Comment. Math. Helv. 97 (2022), 431–456
DOI 10.4171/CMH/537

Commentarii Mathematici Helvetici

Cutoff on Ramanujan complexes and classical groups

Michael Chapman and Ori Parzanchevski

Abstract. The total-variation cutoff phenomenon has been conjectured to hold for simple random walk on all transitive expanders. However, very little is actually known regarding this conjecture, and cutoff on sparse graphs in general. In this paper we establish total-variation cutoff for simple random walk on Ramanujan complexes of type $\widetilde{A}_d$ $(d \geq 1)$. As a result, we obtain explicit generators for the finite classical groups $\mathrm{PGL}_n(\mathbb{F}_q)$ for which the associated Cayley graphs exhibit total-variation cutoff.

Mathematics Subject Classification (2020). 05C48; 05C81, 20E42, 11E95, 37A25.

Keywords. Cutoff, random walks, expanders, Bruhat–Tits buildings, Ramanujan complexes.

1. Introduction

The ε-mixing time of a finite Markov chain is the earliest time at which the distribution on states becomes ε-close to the stationery one, regardless of the starting distribution. There are several distance functions which one may use, and we focus on L^1, or total-variation (see (1.1)). The parameter ε can be thought of as a measure of "standards": for example, in a professional poker tournament one expects the dealer to shuffle the decks longer than in an amateur one. Loosely speaking, a sequence of Markov chains is said to exhibit the *cutoff phenomenon* if it is insensitive to ones' standards. Namely, whether one seeks to be at most 0.01 away from the stationary distribution or at most 0.99 away from it, it will roughly take the same amount of time. In other words, for a long period of time the distribution is at almost maximal distance from stationery, and then over a short period of time it mixes almost completely. This counter-intuitive phenomenon was first demonstrated by Diaconis–Shahshahani and Aldous [1, 6], and was subsequently shown to hold in many naturally occurring Markov chains (see the surveys [5, 23]). Common to all of these examples is that the number of legal moves grows together with the number of states.

The case of a bounded number of legal moves – for example, simple random walk (SRW) on a family of graphs with bounded degrees – turned out to be more resistant, and much less is known about it. In 2004, Peres has conjectured that SRW on every family of transitive bounded degree expanders exhibits the cutoff phenomenon [4], even though at the time *no* family of bounded degree expanders was known to do so. In [16], Lubetzky and Sly used probabilistic methods to show that random regular graphs exhibit cutoff asymptotically almost surely. The next big breakthrough was achieved by Lubetzky and Peres [15], who showed that SRW on all Ramanujan graphs (which are optimal expanders) exhibit cutoff. A main ingredient of [15] is to show first that non-backtracking random walk (NBRW) on Ramanujan graphs exhibits cutoff at an optimal time. The last assertion was generalized in [14] to the context of Ramanujan complexes, which are high-dimensional analogues of Ramanujan graphs, defined in [13, 20]. In the paper [14], Lubetzky, Lubotzky and the second author establish optimal-time cutoff for a large family of *asymmetric* random walks on the cells of these complexes (in the graph case, NBRW is an asymmetric walk on edges). However, the techniques of [14] cannot be applied neither to any symmetric random walk, nor to any walk on vertices.

The goal of the current paper is to establish cutoff for SRW on Ramanujan complexes arising from the group PGL_d over a non-archimedean local field. Interestingly, while SRW on vertices only "sees" the 1-skeleton of the complex, our proof makes use of asymmetric random walks on cells of *all* dimensions of the complex, showing that the high-dimensional geometry can play an important role even when studying walks on graphs. A main motivation to study these complexes is the study of expansion in simple groups: the finite groups $\mathrm{PGL}_2(\mathbb{F}_q)$ admit a Cayley structure of a Ramanujan graph due to Lubotzky, Phillips and Sarnak [18], whereas the groups $\mathrm{PGL}_d(\mathbb{F}_q)$ for general d can be endowed with a Cayley structure which is the 1-skeleton of a Ramanujan complex. Thus, we establish here cutoff for SRW on the groups $\mathrm{PGL}_d(\mathbb{F}_q)$, with respect to the appropriate generators. We remark that the situation in the case $d \geq 3$ is even more striking than in the graph case ($d = 2$): by Kazhdan's property (T), for $d \geq 3$ every generating set of $\mathrm{PGL}_d(\mathbb{Z})$ gives rise to an expander family of Cayley graphs of $\mathrm{PGL}_d(\mathbb{F}_q)$ [21], but – except for the case which we handle in this paper – it is unknown whether these families exhibit cutoff or not.

We now move on to rigorous terms. Let $\mathcal{D}$ be a connected directed graph (*digraph*), which we assume for simplicity to be k-out and k-in regular. Consider random walk on $\mathcal{D}$ starting at a vertex v_0 with uniform transition probabilities, and denote by $\mu^t_{\mathcal{D},v_0}$ its distribution at time t. The ε-mixing time of $\mathcal{D}$ is

$$t_{\mathrm{mix}}(\varepsilon) = t_{\mathrm{mix}}(\varepsilon, \mathcal{D}) = \min\bigl\{t \in \mathbb{N} \mid \forall v_0 \in \mathcal{D}^0,\ \|\mu^t_{\mathcal{D},v_0} - \pi_{\mathcal{D}}\|_{TV} < \varepsilon\bigr\},$$

where $\pi_{\mathcal{D}}$ is the uniform distribution on $\mathcal{D}^0$ (the vertices of $\mathcal{D}$), and $\|\cdot\|_{TV}$ is the total-variation norm

$$\|\mu - \nu\|_{TV} = \max_{A \subseteq \mathcal{D}^0} |\mu(A) - \nu(A)| = \frac{1}{2}\|\mu - \nu\|_1. \tag{1.1}$$

A family of digraphs $\{\mathcal{D}_n\}$ is said to exhibit cutoff if

$$\frac{t_{\text{mix}}(\varepsilon, \mathcal{D}_n)}{t_{\text{mix}}(1-\varepsilon, \mathcal{D}_n)} \xrightarrow{n\to\infty} 1$$

for every $0 < \varepsilon < 1$. The cutoff is said to occur at time $t(n)$, if for every $\varepsilon > 0$ there exists a window of size $w(n,\varepsilon) = o(t(n))$, such that $|t_{\text{mix}}(\varepsilon, \mathcal{D}_n) - t(n)| \le w(n,\varepsilon)$ for n large enough. If $t(n) = \log_k |\mathcal{D}_n|$ we say that the cutoff is *optimal*, since a k-regular walk cannot mix in less steps.

Recall that a connected k-regular graph is called a *Ramanujan graph* if its adjacency spectrum is contained in $[-2\sqrt{k-1}, 2\sqrt{k-1}] \cup \{k\}$.

Theorem ([15]). *The family $\{G_n\}$ of all k-regular Ramanujan graphs exhibits*

(1) *cutoff for SRW at time $\frac{k}{k-2}\log_{k-1}|G_n|$, with a window of size $O(\sqrt{\log|G_n|})$;*

(2) *optimal cutoff for NBRW (at time $\log_{k-1}|G_n|$), with a window of size $O(\log\log|G_n|)$.*

In [15], SRW-cutoff is first reduced to optimal NBRW-cutoff, by studying the distance of SRW on the tree from its starting point. To obtain optimal cutoff for NBRW new spectral techniques are developed for analyzing non-normal operators.

To see how the notion of Ramanujan graphs generalizes to higher dimension, recall that $[-2\sqrt{k-1}, 2\sqrt{k-1}]$ is the L^2-adjacency spectrum of the k-regular tree, which is the universal cover of every k-regular graph [12]. In accordance, *Ramanujan complexes* are roughly defined as finite complexes which spectrally mimic their universal cover; for a precise definition, see Section 2. In [14], a vast generalization of part (2) of the theorem above is proved: say that a walk rule is *collision-free* if two walkers which depart from each other can never cross paths again.

Theorem ([14]). *Let $\mathcal{B}$ be an affine Bruhat–Tits building (see Section 2.1), and fix a collision-free walk rule on cells of $\mathcal{B}$. The family of Ramanujan complexes with universal cover $\mathcal{B}$ exhibit optimal cutoff with respect to the corresponding walk rule on them.*

This recovers NBRW on Ramanujan graphs, since NBRW is indeed collision-free on the edges of the tree. In higher dimension, NBRW is not collision-free anymore, but in [14, Section 5.1] it is shown that collision-free walks do exist, on cells of every dimension, except for vertices. As SRW is not collision-free, the techniques of [14] cannot address it (in fact, they cannot address any operator on vertices; see [22, Remark 3.5 (b)]).

The goal of this paper is to establish cutoff for SRW on Ramanujan complexes. We fix a non-archimedean local field F with residue field of size q, and denote by $\mathcal{B} = \mathcal{B}_{d,F}$ the Bruhat–Tits building associated with $\mathrm{PGL}_d(F)$.

Theorem 1.1 (Main theorem). *The family $\{X_n\}$ of all Ramanujan complexes with universal cover $\mathcal{B}$ exhibits total-variation cutoff for SRW at time $C_{d,q}\log_q|X_n|$, with* *a window of size* $O(\sqrt{\log|X_n|})$*. The constant* $C_{d,q}$ *is determined in* (4.2) *and for each* d*, is a rational function in* q *of magnitude*

$$\frac{1}{\lfloor d/2\rfloor\lceil d/2\rceil}+O\Big(\frac{1}{q}\Big)$$

(see Table 1).

We emphasize that the graphs underlying these walks are not Ramanujan graphs when $d > 2$. For example, in the two-dimensional case (where $d = 3$), the 1-skeleton of X is a $2(q^2+q+1)$-regular graph. If it was a Ramanujan graph, its second largest adjacency eigenvalue would be bounded by

$$2\sqrt{2q^2+2q+1}\approx 2.8q,$$

but in truth this eigenvalue equals $6q - o_n(1)$ (cf. [13, 24]), reflecting the abundance of triangles in X. In this case, the cutoff is achieved at time

$$\frac{q^2+q+1}{2(q^2-1)}\log_q n$$

(see Theorem 3.1).

A motivation for our result, which does not require the notions of Ramanujan complexes or buildings, is the study of expansion in finite simple groups; see [3] for a recent survey. A celebrated result of Lubotzky, Phillips and Sarnak [18] uses the building of $\mathrm{PGL}_2(\mathbb{Q}_p)$ to show that the groups $\mathrm{PSL}_2(\mathbb{F}_q)$ have explicit generators for which the resulting Cayley graphs are Ramanujan, so that [15] yields total-variation cutoff for SRW on these groups. Turning to $\mathrm{PSL}_d(\mathbb{F}_q)$, the work of [15] does not apply anymore, since it is not known whether $\mathrm{PSL}_d(\mathbb{F}_q)$ have generators which yield Ramanujan Cayley graphs. Nevertheless, by considering the building of $\mathrm{PGL}_d(\mathbb{F}_q(\!(t)\!))$, it was shown by Lubotzky, Samuels and Vishne ([19], see also [25]) that the groups $\mathrm{PSL}_d(\mathbb{F}_q)$ have explicit generators, for which the resulting Cayley graph is precisely the 1-skeleton of a Ramanujan complex of type $\widetilde{A}_d$. For $d = 3$, such generators can also be given using the building of $\mathrm{PGL}_3(\mathbb{Q}_p)$ [2, 7]. We thus achieve:

Corollary 1.2. (1) *Fix* $d \geq 3$ *and a prime power* q*. The family* $\{\mathrm{PSL}_d(\mathbb{F}_{q^\ell})\}_{\ell\to\infty}$ *has an explicit symmetric set of* $k = \sum_{j=1}^{d-1}\left[\begin{smallmatrix} d\\ j\end{smallmatrix}\right]_q$ (Gaussian binomial coefficients) *generators exhibiting TV-cutoff.*

(2) *For* $d = 3$ *and infinitely many pairs of primes* $p \neq q$*, the family* $\mathrm{PSL}_3(\mathbb{F}_q)$ *has an explicit symmetric set of* $2(p^2+p+1)$ *generators exhibiting TV-cutoff (at time* $\frac{p^2+p+1}{2(p^2-1)}\log_p|\mathrm{PSL}_3(\mathbb{F}_q)|$*).*

We remark that even though this is a claim on Cayley graphs, the proof makes use of the high-dimensional geometry of their clique complexes!

Let us briefly explain our strategy for proving Theorem 1.1. Given a walk on X, we lift it to a walk on $\mathcal{B}$, and then project it to a sector $\mathcal{S} \subseteq \mathcal{B}$, which can be identified with the quotient of $\mathcal{B}$ by the stabilizer of the starting point of the walk. If X is a k-regular graph then $\mathcal{B}$ is the k-regular tree, and $\mathcal{S}$ is simply an infinite ray which we identify with $\mathbb{N}$: SRW on $\mathcal{B}$ then projects to a $(\frac{1}{k}, \frac{k-1}{k})$-biased walk on this ray, and the projected location $\ell \in \mathbb{N}$ of the walker is precisely its distance from the starting point. On the other hand, this point is also the projection to $\mathcal{S}$ of all terminal vertices of non-backtracking walks of length ℓ; combining this with the optimal cutoff for NBRW is used to establish SRW cutoff in [15].

For the building of dimension d, the so called *Cartan decomposition* gives an isomorphism $\mathcal{S} \cong \mathbb{N}^d$, and the projected walk from $\mathcal{B}$ on $\mathcal{S}$ is an explicit, homogeneous, drifted walk on $\mathbb{N}^d$ (with appropriate boundary conditions). Following the Lubetzky–Peres strategy, we would have liked to use this walk to reduce SRW-cutoff to some collision-free walk from [14], for which optimal cutoff is already established. However, the terminal vertices of the various walks studied in [14] are all located on the special rays in $\mathcal{S}$ which correspond to the standard axes in $\mathbb{N}^d$. This is enough for the graph case (when $d = 1$), but not in general. Our solution combines all the walks from [14], using cells of all positive dimensions. For each point $\alpha \in \mathcal{S}$, we construct a concatenation of collision-free walks on cells of different dimensions, so that the possible paths of the walk terminate in a uniform vertex in the preimage of α in $\mathcal{B}$. The results of [14, 22] are then used to bound the total-variation mixing of the corresponding concatenated walk on a Ramanujan complex.

For the convenience of the reader, we have divided the proof to the two-dimensional case (namely PGL_3) in Section 3, and the general case in Section 4. The case of $d = 3$ is considerably simpler, due to the fact that it has additional symmetry: in this case $\mathrm{PGL}_3(F)$ acts transitively on the cells of every dimension of $\mathcal{B}$; see [9–11] for detailed combinatorial studies of $\mathcal{B}(\mathrm{PGL}_3(F))$. In addition, it is easier to visualize (see Figure 1), and some computations can be made more explicit and give sharper bounds.

Acknowledgements. The authors are grateful to Ori Gurel-Gurevich for his help with proving Proposition 4.2. They also thank Eyal Lubetzky, Alex Lubotzky and Nati Linial for helpful discussions and encouragement. M.C. was supported by ERC grant 339096 of Nati Linial and by ERC, BSF and NSF grants of Alex Lubotzky. O.P. was supported by ISF grant 2990/21.

2. Preliminaries and notations

We briefly recall the notion of Bruhat–Tits buildings of type $\widetilde{A}_d$ and the Ramanujan complexes associated with them. For a more detailed introduction, we refer the reader to [13, 17, 20].

2.1. Bruhat–Tits buildings. Let F be a non-archimedean local field with ring of integers $\mathcal{O}$, uniformizer ϖ, and residue field $\mathcal{O}/\varpi\mathcal{O}$ of size q. The simplest examples are

$$\begin{aligned} F &= \mathbb{Q}_p && \text{with } (\mathcal{O}, \varpi, q) = (\mathbb{Z}_p, p, p), \\ F &= \mathbb{F}_q(\!(t)\!) && \text{with } (\mathcal{O}, \varpi, q) = \big(\mathbb{F}_q[\![t]\!], t, q\big). \end{aligned}$$

Let $G = \mathrm{PGL}_d(F)$ and $K = \mathrm{PGL}_d(\mathcal{O})$, which is a maximal compact subgroup of G. The Bruhat–Tits building $\mathcal{B} = \mathcal{B}(G)$ of type $\widetilde{A}_{d-1}$ associated with G is an infinite, contractible, $(d-1)$-dimensional simplicial complex, on which G acts faithfully. Denoting by $\mathcal{B}^j$ the cells of $\mathcal{B}$ of dimension j, the action of G on $\mathcal{B}$ is transitive both on $\mathcal{B}^0$ and $\mathcal{B}^{d-1}$. Furthermore, there is a vertex, which we denote by ξ, whose stabilizer is K, so that $\mathcal{B}^0$ can be identified with left K-cosets in G. In this manner each vertex $g\xi$ is associated with the F-homothety class of the $\mathcal{O}$-lattice $g\mathcal{O}^d \leq F^d$. A collection of vertices $\{g_i\xi\}_{i=0}^r$ forms an r-cell if, possibly after reordering, there exist scalars $\alpha_i \in F^\times$ such that

$$\varpi g_0\mathcal{O}^d < \alpha_r g_r \mathcal{O}^d < \alpha_{r-1} g_{r-1}\mathcal{O}^d < \cdots < \alpha_1 g_1 \mathcal{O}^d < g_0\mathcal{O}^d.$$

It follows that the link of a vertex in $\mathcal{B}$ can be identified with the spherical building of $\mathrm{PGL}_d(\mathcal{O}/\varpi\mathcal{O}) \cong \mathrm{PGL}_d(\mathbb{F}_q)$, the finite complex whose cells corresponds to flags in $\mathbb{F}_q^d$. In particular, its vertices correspond to non-zero proper subspaces of $\mathbb{F}_q^d$, so that the degree of the vertices in $\mathcal{B}$ is

$$\deg(\xi) = \textstyle\sum_{j=1}^{d-1} \left[{d \atop j} \right]_q, \quad \text{where } \left[{d \atop j} \right]_q \text{ are Gaussian binomial coefficients}$$

(see examples in Table 1). The vertices of $\mathcal{B}$ are colored by the elements of $\mathbb{Z}/d\mathbb{Z}$, via

$$\mathrm{col}(g\xi) = \mathrm{ord}_\varpi \det(g) \in \mathbb{Z}/d\mathbb{Z} \quad (g \in \mathrm{PGL}_d(F)), \tag{2.1}$$

and this coloring makes $\mathcal{B}$ a d-partite complex, namely, every $(d-1)$-cell contains all colors. For $j \geq 1$, we say that an ordered cell $\sigma \in \mathcal{B}^j$ is of *type one* if

$$\mathrm{col}\,\sigma_{i+1} \equiv \mathrm{col}\,\sigma_i + 1 \pmod d \quad \text{for } 0 \leq i < j,$$

and we denote by $\mathcal{B}_1^j$ all j-cells of type one.

2.2. Ramanujan complexes. A *branching operator* on a set Ω is a function

$$T\colon \Omega \to 2^\Omega.$$

By a *geometric* operator T on $\mathcal{B}$ we mean a branching operator on some subset $\mathcal{C}$ of the cells of $\mathcal{B}$ (e.g., all cells of dimension j), which commutes with the action of G. If Γ is a torsion-free lattice in $G = \mathrm{PGL}_d(F)$, then the quotient $X = \Gamma\backslash\mathcal{B}$ is a finite complex, equipped with a covering map $\varphi\colon \mathcal{B} \to X$, and T induces a branching operator $T|_X$ on the cells $\Gamma\backslash\mathcal{C}$ in X, via $T|_X = \varphi T \varphi^{-1}$. A function on X is considered *trivial* if its lift to a function on $\mathcal{B}$ is constant on every orbit of $G' = \mathrm{PSL}_d(F)$, and an eigenvalue of $T|_X$ is called trivial if its eigenfunction is trivial. Denote by $L^2_{\mathrm{col}}(X)$ the space of trivial functions, and by $L^2_0(X)$ its orthogonal complement.

Definition 2.1. The complex $X = \Gamma\backslash\mathcal{B}$ is called a *Ramanujan complex* if for every geometric operator T on $\mathcal{C} \subseteq \mathcal{B}$, the non-trivial spectrum $\mathrm{Spec}(T|_{L^2_0(\Gamma\backslash\mathcal{C})})$ is contained in the spectrum of T acting on $L^2(\mathcal{C})$.

In the above notations, we denote by $\mathcal{D}_T(\mathcal{B})$ the digraph with vertices $\mathcal{C}$, and edges $\{\sigma \to \sigma' \mid \sigma \in \mathcal{C},\ \sigma' \in T(\sigma)\}$, and similarly $\mathcal{D}_T(X)$ for the induced digraph on $\Gamma\backslash\mathcal{C}$.

Theorem 2.2 ([14, Theorem 3 and Proposition 5.3]). *Let T be a k-regular geometric operator on $\mathcal{B}^j_1$. If $\mathcal{D}_T(\mathcal{B})$ is collision-free and $X = \Gamma\backslash\mathcal{B}$ is a Ramanujan complex, then $\mathcal{D}_T(X)$ is a $(d)_j$-normal Ramanujan digraph (where $(d)_j = d!/(d-j)!$).*

This requires some explanation. A k-regular digraph $\mathcal{D}$ is called:

(1) *collision-free* if it has at most one directed path between any two vertices;

(2) *r-normal* if its adjacency matrix $A_{\mathcal{D}}$ is unitarily similar to a block diagonal matrix with blocks of size at most $r \times r$;

(3) a *Ramanujan digraph* if the spectrum of $A_{\mathcal{D}}$ is contained in

$$\{z \in \mathbb{C} \mid |z| = k \text{ or } |z| \leq \sqrt{k}\}.$$

Denoting by $L^2_0(\mathcal{D})$ the orthogonal complement to all $A_{\mathcal{D}}$-eigenfunctions with eigenvalue of absolute value k, we have the following theorem.

Theorem 2.3 ([22, Proposition 4.1]). *If $\mathcal{D}$ is a k-regular r-normal digraph with $\lambda = \max\{|z| \mid z \in \mathrm{Spec}(A_{\mathcal{D}}|_{L^2_0(\mathcal{D})})\}$, then*

$$\|A^\ell_{\mathcal{D}}|_{L^2_0(\mathcal{D})}\|_2 \leq \binom{\ell+r-1}{r-1} k^{r-1}\lambda^{\ell-r+1}.$$

In particular, if $\mathcal{D}$ is a k-regular r-normal Ramanujan digraph, then we have $|\lambda| \leq \sqrt{k}$ for every $\lambda \in \mathrm{Spec}(A_{\mathcal{D}}|_{L^2_0(\mathcal{D})})$, so that

$$\|A^\ell_{\mathcal{D}}|_{L^2_0(\mathcal{D})}\|_2 \leq \binom{\ell+r-1}{r-1} k^{(\ell+r-1)/2} \leq (\ell+r)^r k^{(\ell+r)/2}. \tag{2.2}$$

2.3. Cartan decomposition. With the notations of Section 2.1, the *fundamental apartment* $\mathcal{A} \subseteq \mathcal{B}$ is the subcomplex of $\mathcal{B}$ induced by all translations of ξ by diagonal matrices in G. Geometrically, $\mathcal{A}$ is a simplicial tessellation of the affine space $\mathbb{R}^{d-1}$. The edges in $\mathcal{A}$ are as follow: every vertex $\varpi^\alpha \xi = \operatorname{diag}(\varpi^{\alpha_1}, \dots, \varpi^{\alpha_{d-1}}, \varpi^{\alpha_d})\xi$ is connected to $\varpi^{\alpha+\gamma}\xi$ where γ runs over all non-constant binary vectors, i.e. $\gamma \in \{0,1\}^d \backslash \{\mathbb{0}, \mathbb{1}\}$ (here $\mathbb{0}$ and $\mathbb{1}$ denote the all-zero and all-one vectors in $\{0,1\}^d$, respectively).

We denote by $\mathcal{S} \subseteq \mathcal{A}$ the sector in $\mathcal{A}$ induced by $A\xi$, where

$$A = \left\{\varpi^\alpha = \operatorname{diag}(\varpi^{\alpha_1}, \dots, \varpi^{\alpha_{d-2}}, \varpi^{\alpha_{d-1}}, 1) \mid \alpha_1 \geq \dots \geq \alpha_{d-1} \geq \alpha_d = 0\right\}.$$

It is easy to see that $\mathcal{S}$ is a fundamental domain for the action of $S_d \leq G$ (the so called *spherical Weyl group*) on $\mathcal{A}$. We identify $\mathcal{S}$ with $\mathbb{N}^{d-1}$ via

$$\operatorname{diag}(\varpi^{\alpha_1}, \dots, \varpi^{\alpha_{d-2}}, \varpi^{\alpha_{d-1}}, 1)\xi \mapsto (\alpha_1 - \alpha_2, \dots, \alpha_{d-2} - \alpha_{d-1}, \alpha_{d-1}),$$

thereby giving $\mathbb{N}^{d-1}$ a graph structure. Denote by $\partial\mathcal{S}$ the boundary of $\mathcal{S}$, which corresponds to

$$\partial\mathbb{N}^{d-1} = \left\{\vec{x} \in \mathbb{N}^{d-1} \mid x_i = 0 \text{ for some } i\right\}.$$

Except at $\partial\mathcal{S}$, the edges are the same as in $\mathcal{A}$, parametrized by $\gamma \in \{0,1\}^d \backslash \{\mathbb{0}, \mathbb{1}\}$. For $\varpi^\alpha \in \partial\mathcal{S}$, it might happen that $\varpi^{\alpha+\gamma} \notin \mathcal{S}$, e.g. when $\alpha_i + \gamma_i > \alpha_{i-1} + \gamma_{i-1}$, and one obtains the appropriate terminus of γ by reordering the entries of $\varpi^{\alpha+\gamma}$ in descending order, and then dividing it by its last coordinate if it is not 1. The case of $d = 3$ is depicted in Figure 1.

The Cartan decomposition for PGL_d states that

$$G = \bigsqcup_{a \in A} KaK, \quad \text{or (equivalently) } \mathcal{B}^0 = \bigsqcup_{a \in A} Ka\xi,$$

and the proof is a simple exercise (see e.g. [8, Section 13.2]). It follows that $\mathcal{S}$ can also be identified with the quotient of $\mathcal{B}$ by K, and we denote the obtained projection from $\mathcal{B}$ to $\mathbb{N}^{d-1}$ by Φ. In conclusion, we have identified four complexes:

$$K \backslash \mathcal{B} \cong S_d \backslash \mathcal{A} \cong \mathcal{S} \cong \mathbb{N}^{d-1}.$$

3. The PGL_3 case

In this section, $\mathcal{B} = \mathcal{B}_{3,F}$ is the two-dimensional Bruhat–Tits building of $G = \mathrm{PGL}_3(F)$. The 1-skeleton of $\mathcal{B}$ is a k-regular graph, with

$$k = \deg(\xi) = 2(q^2 + q + 1),$$

where q is the size of the residue field of F.

Theorem 3.1. *Let $X = \Gamma \backslash \mathcal{B}_{3,F}$ be a Ramanujan complex with n vertices. Then SRW on the underlying graph of X has total-variation cutoff at time $\frac{q^2+q+1}{q^2-1} \log_{q^2} n$ with a window of size $O(\sqrt{\log n})$.*

3.1. A lower bound on the mixing time. Throughout this section we fix $\varepsilon > 0$. Denote by $B(\xi, r)$ the r-ball around ξ, i.e. the vertices of graph distance at most r from ξ in $\mathcal{B}$. First, we show that the ball of radius

$$r_0 = \log_{q^2} n - 3 \log_{q^2} \log_{q^2} n$$

can cover only a small fraction of any n-vertex quotient of $\mathcal{B}$:

Proposition 3.2. *For n large enough, $|B(\xi, r_0)| \leq \varepsilon n$.*

Proof. Given $r \geq 1$, the r-sphere $S(\xi, r)$ is shown in [7] to be of size

$$|S(\xi, r)| = (r+1)q^{2r} + 2rq^{2r-1} + 2rq^{2r-2} + (r-1)q^{2r-3}.$$

Thus, one can crudely bound the size of the r-ball by $|B(\xi, r)| \leq 8r^2 q^{2r}$, hence

$$|B(\xi, r_0)| \leq 8r_0^2 q^{2r_0} \leq \frac{8(\log_{q^2} n)^2}{(\log_{q^2} n)^3} n \leq \varepsilon n$$

for n large enough. □

Let $(\mathcal{X}_t)$ be a SRW on $\mathcal{B}$ starting at ξ. We would like to determine until when does the walk remain in the r_0-ball around ξ with high probability. Since the distance from ξ is K-invariant, we have

$$\operatorname{dist}(\zeta, \xi) = \operatorname{dist}(\Phi(\zeta), \Phi(\xi)) = \operatorname{dist}(\Phi(\zeta), (0, 0))$$

for $\zeta \in \mathcal{B}^0$, which leads us to consider the projection of $\mathcal{X}_t$ by Φ. In this manner, we obtain a (non-simple) random walk $(\Phi(\mathcal{X}_t))$ on $\mathbb{N}^2$, and we define

$$\rho(t) = \operatorname{dist}(\Phi(\mathcal{X}_t), (0, 0)) = \operatorname{dist}(\mathcal{X}_t, \xi).$$

Recall that we identified $\mathcal{S} \cong \mathbb{N}^2$ by mapping $\operatorname{diag}(\varpi^\alpha, \varpi^\beta, 1)\xi$ to $(\alpha - \beta, \beta)$, and the edges in $\mathbb{N}^2$ (except at the boundary) are $\pm(1, 0)$, $\pm(0, 1)$, $\pm(1, -1)$; see Figure 1.

Let $\vec{x}$ and $\vec{y}$ be the boundary lines of $\mathcal{S}$ (the x and y axes in Figure 1). The transition probabilities of the projected random walk are as follows: from $(0, 0)$ there is a probability of $\frac{1}{2}$ of moving to $(1, 0)$ and to $(0, 1)$. Outside the boundary, the edge (Δ_x, Δ_y) is taken with probability $q^{\Delta_x + \Delta_y + 1}/k$. On $\vec{x} \setminus (0, 0)$, the edges with $\Delta_y = -1$ are folded back in, giving the probabilities shown in Figure 1, and on $\vec{y}$ the folding is symmetric.

Denote $y(t) = \operatorname{dist}(\Phi(\mathcal{X}_t), \vec{x})$ and $x(t) = \operatorname{dist}(\Phi(\mathcal{X}_t), \vec{y})$, which measure the distance of the projected walk from the boundary. Clearly, $\rho(t) = y(t) + x(t)$. We consider $y(t)$, $x(t)$ and $\rho(t)$ as random walks on $\mathbb{N}$ starting at zero.

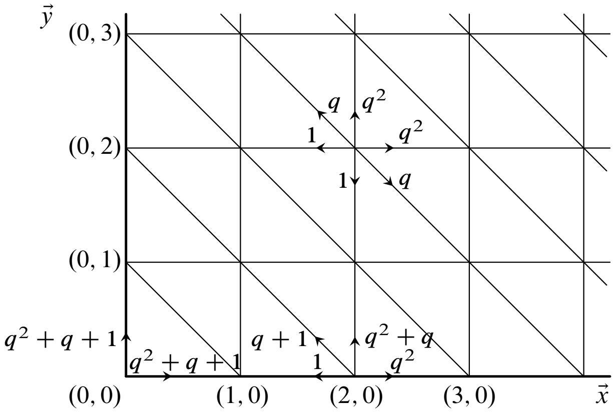

Figure 1. The sector $\mathcal{S} \subseteq \mathcal{B}$ as $\mathbb{N}^2$, and transition probabilities projected from SRW on $\mathcal{B}$, scaled by $2(q^2+q+1)$.

Proposition 3.3. *The walks $x(t)$ and $y(t)$ are transient.*

Proof. We treat only $x(t)$, and the proof for $y(t)$ is analogous. Although the distribution of the random variable $\partial x(t) = x(t) - x(t-1)$ depends on the position of the walk at time $t-1$, there are only four cases to consider: when the walk is at the origin, when the walk is on $\vec{x}$ or $\vec{y}$, and when both $x(t-1)$ and $y(t-1)$ are positive. In all of these cases, we have

$$\mathbb{E}[\partial x(t)] \geq \frac{q^2-q-2}{2(q^2+q+1)}.$$

Thus, for $q > 2$, the value of $x(t)$ is expected to strictly grow at each step and thus $x(t)$ is transient. To cover the case of $q = 2$, one can look "two steps ahead", namely on $\partial^2 x(t) = x(t) - x(t-2)$. There are more cases to check, but explicit computation shows that

$$\mathbb{E}[\partial^2 x(t)] \geq \frac{4q^4+q^3-5q^2-9q-7}{4(q^2+q+1)^2},$$

which is positive for all $q \geq 2$, giving again transience. □

Define

$$S(t) = \sum_{i=1}^{t} Y_i,$$

where $Y_i = \rho(i) - \rho(i-1)$ whenever $x(i-1), y(i-1) > 0$, and otherwise Y_i is a random variable independent of any other, attaining $1, 0, -1$ with respective probabilities $2q^2/k$, $2q/k$, $2/k$. It follows that the Y_i's are i.i.d., and by the central limit theorem, we have

$$\widetilde{S}(t) = \frac{S(t) - \mathcal{E}t}{\sigma\sqrt{t}} \Rightarrow \mathcal{N}(0,1),$$

where

$$\mathcal{E} = \frac{q^2-1}{q^2+q+1}, \quad \sigma = \frac{\sqrt{q^3+4q^2+q}}{q^2+q+1}, \tag{3.1}$$

and $\mathcal{N}(0,1)$ is the standard normal distribution. Let $\mathscr{E}(t) = (\rho(t) - S(t))/\sigma\sqrt{t}$. Since $x(i)$ and $y(i)$ are transient, the difference $\rho(t) - S(t)$ is bounded with probability 1, so that

$$\mathbb{P}[|\mathscr{E}(t)| < C] \xrightarrow{t\to\infty} 1$$

for every $C > 0$. Hence, $\mathscr{E}(t)$ converges to the Dirac measure concentrated at 0, and

$$\Xi(t) = \frac{\rho(t) - \mathcal{E}t}{\sigma\sqrt{t}} = \widetilde{S}(t) + \mathscr{E}(t) \Rightarrow \mathcal{N}(0,1). \tag{3.2}$$

Recall that $\varepsilon > 0$ and r_0 were fixed at the beginning of Section 3.1.

Proposition 3.4. *For n large enough and any $s \geq 0$, at time*

$$t_0 = t_0(s) = \frac{q^2+q+1}{q^2-1}\log_{q^2} n - (s+1)\sqrt{\log_{q^2} n} \tag{3.3}$$

the distance of $\mathcal{X}_t$ from ξ satisfies

$$\mathbb{P}[\rho(t_0) > r_0] < \mathbb{P}[Z > c_q \cdot s] + \varepsilon,$$

where $Z \sim \mathcal{N}(0,1)$ and

$$c_q = \frac{\mathcal{E}^{3/2}}{\sigma} = \sqrt{\frac{(q^2-1)^3}{(q^2+q+1)(q^3+4q^2+q)}}$$

(see (3.1)).

Proof. We note that $\rho(t_0) > r_0$ is equivalent to

$$\Xi(t_0) > \frac{r_0 - \mathcal{E}t_0}{\sigma\sqrt{t_0}} = \frac{(s+1)\mathcal{E}\sqrt{\log_{q^2} n} - 3\log_{q^2}\log_{q^2} n}{\sigma\sqrt{t_0}}.$$

For n large enough we have $\mathcal{E}\sqrt{\log_{q^2} n} \geq 3\log_{q^2}\log_{q^2} n$, and thus

$$\frac{(s+1)\mathcal{E}\sqrt{\log_{q^2} n} - 3\log_{q^2}\log_{q^2} n}{\sigma\sqrt{t_0}} > \frac{s\mathcal{E}\sqrt{\log_{q^2} n}}{\sigma\sqrt{t_0}},$$

and from $\sqrt{t_0} < \mathcal{E}^{-1/2}\sqrt{\log_{q^2} n}$ it follows that

$$\frac{s\mathcal{E}\sqrt{\log_{q^2} n}}{\sigma\sqrt{t_0}} > \frac{s\mathcal{E}^{3/2}}{\sigma} = c_q s.$$

Lastly, since $\Xi(t_0)$ converges to Z in distribution and $t_0 \xrightarrow{n\to\infty} \infty$, for n large enough, we have

$$|\mathbb{P}[\Xi(t_0) > c_q s] - \mathbb{P}[Z > c_q s]| < \varepsilon.$$

All in all, we conclude that

$$\mathbb{P}[\rho(t_0) > r_0] \le \mathbb{P}[\Xi(t_0) > c_q s] \le \mathbb{P}[Z > c_q s] + \varepsilon. \qquad \square$$

Now let X be a quotient of $\mathcal{B}$ with n vertices. For any $v \in X^0$ we can choose the covering map $\varphi\colon \mathcal{B} \to X$ to satisfy $\varphi(\xi) = v$. This map induces a correspondence between paths in X starting at v and paths in $\mathcal{B}$ starting at ξ, and in particular,

$$\varphi(B(\xi, r)) = B(v, r).$$

The projection $X_t = \varphi(\mathcal{X}_t)$ is a SRW on (the 1-skeleton of) X starting at v. We recall that $\mu_X^t = \mu_{X,v}^t$ denotes the distribution of (X_t) and π_X the uniform distribution on X^0.

Proposition 3.5. *There exists $s = s(q, \varepsilon)$ such that for n large enough, the $(1-3\varepsilon)$-mixing time of SRW on X is at least $t_0 = t_0(s)$.*

Proof. Using $\varphi(B(\xi, r)) = B(v, r)$, which implies in particular

$$\mu_{X,v}^t(B(v, r)) \ge \mu_{\mathcal{B},\xi}^t(B(\xi, r)),$$

together with Propositions 3.4 and 3.2, we have for n large enough

$$\begin{aligned}\left\|\mu_{X,v}^{t_0} - \pi_X\right\|_{TV} &\ge \pi_X(X^0 \backslash B(v, r_0)) - \mu_{X,v}^{t_0}(X^0 \backslash B(v, r_0)) \\ &\ge \frac{n - |B(v, r_0)|}{n} - \mu_{\mathcal{B},\xi}^{t_0}(\mathcal{B}^0 \backslash B(\xi, r_0)) \\ &\ge \frac{n - |B(\xi, r_0)|}{n} - \mathbb{P}[Z > c_q s] - \varepsilon \\ &\ge 1 - 2\varepsilon - \mathbb{P}[Z > c_q s].\end{aligned}$$

This implies in particular

$$\max_{v \in X^0} \|\mu_{X,v}^{t_0} - \pi_X\|_{TV} \ge 1 - 2\varepsilon - \mathbb{P}[Z > c_q s],$$

and we can choose s such that $\mathbb{P}[Z > c_q s] < \varepsilon$, and thus $t_{\mathrm{mix}}(1 - 3\varepsilon) > t_0$. $\square$

3.2. An upper bound for the mixing time. Recall from (2.1) that $\mathcal{B}$ is tri-partite via col: $\mathcal{B}^0 \to \mathbb{Z}/3\mathbb{Z}$. The quotient $X = \Gamma\backslash\mathcal{B}$ is tripartite if and only if the map col factors through X^0, which is equivalent to $\mathrm{ord}_\varpi \det \gamma \in 3\mathbb{Z}$ for all $\gamma \in \Gamma$. When this is the case, the trivial functions $L^2_{\mathrm{col}}(X^0)$ (see Section 2.2) are those which are constant on each color, and when X is not tri-partite, $L^2_{\mathrm{col}}(X^0)$ are the constant functions. Denote by $\mathcal{P}_{\mathrm{col}}$ and $\mathcal{P}_0$ the orthogonal projections corresponding to the decomposition $L^2(X^0) = L^2_{\mathrm{col}}(X^0) \oplus L^2_0(X^0)$. For any t, we have

$$\|\mu_X^t - \pi_X\|_{TV} \le \|\mathcal{P}_0(\mu_X^t)\|_{TV} + \|\mathcal{P}_{\mathrm{col}}(\mu_X^t) - \pi_X\|_{TV}. \tag{3.4}$$

We first bound the second term.

Proposition 3.6. *There exist $t_\triangle = t_\triangle(\varepsilon)$ such that $\|\mathcal{P}_{\mathrm{col}}(\mu_X^t) - \pi_X\|_{TV} \le \varepsilon$ for any $t \ge t_\triangle$.*

Proof. If X is non-tripartite then $\mathcal{P}_{\mathrm{col}}(\mu_X^t) = \pi_X$, as both are constant functions of sum one. If X is tripartite, col induces a simplicial map col: $X \to \triangle$, where $\triangle$ is the 2-simplex with vertices $\mathbb{Z}/3\mathbb{Z}$. In this case, $\mathcal{P}_{\mathrm{col}}(\mu_X^t)$ is the pullback of SRW on the 2-simplex starting at $0 = \mathrm{col}(v)$, i.e.

$$\mathcal{P}_{\mathrm{col}}(\mu_X^t)(w) = \frac{3}{n} \cdot \mu_\triangle^t(\mathrm{col}(w)).$$

The triangle is connected and non-bipartite, so there exists a time $t_\triangle$, not depending on n, such that $\|\mu_\triangle^t - \pi_\triangle\|_{TV} < \varepsilon$ for $t > t_\triangle$, hence

$$\|\mathcal{P}_{\mathrm{col}}(\mu_X^t) - \pi_X\|_{TV} = \|\mu_\triangle^{t_1} - \pi_\triangle\|_{TV} \le \varepsilon. \qquad \square$$

Next, we define

$$r_1 = \log_{q^2} n + 16 \log_{q^2} \log_{q^2} n,$$
$$t_1 = \frac{q^2+q+1}{q^2-1} \log_{q^2} n + (s+1)\sqrt{\log_{q^2} n},$$

where s is as in Proposition 3.5. Observe that by time t_1 SRW on $\mathcal{B}$ leaves $B(\xi, r_1)$ with high probability: the same arguments as in Propositions 3.4 and 3.5 give for n large enough

$$\mathbb{P}[\rho(t_1) < r_1] \le \mathbb{P}[Z > c_q s] + \varepsilon < 2\varepsilon. \tag{3.5}$$

It is left to bound $\|\mathcal{P}_0(\mu_X^{t_1})\|_{TV}$, and for this we use for the first time the assumption that X is a Ramanujan complex. We decompose $\mu_X^{t_1}$ by conditioning on the values of ρ, x, y at time t_1: denoting

$$\mu_X^{t,x,y} = \mathbb{P}\big[X_t = \cdot \mid {\scriptstyle x(t)=x, \atop y(t)=y}\big],$$

we have

$$\begin{aligned}\|\mathcal{P}_0(\mu_X^{t_1})\|_{TV} &= \Big\| \mathbb{P}\big[\rho(t_1) < r_1\big] \mathcal{P}_0\big(\mathbb{P}\big[X_t = \cdot \mid \rho(t_1) < r_1\big]\big) \\ &\qquad + \sum_{r_1 \le x+y} \mathbb{P}\big[{\scriptstyle x(t_1)=x, \atop y(t_1)=y}\big] \mathcal{P}_0(\mu_X^{t_1,x,y}) \Big\|_{TV} \\ &\le 2\varepsilon + \max_{r_1 \le x+y \le t_1} \big\|\mathcal{P}_0(\mu_X^{t_1,x,y})\big\|_{TV},\end{aligned} \tag{3.6}$$

using $x(t) + y(t) \le t_1$ and (3.5). To understand the L_0^2-projection of the conditional distribution $\mu_X^{t_1,x,y}$, we turn to study the fiber $\Phi^{-1}(x, y)$, using carefully chosen geometric operators on the cells of $\mathcal{B}$.

Recall the definition of cells of type one from Section 2.1. While $g \in G$ does not preserve colors in $\mathcal{B}^0$ in general, it does preserve the difference between colors, so that the cells of type one in X are well defined (namely, $X_1^j = \Gamma\backslash\mathcal{B}_1^j$). Let T_1 and T_2 be the geodesic edge-flow and triangle-flow operators from [14, Section 5.1]: the operator T_1 acts on $\mathcal{B}_1^1$, taking a (directed) edge vw to all edges wu of type one such that vwu is not a triangle in $\mathcal{B}$. The operator T_2 acts on $\mathcal{B}_1^2$, taking the (ordered) triangle vwu to all triangles wuy with $y \neq v$. We introduce the operators:

$$\begin{aligned}
&T_{01}\colon \mathcal{B}^0 \to \mathcal{B}_1^1, \qquad T_{01}(v) = \{wv \mid w \in \mathcal{B}^0 \text{ (and } wv \text{ is of type one)}\},\\
&T_{12}\colon \mathcal{B}_1^1 \to \mathcal{B}_1^2, \qquad T_{12}(wv) = \{uwv \mid uwv \in \mathcal{B}_1^2\},\\
&T_{20}\colon \mathcal{B}_1^2 \to \mathcal{B}^0, \quad T_{20}(uwv) = \{v\}.
\end{aligned}$$

All of the operators T_i, T_{ij} are regular and geometric.

Proposition 3.7. *For any $(x, y) \in \mathbb{N}^2$, we have $\Phi^{-1}(x, y) = T_{(x,y)}(\xi)$, where*

$$T_{(x,y)} = T_{20} \circ T_2^{2y} \circ T_{12} \circ T_1^x \circ T_{01}\colon \mathcal{B}^0 \to \mathcal{B}^0.$$

Proof. If σ_1, σ_2 are two cells in $\mathcal{B}$ with corresponding G-stabilizers G_{σ_i}, any double coset $G_{\sigma_1} g G_{\sigma_2}$ defines a geometric branching operator from the orbit $G\sigma_1$ to $G\sigma_2$, by

$$(G_{\sigma_1} g G_{\sigma_2})(g'\sigma_1) = g' G_{\sigma_1} g \sigma_2. \tag{3.7}$$

Defining

$$e_1 = \operatorname{diag}(\varpi, \varpi, 1)\xi \to \xi \quad \text{and} \quad \tau_1 = [\operatorname{diag}(\varpi, 1, 1)\xi, \operatorname{diag}(\varpi, \varpi, 1)\xi, \xi],$$

we have orbits $\mathcal{B}^0 = G\xi$, $\mathcal{B}_1^1 = Ge_1$, $\mathcal{B}_1^2 = G\tau_1$, and stabilizers

$$K = G_\xi, \quad P_1 = G_{e_1} = \begin{pmatrix} \mathcal{O} & \mathcal{O} & \mathcal{O} \\ \varpi\mathcal{O} & \mathcal{O} & \mathcal{O} \\ \varpi\mathcal{O} & \mathcal{O} & \mathcal{O} \end{pmatrix} \cap K, \quad P_2 = G_{\tau_1} = \begin{pmatrix} \mathcal{O} & \mathcal{O} & \mathcal{O} \\ \varpi\mathcal{O} & \mathcal{O} & \mathcal{O} \\ \varpi\mathcal{O} & \varpi\mathcal{O} & \mathcal{O} \end{pmatrix} \cap K.$$

The operators we defined arise as

$$\begin{gathered}
T_{01} = KP_1, \quad T_1 = P_1 \begin{pmatrix} \varpi & & \\ & 1 & \\ & & 1 \end{pmatrix} P_1, \quad T_{12} = P_1 P_2,\\
T_2 = P_2 \begin{pmatrix} & 1 & \\ \varpi & & \\ & & 1 \end{pmatrix} P_2, \quad T_{20} = P_2 K.
\end{gathered}$$

Thus, successively applying (3.7) we obtain

$$T_{(x,y)}(\xi) = KP_1 \left(P_1 \begin{pmatrix} \varpi & & \\ & 1 & \\ & & 1 \end{pmatrix} P_1 \right)^x P_1 P_2 \left(P_2 \begin{pmatrix} & 1 & \\ \varpi & & \\ & & 1 \end{pmatrix} P_2 \right)^{2y} P_2 \xi.$$

Explicit computation in [14, Section 5.1] shows that

$$\Big(P_1\Big(\begin{smallmatrix}\varpi & & \\ & 1 & \\ & & 1\end{smallmatrix}\Big)P_1\Big)^x = \Big\{\Big(\begin{smallmatrix}\varpi^x & \alpha & \beta \\ & 1 & \\ & & 1\end{smallmatrix}\Big) \mid \alpha,\beta \in \mathcal{O}/\varpi^x\mathcal{O}\Big\}P_1,$$

and we note that

$$K\Big(\begin{smallmatrix}\varpi^x & \alpha & \beta \\ & 1 & \\ & & 1\end{smallmatrix}\Big) = K\Big(\begin{smallmatrix}\varpi^x & & \\ & 1 & \\ & & 1\end{smallmatrix}\Big)$$

for $\alpha,\beta \in \mathcal{O}$, so that

$$T_{(x,y)}(\xi) = K\Big(\begin{smallmatrix}\varpi^x & & \\ & 1 & \\ & & 1\end{smallmatrix}\Big)P_1P_2\Big(P_2\Big(\begin{smallmatrix} & 1 & \\ \varpi & & \\ & & 1\end{smallmatrix}\Big)P_2\Big)^{2y}\xi$$

(we have used also $P_1, P_2 \le K$). Denoting

$$K_{1,2} = \big\{\big(\begin{smallmatrix}\mu & \\ & A\end{smallmatrix}\big) \mid \mu \in \mathcal{O}^\times,\ A \in GL_2(\mathcal{O})\big\},$$

one can verify that $P_1P_2 \subseteq K_{1,2}P_2$. In fact,

$$P_1P_2 = \Big\{I, \Big(\begin{smallmatrix}1 & & \\ & \mathcal{O} & 1 \\ & 1 & \end{smallmatrix}\Big)\Big\}P_2,$$

and since the elements of $K_{1,2}$ commute with $\Big(\begin{smallmatrix}\varpi^x & & \\ & 1 & \\ & & 1\end{smallmatrix}\Big)$ this implies

$$T_{(x,y)}(\xi) = K\Big(\begin{smallmatrix}\varpi^x & & \\ & 1 & \\ & & 1\end{smallmatrix}\Big)\Big(P_2\Big(\begin{smallmatrix} & 1 & \\ \varpi & & \\ & & 1\end{smallmatrix}\Big)P_2\Big)^{2y}\xi.$$

Finally, explicit computation shows that

$$\Big(P_2\Big(\begin{smallmatrix} & 1 & \\ \varpi & & \\ & & 1\end{smallmatrix}\Big)P_2\Big)^{2y} = \Big\{\Big(\begin{smallmatrix}\varpi^y & & \alpha \\ & \varpi^y & \beta \\ & & 1\end{smallmatrix}\Big) \mid \alpha,\beta \in \mathcal{O}/\varpi^y\mathcal{O}\Big\}P_2,$$

yielding (with α,β ranging over $\mathcal{O}/\varpi^y\mathcal{O}$)

$$T_{(x,y)}(\xi) = K\left(\begin{smallmatrix}\varpi^{x+y} & & \varpi^x\alpha \\ & \varpi^y & \beta \\ & & 1\end{smallmatrix}\right)\xi = K\Big(\begin{smallmatrix}\varpi^{x+y} & & \\ & \varpi^y & \\ & & 1\end{smallmatrix}\Big)\xi = \Phi^{-1}(x,y). \qquad \square$$

Proposition 3.8. *If $r_1 \le x+y \le t_1$, then for n large enough $\|\mathcal{P}_0(\mu_X^{t_1,x,y})\|_{TV} \le \varepsilon$.*

Proof. Recall that φ induces a correspondence between SRW on $\mathcal{B}$ and X, so that $\mu_X^{t,x,y}$ (for any t,x,y) is the pushforward of $\mu_{\mathcal{B},\xi}^{t,x,y}$ by φ:

$$\mu_X^{t,x,y}(w) = \mu_{\mathcal{B},\xi}^{t,x,y}(\varphi^{-1}(w)) = \mathbb{P}\big[\mathcal{X}_t \in \varphi^{-1}(w) \,\big|\, \begin{smallmatrix}x(t)=x,\\ y(t)=y\end{smallmatrix}\big] \quad \forall w \in X^0.$$

It follows from the Cartan decomposition that the distances from $\vec{x}$ and $\vec{y}$ together determine a unique K-orbit in $\mathcal{B}^0$. Since the SRW on $\mathcal{B}$ commutes with K, this implies that for any distance profile $(x,y) \in \mathbb{N}\times\mathbb{N}$ the distribution $\mu_{\mathcal{B},\xi}^{t,x,y}$ is the uniform distribution over $\Phi^{-1}(x,y)$, which we denote by $\pi_{x,y}$. We conclude that

$$\mu_X^{t,x,y} = \pi_{x,y}\circ\varphi^{-1}.$$

For any of the geometric operators $T = T_i, T_{ij}, T_{(x,y)}$, we denote by $\widetilde{T}$ the corresponding stochastic operator on L^2-spaces, e.g.

$$\widetilde{T}_{01}: L^2(\mathcal{B}^0) \to L^2(\mathcal{B}_1^1), \quad (\widetilde{T}_{01} f)(e) = \sum_{w: e \in T_{01}(w)} \frac{f(w)}{|T_{01}(w)|}.$$

By Proposition 3.7, we have

$$\operatorname{supp} \widetilde{T_{(x,y)}}(\mathbb{1}_\xi) \subseteq \Phi^{-1}(x,y).$$

Furthermore, $\widetilde{T_{(x,y)}}(\mathbb{1}_\xi)$ is K-invariant as

$$\widetilde{T_{(x,y)}}(\mathbb{1}_\xi)(k\xi') = \widetilde{T_{(x,y)}}(\mathbb{1}_{k^{-1}\xi})(\xi') = \widetilde{T_{(x,y)}}(\mathbb{1}_\xi)(\xi'),$$

hence

$$\widetilde{T_{(x,y)}}(\mathbb{1}_\xi) = \pi_{x,y}.$$

The stochastic operator $\widetilde{T}|_X$ corresponding to $T|_X = \varphi T \varphi^{-1}$ satisfies $(\widetilde{T}\mu)\circ\varphi^{-1} = \widetilde{T}|_X(\mu\circ\varphi^{-1})$ for any distribution μ on $\mathcal{B}$, so that

$$\mu_X^{t,x,y} = \widetilde{T_{(x,y)}}(\mathbb{1}_\xi) \circ \varphi^{-1} = \widetilde{T_{(x,y)}}\big|_X(\mathbb{1}_v) = \widetilde{T_{20}}\,\widetilde{T_2^{2y}}\,\widetilde{T_{12}}\,\widetilde{T_1^x}\,\widetilde{T_{01}}\big|_X(\mathbb{1}_v). \quad (3.8)$$

It follows from the regularity of incidence relations in X that the operators $\widetilde{T_i}|_X$ and $\widetilde{T_{ij}}|_X$ decompose with respect to the direct sums $L^2 = L^2_{\mathrm{col}} \oplus L^2_0$ of the appropriate cells, and in particular

$$\mathcal{P}_0(\mu_X^{t,x,y}) = \widetilde{T_{(x,y)}}\big|_X(\mathcal{P}_0(\mathbb{1}_v)).$$

The operators T_1 and T_2 are q^2- and q-regular, respectively, and they are shown in [14, Proposition 5.2] to be collision-free. By Theorems 2.2 and 2.3, this implies

$$\|\widetilde{T_1^x}|_{L_0^2(X_1^1)}\|_2 \le \frac{1}{q^{2x}}\binom{x+2}{2} q^4 \cdot q^{x-2} = \binom{x+2}{2} q^{2-x},$$

$$\|\widetilde{T_2^{2y}}|_{L_0^2(X_1^2)}\|_2 \le \frac{1}{q^{2y}}\binom{2y+5}{5} q^5 \cdot \sqrt{q}^{2y-5} = \binom{2y+5}{5} q^{5/2-y}.$$

In addition, we have

$$\|\widetilde{T_{01}}|_X\|_2 = 1/\sqrt{q^2+q+1}, \quad \|\widetilde{T_{12}}|_X\|_2 = 1/\sqrt{q+1},$$
$$\|\widetilde{T_{20}}|_X\|_2 = \sqrt{(q^2+q+1)(q+1)}$$

by degree considerations and evaluation on constant functions. Returning to (3.8), we use $\|\cdot\|_{TV} \le \frac{\sqrt{n}}{2}\|\cdot\|_2$ to conclude that

$$\begin{aligned}\big\|\mathcal{P}_0(\mu_X^{t,x,y})\big\|_{TV} &\le \frac{\sqrt{n}}{2}\big\|\widetilde{T_{(x,y)}}\big|_{L_0^2(X^0)}(\mathcal{P}_0(\mathbb{1}_v))\big\|_2 \\ &\le \frac{\sqrt{n}}{2}\big\|\widetilde{T_{(x,y)}}\big|_{L_0^2(X^0)}\big\|_2 \le \frac{\sqrt{n}\binom{x+2}{2}\binom{2y+5}{5} q^{9/2}}{2q^{x+y}}.\end{aligned}$$

Taking now $r_1 \le x + y \le t_1$, we assume n is large enough that $t_1 \le 3r_1$, hence for n large enough

$$\|\mathcal{P}_0(\mu_X^{t_1,x,y})\|_{TV} \le \frac{\sqrt{n}\binom{3r_1+2}{2}\binom{6r_1+5}{5}q^{9/2}}{2q^{r_1}} \le \frac{\sqrt{n}(7r_1)^7 q^{9/2}}{2q^{r_1}}$$
$$= \frac{(7\log_{q^2} n + 112\log_{q^2}\log_{q^2} n)^7 q^{9/2}}{2(\log_{q^2} n)^8} \le \varepsilon. \qquad \square$$

We come to the proof of the main theorem of this section.

Proof of Theorem 3.1. From (3.4), Proposition 3.6 (which applies once $t_1 \ge t_\triangle$), (3.6), and Proposition 3.8, we conclude that

$$\|\mu_X^{t_1} - \pi_X\|_{TV} \le 3\varepsilon + \max_{r_1 \le x+y \le t_1} \|\mathcal{P}_0(\mu_X^{t_1,x,y})\|_{TV} \le 4\varepsilon,$$

so that $t_{\mathrm{mix}}(4\varepsilon) \le t_1$. Together with Proposition 3.5, this implies the cutoff phenomenon at time $\frac{q^2+q+1}{q^2-1}\log_{q^2} n$, with a window of size $O(\sqrt{\log_{q^2} n})$. $\square$

4. The case of PGL_d for all $d \ge 2$

The main difference between PGL_3 and the general case is that PGL_3 acts transitively on $\mathcal{B}^j$ for all j, but the same does not happen for general d. As a result, the projected walk on the sector $\mathcal{S} = K\backslash\mathcal{B}$ is no longer isotropic – some directions are more likely to be chosen than others. Our approach is to define a suitable metric on $\mathcal{S}$ and $\mathcal{B}$, which takes this asymmetry into account. Albeit, PGL_d still acts transitively on $\mathcal{B}^0$, so the 1-skeleton of $\mathcal{B}$ is a regular graph.

4.1. The projected walk on $\mathcal{S}$.

As in Section 3, we consider a SRW $(\mathcal{X}_t)$ on $\mathcal{B}$ starting from ξ, which projects modulo K to a weighted random walk on $\mathcal{S}$. Recalling the identification $\mathcal{S} \cong \mathbb{N}^{d-1}$, we define $x_i(t)$ to be the i-th index of the projected walk $\Phi(\mathcal{X}_t)$, so that

$$\Phi(\mathcal{X}_t) = \vec{x}(t) = (x_1(t), \dots, x_{d-1}(t)).$$

We consider $\mathcal{S}$ as a weighted directed graph, with the weight of an edge being the probability that the projected walk chooses this edge. The weights are easier to describe outside the boundary: it follows from the identification of the link of a vertex as the flag complex of $\mathbb{F}_q^d$ that for every $\gamma \in \{0,1\}^d \backslash \{\mathbb{0}, \mathbb{1}\}$ (see Section 2.3) and $\varpi^\alpha \xi \in \mathcal{S}\backslash\partial\mathcal{S}$, the probability of moving from $\varpi^\alpha\xi$ to $\varpi^{\alpha+\gamma}\xi$ is

$$\mathbb{P}[\varpi^\alpha\xi \to \varpi^{\alpha+\gamma}\xi] = \frac{q^{Z_\gamma}}{\deg \xi}, \quad \text{where } Z_\gamma = \#\{(i,j) \mid i < j,\ \gamma_i = 1,\ \gamma_j = 0\}.$$

At the boundary, the only difference is that γ which leads outside of $\mathcal{S}$ is folded back into it by the action of the spherical Weyl group S_d.

Claim 4.1. If $\Phi(\mathcal{X}_{t-1}) \notin \partial\mathbb{N}^{d-1}$, then

$$\mathbb{P}[x_i(t) - x_i(t-1) = 1] = q \cdot \mathbb{P}[x_i(t) - x_i(t-1) = -1] \quad (1 \le i \le d-1).$$

Proof. When moving along γ (except at the boundary) the i-th index x_i of the projected walk changes by $\gamma_{i-1} - \gamma_i$. The permutation τ on $\{0,1\}^d$, which transposes the i-th and $(i-1)$-th indices, induces an involution on $\{0,1\}^d \backslash \{\mathbb{0}, \mathbb{1}\}$ that reverses the change in x_i. If $\gamma_{i-1} = 1$ and $\gamma_i = 0$, then $Z_\gamma = \frac{1}{q} Z_{\tau(\gamma)}$, so for every edge that decreases x_i there is an edge which increases it whose weight is q times larger. □

In what follows, for $\gamma \in \mathbb{Z}^d$ we denote by γ' the difference vector

$$\gamma' = (\gamma_1 - \gamma_2, \gamma_2 - \gamma_3, \ldots, \gamma_{d-1} - \gamma_d) \in \mathbb{Z}^{d-1}.$$

Proposition 4.2. *The projected walk $\Phi(\mathcal{X}_t)$ visits $\partial\mathbb{N}^{d-1}$ only a finite number of times with probability one.*

Proof. In essence this follows from the fact that the boundary is sink-less, and on its complement the walk is a positively-drifted walk on $\mathbb{Z}^{d-1}$. Namely, from any point in $\mathbb{N}^{d-1}$ the probability to enter

$$D = \{\alpha \in \mathbb{N}^{d-1} \mid \forall i, \alpha_i \ge d-1\}$$

in $d(d-1)$ steps is at least $(\deg \xi)^{-d(d-1)}$. Let $P_{\vec{x}}$ be the probability that a walk which starts from $\vec{x} \in \mathbb{N}^{d-1}$ ever touches the boundary. This is the same as the probability of the walk on $\mathbb{Z}^{d-1}$ with transition probability $q^{Z_\gamma}/\deg \xi$ of moving along γ', where $\gamma \in \{0,1\}^d \backslash \{\mathbb{0}, \mathbb{1}\}$, to ever reach from $\vec{x}$ to a point with a zero coordinate. This is bounded by

$$\sum_{i=1}^{d-1} P_{\vec{x},i},$$

where $P_{\vec{x},i}$ is the probability that the i-th coordinate ever vanishes. But on $\mathbb{Z}^{d-1}$ each coordinate is a drifted walk as in Claim 4.1, hence it follows from standard arguments that $P_{\vec{x},i} \le 1/q^{x_i}$. Thus, for $\vec{x} \in D$ we obtain $P_{\vec{x}} \le d/q^d < 1$, and it follows that the expected number of visits to the boundary is bounded by

$$\sum_{i=0}^{\infty} d^i (\deg \xi)^{d(d-1)}/q^{di} < \infty. \qquad \square$$

4.2. Geometric operators on $\mathcal{B}(\mathrm{PGL}_d)$. For $1 \le j < d$, the geodesic j-flow T_j defined in [14] is a q^{d-j}-regular branching operator on $\mathcal{B}_1^j$ (the j-cells of type one), which takes the (ordered) cell $[v_0, \ldots, v_j]$ to all cells $[v_1, \ldots, v_j, w] \in \mathcal{B}_1^j$ such that $\{v_0, \ldots, v_j, w\} \notin \mathcal{B}$. Defining

$$\xi_i = \mathrm{diag}(\varpi^{\times(d-i)}, 1^{\times i})\xi \quad \text{and} \quad \sigma^j = [\xi_j, \xi_{j-1}, \ldots, \xi_0],$$

we have $\mathcal{B}_1^j = G\sigma^j$, and the operator T_j corresponds to the double coset $P_j w_j P_j$, where

$$P_j := G_{\sigma^j} = \{g \in K \mid g_{r,c} \in \varpi\mathcal{O} \text{ for } c \leq \min(j, r-1)\},$$

$$w_j = \left(\begin{array}{c|c|c} & I_{j-1} & \\ \hline \varpi & & 0 \cdots 0 \\ \hline & & I_{d-j} \end{array}\right)$$

(note $\mathcal{B}_1^0 = \mathcal{B}^0 = G\sigma^0$ and $P_0 = K$, though there is no 0-flow). Each double coset $P_j P_{j+1}$ $(0 \leq j \leq d-2)$ gives via (3.7) an operator

$$T_{j,j+1} \colon \mathcal{B}_1^j \to \mathcal{B}_1^{j+1},$$

which takes $\sigma \in \mathcal{B}_1^j$ to all $v\sigma \in \mathcal{B}_1^{j+1}$ $(v \in \mathcal{B}^0)$. In addition, $P_{d-1}P_0$ yields

$$T_{d-1,0} \colon \mathcal{B}_1^{d-1} \to \mathcal{B}^0,$$

which returns the last vertex of a cell.

Proposition 4.3. *For any $\vec{x} \in \mathbb{N}^{d-1}$, the fiber $\Phi^{-1}(\vec{x})$ equals $T_{\vec{x}}(\xi)$, where*

$$T_{\vec{x}} := T_{d-1,0} \prod_{j=d-1}^{1} T_j^{jx_j} T_{j-1,j} \colon \mathcal{B}^0 \to \mathcal{B}^0.$$

Proof. Denoting $g_t = \operatorname{diag}(\varpi^{x_1+\cdots+x_t}, \varpi^{x_2+\cdots+x_t}, \ldots, \varpi^{x_t}, 1, \ldots, 1)$, we claim that

$$T_{\vec{x}}(\xi) = Kg_{t-1}\Big[\prod_{j=t}^{d-1} P_j(w_j P_j)^{jx_j}\Big]\xi \quad \text{for } 1 \leq t \leq d. \tag{4.1}$$

For $t = 1$, the definitions of $T_{\vec{x}}$ and the operators $T_i, T_{i,j}$ indeed give

$$T_{\vec{x}}(\xi) = \Big[\prod_{j=1}^{d-1} P_{j-1}P_j(P_j w_j P_j)^{jx_j}\Big]P_{d-1}P_0\xi = Kg_0\Big[\prod_{j=1}^{d-1} P_j(w_j P_j)^{jx_j}\Big]\xi.$$

Assume that (4.1) holds for some $1 \leq t \leq d-1$. Explicit computation as in [14, Section 5.1] gives

$$P_t(w_t P_t)^{tx_t} = \left(\begin{array}{c|c} \varpi^{x_t} I_t & M_{t\times d-t}(\mathcal{O}) \\ \hline 0 & I_{d-t} \end{array}\right) P_t,$$

and using K to perform row elimination we obtain

$$\begin{aligned} T_{\vec{x}}(\xi) &= Kg_{t-1}\left(\begin{array}{c|c} \varpi^{x_t} I_t & M_{t\times d-t}(\mathcal{O}) \\ \hline 0 & I_{d-t} \end{array}\right) P_t\Bigg[\prod_{j=t+1}^{d-1} P_j(w_j P_j)^{jx_j}\Bigg]\xi \\ &= Kg_t P_t P_{t+1}\Bigg[\prod_{j=t+1}^{d-1} P_j(w_j P_j)^{jx_j}\Bigg]\xi. \end{aligned}$$

Next, observe that $P_t P_{t+1}$ decomposes as $S_t P_{t+1}$ when $S_t \subseteq K$ is any set which takes σ^t to all $(t+1)$-cells containing it. There are $(q^{d-t}-1)/(q-1)$ such cells, as in the spherical building σ^t corresponds to a t-dimensional subspace of $\mathbb{F}_q^d$, and these cells to the minimal subspaces containing it. This also shows how to compute such a transversal S_t, and

$$S_t = \bigsqcup_{j=1}^{d-t} \operatorname{diag}(I_t, Q_j, I_{d-t-j}), \quad \text{where } Q_j = \begin{pmatrix} \mathbb{F}_q & 1 & & \\ \vdots & & \ddots & \\ \mathbb{F}_q & & & 1 \\ 1 & 0 & \cdots & 0 \end{pmatrix} \subseteq GL_j(\mathcal{O}),$$

is one option. Since the matrices in S_t above commute with g_t (and lie in K), this shows that

$$Kg_t P_t P_{t+1}\Bigg[\prod_{j=t+1}^{d-1} P_j(w_j P_j)^{jx_j}\Bigg]\xi = Kg_t\Bigg[\prod_{j=t+1}^{d-1} P_j(w_j P_j)^{jx_j}\Bigg]\xi,$$

establishing (4.1) for $t+1$. Taking $t = d$ in (4.1), we obtain

$$T_{\vec{x}}(\xi) = Kg_{d-1}\xi = \Phi^{-1}(x_1,\dots,x_{d-1}). \qquad \square$$

The decomposition of $\Phi^{-1}(\vec{x})$ suggests the metric to impose on $\mathcal{S}$.

Definition 4.4. The R-norm on $\mathbb{N}^{d-1} \cong \mathcal{S}$ is

$$R(x_1,\dots,x_{d-1}) = \sum_{j=1}^{d-1} j(d-j)x_j.$$

In addition, we obtain a bound on the size of the fiber above $\vec{x}$.

Corollary 4.5. *For $\vec{x} \in \mathbb{N}^{d-1}$, the size of the fiber $\Phi^{-1}(\vec{x})$ is bounded by*

$$|\Phi^{-1}(\vec{x})| \le \prod_{j=1}^{d-1} \frac{q^{j+1}-1}{q-1} \cdot q^{R(\vec{x})} \le d!\, q^{\binom{d}{2}+R(\vec{x})}.$$

Proof. This follows from Proposition 4.3, since T_j is q^{d-j}-regular, $T_{j,j+1}$ is $(q^{d-j}-1)/(q-1)$-regular (see proof of Proposition 4.3), and $T_{d-1,0}$ is 1-regular. $\square$

Later, we will be interested in the long-term behavior of the R-distance of the random walk on $\mathcal{S}$ from ξ. The change in the R-distance distributes in the same manner whenever the walk is not at the boundary $\partial\mathcal{S}$. We denote this distribution by $\mathcal{D}$:

$$\mathbb{P}[\mathcal{D} = j] = \sum_{\gamma\in\{0,1\}^d\setminus\{\mathbb{0},\mathbb{1}\}:R(\gamma')=j} \frac{q^{Z_\gamma}}{\deg(\xi)},$$

and define $\mathcal{E}_d = \mathbb{E}[\mathcal{D}]$ and $\sigma_d = \sqrt{\mathrm{Var}[\mathcal{D}]}$ (note that $\mathcal{E}$ from Section 3 is $\mathcal{E}_3/2$). The reciprocal of $\mathcal{E}_d$ is the constant $C_{d,q}$, which appears in Theorem 1.1:

$$C_{d,q} := \frac{1}{\mathcal{E}_d} = \Bigg[\sum_{\gamma\in\{0,1\}^d\setminus\{\mathbb{0},\mathbb{1}\}} \frac{R(\gamma')q^{Z_\gamma}}{\deg(\xi)}\Bigg]^{-1}. \tag{4.2}$$

Proposition 4.6. *We have* $\mathcal{E}_d = \lfloor\frac{d}{2}\rfloor\lceil\frac{d}{2}\rceil + O(\frac{1}{q})$.

Proof. Recall that

$$\begin{bmatrix} d \\ j \end{bmatrix}_q = \prod_{i=1}^{j} \frac{q^{d-i+1}-1}{q^i-1}.$$

Writing $f \approx g$ for $f(q) = g(q)(1 + O(1/q))$, this implies $[{d \atop j}]_q \approx q^{j(d-j)}$, and thus

$$\deg(\xi) \approx \sum_{j=\lfloor d/2\rfloor}^{\lceil d/2\rceil} \begin{bmatrix} d \\ j \end{bmatrix}_q \approx \frac{3-(-1)^d}{2} q^{\lfloor d/2\rfloor\lceil d/2\rceil}.$$

Similarly, Z_γ is largest when γ is a sequence of $\lfloor d/2\rfloor$ or $\lceil d/2\rceil$ ones, followed by zeros. In this case, we have $Z_\gamma = \lfloor d/2\rfloor\lceil d/2\rceil$, and also $R(\gamma') = \lfloor d/2\rfloor\lceil d/2\rceil$, hence it follows from (4.2) that $\mathcal{E}_d \approx \lfloor d/2\rfloor\lceil d/2\rceil$. □

We demonstrate the first few cases of $\mathcal{E}_d$ and $\deg(\xi)$ in Table 1.

4.3. Cutoff on $\mathcal{B}(\mathrm{PGL}_d)$. Fix $\varepsilon > 0$. For $r \geq 0$ we define $B^R(\xi, r)$, the R-normalized r-ball around ξ, to be the set of vertices $\zeta \in \mathcal{B}^0$ satisfying $R(\Phi(\zeta))) \leq r$. From Corollary 4.5, we obtain the bound

$$|B^R(\xi,r)| \leq |\{\vec{x} \mid R(\vec{x}) \leq r\}| d!q^{\binom{d}{2}+r} \leq d!q^{\binom{d}{2}} \cdot r^{d-1}q^r. \tag{4.3}$$

Defining

$$r_0 = \log_q n - d\log_q \log_q n,$$

we obtain, for n large enough, that

$$|B^R(\xi,r_0)| \leq d!q^{\binom{d}{2}} \frac{\log_q^{d-1}(n)\cdot n}{\log_q^d(n)} < \varepsilon n. \tag{4.4}$$

d	$\mathcal{E}_d \cdot \deg(\xi)$	$\deg(\xi) = \sum_{j=1}^{d-1} \left[{d \atop j} \right]_q$
2	$q-1$	$q+1$
3	$4q^2-4$	$2q^2+2q+2$
4	$4q^4+8q^3+2q^2-4q-10$	$q^4+3q^3+4q^2+3q+3$
5	$12q^6+8q^5+16q^4+4q^3$ $-8q^2-12q-20$	$2q^6+2q^5+6q^4+6q^3$ $+6q^2+4q+4$
6	$9q^9+23q^8+22q^7+25q^6$ $+21q^5+3q^4-15q^3$ $-28q^2-25q-35$	$q^9+3q^8+4q^7+7q^6$ $+9q^5+11q^4+9q^3$ $+8q^2+5q+5$
7	$24q^{12}+20q^{11}+52q^{10}+52q^9$ $+56q^8+32q^7+24q^6$ $-8q^5-40q^4-52q^3$ $-60q^2-44q-56$	$2q^{12}+2q^{11}+6q^{10}+18q^6$ $+16q^5+8q^9+12q^8$ $+12q^7+16q^4+12q^3$ $+10q^2+6q+6$

Table 1. The polynomials which arise in the computation of $\deg(\xi)$, $\mathcal{E}_d$, $C_{d,q}$.

For a finite quotient X of $\mathcal{B}$ and $v \in X^0$, we choose a covering map $\varphi\colon \mathcal{B} \to X$ with $\varphi(\xi) = v$ as before, and define $B^R(v,r) = \varphi(B^R(\xi,r))$ (this is independent of the choice of φ as R is K-invariant). As in Section 3, $(X_t) = \varphi(\mathcal{X}_t)$ is a SRW on X starting from v. We define

$$\rho(t) = R(\Phi(\mathcal{X}_t)) = \sum_{j=1}^{d-1} j(d-j)x_j(t),$$

and recall that $\rho(t) - \rho(t-1) \sim \mathcal{D}$ when $\Phi(\mathcal{X}_{t-1}) \notin \partial\mathbb{N}^{d-1}$. By the same arguments as in PGL_3 (with Proposition 4.2 replacing Proposition 3.3), we see that

$$(\rho(t) - \mathcal{E}_d t)/(\sigma_d \sqrt{t}) \Rightarrow \mathcal{N}(0,1).$$

Proposition 4.7. *There exists $s = s(q,\varepsilon)$ such that $t_{\mathrm{mix}}(1-3\varepsilon) > t_0$ for large enough X, where*

$$t_0 = \frac{1}{\mathcal{E}_d}\log_q n - (s+1)\sqrt{\log_q n}.$$

Proof. Using a similar computation to the one in Proposition 3.4, we obtain

$$\mathbb{P}[\rho(t_0) > r_0] < \mathbb{P}[Z > cs] + \varepsilon$$

for $Z \sim \mathcal{N}(0,1)$ and $c = c(q,d) = \mathcal{E}_d^{3/2}/\sigma_d$. Combining this with (4.4), the proof continues as that of Proposition 3.5, with $B^R(v,r_0)$ replacing $B(v,r_0)$. □

We turn to the upper bound, starting again with the trivial spectrum. For $X = \Gamma\backslash\mathcal{B}$, we have

$$\{\mathrm{ord}_\varpi \det\gamma \mid \gamma \in \Gamma\} = m\mathbb{Z}$$

for a unique $m \mid d$, and we say that X is m-partite. We obtain a map col: $X^0 \to \mathbb{Z}/m\mathbb{Z}$, which we again consider as a simplicial map from X to $\triangle_{m-1}$, the $(m-1)$-dimensional simplex. We have

$$L^2_{\mathrm{col}}(X^0) = \mathrm{col}^{-1}(L^2(\triangle^0_{m-1})),$$

and $L^2_0(X^0)$, $\mathcal{P}_{\mathrm{col}}$, $\mathcal{P}_0$ are defined as before. The walk induced from X on $\triangle_{m-1}$ is not simple, but every edge is taken with positive probability. Furthermore, unless $d = m = 2$, the walk is aperiodic, since even if $m = 2$ there are loops at the vertices of $\triangle_1$ when $d \geq 3$. The case $d = m = 2$ is that of bipartite Ramanujan graphs, on which SRW does not mix, and for the rest of the paper we exclude this case. We conclude as before that there exists $t_\triangle = t_\triangle(\varepsilon)$ with $\|\mathcal{P}_{\mathrm{col}}(\mu_X^t) - \pi_X\|_{TV} \leq \varepsilon$ for any $t \geq t_\triangle$, hence

$$\begin{aligned}\|\mu_X^t - \pi_X\|_{TV} &\leq \|\mathcal{P}_0(\mu_X^t)\|_{TV} + \|\mathcal{P}_{\mathrm{col}}(\mu_X^t) - \pi_X\|_{TV}\\ &\leq \|\mathcal{P}_0(\mu_X^t)\|_{TV} + \varepsilon.\end{aligned} \tag{4.5}$$

We now choose

$$\begin{aligned}r_1 &= \log_q n + 4\,d!\log_q\log_q n,\\ t_1 &= \frac{1}{\mathcal{E}_d}\log_q n + (s+1)\sqrt{\log_q n},\end{aligned}$$

and the same c and s as in Proposition 4.7 give for n large enough

$$\mathbb{P}[\rho(t_1) < r_1] \leq \mathbb{P}[Z > cs] + \varepsilon < 2\varepsilon.$$

Denoting $\mu_X^{t,\vec{x}} = \mathbb{P}[X_t = \cdot \mid \vec{x}(t) = \vec{x}]$ and

$$S = \Big\{\vec{x} \in \mathbb{N}^{d-1} \mid \sum_{i=1}^{d-1} x_i \leq t_1 \text{ and } R(\vec{x}) \geq r_1\Big\},$$

we obtain

$$\begin{aligned}\big\|\mathcal{P}_0(\mu_X^{t_1})\big\|_{TV} &= \Big\|\mathbb{P}\big[\rho(t_1) < r_1\big]\mathcal{P}_0\big(\mathbb{P}[X_t = \cdot \mid \rho(t_1) < r_1]\big)\\ &\qquad + \sum_{\vec{x}: r_1 \leq R(\vec{x})} \mathbb{P}\big[\vec{x}(t_1) = \vec{x}\big]\mathcal{P}_0(\mu_X^{t_1,\vec{x}})\Big\|_{TV}\\ &\leq 2\varepsilon + \max_{\vec{x}\in S}\big\|\mathcal{P}_0(\mu_X^{t_1,\vec{x}})\big\|_{TV}.\end{aligned} \tag{4.6}$$

Proposition 4.8. *If* $\sum_{i=1}^{d-1} x_i \leq t_1$ *and* $R(\vec{x}) \geq r_1$, *then for* n *large enough, we have*

$$\left\| \mathcal{P}_0(\mu_X^{t_1,\vec{x}}) \right\|_{TV} \leq \varepsilon.$$

Proof. Denote by $\pi_{\vec{x}}$ the uniform distribution on $\Phi^{-1}(\vec{x})$. Using the same argument as in Proposition 3.8, with Proposition 4.3 replacing Proposition 3.7, for any t, we obtain

$$\begin{aligned} \mu_X^{t,\vec{x}} &= \pi_{\vec{x}} \circ \varphi^{-1} = \widetilde{T_{\vec{x}}}(\mathbb{1}_\xi) \circ \varphi^{-1} \\ &= \widetilde{T_{\vec{x}}}\big|_X(\mathbb{1}_v) = \widetilde{T_{d-1,0}} \prod_{j=d-1}^{1} \widetilde{T_j}^{jx_j} \widetilde{T_{j-1,j}}\big|_X(\mathbb{1}_v). \end{aligned}$$

Again the operators $\widetilde{T_i}|_X$ and $\widetilde{T_{ij}}|_X$ decompose with respect to $L^2 = L^2_{\text{col}} \oplus L^2_0$, so that

$$\mathcal{P}_0(\mu_X^{t,\vec{x}}) = \widetilde{T_{\vec{x}}}\big|_X(\mathcal{P}_0(\mathbb{1}_v)).$$

By [14, Section 5.1], the j-flow operator T_j is q^{d-j}-regular and collision-free, and using Theorem 2.2 and (2.2), we obtain

$$\begin{aligned} \|\widetilde{T_j^{jx_j}}|_{L^2_0(X_1^j)}\|_2 &\leq \frac{1}{q^{(d-j)jx_j}} \cdot \big(jx_j + (d)_j\big)^{(d)_j} (q^{d-j})^{\frac{jx_j+(d)_j}{2}} \\ &= \frac{\big(q^{(d-j)/2}(jx_j + (d)_j)\big)^{(d)_j}}{q^{(d-j)jx_j/2}} \overset{(*)}{\leq} \frac{(2q^d dt_1)^{(d)_j}}{q^{(d-j)jx_j/2}}, \end{aligned}$$

where $(*)$ assumes n is large enough. If T is a branching operator of out-degree d_o and in-degree d_i, then $\|\widetilde{T}\|_2 = \sqrt{d_i/d_o}$, so that

$$\|\widetilde{T_{d-1,0}}\|_2 \prod_{j=d-1}^{1} \|\widetilde{T_{j-1,j}}\|_2 = 1.$$

Using $\sum_{j=1}^{d-1} (d)_j < 2 \cdot d!$, we obtain for n large enough that

$$\begin{aligned} \|\widetilde{T_{\vec{x}}}|_{L^2_0(X^0)}\|_2 &\leq \prod_{j=1}^{d-1} \|\widetilde{T_j^{jx_j}}|_{L^2_0(X_1^j)}\|_2 \\ &\leq \frac{1}{q^{R(\vec{x})/2}} \prod_{j=1}^{d-1} (2q^d dt_1)^{(d)_j} \leq \frac{(2q^d dt_1)^{2d!-1}}{q^{R(\vec{x})/2}}. \end{aligned}$$

Taking n large enough such that $t_1 \leq \frac{2}{\varepsilon_d} \log_q n$, we obtain from $R(\vec{x}) \geq r_1$ that for n large enough

$$\begin{aligned} \left\| \mathcal{P}_0(\mu_X^{t,\vec{x}}) \right\|_{TV} &= \left\| \widetilde{T_{\vec{x}}}\big|_X(\mathcal{P}_0(\mathbb{1}_v)) \right\|_{TV} \leq \frac{\sqrt{n}}{2} \left\| \widetilde{T_{\vec{x}}}|_{L^2_0(X^0)} \right\|_2 \\ &\leq \frac{(4q^d d\,\varepsilon_d^{-1} \log_q n)^{2d!-1}}{2(\log_q n)^{2d!}} < \varepsilon. \end{aligned}$$

□

We conclude with the proof of the main theorem:

Proof of Theorem 1.1. From (4.5), (4.6), and Proposition 4.8, we conclude that

$$t_{\text{mix}}(4\varepsilon) \le t_1$$

for $n = |X^0|$ large enough. Together with Proposition 4.7, this implies the cutoff phenomenon at time $\frac{1}{\mathcal{E}_d} \log_q n$, with a window of size $O(\sqrt{\log n})$, and

$$C_{d,q} = \frac{1}{\mathcal{E}_d} = \frac{1}{\lfloor d/2 \rfloor \lceil d/2 \rceil} + O\Big(\frac{1}{q}\Big)$$

by Proposition 4.6. □

References

[1] D. Aldous, Random walks on finite groups and rapidly mixing Markov chains. In *Seminar on probability XVII*, pp. 243–297, Lecture Notes in Math. 986, Springer, Berlin, 1983 Zbl 0514.60067 MR 770418

[2] C. Ballantine, S. Evra, B. Feigon, K. Maurischat, and O. Parzanchevski, Explicit Cayley-like Ramanujan bigraphs. In preparation

[3] E. Breuillard and A. Lubotzky, Expansion in simple groups. In *Dynamics, geometry, number theory. The impact of Margulis on modern mathematics*, pp. 246–275, Univ. Chicago Press, Chicago, IL, 2022 MR 4422056

[4] A. Dembo, Y. Peres, and D. Revelle, Sharp thresholds for mixing times. Workshop, American Institute of Mathematics, San Jose, CA, 2004, `https://aimath.org/pastworkshops/mixingtimes.html`

[5] P. Diaconis, The cutoff phenomenon in finite Markov chains. *Proc. Nat. Acad. Sci. U.S.A.* **93** (1996), no. 4, 1659–1664 Zbl 0849.60070 MR 1374011

[6] P. Diaconis and M. Shahshahani, Generating a random permutation with random transpositions. *Z. Wahrsch. Verw. Gebiete* **57** (1981), no. 2, 159–179 Zbl 0485.60006 MR 626813

[7] S. Evra and O. Parzanchevski, Ramanujan complexes and golden gates in $PU(3)$. *Geom. Funct. Anal.* **32** (2022), no. 2, 193–235 Zbl 07511027 MR 4408431

[8] D. Goldfeld and J. Hundley, *Automorphic representations and L-functions for the general linear group. Vol. II*. Cambridge Stud. Adv. Math. 130, Cambridge Univ. Press, Cambridge, 2011 Zbl 1253.11060 MR 2808915

[9] K. Golubev and O. Parzanchevski, Spectrum and combinatorics of two-dimensional Ramanujan complexes. *Israel J. Math.* **230** (2019), no. 2, 583–612 Zbl 1410.05126 MR 3940429

[10] M.-H. Kang and W.-C. W. Li, Zeta functions of complexes arising from PGL(3). *Adv. Math.* **256** (2014), 46–103 Zbl 1328.22008 MR 3177290

[11] T. Kaufman and O. Parzanchevski, Free flags over local rings and powering of high-dimensional expanders. *Int. Math. Res. Not. IMRN* (2021), Art. ID. rnab126

[12] H. Kesten, Symmetric random walks on groups. *Trans. Amer. Math. Soc.* **92** (1959), 336–354 Zbl 0092.33503 MR 109367

[13] W.-C. W. Li, Ramanujan hypergraphs. *Geom. Funct. Anal.* **14** (2004), no. 2, 380–399 Zbl 1084.05047 MR 2060199

[14] E. Lubetzky, A. Lubotzky, and O. Parzanchevski, Random walks on Ramanujan complexes and digraphs. *J. Eur. Math. Soc. (JEMS)* **22** (2020), no. 11, 3441–3466 Zbl 1454.05114 MR 4167012

[15] E. Lubetzky and Y. Peres, Cutoff on all Ramanujan graphs. *Geom. Funct. Anal.* **26** (2016), no. 4, 1190–1216 Zbl 1351.05208 MR 3558308

[16] E. Lubetzky and A. Sly, Cutoff phenomena for random walks on random regular graphs. *Duke Math. J.* **153** (2010), no. 3, 475–510 Zbl 1202.60012 MR 2667423

[17] A. Lubotzky, Ramanujan complexes and high dimensional expanders. *Jpn. J. Math.* **9** (2014), no. 2, 137–169 Zbl 1302.05095 MR 3258617

[18] A. Lubotzky, R. Phillips, and P. Sarnak, Ramanujan graphs. *Combinatorica* **8** (1988), no. 3, 261–277 Zbl 0661.05035 MR 963118

[19] A. Lubotzky, B. Samuels, and U. Vishne, Explicit constructions of Ramanujan complexes of type $\widetilde{A}_d$. *European J. Combin.* **26** (2005), no. 6, 965–993 Zbl 1135.05038 MR 2143204

[20] A. Lubotzky, B. Samuels, and U. Vishne, Ramanujan complexes of type $\widetilde{A}_d$. *Israel J. Math.* **149** (2005), 267–299 Zbl 1087.05036 MR 2191217

[21] G. A. Margulis, Explicit constructions of expanders. *Probl. Peredaci Inform.* **9** (1973), no. 4, 71–80 Zbl 0312.22011 MR 0484767

[22] O. Parzanchevski, Ramanujan graphs and digraphs. In *Analysis and geometry on graphs and manifolds*, pp. 344–367, London Math. Soc. Lecture Note Ser. 461, Cambridge Univ. Press, Cambridge, 2020 Zbl 1473.05113 MR 4412979

[23] L. Saloff-Coste, Random walks on finite groups. In *Probability on discrete structures*, pp. 263–346, Encyclopaedia Math. Sci. 110, Springer, Berlin, 2004 Zbl 1049.60006 MR 2023654

[24] L. Saloff-Coste and W. Woess, Transition operators, groups, norms, and spectral radii. *Pacific J. Math.* **180** (1997), no. 2, 333–367 Zbl 0899.60005 MR 1487568

[25] A. Sarveniazi, Explicit construction of a Ramanujan $(n_1, n_2, \ldots, n_{d-1})$-regular hypergraph. *Duke Math. J.* **139** (2007), no. 1, 141–171 Zbl 1180.11016 MR 2322678

Received 16 April 2021

M. Chapman, Einstein Institute of Mathematics, The Hebrew University of Jerusalem, Givat Ram, Jerusalem 9190401, Israel

E-mail: michael.chapman@mail.huji.ac.il

O. Parzanchevski, Einstein Institute of Mathematics, The Hebrew University of Jerusalem, Givat Ram, Jerusalem 9190401, Israel

E-mail: parzan@math.huji.ac.il