

Expander Graphs – Both Local and Global

Michael Chapman

*Einstein Institute of Mathematics
The Hebrew University of Jerusalem
Jerusalem, Israel*

Email: michael.chapman@mail.huji.ac.il

Nati Linial

*Department of Computer Science
The Hebrew University of Jerusalem
Jerusalem, Israel*

Email: nati@cs.huji.ac.il

Yuval Peled

*Courant Institute of Mathematical Sciences
NYU
New York, USA*

Email: yuval.peled@courant.nyu.edu

Abstract—Let $G = (V, E)$ be a finite graph. For $v \in V$ we denote by G_v the subgraph of G that is induced by v 's neighbor set. We say that G is (a, b) -regular for $a > b > 0$ integers, if G is a -regular and G_v is b -regular for every $v \in V$. Recent advances in PCP theory call for the construction of infinitely many (a, b) -regular expander graphs G that are expanders also locally. Namely, all the graphs $\{G_v | v \in V\}$ should be expanders as well. While random regular graphs are expanders with high probability, they almost surely fail to expand locally. Here we construct two families of (a, b) -regular graphs that expand both locally and globally. We also analyze the possible local and global spectral gaps of (a, b) -regular graphs. In addition, we examine our constructions vis-a-vis properties which are considered characteristic of high-dimensional expanders.

Keywords-Expander-Graphs, High-dimensional combinatorics

I. INTRODUCTION

It is hard to overstate the significance of expander graphs in theoretical computer science and the impact their study has had on a number of mathematical areas. Dinur's proof of the PCP Theorem, (e.g., [RS07]) is a prime example of their role in TCS. As the field develops, further refinements and extensions of the theory of expanders are called for. Thus, the recent breakthrough sampling algorithm of bases in a matroid by Anari, Liu, Oveis-Gharan and Vinzant [ALOGV19] builds on the newly emerging field of high-dimensional expanders. The present paper is motivated by Dinur and Kaufman's recent work in PCP theory [DK17], which needs graphs that are expanders both locally and globally. If v is a vertex in a graph G , the *link* of v in G , denoted G_v , is the subgraph of G that is induced by v 's neighbors. We seek large regular expanders G such that G_v is an expander for every $v \in V(G)$.

As observed already at the very early days of this area (e.g., [Pin73]), for every $d \geq 3$ asymptotically almost every d -regular graph is a very good expander. However, it is easy to verify that almost every d -regular graph is very far from satisfying the above requirement, as G_v is typically an anticlique. So, we ask: Given positive integers $a > b$ do there exist arbitrarily large (a, b) -expanders? Namely, a -regular expander graphs G such that every G_v is a b -regular expander. If so, how good can

the expansion properties (edge expansion, spectral gap) of G and the graphs G_v be? This last question can be made concrete in several ways, at an increasing order of difficulty, for given $a > b > 0$:

- 1) Can you construct infinitely many connected (a, b) -graphs?
- 2) Can you guarantee that all G_v be connected and that G is an expander?
- 3) Can you supply a bound on the spectral gap of all G_v ? These estimates should be bounded away from zero and expressed only in terms of a, b .
- 4) Can you, moreover, construct connected (a, b) -graphs where the second largest eigenvalue of every G_v is strictly smaller than $b/2$?

We have positive answers to the first three questions but not to the fourth one (see the open problem section). This is particularly intriguing because in case (4) *Garland's method* [Gar73] as elucidated by Oppenheim [Opp17] yields a spectral gap for G .

Theorem I.1 ([Opp17]). *Let G be a connected (a, b) -regular graph in which the second eigenvalue of every link is at most ϵb . Then G 's second eigenvalue is at most $\frac{\epsilon}{1-\epsilon}a$.*

Here we provide a purely combinatorial construction of graphs satisfying (3) that cannot satisfy (4). We also complement the Garland-type results, by proving (Theorem I.4) a tight Alon-Boppana bound on the second eigenvalue of any (a, b) -regular graph. Furthermore, we prove (Theorem I.5) bounds on the possible relations between the local and global spectral gaps.

Our investigations are closely related to the recently emerging study of high-dimensional expansion. The high-dimensional realm is not as well-behaved as the one-dimensional situation, where vertex-expansion, edge-expansion, spectral gaps and the convergence rate of the random walk mutually control each other quite tightly. Therefore, a number of inherently different ways to quantify high-dimensional expansion were suggested. We explore our

new (a,b) -expanders in light of these different measures of high-dimensional expansion.

Preliminaries, main results and organization

Let G be a graph and $v \in V(G)$. The link of v denoted G_v is the subgraph of G that is induced by the vertex set $\{u \in V \mid uv \in E\}$.

Definition I.2. Let $a > b \geq 0$ be integers. An (a,b) -regular graph G is an a -regular graph, where for every vertex $v \in V(G)$ the link G_v is b -regular.

We recall some basic terminology of expander graphs. Let G be a d -regular graph with adjacency matrix A_G , and let $d = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ be its eigenvalues. We say that G is an ε -spectral expander if its normalized spectral gap is at least ε , i.e., $1 - \frac{\lambda_2}{d} \geq \varepsilon$. We say that G is a δ -edge expander if $|E(U, V \setminus U)| \geq \delta \cdot \min(|U|, |V \setminus U|)$ for every $U \subseteq V(G)$, where $E(A, B)$ is the set of edges with one vertex in A and one in B . The largest such δ is called the edge-expansion (or Cheeger constant) of G . When we say below that G is an expander or an (a,b) -expander, we mean that G has some non-trivial but unspecified spectral gap resp. that G is (a,b) -regular and has some non-trivial but unspecified local and global spectral gaps. In the more technical parts of the paper we avoid such loose language and specify the relevant parameters.

Examples I.3. Here are some known families of (a,b) -regular graphs:

- 1) An n -clique is $(n-1, n-2)$ -regular and has good expansion properties, but a large $(n-1, n-2)$ -regular graph is the disjoint union of n -cliques, and clearly not a global expander.
- 2) The 2-dimensional Ramanujan complexes coming from $PGL_3(\mathbb{Q}_p)$ for prime p , are $(2p^2 + 2p + 2, p+1)$ -regular and are an example for a construction satisfying the Garland's method requirements (See [LSV05] and [GP19]). These graphs have many high dimensional expansion properties, see e.g., [DK17] and [EK16]. Note, however, that this is the **only** family of Ramanujan complexes whose 1-skeleton is (a,b) -regular for any a and b . Thus, additional constructions must be sought elsewhere, e.g., using combinatorial arguments as we do here.
- 3) The 1-skeleton of a non-singular, a -regular triangulation of a surface is $(a, 2)$ -regular. See Section V for more on this.
- 4) A construction due to Kaufman and Oppenheim of (a,b) -regular graphs that satisfy the Garland's method requirements can be found in [KO17].
- 5) Conlon's hypergraph expanders [Con17] are also (a,b) -expanders, but his a and b are unbounded.

In Section II we consider the largest possible spectral gap of an (a,b) -regular graph. We prove the following Alon-Boppana type bound. This bound is tight, but it makes no reference to local expansion.

Theorem I.4. The second eigenvalue of an (a,b) -regular graph satisfies

$$\lambda_2 \geq b + 2\sqrt{a-b-1} - o_n(1).$$

The bound is tight.

The full proof is in Section VII. Graphs for which Theorem I.4 holds with equality have disconnected links. Indeed, the situation changes when the links are expanders, or at least connected. As the next theorem shows, there is a tradeoff between local and global expansion. Recall the definition of the link G_v of v in G appearing in the first paragraph of the paper.

Theorem I.5. Let G be an (a,b) -regular graph where every link has edge expansion at least $\delta > 0$. Then G 's second eigenvalue satisfies:

$$\lambda_2 \geq \left(b + 2\sqrt{a-b-1}\right)(1 + \varepsilon) - o_n(1),$$

where $\varepsilon = \varepsilon(a,b,\delta) > 0$ is given explicitly. For $a \geq b^2 + O(b)$, ε strictly increases with δ . The same holds for all values of a and b , provided that δ is small enough.

In Section III we introduce our *Polygraph* constructions, which can be viewed as a new kind of graph products. These constructions transform a high-girth regular expander into an (a,b) -expander. For example, let $q > p \geq 0$ be integers, let G be a graph with distance function ρ and $\text{girth}(G) > 3p + 3q$. The vertex set of the *polygraph* G_S is $V(G)^3$ and (x_1, x_2, x_3) is a neighbor of (y_1, y_2, y_3) iff the multiset of three distances $[\rho(x_i, y_i) \mid i = 1, 2, 3]$ coincides with the multiset $[p, q, p+q]$.

For illustration, here is what this Polygraph looks like when $p = 0$ and $q = 1$. Take three copies of a d -regular triangle-free graph G and have a token move on each of them. At every step two of the tokens move to a neighboring vertex and the third token stays put. Any configuration of tokens is a vertex of the polygraph and the above process defines its adjacency relation.

Theorem I.6. Let $q > p \geq 0$ be even integers. If G is a regular, connected, non-bipartite graph of girth bigger than $3p + 3q$, then G_S is an (a,b) -regular local ε -spectral expander and global ε' -spectral expander. Here a and b depend on p, q and G 's regularity, ε can be bounded from below in terms of only p and q , and ε' has a closed formula that involves p, q and the spectral gap of G .

In Section III-D we investigate in detail the regularity and local spectral gaps of two specific Polygraph constructions.

In Section IV we examine Polygraphs from the perspective of high-dimensional expanders. We discuss properties such as geometric overlap, discrepancy, coboundary expansion and mixing of the edge-triangle-edge random walks. In Section V we provide some additional constructions of (a,b) -regular graphs, based on regular triangulations of surfaces and tensor products of graphs. We conclude the paper with some open questions related to this study.

II. THE SECOND EIGENVALUE OF (a,b) -REGULAR GRAPHS - LOWER BOUNDS

Proof sketch for Theorem I.4 and Theorem I.5: The proof relies on the moment method. An example of this approach is the proof of the Alon-Boppana bound in [HLW06] Section 5.2, which we now recall. Let G be a d -regular graph with adjacency matrix $A_G = A$ and eigenvalues $d = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$. For t a positive integer we note that

$$\text{trace}(A^t) = \sum \lambda_i^t \leq d^t + (n-1) \cdot \Lambda^t, \quad (\text{II.1})$$

where $\Lambda = \lambda(G) := \max\{\lambda_2, -\lambda_n\}$. On the other hand, $\text{trace}(A^t)$ is the number of closed walks of length t in G . This number can be bounded from below by counting length- t walks that start and end at some given *origin* vertex in G 's universal cover \mathbb{T}_d , the (infinite) d -regular tree. Associated with such a walk is a word in $\{F, B\}^t$, where F (resp. B) stands for a *forward* step away from the origin (*backward* step toward it). This word satisfies the *Catalan condition*, i.e., it has an equal number of B 's and F 's, and every initial segment has at least as many F 's as B 's. Also, B -steps are uniquely defined whereas every F -step can be realized in $d-1$ ways. By working out the number of such words, the Alon-Boppana bound $\Lambda \geq 2\sqrt{d-1} - o_n(1)$ is obtained.

We want to show that if G is (a,b) -regular, there are many more closed paths. To this end we introduce a *local system of coordinates* and keep track of a trail back to the origin (See Figure 1, where (z,x) is the last edge in this trail) that mimics the above association between words and closed paths. We distinguish between "sideways" steps to common neighbors of x and z and other "forward" steps, so our alphabet is $\{B, F, S\}$. To every word with x occurrences of F and B , and $t-2x$ occurrences of S , which satisfies the Catalan condition upon omitting the S 's, we can associate $(a-b-1)^x b^{t-2x}$ distinct closed paths of length $t+2$ in G . By optimizing over the choice of t and x , we conclude that $\lambda_2 \geq b + 2\sqrt{a-b-1} - o_n(1)$.

If the graph is not only (a,b) -regular, but also locally expanding, we can significantly increase the census of closed paths by finding many more local backwards steps. Whereas previously backward steps were uniquely defined, in the locally expanding case there are more FB combinations than before, which yield the required bound. ■

III. THE POLYGRAPH

In this section we construct infinite families of (a,b) -regular graphs with good local and global expansion properties. This construction is strictly combinatorial and is reminiscent of certain graph product operations. Theorem I.6 is proved at the end of this section.

Some terminology: If S is a multiset of k distinct non-negative integers that appear with multiplicities m_1, \dots, m_k and thus has a total of $m = \sum m_i$ members, we denote $S = [l_1, \dots, l_m]$. Let G be a d -regular graph with $d \geq 3$ and $\text{girth}(G) > 3 \max(S)$, and let ρ be its graph metric.

Definition III.1. *The Polygraph* $G_S = (V_S, E_S)$ has vertex set $V_S = V(G)^m$, and two vertices $\bar{x} = (x_1, \dots, x_m)$ and $\bar{y} = (y_1, \dots, y_m)$ in V_S are neighbors if and only if $[\rho(x_1, y_1), \dots, \rho(x_m, y_m)] = S$ as multisets.

The *distance profile* of any two vertices \bar{x} and \bar{y} in V_S is

$$\bar{\rho}(\bar{x}, \bar{y}) = (\rho(x_1, y_1), \dots, \rho(x_m, y_m)).$$

Thus, $\bar{x}\bar{y} \in E_S$ given $\bar{\rho}(\bar{x}, \bar{y}) = S$ as multisets. Conversely, if $\bar{\rho}(\bar{x}, \bar{y}) = (d_1, \dots, d_m)$, then $\bar{x}\bar{y} \in E_Z$, where $Z = [d_1, \dots, d_m]$.

Remark III.2. When $S' = [1, 0, \dots, 0]$, the polygraph $G_{S'}$ coincides with $G^{\square m}$, the m -th Cartesian power of G . If $S = [l_1, \dots, l_m]$ and $N = \sum_{i=1}^m l_i$, then G_S embeds in $G^{\square m}$, where every edge of G_S can be mapped to a (non-unique) length N non-backtracking path. Some of the proofs below involve polygraphs with the same graph G and two distinct multisets S_1, S_2 . In such situations, it is useful to embed both G_{S_1} and G_{S_2} into the same $G^{\square m}$.

Claim III.3. *The polygraph G_S is (a_S, b_S) -regular, where a_S and b_S depend only on S .*

A closed formula for b_S appears in the appendix, and is based on the arithmetic properties of S .

Claim III.4. *With the above notation, $b_S > 0$ iff there is a $3 \times m$ matrix, every row of which is comprised of the integers l_1, \dots, l_m in some order, where every column has an even sum and satisfies the triangle inequality.*

Proof: Necessity: Assume $\bar{x}, \bar{y}, \bar{z}$ is a triangle in G_S . Since G 's girth is large, the three geodesic paths connecting x_i, y_i and z_i form a tree, whence the sum of their lengths is even and the three lengths satisfy the triangle inequality. Hence, one can construct a matrix whose rows are the distance profiles of the edges of $\bar{x}\bar{y}, \bar{x}\bar{z}, \bar{y}\bar{z}$, meaning $\bar{\rho}(\bar{x}, \bar{y}), \bar{\rho}(\bar{x}, \bar{z})$ and $\bar{\rho}(\bar{y}, \bar{z})$. Sufficiency is not hard either: given three integers smaller than $\text{girth}(G)$ with even sum that satisfy the triangle inequality, there are three vertices in G the distances between which are these three integers. This allows us to construct $\bar{x}, \bar{y}, \bar{z}$ one coordinate at a time. ■

Remark III.5. Since our main interest is in connected polygraphs G_S with $b_S > 0$, we mostly restrict ourselves to the case where $N = \sum_{i=1}^m l_i$ is even and G is not bipartite.

We do not know how to efficiently test the condition in Claim III.4 for a given S , and suspect that it is hard in worst case. This is clearly no problem for small m . Here is, e.g., the solution for $m = 3$:

Claim III.6. Let S be the multiset of integers $p, q, r \geq 0$. Then $b_S \neq 0$ if and only if (i) p, q, r are all even, or (ii) Their sum is even and they satisfy the triangle inequality.

A. Non-backtracking paths

Let G be a graph with adjacency matrix A_G and let $A_G^{(t)}$ be the matrix whose (i, j) -th entry is the number of length- t non-backtracking paths between vertices i and j in G . We also view $A_G^{(t)}$ as the adjacency matrix of a multigraph $G^{(t)}$. When G is d -regular, $G^{(t)}$ is $d(d-1)^{t-1}$ -regular. In this case these matrices satisfy the following recursion:

$$\begin{aligned} A_G^{(1)} &= A_G ; A_G^{(2)} = A_G^2 - dI_n \\ A_G^{(t+1)} &= A_G A_G^{(t)} - (d-1)A_G^{(t-1)}. \end{aligned}$$

The polynomial $p^{(t)}(\cdot)$ that satisfies $p^{(t)}(A_G) = A_G^{(t)}$ is sometimes called the t -th *Geronimus polynomial*. For more information see [Sol92], [ABLS07] or [DSV03].

Lemma III.7. If G is a connected, non bipartite graph with minimum vertex degree at least 3, then for every t , $G^{(t)}$ is a connected non bipartite graph.

B. Connectivity and Spectral analysis of G_S

Let $\Omega = \binom{[m]}{m_1, \dots, m_k}$ be the set of rearrangements $\omega = (\omega_1, \dots, \omega_m)$ of $S = [l_1, \dots, l_m]$. Thus

$$A_{G_S} = \sum_{\omega \in \Omega} \bigotimes_{j=1}^m A_G^{(\omega_j)}$$

where \otimes is the Kronecker tensor product.

If v is an eigenvector of A with eigenvalue λ , likewise for v', A' and λ' , then $v \otimes v'$ is an eigenvector of $A \otimes A'$ with eigenvalue $\lambda\lambda'$. Also, $A_G^{(t)}$ and A_G have the same eigenvectors, since $A_G^{(t)}$ is a polynomial in A_G . It follows that every eigenvector of A_{G_S} has the form $v_1 \otimes \dots \otimes v_m$ where each v_i is an eigenvector of A_G . Moreover, by going through all such choices of v_1, \dots, v_m we obtain the full list of eigenvectors. The eigenvalue of $v_1 \otimes \dots \otimes v_m$ is

$$\chi(\lambda_1, \dots, \lambda_m) = \chi_S(\lambda_1, \dots, \lambda_m) = \sum_{\omega \in \Omega} \prod_{j=1}^m p^{(\omega_j)}(\lambda_j)$$

where λ_i the eigenvalue of v_i , and $p^{(t)}(x)$ is the t -th Geronimus polynomial mentioned above. We obtain all the eigenvalues of

A_{G_S} by evaluating the symmetric polynomial χ_S on all m -tuples of eigenvalues $(\lambda_1, \dots, \lambda_m)$.

Claim III.8. Let G be a connected non bipartite d -regular graph and S a multiset of non-negative integers, not all zero. Then G_S is connected and non bipartite.

We now provide stronger bounds on G_S 's spectral gap under appropriate assumptions on G :

Lemma III.9. Let $\alpha \in (-2\sqrt{d-1}, 2\sqrt{d-1})$ and $|\beta| > 2\sqrt{2}\sqrt{d-1}$. Then $|p^{(t)}(\beta)| > |p^{(t)}(\alpha)|$, where $p^{(t)}$ is the t -th Geronimus Polynomial. Also, if $|x| \geq 2\sqrt{d-1}$, then $|p^{(t)}(x)| < |x|^t$.

Therefore, if $\lambda(G) \geq 2\sqrt{2}\sqrt{d-1}$, then we have

$$\begin{aligned} \lambda(G_S) &= \chi(\lambda(G), d, \dots, d) \\ &= \sum_{\omega \in \Omega} p^{(\omega_1)}(\lambda(G)) \prod_{j=2}^m p^{(\omega_j)}(d) \\ &= \sum_{i=1}^m p^{(l_i)}(\lambda(G)) \cdot \frac{(m-1)! \cdot m_i}{m_1! \dots m_k!} \prod_{j \neq i} [d(d-1)^{l_j-1}]. \end{aligned}$$

This formula is nice, but we are interested in a more practical bound on $\lambda(G)$. Thus the following:

Proposition III.10. Let G be a d -regular graph and let S be a multiset of non-negative integers, the smallest of which is s , whose sum is N . Then, letting $\mu = \max(\lambda(G), 2\sqrt{d-1})$, we have

$$\lambda(G_S) \leq \binom{m}{m_1, \dots, m_k} \mu^s d^{k-1} (d-1)^{N-k-s+1}.$$

C. Local Connectivity of G_S

We start with a necessary condition for $L = (G_S)_v \cong ((\mathbb{T}_d)_S)_{(\xi, \dots, \xi)}$ to be connected, where ξ is some fixed vertex in the d -regular tree \mathbb{T}_d .

Lemma III.11. If L is connected then either (i) $0 \in S$, or (ii) there is a positive $s \in S$ such that $2s \in S$ as well, or (iii) there are three distinct $s, s', s'' \in S$ satisfying $s'' = s + s'$.

The following two claims give a necessary and sufficient condition for connectivity for $m = 2, 3$.

Claim III.12. Let $S = [p, q]$, where $q \geq p$ and $q > 0$. Then, L is connected if and only if p is even and $q = 2p$.

Claim III.13. Let $S = [p, q, r]$, where $p \leq q \leq r$ and $0 < r$. Then L is connected if and only if $p + q + r$ is even and one of the following holds:

- 1) $r = p + q$ and either: p is even, or q is even and $2p \geq q$, or $p = q$;
- 2) $q = 2p$ and $r \leq p + q$, or $r = 2p$;
- 3) p, q and r are even, $4p \geq 2q \geq r$ and either $r = 2p$ or $r = 2q$ or $q = 2p$.

There is also a criterion for the connectivity of L when all members of S are even. We omit the (very technical) details. One can deduce Theorem 1.6 from this entire section, in particular Claim III.13 and Lemma III.10.

D. Some Concrete Explicit Constructions

1) *The case $S = [1, 1, 0]$:* Recall the concrete description of G_S from the introduction: Take three copies of a d -regular non-bipartite triangle-free graph G and have a token move on each of them. At every step two of the tokens move to a neighboring vertex and the third token stays put. The resulting graph G_S is $(3d^2, 2d)$ -regular, it is connected (Claim III.8) and has connected links (Claim III.13). For $v_0 \in V(G)$, we turn to study the spectrum of L – the link of (v_0, v_0, v_0) in G_S . The graph L is tripartite with parts V_1, V_2, V_3 , where $V_1 = \{(v_0, v_i, v_j) \mid 1 \leq i, j \leq d\}$, and likewise for V_2, V_3 . Edges between V_2 and V_3 are defined via $(v_i, v_j, v_0) \sim (v_{i'}, v_0, v_{j'})$ iff $i' = i$. The two other adjacency conditions are similarly defined.

Lemma III.14. *The eigenvalues of L are $2d, d, 0, -d$ with corresponding multiplicities $1, 3(d-1), 3(d-1)^2, 3d-1$.*

2) *The case $S = [1, 2, 3]$:* Here G_S is $(6d^3(d-1)^3, 2(d-1)^2(4d-7))$ -regular. It is connected and so are its links. It also has interesting spectral properties, since by Proposition III.10, if G is Ramanujan, then $\lambda(G_S) \leq 12d^2(d-1)^{7/2}$. Actually, a similar conclusion can be drawn whenever G has a substantial spectral gap and S contains no zeros.

IV. THE PERSPECTIVE OF HIGH DIMENSIONAL EXPANSION

The study of (a, b) -regular graphs can be cast in the language of simplicial complexes. Let us recall some basic facts from that theory. Let X be a simplicial complex, i.e., a collection of finite sets that is closed under inclusion. Namely if $\sigma \in X$ and $\tau \subseteq \sigma$, then $\tau \in X$. The singletons in X are called vertices. An element $\sigma \in X$ is called a k -dimensional face if $|\sigma| = k + 1$. The *link* of a face σ in X is the following simplicial complex:

$$X_\sigma = \{\tau \in X \mid \sigma \cap \tau = \emptyset, \sigma \cup \tau \in X\}.$$

The i -th *skeleton* $X^{(i)}$ of X is the simplicial complex that is comprised of all faces of X of dimension less or equal to i . Associated with a graph $G = (V, E)$ is its *clique complex* \mathcal{C}_G , whose vertex set is V and $S \subseteq V$ is a face of it if and only if S spans a clique in G . Hence G is (a, b) -regular if and only if the 1-skeleton $\mathcal{C}_G^{(1)}$ is a -regular and the link of every vertex $v \in \mathcal{C}_G^{(2)}$ is a b -regular graph. If G is an (a, b) -regular graph, then the number of 2-dimensional faces in $\mathcal{C}_G^{(2)}$ is $\frac{abn}{6}$.

This section contains both negative and positive results. The negative results are mainly about the $[1, 1, 0]$ -polygraphs and the positive ones are about $[1, 2, 3]$ -polygraphs.

There is a considerable body of research, mostly quite recent on expansion in high dimensional simplicial complexes.

Several different ways were proposed to quantify this notion. For the definitions of cosystolic and coboundary expansion, see e.g., [EK16].

A. Discrepancy

Every $[1, 1, 0]$ -polygraph $G_{[1,1,0]}$ has poor discrepancy, since it has two sets, each with $\frac{1}{8}$ of its vertices and no edges between them. Namely, if $A \subseteq V(G)$ contains a half of G 's vertices, then no edge connects A^3 and $(A^c)^3$. A similar construction can be given whenever S contains a zero. In contrast, $[1, 2, 3]$ -polygraphs exhibit better discrepancy properties, and in particular have the geometric overlap property (see IV-C and IX-A below).

B. Coboundary expansion

This part is inspired by ideas of Luria, Gundert and Rosenthal (e.g., Section 3 of [Cha]). As they show, Conlon's hypergraph [Con17] contains small non-trivial cocycles and thus is not a cosystolic expanders and a fortiori not a coboundary expander either.

Again let G be non-bipartite triangle-free and d -regular, and $\Gamma = G_{[1,1,0]}$. We exhibit a set $A \subseteq E(\Gamma)$ such that:

- 1) Every triangle in Γ has exactly two edges from A ;
- 2) A is not a cut in Γ .

It follows that the 2-skeleton of Γ 's clique complex¹ has a non-trivial first \mathbb{F}_2 -cohomology and is thus not a coboundary expander.

The distance profile of every edge in Γ is one of three: $(0, 1, 1)$, $(1, 1, 0)$ or $(1, 0, 1)$, and every triangle has exactly one edge of each kind. The set A of those edges with profile $(1, 1, 0)$ or $(1, 0, 1)$ clearly satisfies condition (1). To show condition (2) we find an odd cycle in the graph $(V(\Gamma), A)$. Since G is non-bipartite, it has an odd cycle, say v_1, \dots, v_ℓ, v_1 . But then

$$(v_1, v_1, v_1), (v_2, v_2, v_1), \dots, (v_\ell, v_\ell, v_1), (v_1, v_1, v_1)$$

is an odd cycle in $(V(\Gamma), A)$.

We note that this argument fails for $S = [1, 2, 3]$. On the other hand, this argument does work for $S = [1, 1, 2]$, showing that even a zero-free S need not yield coboundary expansion.

C. Geometric overlap property

We say that a 2-dimensional simplicial complex X has the α -geometric overlap property if the following holds: For every embedding of $V(X) \rightarrow \mathbb{R}^2$ with the induced affine extension to X 's edges and faces there is a point in \mathbb{R}^2 that meets at least an α -fraction of the images of X 's 2-faces. Work in this section and the following parallels [Con17].

¹The clique complex of a graph $G = (V, E)$ has vertex set V and $S \subseteq V$ is a face iff it spans a clique in G .

Theorem IV.1. *Let G be a d -regular graph with $\text{girth}(G) > 9$, $d > d_0$ and $\lambda(G) < \varepsilon_0 d$. Then $G_{[1,2,3]}$ has the α_0 -geometric overlap property. Here $d_0, \varepsilon_0, \alpha_0 > 0$ are absolute constants.*

D. Mixing of the edge-triangle-edge random walk

Random walks play a key role in the study of expander graphs, and similar questions are being studied in the high-dimensional realm as well, e.g., [KM16], [LLP17], [DK17] and [Con17]. Consider the following random walk on the one-dimensional faces (i.e., edges) of a 2-dimensional simplicial complex X . We move from an edge $e \in X^{(1)}$ to an edge that is chosen uniformly among all edges $f \in X^{(1)}$ with $f \cup e \in X^{(2)}$ (a triangle in X). We wish to decide for simplicial complexes of interest if this walk mixes rapidly. Differently stated, this is a walk on $\text{Aux}(X)$, a graph with vertex set $X^{(1)}$, where ef is an edge iff $e \cup f \in X^{(2)}$. Given a multiset S , we consider $\text{Aux}(X)$ for $X = \mathcal{C}_{G_S}$, the clique complex of G_S .

Theorem IV.2. *The edge-triangle-edge random walk on $\mathcal{C}_{G_{[1,1,0]}}$ mixes rapidly, i.e. in $O(\log |V(G)|)$ time.*

V. FOR WHICH (a, b) DO LARGE (a, b) -REGULAR GRAPHS EXIST?

This section provides a partial answer to the question in its title. If $b = 2$ and the links are connected, then every link is a cycle. So, the graph in question is the 1-skeleton of a triangulated 2-manifold. It is instructive to consider the case $a = 6$. The Cayley graph of \mathbb{Z}^2 with generators $(\pm 1, 0), (0, \pm 1), \pm(1, 1)$ is the planar triangular grid. The quotient of this graph mod $m\mathbb{Z} \times n\mathbb{Z}$, is a $(6, 2)$ -regular finite triangulation of the torus whose links are connected. We ask for which values of a there exist infinitely many such graphs.

This discussion is closely related to the study of equivelar polyhedral 2-manifolds and non-singular $\{p, q\}$ -patterns on surfaces, a subject on which there exists a considerable body of literature. We only mention [MSW82] and [MSW83] where infinitely many such graphs for $a \geq 7$ are constructed. Some of these constructions are inductive and start from the above triangulations of the torus. Other constructions are iterative and use snub polyhedra of prisms.

We recall that the *tensor product* $G \otimes H$ of two graphs G and H , is a graph with vertex set $V(G) \times V(H)$ where (u, v) and (u', v') are neighbors when both $uu' \in E(G)$ and $vv' \in E(H)$. Therefore its adjacency matrix is the Kronecker tensor product $A_G \otimes A_H$. Note that $G \otimes G$ is isomorphic to $G_{[1,1]}$. Also, if G is (a, b) -regular and H is (a', b') -regular, then $G \otimes H$ is (aa', bb') -regular.

Thus, if G is $(k, 2)$ -regular, then $G \otimes K_m$ is $(k(m-1), 2(m-2))$ -regular. This yields arbitrarily large $(k(m-1), 2(m-2))$ -regular graphs with connected links. This means that the

question in the title is answered positively for many $a > b$ and for all asymptotic relations between a and b .

VI. OPEN QUESTIONS

Countless questions suggest themselves in this new domain of research. We mention below three which we view as the most attractive.

- 1) **A randomized model:** One of the earliest discoveries in the study of expander graphs is that in essentially every reasonable model of random graphs, and in particular for random d -regular graphs, almost all graphs are expanders. It would be very interesting to find a randomized model of (a, b) -regular graphs and in particular one where most members are expanders both locally and globally.
- 2) **Higher-dimensional constructions:** We have touched upon the connections of our subject with the study of expansion in higher-dimensional simplicial complexes. Clearly, (a, b) -regularity is a two-dimensional condition, and we know essentially nothing for higher dimensions. Concretely: do (a, b, c) -regular graphs exist? Namely for fixed $a > b > c > 1$, we ask whether there exist arbitrarily large a -regular graphs where the link of every vertex is b -regular, and the link of every edge is c -regular. We want, moreover, that the whole graph, every vertex link and every edge link be expanders. We stress that no such constructions based on Ramanujan complexes [LSV05] are presently known. There are indications that the situation in dimension two is less rigid than in higher dimensions. Does this translate to some non-existence theorems?
- 3) **Garland's method** [Gar73] is a powerful tool in the study of high-dimensional expansion, e.g., [Opp17], [Pap16]. In order to apply the method for an (a, b) -regular graph G , it needs to have the property that the spectrum of every vertex link is contained in $\{-b, b\} \cup [-\beta, \beta]$ for some $\beta < \frac{b}{2}$. In such case, Garland's method asserts that G is also a global expander. However, some substantial new ideas will be needed to construct such an (a, b) -regular graph using only combinatorial arguments. For instance, polygraphs cannot have this property. Indeed, compare what happens when we start from a d -regular graph G that is a very good expander vs. a very bad one. While G_S inherits G 's expansion quality, the links of the two graphs are identical.
- 4) **Trade off:** The lower bound on λ_2 from Theorem I.5 is an increasing function of δ . However, we do not know how tight this bound is and whether the *best possible* lower bound on λ_2 increases with δ . Theorem I.4 is tight, so the best bounds for $\delta = 0$ and $\delta > 0$ differ.

But whether the same holds as $\delta > 0$ increases, we do not know.

5) Is there an efficient way of checking whether $b_S \neq 0$ for every S ?

VII. APPENDIX A: COMPLETE PROOFS FOR THEOREM I.4 AND THEOREM I.5

Proof of Theorem I.4: This theorem is reminiscent of the Alon-Boppana Theorem. We are inspired by the proof of that theorem via the moment method (e.g., [HLW06] Section 5.2). Let G be a d -regular graph with adjacency matrix $A_G = A$ and eigenvalues $d = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$. For t a positive integer we note that

$$\text{trace}(A^t) = \sum \lambda_i^t \leq d^t + (n-1) \cdot \Lambda^t, \quad (\text{VII.1})$$

where $\Lambda = \lambda(G) := \max\{\lambda_2, -\lambda_n\}$. On the other hand, $\text{trace}(A^t)$ is the number of closed walks of length t in G . This number can be bounded from below by counting length- t walks that start and end at some given *origin* vertex in G 's universal cover \mathbb{T}_d , the (infinite) d -regular tree. Associated with such a walk is a word in $\{F, B\}^t$, where F (resp. B) stands for a *forward* step away from the origin (*backward* step toward it). This word satisfies the *Catalan condition*, i.e., it has an equal number of B 's and F 's, and every initial segment has at least as many F 's as B 's. Also, B -steps are uniquely defined whereas every F -step can be realized in $d-1$ ways. By working out the number of such words, the Alon-Boppana bound $\Lambda \geq 2\sqrt{d-1} - o_n(1)$ is obtained.

En route to a proof there are two obvious obstacles:

At the origin, there are d possible F steps. (VII.2)

We soon address this point.

A closed walk's length in \mathbb{T}_d is even, but we use an odd t . (VII.3)

The advantage of having an odd t is that the term Λ^t in (VII.1) can be replaced by the possibly smaller λ_2^t .

Our proof needs a modified notion of forward and backward steps and also allow for *sideways* step. We consider length- t *Catalan words* in the alphabet $\Sigma = \{F_j | j = 1, \dots, a-b-1\} \cup \{B\} \cup \{S_i | i = 1, \dots, b\}$. Namely, a word with an equal number of F 's and B 's where $\#F \geq \#B$ in each initial segment. We wish to injectively associate to each such word a closed walk in our graph. Roughly speaking, when the next letter in the word is F_j we should move to the j -th *forward neighbor* of our current position, likewise move to the i -th *sideways neighbor* upon reading S_i , and finally moving one step backward on a B .

In \mathbb{T}_d it is perfectly clear what forward and backward mean and sideways does not exist. As we explain next, we navigate a general graph using a *local system of coordinates*. To this end we use a stack X in which we store vertices, where every




Figure 1: The local system of coordinates of x , at which the walk currently resides, with respect to its neighbor z , that is currently at the top of the stack.

two consecutive entries in X are two adjacent vertices in G . An invariant that we maintain is that x , the vertex at which the walk currently resides is always a neighbor of **top**, the top entry of X . Suppose that we move next from x to a neighbor y .

- If y is not a neighbor of **top**, this is a *forward* step, and we *push* x .
- If $y = \mathbf{top}$, this is a *backward* step and we *pop*.
- If y is a neighbor of **top**, this is a *sideways* step, and the stack stays unchanged.

It remains to define which is x 's j -th forward (resp. sideways) neighbor. This choice is not absolute, but rather depends on the current **top**: Consider two neighbors x and z in G , where we think of z as the current **top**, and x as our current position. We fix some ordering on the set $\Phi_{x,z}$ of the $a-b-1$ neighbors of x that are not neighbors of z and an ordering on the set $\Psi_{x,z}$ of the b joint neighbors of x and z . Thus, if we are currently at x , and z is at the top of the stack, we interpret the symbol F_j as “move to the j -th vertex in $\Phi_{x,z}$ ”. Likewise S_i means “move from x to the i -th vertex in $\Psi_{x,z}$ ” and B means “step from x to z ”. In other words, $\Psi_{x,z} = G_x \cap G_z$, whereas $\Phi_{x,z} = G_x \setminus (G_z \cup \{z\})$. (See Figure 1).

Given a starting vertex v and a length- t Catalan-word ω over the alphabet Σ , we will specify a closed walk of length $t+2$ that starts and ends at v . Before we do that we need to deal with the issue raised in (VII.2) above. We associate with every vertex v one of its neighbors ϕ_v . We start our walk at v , then move to ϕ_v and push v on the stack. Henceforth we follow the transitions that are dictated by ω and the push/pop rules described above. Since $\#F_\omega = \#B_\omega$, when we are done reading ω the stack contains only the symbol v , and we therefore reside at a neighbor of v . We now empty the stack and move to v . This clearly associates injectively a closed path as described

with every pair (v, ω) for v a vertex and ω a Catalan word.

We have thus shown that the number of closed walks of length $t+2$ in G is at least

$$n \sum_{0 \leq k < \frac{t}{2}} \binom{t}{k, k, t-2k} \frac{1}{k+1} b^{t-2k} (a-b-1)^k. \quad (\text{VII.4})$$

Here n counts the choices of the starting vertex v . The trinomial coefficient counts words ω with $\#F_\omega = \#B_\omega = k$ and $\#S_\omega = t-2k$. The term $\frac{1}{k+1}$ accounts for the probability that the Catalan Condition holds. Finally every F -step can be indexed in $a-b-1$ ways and every S -step in b ways.

There are only $O(t)$ terms in this sum whereas the largest term is exponential in t , so it suffices to determine the largest term in the sum. To this end we express $k = \alpha t$, and then we need to find the α that maximizes the expression²

$$\mathcal{H}(\alpha, \alpha, 1-2\alpha) + \alpha \log(a-b-1) + (1-2\alpha) \log b, \quad (\text{VII.5})$$

where \mathcal{H} is the entropy function with base 2. Straightforward calculation yields that the maximum is $\log(b+2\sqrt{a-b-1})$ which is attained for $\alpha = \frac{\sqrt{a-b-1}}{b+2\sqrt{a-b-1}}$. When we return to (VII.1), the best lower bound on λ_2 is attained for $t \approx \frac{\log n}{\log d - \log(\lambda_2)}$ and yields

$$\lambda_2 \geq b + 2\sqrt{a-b-1} - o_t(1)$$

as claimed.

We now prove that this bound is tight. Let $H = (L, R, F)$ be a connected, bipartite, left c -regular and right d -regular graph with $\text{girth}(H) \geq 8$. Associated with H is the $c(d-1)$ -regular graph $G = (L, E)$, where $xy \in E$ if and only if there is a vertex $z \in R$ such that $xz, yz \in F$. Note that every link in G is a $(d-2)$ -regular graph which is the disjoint union of c graphs each of which a $(d-1)$ -clique. All told, this is a construction of $(c(d-1), d-2)$ -regular graphs. For example, here is a concrete family of bipartite graphs H as above with $c=2$. Let Γ be a d -regular triangle-free graph, and let $L = E(\Gamma), R = V(\Gamma)$ and F the vertex-edge incidence relation of Γ .

Of course, the links in this graph are not expanders - they are not even connected. It is easy to see that the adjacency matrix of G is a block in $A_H^2 - cI$. If H is a (c, d) -biregular bipartite Ramanujan graph (see Section 2.3 in [MSS15]), then $\lambda(G) \leq \sqrt{c-1} + \sqrt{d-1}$. Thus

$$\begin{aligned} \lambda(G) &\leq (\sqrt{c-1} + \sqrt{d-1})^2 - c \\ &= d - 2 + 2\sqrt{c(d-1)} - (d-2) - 1 \end{aligned}$$

showing that the bound is tight. \blacksquare

Proof of Theorem 1.5:

To start, we improve the lower bound on λ_2 in Theorem 1.4 when each G_v is *connected*. In those cases where Theorem 1.4 gives a tight bound, our census of closed walks is complete.

²Logarithms here are to base 2, unless otherwise stated.




Figure 2: The assumption that G_x is connected guarantees that there is an edge uy between $\Phi_{x,z}$ and $\Psi_{x,z}$. The walk $x \rightarrow u \rightarrow y$ is now considered an FB move rather than an FS move as in Theorem 1.4. As a result, the vertex z , and not x , is on the top of the stack when the walk reaches y . The better an expander graph G_x is, the more edges there are between $\Phi_{x,z}$ and $\Psi_{x,z}$.

However, as we soon observe, when all the links are connected, many additional closed walks emerge. To maintain the overall structure of the proof F, B and S steps still go with push, pop and no change to the stack, but they need no longer reflect the distance from the origin.

Given an initial vertex v and a word in B, F, S (with appropriate indices) again we associate to these data a walk in G that starts and ends in v . However, the correspondence is now somewhat different. Suppose that the walk currently resides at the vertex x , its neighbor z is at the top of the stack, and the coming two letters are FB in this order. Because the link of x is connected, there must be an edge between some vertex $u \in \Phi_{x,z}$ and some vertex $y \in \Psi_{x,z}$. Say that we realize the F -step by moving from x to u . After this move $\mathbf{top} = x$ and the penultimate entry in the stack is z . In the proof of Theorem 1.4 the coming B -step is realized now by moving back to x , and popping x , making $\mathbf{top} = z$. But because u has a neighbor $y \in \Psi_{x,z}$, we can also move from u to y and pop x while respecting the structure of the proof. In other words, now we can and will consider the transitions $x \rightarrow u \rightarrow y$ as realizing the subword FB rather than FS (see Figure 2).

To complete the details, we place u first in the ordering of $\Phi_{x,z}$, and y first in $\Psi_{u,x}$. We interpret each subword $F_i B$ for $i = 1, \dots, a-b-1$ (including $F_1 B$) as before. However, we allow as well the subword $F_1 B$ to which we associate the transitions $x \rightarrow u \rightarrow y$. The same applies to subwords $F_i S_j$ which we interpret as usual. However, we forbid the subword $F_1 S_1$ to avoid overcounting the walk $x \rightarrow u \rightarrow y$.

This change affects the census in Theorem 1.4. A subword FB has now $a-b$ rather $a-b-1$ realizations, whereas for

FS the count goes down from $(a-b-1)b$ to $(a-b-1)b-1$. As the next calculation shows, the gain outweighs the loss, yielding a better lower bound on λ_2 .

Clearly there are $(\beta^{2\beta}(1-2\beta)^{1-2\beta})^{-t(1+o_t(1))}$ length t words in the alphabet $\{F, B, S\}$ with βt letters F and B and $(1-2\beta)t$ letters S . Standard concentration-of-measure inequalities show that with a proper choice of the $o_t(1)$ terms, the same asymptotic counts remains even if we insist that:

- The Catalan condition for F and B holds
- Every pair of consecutive letters appears the "right" number of times. E.g., the number of FB, FS, SS subwords is $(1+o_t(1))\beta^2 t$, resp. $(1+o_t(1))\beta(1-2\beta)t$, and $(1+o_t(1))(1-2\beta)^2 t$, etc.

For every such word, we compute the number of permissible ways to index the F -steps and the S -steps as following:

- A letter F that is followed by an F can be indexed in $a-b-1$ ways.
- A letter S that is not preceded by an F can be indexed in b ways.
- A pair of consecutive letters FB can be indexed in $a-b$ ways.
- A pair of consecutive letters FS can be indexed in $(a-b-1)b-1$ ways.

In summary, we seek to maximize

$$\begin{aligned} \mathcal{H}(\beta, \beta, 1-2\beta) + \beta^2 \log(a-b-1) \\ + (1-2\beta)(1-\beta) \log b + \beta^2 \log(a-b) \\ + (\beta-2\beta^2) \log(b(a-b-1)-1). \end{aligned}$$

Write $\log(b(a-b-1)-1) = \log b + \log(a-b-1) + \log(1 - \frac{1}{b(a-b-1)})$ and $\log(a-b) = \log(a-b-1) + \log(1 + \frac{1}{a-b-1})$ to conclude that instead of the analysis of Equation (VII.5) we now seek β that maximizes

$$\begin{aligned} S(\beta, a, b) = \mathcal{H}(\beta, \beta, 1-2\beta) \\ + \beta \log(a-b-1) + (1-2\beta) \log b + \Delta, \end{aligned} \quad (\text{VII.6})$$

where

$$\begin{aligned} \Delta = \beta^2 \log \left(1 + \frac{1}{a-b-1} \right) \\ + \beta(1-2\beta) \log \left(1 - \frac{1}{b(a-b-1)} \right). \end{aligned}$$

We now prove that $\max_{\beta} S(\beta, a, b) > \log(b + 2\sqrt{a-b-1})$ whenever $a-b \geq 3$ and $b \geq 2$. The proof for complementary parameter range follows from by observing that G is necessarily comprised of disjoint copies of the same graph H and is therefore not even connected. If $b = 0, 1$ and $\delta > 0$, then H is a triangle. When $a-b = 1, 2$ the same holds with $H = K_{a+1}$ and $H = (K_{a+2} \text{ minus a perfect matching})$ respectively.

We denote $c = \sqrt{a-b-1}$. The values of

$$\frac{(b+2c)^2}{\log(e)} \left(\max_{\beta} S(\beta, a, b) - \log(b+2c) \right)$$

for small a and b 's are shown in the following table:

$a-b-1 \setminus b$	2	3	4	5	6
2	0.062	0.08	0.088	0.092	0.094
3	0.281	0.287	0.29	0.29	0.291
4	0.397	0.401	0.402	0.402	0.402
5	0.472	0.474	0.475	0.475	0.475
6	0.525	0.527	0.527	0.527	0.527
7	0.565	0.567	0.567	0.567	0.567
8	0.597	0.598	0.599	0.599	0.598

Consequently, in proving our statement we can ignore the case where both b and c are small. We do not find a closed form expression for β that maximizes (VII.6). Instead we let $\beta := \frac{c}{b+2c}$, and show that $S(\frac{c}{b+2c}, a, b) > \log(b+2c)$. With this choice of β there holds $\frac{1-2\beta}{\beta} = \frac{b}{c}$, so that

$$\Delta = \frac{\beta^2}{c} \left(c \cdot \log \left(1 + \frac{1}{c^2} \right) + b \cdot \log \left(1 - \frac{1}{b \cdot c^2} \right) \right).$$

It is easily seen that for $c > 0$ fixed, this expression is an increasing function of b , whence it suffices to verify that $\Delta > 0$ when $b = 2$. Using Taylor expansion it is easily verified that this inequality holds already for $c > 1.5$. The same analysis yields that for every $b \geq 2$ and large c there holds $\Delta = \frac{\log(e) - o_c(1)}{(b+2c)^2}$. It follows that if all G_v are connected then

$$\lambda_2(G) \geq \left(b + 2\sqrt{a-b-1} \right) \left(1 + \Omega \left(\frac{\log e}{(b+2c)^2} \right) \right) - o_n(1)$$

as claimed.

We turn to consider what happens when the graphs G_v expand. In this case, for every two adjacent vertices x, z there are some edges between the sets $\Psi_{x,z}$ and $\Phi_{x,z}$, where, as above, $\Psi_{x,z} := G_x \cap G_z$, and $\Phi_{x,z} := G_x \setminus (G_z \cup \{z\})$. Let R be the least number of such edges over all $xz \in E(G)$. Hence, by the definition of edge expansion, $R \geq \min(b+1, a-b-1) \cdot \delta$. Under the assumption that all G_v are connected we pick one edge uv with $u \in \Phi_{x,z}$ and $v \in \Psi_{x,z}$ and create a special forward step denoted by F_{1*} . We interpret the subword $F_{1*}B$ as an instruction to move $x \rightarrow u \rightarrow y$ and maintaining z on the top of the stack. In addition, we forbid the subword $F_i S_j$ where u is the i -th vertex in $\Phi_{x,z}$ and y is the j -th vertex in $\Psi_{u,x}$. In the present context we can likewise consider some $r \leq R$ edges $u_k v_k$, $k = 1, \dots, r$, with $u_k \in \Phi_{x,z}, v_k \in \Psi_{x,z}$. Associated with them we create r types of forward steps called $F_{1*}, F_{2*}, \dots, F_{r*}$ and associate with the subword $F_{k*}B$ the move $x \rightarrow u_k \rightarrow y_k$, while z

stays on the top of the stack. In addition, we forbid subwords of the form $F_{i_k}S_{j_k}$, where u_k is the i_k -th vertex in $\Phi_{x,z}$ and y_k is the j_k -th vertex in $\Psi_{u_k,x}$. This works for any choice of $r \leq R$ such edges. Now a pair of consecutive letters FB can be indexed in $a-b+r-1$ ways but a pair of consecutive letters FS can be indexed only in $(a-b-1)b-r$ ways. This yields the same maximization problem of (VII.6) with the correction term

$$\Delta = \beta^2 \log \left(1 + \frac{r}{a-b-1} \right) + \beta(1-2\beta) \log \left(1 - \frac{r}{b(a-b-1)} \right).$$

By letting $\beta := c/(b+2c)$ as before this reformulates as

$$\Delta = \frac{\beta^2}{c} \left(c \cdot \log \left(1 + \frac{r}{c^2} \right) + b \cdot \log \left(1 - \frac{r}{bc^2} \right) \right).$$

Straightforward calculations show that the value of r that maximizes this expression is $\frac{b(c^3-c^2)}{b+c}$. So, we let $r := \min \left(R, \lceil \frac{b(c^3-c^2)}{b+c} \rceil \right)$. In this case, when c is large, we get $\Delta = \frac{r \log(e) - o_c(1)}{(b+2c)^2}$ and

$$\lambda_2(G) \geq \left(b + 2\sqrt{a-b-1} \right) \left(1 + \Omega \left(\frac{r \log(e)}{(b+2c)^2} \right) \right) - o_n(1)$$

completing the proof. \blacksquare

Remark VII.1. If $\delta < \frac{b(a-b-1)(\sqrt{a-b-1}-1)}{(b+\sqrt{a-b-1})\min(b+1,a-b-1)}$, then ε increases with δ . Note that if $a \geq b^2 + 5b + 5$, this restriction on δ is vacuous and ε is always increasing, since the edge expansion of a b -regular graph cannot exceed $\frac{b}{2}$.

VIII. APPENDIX B: REGULARITY OF THE LINKS

Recall that Ω is its set of all the arrangements of the multiset $S = [l_1, \dots, l_m]$. For a positive integer i , define $f_i: \mathbb{Z}_{\geq 0}^2 \rightarrow \mathbb{Z}_{\geq 0}$ as follows:

$$f_i(j, k) = \begin{cases} 0, & \text{for } i+j+k \text{ odd} \\ 0, & \text{for } \frac{i+j+k}{2} < \max\{i, j, k\} \\ (d-1)^{\frac{j+k-i}{2}}, & \text{for } \frac{i+j+k}{2} = \max\{i, j, k\} \\ (d-2)(d-1)^{\frac{i+k-i}{2}-1}, & \text{otherwise,} \end{cases}$$

and for $i = 0$,

$$f_0(j, k) = \begin{cases} 0, & j \neq k \\ \lfloor d(d-1)^{j-1} \rfloor, & j = k. \end{cases}$$

Thus, by Claim III.4, we conclude that

$$b_S = \sum_{\omega, \omega' \in \Omega} \prod_{j=1}^m f_j(\omega_j, \omega'_j).$$

IX. APPENDIX C: PROOFS FROM SECTION IV

A. Theorem IV.1

Proof of Theorem IV.1: The first ingredient of our argument comes from Bukh's proof of the Boros-Füredi theorem [Buk06]. A fan of three lines in the plane that pass through a point x splits \mathbb{R}^2 into 6 sectors. For every finite $X \subset \mathbb{R}^2$, there is such a fan where each sector contains at least $\lfloor \frac{|X|}{6} \rfloor$ points of X . But then x resides in every triangle whose three vertices come from non-contiguous sectors of the fan. Thus it suffices to show that if $A, B, C \subset V(G_{[1,2,3]})$ are disjoint subsets of size $\lfloor \frac{|V(G_{[1,2,3]})|}{6} \rfloor = \lfloor \frac{n^3}{6} \rfloor$ each, then a constant fraction of the triangles in $G_{[1,2,3]}$ are in $T(A, B, C)$, i.e., they meet A, B and C .

Using the expander mixing lemma (=EML), we derive an estimate of $|E_{[1,2,3]}(A, B)|$ and show that the density of the A, B edges is very close to the overall edge density of $G_{[1,2,3]}$. We assign a midpoint to every directed $A \rightarrow B$ edge $u \rightarrow v$. The crucial property of this midpoint is that $d^3 - O(d^2)$ of its d^3 neighbors in $G_{[1,1,1]}$ form together with u, v a triangle in $G_{[1,2,3]}$. This gives us a good lower bound on the number of triangles in $G_{[1,2,3]}$ that have exactly one vertex in A and one in B . Next we need to show that however we choose C , many of these triangles have a vertex also in C . To this end, we apply the EML to M and C in $G_{[1,1,1]}$, where M is the multiset of all midpoints created as above. Here C is an arbitrary set of $\lfloor \frac{n^3}{6} \rfloor$ vertices outside $A \cup B$.

We turn to carry out this plan. By Proposition III.10, $\lambda(G_{[1,2,3]}) \leq 6\mu d^2(d-1)^3$, where $\mu = \max(\lambda(G), 2\sqrt{d-1})$, and the EML yields:

$$\frac{n^3}{6}(d-1)^3 d^2(d+6\mu) \geq |E_{[1,2,3]}(A, B)| \geq \frac{n^3}{6}(d-1)^3 d^2(d-6\mu)$$

Let $u = (u_1, u_2, u_3) \in A$, $v = (v_1, v_2, v_3) \in B$ be neighbors in $G_{[1,2,3]}$ and suppose that their distance profile is $(1, 2, 3)$, in this order. The vertices u_2, v_2 have a unique common neighbor in G , called w . Also, let z_1, z_2 be the vertices on the shortest path from u_3 to v_3 . Then $m = (u_1, v_2, z_1)$ is a midpoint of the directed edge $u \rightarrow v$. Let $x = (x_1, x_2, x_3)$ be a neighbor of m in $G_{[1,1,1]}$, i.e., $x_1 u_1, x_2 v_2, x_3 z_1 \in E(G)$. It is easily verified that if in addition $x_1 \neq v_1$, $x_2 \neq w$, and $x_3 \neq u_3, z_2$, then uvx is a triangle in $G_{[1,2,3]}$. Clearly $(d-1)^2(d-2) = d^3 - O(d^2)$ of the d^3 neighbors of m satisfy these additional conditions. Figure 3 provides a local view of the three factors of $G_{[1,2,3]}$:




Figure 3: A triangle in $G_{[1,2,3]}$ as viewed in $G^{\square 3}$.

Clearly the midpoint we chose for $v \rightarrow u$ differs from the one we choose for $u \rightarrow v$. Let $M = M(A, B)$ be the multiset of all such midpoints (for both $A \rightarrow B$ and $B \rightarrow A$ edges). Let C be a set of $\lfloor \frac{n^3}{6} \rfloor$ vertices outside $A \cup B$. To count M, C edges we need a version of the EML that applies as well to multisets of vertices.

Lemma IX.1 ([Con17]). *Let P, Q be two multisets of vertices in a D -regular N -vertex graph H . Then:*

$$\left| E(P, Q) - \frac{D}{N} |P| |Q| \right| \leq 1\lambda(H) \sqrt{\left(\sum_{x \in P} w_x^2 - \frac{|P|^2}{N} \right) \left(\sum_{y \in Q} w_y^2 - \frac{|Q|^2}{N} \right)}$$

where w_x, w_y is the multiplicity of $x \in P$ resp. of $y \in Q$.

By Proposition III.10, $\lambda(G_{[1,1,1]}) \leq \mu d^2$, hence

$$|E_{[1,1,1]}(M, C)| \geq \frac{d^3}{n^3} |M| |C| - \mu d^2 \sqrt{|C| \sum_{y \in M} w_y^2}.$$

As noted before, if uv is an edge in $G_{[1,2,3]}$ and m is the midpoint we chose, then out of the d^3 neighbors that m has in $G_{[1,1,1]}$, at least $(d-1)^2(d-2)$ form a $G_{[1,2,3]}$ -triangle with u, v . Therefore,

$$|T(A, B, C)| \geq |E_{[1,1,1]}(M, C)| - (d^3 - (d-1)^2(d-2)) |M|$$

A long but routine calculation shows that the theorem holds, e.g., with $d_0 = 1600$, $\varepsilon_0 = 1/20$ and $\alpha_0 = 1/100$. ■

Remark IX.2. So far we have provided no systematic explanation for the connection between the multiset $S = [1, 2, 3]$ and $[1, 1, 1]$. We turn to discuss this issue. Our argument utilizes two properties of S :

- 1) All its elements are positive;
- 2) Given a list of all the potential length profiles of triangles in G_S , one should check whether there is a triangles that has a center with a distance profile from one of the triangle's vertices which is zero-free.

These are the only conditions we used about G_S . Thus, $[1, 1, 0]$ fails condition (1), and $[1, 1, 2]$ does not satisfy (2), but $[1, 2, 3]$ has such triangles, namely those with our specific choice of midpoints as centers. There are many other examples such as

$[2, 2, 2]$. A curious aspect of the latter example is that for large enough d it has the geometric overlap property even though its link is not even connected.

B. Theorem IV.2

We establish a two sided spectral gap for $\text{Aux}(\mathcal{C}_{G_{[1,1,0]}}^{(2)})$ (and hence rapid mixing of the walk). A similar, but slightly harder argument applies as well to $S = [1, 2, 3]$.

Lemma IX.3. *Let G be a d -regular triangle-free graph with n vertices, where $\lambda(G) = (1 - \varepsilon)d$, and let $\Gamma = \mathcal{C}_{G_{[1,1,0]}}^{(2)}$. Then $\lambda_2(\text{Aux}(\Gamma)) \leq (1 - \Omega(\varepsilon^4))4d$.*

Proof: Let $F \subseteq E_{[1,1,0]}$ be a set of at most a half of $G_{[1,1,0]}$'s edges and consider the triangles of $G_{[1,1,0]}$ that have edges from F . Proving a lower bound on $\text{Aux}(\Gamma)$'s edge expansion entails showing that for every such F , a constant fraction of these triangles are not contained in F .

Let Γ denote the set of triangles in \mathcal{T} . We freely refer to them either as triangles or as 2-faces of a complex. Each triangle is associated to its center, so that $\mathcal{T} = \bigsqcup_{x \in \Gamma^{(0)}} \mathcal{T}_x$, where \mathcal{T}_x is the set of those triangles in Γ whose center is the vertex x .

Every edge xy in $G_{[1,1,0]}$ has two midpoints, call them x' and y' . It is easy to verify that x' and y' are neighbors in $G_{[1,1,0]}$ and the midpoints of the edge $x'y'$ are x and y . This yields a natural duality $xy \longleftrightarrow x'y'$ among the edges of $G_{[1,1,0]}$. Let F_x be the set of those edges in F that belong to a triangle in \mathcal{T}_x . Note that $2|F| = \sum_{x \in \Gamma^{(0)}} |F_x|$, since every edge of $G_{[1,1,0]}$ has exactly two midpoints, each of which is a center of $3d^2$ triangles.

Let X be the set of those vertices x of Γ such that the vast majority of edges in \mathcal{T}_x belong to F , i.e., $X := \{x \in \Gamma^{(0)} \mid |F_x| \geq (1 - \delta)3d^2\}$. We intend to show that with a proper choice of $\delta > 0$ there holds

$$\sum_{x \notin X} |F_x| \geq \frac{\varepsilon}{12} |F|. \quad (\text{IX.1})$$

We can assume that $\sum_{x \in X} |F_x| \geq |F|$, for otherwise $\sum_{x \notin X} |F_x| \geq |F|$ and (IX.1) clearly holds. Consequently, $|X| \geq \frac{|F|}{3d^2}$, since $|F_x| \leq 3d^2$ for every x . We can and will take $\delta < \frac{1}{4}$, so that:

$$\frac{3}{4} |X| \cdot 3d^2 \leq |X| (1 - \delta) 3d^2 \leq \sum_{x \in X} |F_x| \leq 2|F| \leq \frac{3d^2 n^3}{2}.$$

Hence $|X| \leq \frac{2n^3}{3}$ and $|X^c| \geq \frac{1}{2}|X|$. It is well known that a k -regular graph whose second eigenvalue is μ has edge expansion $\geq \frac{k-\mu}{2}$. Since $\lambda(G_{[1,1,0]}) = d^2 + 2d\lambda(G) = 3d^2 - 2\varepsilon d^2$, we conclude that

$$|E(X, X^c)| \geq \varepsilon d^2 \min(|X|, |X^c|) \geq \frac{\varepsilon d^2}{2} |X|.$$

To derive an upper bound on $|E(X, X^c)|$, let $xy \in E(X, X^c)$, with $x \in X, y \notin X$ and let e be the edge that is dual to xy . Clearly e is in both \mathcal{T}_x and \mathcal{T}_y , so either $e \in F$ and $e \in F_y$ or $e \notin F$ and $e \in \mathcal{T}_x \setminus F_x$. Therefore

$$\sum_{y \in X^c} |F_y| + \sum_{x \in X} |\mathcal{T}_x \setminus F_x| \geq |E(X, X^c)|.$$

But if $x \in X$, then $|\mathcal{T}_x \setminus F_x| \leq 3\delta d^2$. We sum the above inequalities and conclude that

$$\sum_{y \in X^c} |F_y| \geq \frac{\varepsilon d^2}{2} |X| - 3\delta d^2 |X| \geq \left(\frac{\varepsilon}{6} - \delta\right) |F|.$$

By choosing $\delta = \frac{\varepsilon}{12}$ we obtain (IX.1).

We proceed to prove the main statement. The fact that $\text{Aux}(\mathcal{T}_x)$ is isomorphic to $((\mathbb{T}_d)_S)_{(\xi, \xi, \xi)}$ together with the spectral information in Lemma III.14, imply that $\text{Aux}(\mathcal{T}_x)$ has edge expansion $\geq \frac{d}{2}$. Hence if $x \in X^c$, then

$$|E_{\text{Aux}(\Gamma)}(F_x, \mathcal{T}_x \setminus F_x)| \geq \frac{d}{2} \min(|F_x|, |\mathcal{T}_x \setminus F_x|) \geq \frac{\varepsilon d}{24} |F_x|.$$

Therefore,

$$\begin{aligned} |E_{\text{Aux}(\Gamma)}(F, F^c)| &= \sum_{x \in \Gamma^{(0)}} |E_{\text{Aux}(\Gamma)}(F_x, \mathcal{T}_x \setminus F_x)| \\ &\geq \sum_{x \in X^c} |E_{\text{Aux}(\Gamma)}(F_x, \mathcal{T}_x \setminus F_x)| \\ &\geq \sum_{x \in X^c} \frac{\varepsilon d}{24} |F_x| \geq \frac{\varepsilon^2}{288} d |F|, \end{aligned}$$

where the last step uses Inequality IX.1. In other words, $\text{Aux}(\Gamma)$ has edge expansion $\geq \frac{\varepsilon^2}{288} d$. But the second eigenvalue of a k -regular graph with edge-expansion h is at most $\sqrt{k^2 - h^2}$ (see Appendix B in [RS07]). Since $\text{Aux}(\Gamma)$ is $4d$ -regular, this yields

$$\lambda_2(\text{Aux}(\Gamma)) \leq \left(1 - \frac{\varepsilon^4}{3 \cdot 10^6}\right) 4d. \quad \blacksquare$$

In order to control the low end of $\text{Aux}(\Gamma)$'s spectrum we recall the following:

Lemma IX.4 ([DR94]). *Let $G = (V, E)$ be an N -vertex D -regular graph with eigenvalues $\lambda_1 \geq \dots \geq \lambda_N$. For $U \subseteq V$ let $b(U)$ denote the least number of edges that must be removed to make subgraph induced by U bipartite. Then*

$$\lambda_N \geq -D + \frac{\Psi^2}{4D},$$

where

$$\Psi = \min_{U \neq \emptyset} \frac{b(U) + |E(U, U^c)|}{|U|}.$$

We can now establish a gap at the bottom of $\text{Aux}(\Gamma)$'s spectrum. We consider U either as a set of vertices in $\text{Aux}(\Gamma)$, or a set of edges in $G_{[1,1,0]}$. We separate the proof into two cases:

- When U is very large, and therefore contains many triangles;
- When U is not very large in which case we can apply Claim IX.3.

We need the following

Claim IX.5. *A set W of w edges in $K_{d,d,d}$ contains at least $d(w - 2d^2)^+$ triangles. The bound is tight.*

Proof: Tightness is easy. If $w \leq 2d^2$, we can have W completely avoid one of the three $K_{d,d}$'s, and therefore be triangle free. When $w > 2d^2$, have W contain two of the $K_{d,d}$'s. Every edge in the third $K_{d,d}$ is in exactly d triangles so W has exactly $d(w - 2d^2)$ triangles. The proof of the bound is very similar: Start with any set W of $w > 2d^2$ edges and sequentially add to W every remaining edge in $K_{d,d,d}$. The addition of a new edge creates at most d new triangles, and eventually we reach the whole of $K_{d,d,d}$ with its d^3 triangles. It follows that we must have started with at least $d(w - 2d^2)$ triangles, as claimed. \blacksquare

We maintain the same notations: U_x is the set of edges in U that belong to a triangle in \mathcal{T}_x , the set of triangles with center x . A triangle is associated to its center, and we partition the triangles contained in U according to their various centers. We also recall that the 1-skeleton of \mathcal{T}_x is a complete tripartite graph $K_{d,d,d}$.

Call vertex x *heavy* if $|U_x| \geq \frac{5}{6} |E(K_{d,d,d})| = \frac{5d^2}{2}$, and note that by the above claim, in this case U_x must contain at least $\frac{d^3}{2}$ triangles, which is also a lower bound on the number of triangles in $\text{Aux}(U_x)$. But all triangles in $\text{Aux}(U_x)$ are edge disjoint, so we must remove at least $\frac{d^3}{2}$ edges from $\text{Aux}(U_x)$ to make it bipartite.

Also recall that every edge in $E_{[1,1,0]}$ belongs to exactly two triangles. Consequently, if $|U| \geq \frac{71}{72} |E_{[1,1,0]}|$, then at least $\frac{5}{6}$ of the vertices are heavy. Therefore, in this case we must remove at least $\frac{5n^3}{6} \cdot \frac{d^3}{2} = \frac{5d^3 n^3}{12}$ edges to make the induced graph on U bipartite. Therefore

$$\frac{b(U)}{|U|} \geq \frac{5d^3 n^3}{12} \cdot \frac{2}{3d^2 n^3} = \frac{5d}{18}.$$

On the other hand, if $|U| \leq \frac{71}{72} |E_{[1,1,0]}|$, then

$$|E_{\text{Aux}(\Gamma)}(U, U^c)| \geq \frac{d\varepsilon^2}{192} \min(|U|, |U^c|) \geq \frac{d\varepsilon^2}{71 \cdot 192} |U|$$

and therefore $\frac{|E_{\text{Aux}(\Gamma)}(U, U^c)|}{|U|} \geq \frac{d\varepsilon^2}{2 \cdot 10^4}$. We conclude that $\Psi \geq \frac{d\varepsilon^2}{2 \cdot 10^4}$, and by Lemma IX.4, $\lambda_N \geq -4d + \frac{\varepsilon^4 d}{32 \cdot 10^8}$. Since we established an additive gap of size $O(d\varepsilon^4)$ both from above and from below for $\text{Aux}(\Gamma)$, it follows that the edge-triangle-edge random walk mixes rapidly.

ACKNOWLEDGEMENTS:

We thank Irit Dinur, whose questions prompted us to work on these problems. Our conversations with Irit took place during a special year on high dimensional combinatorics held in the Israel Institute of Advanced Studies. We thank the IIAS and the organizers of this beautiful year.

REFERENCES

- [ABLS07] N. Alon, I. Benjamini, E. Lubetzky, and S. Sodin, *Non-backtracking random walks mix faster*, Communications in Contemporary Mathematics **9** (2007), no. 4, 585–603.
- [ALOGV19] N. Anari, K. Liu, S. Oveis-Gharan, and C. Vinzant, *Log-Concave Polynomials II: High-Dimensional Walks and an FPRAS for Counting Bases of a Matroid* (2019). Available at <https://arxiv.org/abs/1811.01816>.
- [Buk06] B. Bukh, *A Point in Many Triangles*, Electr. J. Comb. **13** (2006), no. 1.
- [Cha] M. Chapman, *Conlon's construction of hypergraph expanders*. Available at https://cs.huji.ac.il/~michaelchapman/Lecture_11-6-18_HDE.pdf.
- [Con17] D. Conlon, *Hypergraph expanders from Cayley graphs* (2017). Available at <https://arxiv.org/abs/1709.10006>.
- [DK17] I. Dinur and T. Kaufman, *High dimensional expanders imply agreement expanders*, IEEE 58th Annual Symposium on Foundations of Computer Science (2017).
- [DSV03] G. Davidoff, P. Sarnak, and A. Valette, *Elementary Number Theory, Group Theory and Ramanujan Graphs*, Cambridge university press, 2003.
- [DR94] M. Desai and V. Rao, *A characterization of the smallest eigenvalue of a graph*, J. Graph Theory **18–2** (1994), 181–194.
- [EK16] S. Evra and T. Kaufman, *Bounded degree cosystolic expanders of every dimension*, In Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing (2016).
- [Gar73] H. Garland, *p-adic curvature and the cohomology of discrete subgroups of p-adic groups*, Annals of Mathematics **97** (1973), no. 3, 375–423.
- [KO17] T. Kaufman and I. Oppenheim, *Simplicial complexes arising from elementary matrix groups and high dimensional expanders* (2017). Available at <https://arxiv.org/abs/1710.05304>.
- [GP19] K. Golubev and O. Parzanchevski, *Spectrum and combinatorics of Ramanujan triangle complexes*, Israel Journal of Mathematics **230** (2019), 583–612.
- [HLW06] S. Hoory, N. Linial, and A. Wigderson, *Expander Graphs and Their Applications*, Bull. Amer. Math. Soc. **43** (2006), no. 04, 439–562, DOI 10.1090/s0273-0979-06-01126-8.
- [KM16] T. Kaufman and D. Mass, *High Dimensional Combinatorial Random Walks and Colorful Expansion*, CoRR **abs/1604.02947** (2016).
- [LMN02] N. Linial, A. Magen, and A. Naor, *Girth and Euclidean distortion*, Geometric & Functional Analysis GAFA **12** (2002), no. 2, 380–394.
- [LLP17] E. Lubetzky, A. Lubotzky, and O. Parzanchevski, *Random walks on Ramanujan complexes and digraphs* (2017). Available at <https://arxiv.org/abs/1702.05452>.
- [LSV05] A. Lubotzky, B. Samuels, and U. Vishne, *Ramanujan complexes of type \tilde{A}_{ds}* , Israel journal of Mathematics **149** (2005), 267–299.
- [MSS15] A. W. Marcus, D. A. Spielman, and N. Srivastava, *Interlacing families I: bipartite Ramanujan graphs of all degrees*, Ann. of Math. **182-1** (2015), 307–325.
- [MSW82] P. McMullen, C. Schulz, and J. M. Wills, *Equivelar polyhedral manifolds in E^3* , Israel J. Math. **41** (1982), pp 331–346.
- [MSW83] P. McMullen, C. Schulz, and J. M. Wills, *Polyhedral 2-manifolds in E^3 with unusually large genus*, Israel J. Math. **46** (1983), pp 127–144.
- [Opp17] I. Oppenheim, *Vanishing of cohomology with coefficients in representations on Banach spaces of groups acting on buildings*, Comment. Math. Helv. **92** (2017), pp 389–428.
- [Pap16] M. Papikian, *On Garlands vanishing theorem for SL_n* , European Journal of Mathematics **2** (2016), pp 579–613.
- [Pin73] M. Pinsker, *On the complexity of a concentrator*, 7th annual teletraffic conference (1973).
- [RS07] J. Radhakrishnan and M. Sudan, *On Dinurs Proof of the PCP Theorem*, Bulletin of th AMS **44** (2007), pp 19–61.
- [Sol92] P. Solé, *The second eigenvalue of regular graphs of given girth*, J. Combin. Theory Ser. B **56** (1992), no. 2, 239–249.
- [Wei62] P. M. Weichsel, *The Kronecker product of graphs*, Proceedings of the American Mathematical Society **13** (1962), 47–52.