta L \frac{|\partial X''|}{|X''|} \leq \Delta^2 L\ \mathrm{ch}(G)\,. \qquad \blacksquare$$

In light of the above claim we have $\mathrm{ch}(G) \geq \mathrm{ch}(F)/25L$ where the graph F is the result of taking a complete 5-regular tree of height $3h+2$ levels and connecting its $5 \cdot 4^{3h+1}$ leaves, denoted by F', via the 4-regular expander H_2. It therefore remains to show that $\mathrm{ch}(F) \geq \kappa$.

Let S be a set of size $s \leq |F|/2$ vertices that achieves $\mathrm{ch}(F)$. Define its subset of the leaves $S' = S \cap F'$ and set $s' = |S'|$. Since $|F'| \geq \frac{3}{4}|F|$ we clearly have $s' \leq s \leq \frac{2}{3}|F'|$ hence $\big(|S'| \wedge |F' \setminus S'|\big) \geq s'/2$. We thus have the following two options:

1. $s' \geq \frac{2}{3}s$: In this case
$$|\partial_F S| \geq |\partial_{F'} S'| \geq \mathrm{ch}(H_2)s'/2 \geq \mathrm{ch}(H_2)s/3\,.$$
2. $s' < \frac{2}{3}s$: Letting $\mathbb{T}_5$ denote the infinite 5-regular tree (whose Cheeger constant equals 3) we get
$$|\partial_F S| \geq \mathrm{ch}(\mathbb{T}_5)(s - s') - s' = 3s - 4s' > s/3\,.$$

Altogether we deduce that

$$\mathrm{ch}(G) \geq \mathrm{ch}(F) \geq (\mathrm{ch}(H_2) \wedge 1)/3 = \kappa\,.$$

The second part of the lemma (the statement on the subgraph $\tilde{G}$) follows from essentially the same argument given above for $\mathrm{ch}(F)$, as $\tilde{G}$ is precisely a forest of 5-regular trees of height h where all the leaves are connected via the expander H_2. Again we get $\mathrm{ch}(\tilde{G}) \geq \kappa$, completing the proof. $\blacksquare$

Consider the SRW started $k \le h$ levels above the bottom (i.e. at level $3h+2-k$) of the graph G. The height of the walk is then a one-dimensional biased random walk with positive speed $\frac{3}{5}$, implying that it would reach the bottom after $\frac{5}{3}k + o(h)$ steps with high probability.

On the other hand, if the SRW is started closer to the root, i.e. at level $2h+2-k$, then the one-dimensional random walk is delayed by two factors, horizontal (cross-edges) and vertical (stretching the edges into paths). Until reaching level $2h+2$ (after which the previous analysis applies), these delays are encountered along $\frac{5}{3}k + o(h)$ stretched edges with the following effect:

- The former incurs a laziness delay with probability $\frac{3}{5}$ whenever the walk is positioned on an interior vertex of a 2-path.
- The latter delays the walk by the passage time of a SRW through an L-long 2-path, where the walk leaves the origin with probability 1.

It is well-known (and easy to derive) that the expected passage time of the one-dimensional SRW from 0 to $\pm L$ is precisely L^2 and the expected number of visits to the origin by then (including the starting position) is exactly L. It thus follows that the expected delay of the one-dimensional walk representing our height in the tree along a single stretched edge is

$$\tfrac{5}{2}(L^2 - L) + L = \tfrac{1}{2}L(5L-3)\,.$$

Combining the above cases we arrive at the following conclusion:

Claim 2.3. *Consider the* SRW *on the graph G started at a vertex on level $s \in \{0,\dots,3h+1\}$. Set $\alpha = s/h$ and let τ_ℓ be the hitting time of the walk to the leaves (i.e. to level $3h+2$). Then w.h.p.*

$$\tau_\ell = \begin{cases} (\frac{5}{3} + o(1)) \left[L(5L-3)(1-\frac{\alpha}{2}) + 1 \right] h & \text{If } 0 \le \alpha \le 2\,, \\ \frac{5}{3}(3-\alpha)h + o(h) & \text{If } \alpha \ge 2\,. \end{cases}$$

The next lemma relates τ_ℓ, the hitting time to the leaves (addressed by the above claim), and the mixing of the SRW on the graph.

Lemma 2.4. *Let $\varepsilon > 0$, let s_0 be some vertex on level $l_0 \in \{0,\dots,h+2\}$ and $T = (1+\delta)\mathbb{E}_{s_0}\tau_\ell$ for $\delta > 0$ fixed, where τ_ℓ is the hitting time of the* SRW *to the leaves. Then $\|\mathbb{P}_{s_0}(S_T \in \cdot) - \pi\|_{\mathrm{TV}} < \varepsilon$ for any sufficiently large h.*

Proof. Let (S_t) denote the SRW started at some vertex s_0 in level $l_0 \le h+2$ and π be the uniform distribution on $V(G)$. Let $(\tilde{S}_t)$ be a random walk started from the uniform distribution $\tilde{S}_0 \sim \pi$. Write $\mathscr{L}_i$ for $i \in \{0,\dots,3h+2\}$ for the vertices of level i in G (accounting for all the vertices except interior ones along the 2-paths of length L corresponding to stretched edges) and let $\psi : G \to \{0,\dots,3h+2\}$ map vertices in the graph to their level (while mapping interior vertices of 2-paths to the lower of their endpoint levels).

Further let $\Omega = \{2h+3,\dots,3h+2\}$. Clearly, for large enough h we have

$$\mathbb{P}\big(\psi(\tilde{S}_0) \notin \Omega\big) < \frac{\varepsilon}{10}\,.$$

Furthermore, due to the bias of the SRW towards the leaves and the fact that $\tau_\ell = (1+o(1))\mathbb{E}\tau_\ell \asymp h$ (recall Claim 2.3 and that L is fixed) we deduce that $\psi(S_T) > \frac{5}{2}h$ except with probability exponentially small in h, and in particular for any sufficiently large h

$$\mathbb{P}\big(\psi(S_T) \notin \Omega\big) < \frac{\varepsilon}{10}\,.$$

Therefore, an elementary calculation shows that if $\psi(S_T)$ and $\psi(\tilde{S}_T)$ are close in total-variation and so are S_T given $\psi(S_T)=i$ and $\tilde{S}_T$ given $\psi(\tilde{S}_T)=i$ for all $i\in\Omega$, then the required statement on $\|\mathbb{P}(S_T\in\cdot)-\pi\|$ would follow. Namely, if we should show that at time $t=T$ we have

$$\left\|\mathbb{P}\left(S_t\in\cdot\mid\psi(S_t)=i\right)-\mathbb{P}\left(\tilde{S}_t\in\cdot\mid\psi(\tilde{S}_t)=i\right)\right\|_{\mathrm{TV}}<\frac{\varepsilon}{5}\text{ for all }i\in\Omega\,, \tag{2.2}$$

$$\left\|\mathbb{P}\left(\psi(S_t)\in\cdot\right)-\mathbb{P}\left(\psi(\tilde{S}_t)\in\cdot\right)\right\|_{\mathrm{TV}}<\frac{\varepsilon}{5}\,, \tag{2.3}$$

then we would get that $\|\mathbb{P}(S_t\in\cdot)-\pi\|_{\mathrm{TV}}<\varepsilon$ (with room to spare).

Examine the period spent by (S_t) in levels $\{h+2,\ldots,2h+2\}$. The graph in these levels is essentially a product of a 4-ary tree whose edges are stretched into L-long 2-paths and the expander H_1. Let $\varphi: G\to H_1$ map the vertices in these levels to their corresponding vertices in H_1, and let τ_0,τ_1 be the hitting times of (S_t) to levels $\frac{3}{2}h$ and $2h+2$ respectively.

As argued above, the SRW started at level $2h+2-k$ is a one-dimensional biased random walk that w.h.p. passes through $(1+o(1))\frac{5}{3}k+o(h)$ stretched edges until reaching level $2h+2$ for the first time. In particular, between times τ_0,τ_1 the walk w.h.p. passes through $(1+o(1))\frac{5}{6}h$ stretched edges.

Along each stretched edge among levels $\{h+2,\ldots,2h+2\}$, the walk traverses a cross-edge in H_1 (that is, $\varphi(S_{t+1})$ is uniformly distributed over the neighbors of $\varphi(S_t)$ in H_1) with probability $\frac{3}{5}$ whenever it is in an interior vertex in the 2-path, for a total expected number of $\frac{3}{2}(L^2-L)$ such moves.

Finally, due to its bias towards the leaves, with high probability the SRW from level $\frac{3}{2}h$ reaches level $2h+2$ (the vertices B) before hitting level $h+2$. Applying CLT we conclude that the SRW w.h.p. traverses

$$(\tfrac{5}{4}+o(1))L(L-1)h>L^2h$$

cross-edges (each corresponding to a single step of the SRW on H_1) between times τ_0,τ_1, where the last inequality holds for $L>5$ and large enough h. This amounts to at least L^2h consecutive steps of a SRW along H_1.

Aiming for a bound on the total-variation mixing, we may clearly condition on events that occur with high probability: Condition therefore throughout the proof that indeed the above statement holds.

In particular, letting (X_t) be the SRW on the expander H_1 and recalling that $|A|=|H_1|=20\cdot 2^{2h}$ and $r=L^2h$ it follows that

$$\begin{aligned}\max_{s_0}\left\|\mathbb{P}_{s_0}\left(\varphi(S_{\tau_1})\in\cdot\right)-|A|^{-1}\right\|_{\mathrm{TV}}&\le\max_{x_0}\left\|\mathbb{P}_{x_0}\left(X_r\in\cdot\right)-|A|^{-1}\right\|_{\mathrm{TV}}\\&\le\frac{1}{2}\max_{x_0}\left\|\mathbb{P}_{x_0}\left(X_r\in\cdot\right)-|A|^{-1}\right\|_2\,.\end{aligned}$$

Recalling that H_1 is a 3-regular with second largest (in absolute value) eigenvalue $\lambda(H_1)$ and writing $\gamma\triangleq 1-\frac{\lambda(H_1)}{3}$,

$$\begin{aligned}\max_{x_0}\left\|\mathbb{P}_{x_0}\left(X_r\in\cdot\right)-|A|^{-1}\right\|_2&\le\sqrt{|A|}\exp\left(-\gamma r\right)\\&<5\exp\left[-(\gamma L^2-\log 2)h\right]<|A|^{-2}\,,\end{aligned}$$

with the last inequality justified for any sufficiently large h provided that

$$L > \frac{\sqrt{5\log 2}}{\sqrt{1-\lambda(H_1)/3}}, \tag{2.4}$$

a fact inferred from the choice of L in (2.1). In this case, for any s_0

$$\mathbb{P}_{s_0}\left(S_{\tau_1} = u\right) = \frac{1+o(1)}{|A|} \text{ for every } u \in A.$$

By symmetry we now conclude that for any $i \in \Omega$ and $t \geq \tau_1$ we have

$$\mathbb{P}_{s_0}\left(S_t \in \cdot \mid \psi(S_t) = i\right) = \frac{1+o(1)}{|\mathscr{L}_i|} \text{ for every } i \in \{2h+2, \ldots, 3h+2\}\ .$$

This immediately establishes (2.2).

To obtain (2.3), note that $\tilde{S}_t$ for $t = \tau_1$ w.h.p. satisfies $\psi(\tilde{S}_0) > \frac{5}{2}h$. Conditioned on this event we can apply a monotone-coupling to successfully couple (S_t) and $(\tilde{S}_t)$ such that $\psi(S_{\tau_\ell}) = \psi(\tilde{S}_{\tau_\ell})$, yielding (2.3).

The concentration of τ_ℓ established in Claim 2.3 carries the above two bounds to time T, thus completing the proof. ■

Let $t_{\text{MIX}}(\varepsilon; x)$ denote the total-variation mixing time from a given starting position x. That is, if (X_t) is an ergodic Markov chain on a finite state space Ω with stationary distribution π then

$$t_{\text{MIX}}(\varepsilon; x) \stackrel{\triangle}{=} \min\left\{t : \|\mathbb{P}_x(X_t \in \cdot) - \pi\|_{\text{TV}} < \varepsilon\right\}.$$

The above lemma gives an upper bound on this quantity for the SRW started at one of the levels $\{0, \ldots, h+2\}$, which we now claim is asymptotically tight:

Corollary 2.5. *Consider the* SRW *on G started at some vertex s_0 on level $l_0 \in \{0, \ldots, h+2\}$ and let τ_ℓ be the hitting time of the walk to the leaves. Then for any fixed $0 < \varepsilon < 1$ we have $t_{\text{MIX}}(\varepsilon; s_0) = (1+o(1))\mathbb{E}_{s_0}\tau_\ell$.*

Proof. The upper bound on $t_{\text{MIX}}(\varepsilon; s_0)$ was established in Lemma 2.4.

For a matching lower bound on $t_{\text{MIX}}(1-\varepsilon; s_0)$ choose some fixed integer $K = O(\log(L/\varepsilon))$ such that the bottom K levels of the graph comprise at least a $(1-\varepsilon)$-fraction of the vertices of G, i.e.

$$\sum_{i > 3h+2-K} \left|\mathscr{L}_i\right| > (1-\varepsilon)|G|.$$

The lower bound now follows from observing that, by the same arguments that established Claim 2.3, the hitting time from level l_0 to level $3h+2-K$ is w.h.p. $(1-o(1))\mathbb{E}\tau_\ell$ for any sufficiently large h. ■

Having established the asymptotic mixing time of the SRW started at the top h levels, we next wish to show that from all other vertices the mixing time is faster.

Claim 2.6. *Let (S_t) be the* SRW *started at some vertex x in level $\psi(x) > h$. For every $0 < \varepsilon < 1$ and any sufficiently large h we have $t_{\text{MIX}}(\varepsilon; x) < 6L^2 h$.*

Proof. Let G' denote the induced subgraph on the bottom $h+1$ levels of the graph (i.e., levels $2h+2, \ldots, 3h+2$). Since $|G'| = (1-o(1))|G|$ it clearly suffices to show that (S_t) mixes within total-variation distance ε on G'.

By Claim 2.3 (taking the worst case $\alpha = 1$ corresponding to $\psi(x) = h$) with high probability we have

$$\tau_\ell \leq (\tfrac{5}{6} + o(1))[L(5L-3)+2]h < 5L^2 h \triangleq t_1\,,$$

where the above strict inequality holds for any sufficiently large h (as $L \geq 1$).

Recall that $\mathrm{ch}(G') \geq \kappa = (\mathrm{ch}(H_2) \wedge 1)/3$ by Lemma 2.1, where H_2 is the explicit 4-regular expander with second largest (in absolute value) eigenvalue $\lambda(H_2)$. Further consider the graph G'' obtained by adding to G' a perfect matching on level $2h+2$, thus making it 5-regular. Clearly, adding edges can only increase the Cheeger constant and so $\mathrm{ch}(G'') \geq \kappa$ as well. Moreover, the discrete form of Cheeger's inequality ([5, 6, 12, 20]), which for a d-regular expander H with second largest (in absolute value) eigenvalue λ states that

$$\frac{d-\lambda}{2} \leq \mathrm{ch}(H) \leq \sqrt{2d(d-\lambda)}\,,$$

here gives the following:

$$\frac{4-\lambda(H_2)}{6} \wedge \frac{1}{3} \leq \kappa \leq \mathrm{ch}(G'') \leq \sqrt{10(5-\lambda(G''))}\,.$$

In particular we obtain that $\gamma \triangleq 1 - \frac{\lambda(G'')}{5}$ satisfies

$$\gamma > \left(\frac{4-(\lambda(H_2) \vee 2)}{6\sqrt{50}}\right)^2 \tag{2.5}$$

while a simple random walk (X_t) on G'' is well-known to satisfy

$$\max_{x_0} \left\| \mathbb{P}_{x_0}\left(X_t \in \cdot\right) - |G''|^{-1} \right\|_2 \leq \sqrt{|G''|} \exp\left(-\gamma t\right)\,.$$

As $|G''| \asymp 2^{6h}$ we infer that after

$$\frac{3\log 2}{\gamma} h + \frac{O(\log(1/\varepsilon))}{\gamma} < \frac{9}{4\gamma} h \triangleq t_2$$

(the strict inequality holding for large enough h) steps we have

$$\max_{x_0} \left\| \mathbb{P}_{x_0}\left(X_{t_2} \in \cdot\right) - |G''|^{-1} \right\|_{\text{TV}} \leq \varepsilon/2\,.$$

Recall that our choice of t_1 is such that $\tau_\ell < t_1$ w.h.p., i.e. the SRW reaches level $3h+2$ by that time, and thereafter (due to its bias towards the leaves) it does not revisit level $2h+2$ until time $t_1 + t_2$ except with a probability that is exponentially small in h. Since G and G'' are identical on levels $2h+3, \ldots, 3h+2$ we deduce that w.h.p. the SRW performs at least t_2 consecutive steps

in G'' following τ_ℓ. Altogether, for large enough h we have $\left\|\mathbb{P}_x\left(S_{t_1+t_2}\in\cdot\right)-\pi\right\|_{\text{TV}}<\varepsilon$ and so $t_{\text{MIX}}(\varepsilon;x)<t_1+t_2$.

Finally, bearing (2.5) and the choice of L in (2.1) in mind,

$$L\geq\frac{64}{4-(\lambda(H_2)\vee 2)}>\frac{\sqrt{\frac{9}{4}}\cdot 6\sqrt{50}}{4-(\lambda(H_2)\vee 2)}, \tag{2.6}$$

hence $L^2>9/4\gamma$ and so $t_2\leq L^2h$ and $t_1+t_2\leq 6L^2h$, as required. ■

We now claim that the worst-case mixing time within any $0<\varepsilon<1$ is attained by an initial vertex at distance $o(h)$ from the root. Fix $0<\varepsilon<1$, let x be the initial vertex maximizing $t_{\text{MIX}}(\varepsilon;x)$ and recall that $\psi(x)$ denotes its level in the graph. The combination of Claim 2.3 and Corollary 2.5 ensures that if $\psi(x)\leq h$ then necessarily $\psi(x)=o(h)$, in which case

$$t_{\text{MIX}}(\varepsilon;x)=\left(\tfrac{5}{3}+o(1)\right)(5L^2-3L+1)h\stackrel{\triangle}{=}t^\star. \tag{2.7}$$

An immediate consequence of the requirement 2.6 on L is that $L\geq 50$, hence $t^\star>8L^2h$ for any sufficiently large h. Therefore, we cannot have $\psi(x)>h$ since by Claim 2.6 that would imply that $t_{\text{MIX}}(\varepsilon;x)<6L^2h$ contradicting the fact that x achieves the worst-case mixing time.

Overall we deduce that for any $0<\varepsilon<1$ we have

$$t_{\text{MIX}}(\varepsilon)=\max_x t_{\text{MIX}}(\varepsilon;x)=(1+o(1))t^\star,$$

thus confirming that the SRW on the above constructed family of 5-regular expanders exhibits total-variation cutoff from a worst starting location.

It remains to describe how our construction can be (relatively easily) modified to be 3-regular rather than 5-regular.

The immediate step is to use binary trees instead of 4-ary trees, after which we are left with the problem of embedding the explicit expanders H_1 and H_2 without increasing the degree. This will be achieved via the line-graphs of these expanders, hence our explicit expanders will now have slightly different parameters:

- H_1 : An explicit 3-regular expander on 2^{h+1} vertices.
- H_2 : An explicit 3-regular expander on 2^{3h+1} vertices.

Recall that given a tree rooted at some vertex u, denoted by $\mathscr{T}_u$, its edge-stretched version is obtained by replacing each edge by a 2-path of length L, and the collection of all new interior vertices (due to subdivision of edges) is denoted by $\mathscr{T}_u^*$. The modified construction is as follows:

1. Levels 0,1,2: First levels of a binary tree.
   - Denote by $U=\{u_1,\ldots,u_6\}$ the vertices in level 2.
2. Levels $3,\ldots,h+2$: Stretched binary trees rooted at U:

- Connect vertices from $\mathcal{T}^*_{u_i}$ (interior vertices along 2-paths) to the corresponding (isomorphic) vertices in $\mathcal{T}^*_{u_{2i}}$, i.e. inter-connect the interior vertices via perfect matchings.
- Denote by A the $6 \cdot 2^h$ vertices in level $h+2$.

3. Levels $h+3, \ldots, 2h+2$: Edge-stretched binary trees rooted at A and inter-connected via the line-graph of H_1 using auxiliary vertices:
   - Associate each binary tree $\mathcal{T}_{a_i}$ rooted at A to an *edge* of H_1.
   - For each $x \in \mathcal{T}^*_{a_i}$ (interior vertex on a 2-path) we connect it to a new auxiliary vertex x' and associate x' with a unique edge of H_1.
   - We say that $x \in \mathcal{T}^*_{a_i}$ and $y \in \mathcal{T}^*_{a_j}$ are isomorphic if the isomorphism from $\mathcal{T}_{a_i}$ to $\mathcal{T}_{a_j}$ maps x to y. Add $|H_1|$ new auxiliary vertices per equivalence class of $|A|$ such isomorphic vertices, identify them with the vertices of H_1 and connect every new vertex v to the auxiliary vertices x', y', z' representing the edges incident to it.
4. Levels $2h+3, \ldots, 3h+2$: A forest of binary trees.
5. Last level: leaves are inter-connected via the line-graph of H_2:
   - Associate the $6 \cdot 2^{3h}$ vertices with the edges of H_2.
   - Add $|H_2|$ new auxiliary vertices, each connected to the leaves corresponding to edges that are incident to it in H_2.

It is easy to verify that the walk along the cross-edges of the $\mathcal{T}^*_{a_i}$'s now corresponds to a lazy (unbiased) random walk on the edges of H_1. Similarly, the walk along the cross-edges connecting the leaves corresponds to the SRW on the edges of H_2. Hence, all of the original arguments remain valid in this modified setting for an appropriately chosen fixed L. This completes the proof of Theorem 1. ■

2.2 Explicit expanders without total-variation cutoff

The explicit cubic expanders with cutoff constructed in the previous section (illustrated in Fig. 1) can be easily modified so that the SRW on them from a worst starting position would *not* exhibit total-variation cutoff.

To do so, recall that in the above-described family of graphs, each vertex of the subset U was the root of a regular tree of height h whose edges were stretched into L-long 2-paths (see Item 2 of the construction). We now tweak this construction by stretching some of the edges into 2-paths of length L'. Namely, for subtrees rooted at the *odd* vertices in level $h/2$ of these trees we stretch the edges into paths of length $L' > L$. Under this modified stretching the trees $\mathcal{T}_{u_i}$ are clearly still isomorphic, hence the cross edges are inter-connecting 2-paths of the same lengths.

By the arguments above, starting from any level $\ell > h/2$ the mixing is faster compared to the root, and if L'/L is sufficiently small then the root remains the asymptotically worst starting position. However, starting from the root (and in fact, starting from any level $\ell \leq h/2$) the hitting time to the set A is no longer concentrated due to the odd/even choice of subtree at level $h/2$. Therefore, from the worst starting position we have that the hitting time to the leaves is concentrated on two

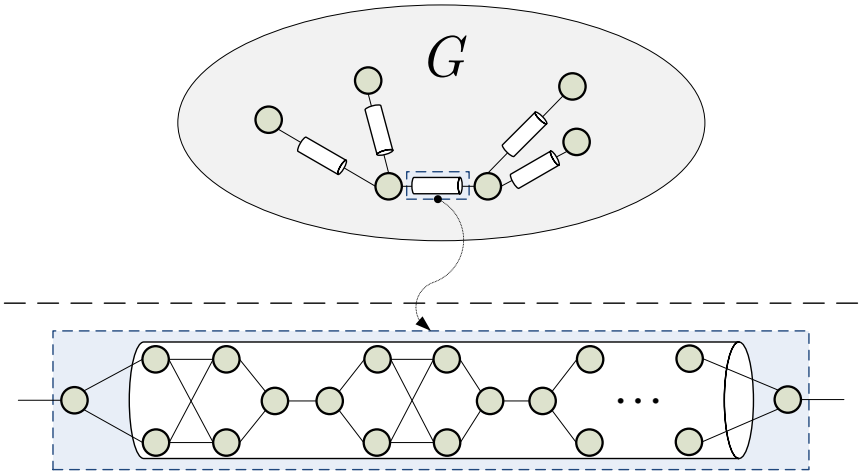

Figure 2: Extension of the explicit expanders with cutoff to graphs with prescribed order of cutoff location.

distinct values (differing by a fixed multiplicative constant), each with probability $\frac{1}{2} - o(1)$. This implies that the ratio $t_{\text{MIX}}(\frac{1}{4})/t_{\text{MIX}}(\frac{3}{4})$ is bounded away from 1 and in particular this explicit family of expanders does not have total-variation cutoff.

2.3 Proof of Theorem 2: cutoff at any prescribed location order

Suppose H is an explicit 3-regular expander on m vertices provided by Theorem 1, and recall that the SRW on this graph exhibits cutoff at $C \log m$ where $C > 0$ is some absolute constant. Our graph G will be the result of replacing every edge of H by the 3-regular analogue of a 2-path, which we refer to as a "cylinder", illustrated in Fig. 2. The length of each cylinder is set to be $L = L(m)$ satisfying $L \equiv 1 \pmod 4$. Notice that the total number of vertices in G is

$$n = |V(H)| + |E(H)|\tfrac{3}{2}(L-1) = \left(1 + \tfrac{9}{4}(L-1)\right) m\,. \tag{2.8}$$

Since $m \to \infty$ and the SRW on H, started at a worst-case starting position, traverses $(C+o(1))\log m$ edges until mixing, we infer from CLT, as well as the fact that the expected passage-time through an L-long cylinder is L^2, that the analogous random walk on G has cutoff at

$$t_{\text{MIX}} = CL^2 \log m = (C + o(1))L^2 \log(n/L)\,. \tag{2.9}$$

When $L(m) = O(1)$ we have $t_{\text{MIX}} \asymp \log n$. On the other extreme end, when L grows arbitrarily fast as a function of m we obtain that it approaches n arbitrarily closely but we must still having $L = o(n)$ since $n/L \asymp m \to \infty$. In that case t_{MIX} approaches n^2 arbitrarily closely while having a strictly smaller order.

To complete the construction it remains to observe that we may choose L and m so that $|G_n| \asymp n$ and $t_{\text{MIX}} \asymp t_n$. This can be achieved by first selecting L so that $t_n \asymp (C + o(1))L^2 \log(n/L)$ and then selecting a graph constructed through Theorem 1 on m vertices for some $m \asymp n/L$ (note that in the theorem we construct graphs of size essentially $(c + o(1))2^{3h}$ so this is always possible).

To show that there is no cutoff whenever $t_{\text{MIX}} \asymp n^2$ we will argue that in that case we have $\mathtt{gap} = O(n^{-2})$, in contrast to the necessary condition for cutoff $\mathtt{gap}^{-1} = o(t_{\text{MIX}})$ due to Peres (cf. [16]), discussed in the introduction.

Lemma 2.7. *Let G be a graph on n vertices with degrees bounded by some Δ fixed on which the* SRW *has $t_{\text{MIX}} \asymp n^2$. Then the spectral-gap of the walk satisfies $\mathtt{gap} \asymp n^{-2}$. In particular, the* SRW *on G does not exhibit cutoff.*

Proof. Observe that the above graph must satisfy $\text{diam}(G) \geq cn$ for some fixed $c > 0$ as it is well-known (cf., e.g., [4]) that the lazy walk on any graph H has $t_{\text{MIX}} = O(\text{diam}(H)\text{vol}(H))$ and in our case $\text{vol}(G) \leq \Delta n = O(n)$.

Let $x, y \in V(G)$ be two vertices whose distance in G is

$$N \stackrel{\triangle}{=} \text{dist}_G(x,y) = \text{diam}(G) \geq cn\,.$$

Our lower bound on the gap will be derived from its representation via the Dirichlet form, according to which

$$\mathtt{gap} = \inf_f \frac{\mathscr{E}(f)}{\text{Var}(f)} = \inf_f \frac{\frac{1}{2}\sum_{x,y\in\Omega}\left[f(x)-f(y)\right]^2 \pi(x)P(x,y)}{\text{Var}_\pi f}\,, \tag{2.10}$$

where π is the (uniform) stationary measure and P is the transition kernel of the SRW. As a test-function $f : V(G) \to \mathbb{R}$ in the above form choose

$$f(v) \stackrel{\triangle}{=} \text{dist}_G(x,v)\,.$$

Clearly we have $\mathscr{E}(f) \leq 1$ and a lower bound of order n^2 on the variance follows from the fact that two sets of linear size each have a linear discrepancy according to f. Namely,

$$\pi\left(f^{-1}(\{0,\ldots,\lfloor N/4\rfloor\})\right) \geq \frac{c}{4}\,,\ \pi\left(f^{-1}(\{\lceil 3N/4\rceil,\ldots,N\})\right) \geq \frac{c}{4}\,,$$

as a result of which

$$\text{Var}(f) \geq (c/4)(N/4)^2 > c'n^2\,.$$

We conclude that $\mathtt{gap} = O(n^{-2})$, thus completing the proofs of Lemma 2.7 and Theorem 2. ∎

3 Concluding remarks and open problems

- Recent results in [15] showed that almost every regular expander graph has total-variation cutoff (prior to that there were no known examples for bounded-degree graphs with this phenomenon); here we provided a *first explicit* construction for bounded-degree expanders with cutoff.
- The expanders constructed in this work are non-transitive. Moreover, our proof exploits their highly asymmetric structure in order to control the mixing time of the random walk from various starting locations. It would be interesting to obtain an explicit construction of transitive expanders with total-variation cutoff.

- A slight variant of our construction gives an example of a family of expanders where the SRW does *not* exhibit cutoff, thereby disagreeing with Peres' cutoff-criterion. Both here and in another such example due to Peres and Wilson [17] the expanders are non-transitive (hence the restriction to transitive graphs in Peres' conjecture stated next).
- While it is conjectured by Peres that the random walk on any family of transitive bounded-degree expanders exhibits total-variation cutoff, there is not even a single example of such a transitive family where cutoff was proved (or disproved).
- For general (not necessarily expanding) bounded-degree graphs on n vertices it is well-known that $t_{\text{MIX}} = O(n^2)$. Here we showed that cutoff can occur essentially anywhere up to $o(n^2)$ by constructing cubic graphs with cutoff at any such prescribed location. Furthermore, this is tight as we prove that if $t_{\text{MIX}} \asymp n^2$ then cutoff cannot occur.

Acknowledgment

We would like to thank Yuval Peres for suggesting the problem, his encouragement and useful discussions.

References

[1] M. Ajtai, *Recursive construction for 3-regular expanders*, Combinatorica **14** (1994), no. 4, 379–416.

[2] D. Aldous, *Random walks on finite groups and rapidly mixing Markov chains*, Seminar on probability, XVII, 1983, pp. 243–297.

[3] D. Aldous and P. Diaconis, *Shuffling cards and stopping times*, Amer. Math. Monthly **93** (1986), 333–348.

[4] D. Aldous and J. A. Fill, *Reversible Markov Chains and Random Walks on Graphs*. In preparation, `http://www.stat.berkeley.edu/~aldous/RWG/book.html`.

[5] N. Alon, *Eigenvalues and expanders*, Combinatorica **6** (1986), no. 2, 83–96.MR875835

[6] N. Alon and V. D. Milman, λ_1, *isoperimetric inequalities for graphs, and superconcentrators*, J. Combin. Theory Ser. B **38** (1985), no. 1, 73–88.MR782626

[7] N. Berestycki, *Phase transitions for the distance of random walks and applications to genome rearrangement*, Ph.D. dissertation, Cornell University (2005).

[8] G.-Y. Chen and L. Saloff-Coste, *The cutoff phenomenon for ergodic Markov processes*, Electronic Journal of Probability **13** (2008), 26–78.

[9] P. Diaconis, *The cutoff phenomenon in finite Markov chains*, Proc. Nat. Acad. Sci. U.S.A. **93** (1996), no. 4, 1659–1664.

[10] P. Diaconis and M. Shahshahani, *Generating a random permutation with random transpositions*, Z. Wahrsch. Verw. Gebiete **57** (1981), no. 2, 159–179.

[11] J. Ding, E. Lubetzky, and Y. Peres, *Total-variation cutoff in birth-and-death chains*, Probab. Theory Related Fields **146** (2010), no. 1, 61–85.

[12] J. Dodziuk, *Difference equations, isoperimetric inequality and transience of certain random walks*, Trans. Amer. Math. Soc. **284** (1984), no. 2, 787–794.MR743744

[13] R. Durrett, *Random graph dynamics*, Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, Cambridge, 2007.

[14] S. Hoory, N. Linial, and A. Wigderson, *Expander graphs and their applications*, Bull. Amer. Math. Soc. **43** (2006), no. 4, 439–561.MR2279667

[15] E. Lubetzky and A. Sly, *Cutoff phenomena for random walks on random regular graphs*, Duke Math. J. **153** (2010), no. 3, 475–510.

[16] Y. Peres, American Institute of Mathematics (AIM) research workshop “Sharp Thresholds for Mixing Times” (Palo Alto, December 2004).

[17] Y. Peres and D. B. Wilson. Private communication.

[18] O. Reingold, S. Vadhan, and A. Wigderson, *Entropy waves, the zig-zag graph product, and new constant-degree expanders*, Ann. of Math. (2) **155** (2002), no. 1, 157–187.

[19] L. Saloff-Coste, *Random walks on finite groups*, Probability on discrete structures, 2004, pp. 263–346.

[20] A. Sinclair and M. Jerrum, *Approximate counting, uniform generation and rapidly mixing Markov chains*, Inform. and Comput. **82** (1989), no. 1, 93–133.MR1003059