Séminaire BOURBAKI Novembre 2021
73e année, 2021–2022, no 1187, p. 281 à 294
doi: 10.24033/ast.1188

HIGH-DIMENSIONAL EXPANDERS
[*after* Gromov, Kaufman–Kazhdan–Lubotzky, and others]

by **Uli Wagner**

1. Introduction

Informally speaking, expander graphs combine two seemingly contradictory properties: they are very sparse yet at the same time highly connected. There are several different ways of quantifying mathematically what it means for a graph to be "highly connected", leading to different definitions of *expansion* (which, however, turn out to be essentially equivalent). Arguably the most elementary one is *edge expansion*:

Definition 1.1 (Edge Expansion). Let $X = (V, E)$ be a graph.[(1)] For disjoint subsets $S, T \subset V$, let $E(S, T)$ denote the set of edges of X between S and T. We say that X is η-edge expanding, for some $\eta \geqslant 0$, if

$$\frac{|E(S, V \setminus S)|}{|E|} \geqslant \eta \cdot \frac{\min\{|S|, |V \setminus S|}{|V|} \qquad (\forall S \subset V, S \neq \emptyset, V) \tag{1}$$

The *edge expansion* of X (also called *Cheeger constant*) is defined as the optimal η such that (1) holds, i.e.,

$$h(X) := \min_{S\colon \emptyset \neq S \subsetneq V} \frac{|E(S, V \setminus S)|}{\min\{|S|, |V \setminus S|\}} \cdot \frac{|V|}{|E|} \tag{2}$$

By definition, we have $h(X) > 0$ if and only if X is connected.

As a trivial example (which, however, will play an important role later on, for generalizations to higher dimensions), the complete graph K_n on n vertices satisfies

$$h(K_n) = 1 + o(1)$$

(1)Throughout we will assume all graphs to be finite, simple (no loops or multiple edges) and undirected, unless explicitly stated otherwise. For disjoint subsets S, T of V, we will denote by $E(S, T) := \{vw \in E\colon v \in S, w \in T\}$ the set of edges between S and T, and for a vertex a vertex $v \in V$, we denote by $\deg(v) = |\{w \in V \mid vw \in E\}|$ the *degree* (also called *valency*) of v in X.

Definition 1.2. An infinite family of finite graphs X_n, $n \in \mathbb{N}$, is called a family of (bounded-degree) *expander graphs* if the graphs are of uniformly bounded degree and their edge expansion is uniformly bounded away from zero, i.e., there are $\eta > 0$ and $k \in \mathbb{N}$ such that $h(X_n) \geqslant \eta$ and $\deg_{X_n}(v) \leqslant k$ for all vertices v of X_n and all $n \in \mathbb{N}$.

Families of expander graphs were shown to exist by probabilistic arguments by Kolmogorov and Barzdin (1993) and Pinsker (1973). The first explicit construction of a family expander graphs was given by Margulis (1973) (using Kazhdan's Property (T)), and by now, many different constructions are known. Expansion and expander graphs play an important role in many different areas of mathematics and computer science and are the source of deep connections between them, see for instance the surveys by Hoory, Linial, and Wigderson (2006) or Lubotzky (2012).

The goal of this *exposé* is to offer a glimpse of the emerging theory of *high-dimensional expanders*, which is still in a formative stage, but has already led to a number of striking results and applications (see, e.g., Lubotzky (2018) for a recent survey, including many topics that we will neglect). One interesting aspect is that even the definition of higher-dimensional expansion is not at all obvious and that, unlike in the case of graphs, there is a rich array of mutually non-equivalent notions of high-dimensional expansion, each of interest in its own right and with its own applications.

Here we will mainly focus on three notions of high-dimensional expansion that have a strong topological flavor and that have played an important role in the development of the field in the last decade: the first is *topological expansion* (also called the *topological overlap property*), which is defined in terms of maps from a d-dimensional simplicial complex to $\mathbb{R}^d$; the second is *coboundary expansion*, which generalizes edge-expansion of graphs and provides a quantitative version of vanishing $\mathbb{F}_2$-cohomology of a complex in higher dimensions; the third is *cosystolic expansion*, which is a weakening of coboundary expansion that allows for non-vanishing $\mathbb{F}_2$-cohomology. Infomally, these notions are related by the following series of implications:

$$\text{coboundary expansion} \quad \Rightarrow \quad \text{cosystolic expansion} \quad \Rightarrow \quad \text{topological expansion}$$

2. Topological Overlap and Topological Expanders

As a starting point, let us consider the following classical result in discrete geometry, due to Boros and Füredi (1984) (for $d = 2$) and Bárány (1982) (for general d), which at first may seem to have little to do with to expansion:

Theorem 2.1. *Let P be a set of n points in $\mathbb{R}^2$. Then there exists a point $\mathbb{R}^2$ that is contained in at least*

$$\left(\frac{2}{9}+o(1)\right)\binom{n}{3}$$

of the triangles (convex hulls of triples of points) spanned by the points in P.

More generally, for every set P of n points in $\mathbb{R}^d$, there exists a point $\mathbb{R}^d$ that is contained in at least

$$(c_d+o(1))\binom{n}{d+1}$$

of the affine d-simplices (convex hulls of $d+1$ points) spanned by the points in P, where $c_d>0$ is a constant that depends only on d.

Theorem 2.1 has led to a host of related results and applications, see Matoušek (2002, Ch. 9). Determining the optimal value of the constant c_d is a well-known open problem. It is known that $c_2 = 2/9$ is optimal, and an analogous construction in higher dimensions shows $c_d \leqslant \frac{(d+1)!}{(d+1)^{d+1}} = e^{-\Theta(d)}$ (Bukh, Matoušek, and Nivasch, 2010). On the other hand, Bárány's proof yields $c_d \geqslant (d+1)^{-d}$, and despite several later improvements, the best known lower bound is still of the form $e^{-\Theta(d\log d)}$.

Theorem 2.1 can be restated as follows. Let Δ_n^d denote the complete d-dimensional simplicial complex on n vertices (in other words, the d-dimensional skeleton of the $(n-1)$-dimensional simplex). Then, for every *affine* map $F\colon \Delta_n^d \to \mathbb{R}^d$, there is a point $p \in \mathbb{R}^d$ that is contained in the F-images of at least a $(c_d+o(1))$-fraction of the d-dimensional faces of Δ_n^d.

Gromov (2010) showed that this remains true for arbitrary continuous maps:

Theorem 2.2 (Gromov). *For every continuous map $F\colon \Delta_n^d \to \mathbb{R}^d$, there is a point $p \in \mathbb{R}^d$ that is contained in the F-images of at least a $(c_d^{\mathrm{top}}+o(1))$-fraction of the d-dimensional faces of Δ_n^d, where c_d^{top} is a constant depending only on d.*

Gromov's argument yields a lower bound of $c_d \geqslant c_d^{\mathrm{top}} \geqslant \frac{2d}{(d+1)!(d+1)}$, recovering the optimal constant $c_2 = c_2^{\mathrm{top}} = 2/9$ in the plane, and improving on the previously known bounds for c_d by a factor exponential in d for general dimensions; however, the lower bound is still of the form $e^{-\Theta(d\log d)}$ and thus far from the upper bound.

One aspect that makes Theorem 2.2 interesting is that for $d \geqslant 2$ and an arbitrary continuous map $F\colon \Delta_n^d \to \mathbb{R}^d$, there is no obvious candidate for the point p. (By contrast, for $d=1$, we can simply take p to be the median of the images of the vertices; moreover, for *affine* maps, as in Theorem 2.1, one can show that the *centerpoint* of the vertex images, a generalization of the median, works in any dimension d, albeit leading to a non-optimal constant, see Bukh, Matoušek, and Nivasch (2010).)

Gromov's argument[2] for the existence of a suitable point p relies on a certain higher-dimensional expansion property of Δ_n^d, *coboundary expansion*, which generalizes edge-expansion of graphs (corresponding to 1-dimensional coboundary expansion); the formal definition will be given in Section 3 below. Interestingly, the notion of coboundary expansion also arose independently (and earlier) in a different context, in the work of Linial and Meshulam (2006) on *random complexes*.

Gromov's proof of Theorem 2.2 is remarkably robust and yields a much more general result as well as a whole new circle of questions:

Definition 2.3. Let X be a finite d-dimensional simplicial complex.

1. We say that X has the *ε-topological overlap property*, for some real parameter $\varepsilon > 0$, if for every continuous map $F\colon X \to \mathbb{R}^d$, there exists a point $p \in \mathbb{R}^d$ that is contained in at least an ε-fraction of the F-images of d-dimensional faces of X.

2. An infinite family of d-dimensional complexes is a family of *topological expanders* if all the complexes in the family have the ε-topological overlap property, for a uniform $\varepsilon > 0$.

In this language, Theorem 2.1 says that for every d, the complete complexes Δ_n^d form a family of geometric expanders (cf., Remark 2.7), and Theorem 2.2 asserts that they form a family of topological expanders. As remarked above, Gromov's proof leads to a more general result, which can be informally summarized as follows (see Theorem 4.2 below for the formal statement): *every d-dimensional complex that has the coboundary expansion property in dimensions* $1, \ldots, d$ *satisfies the topological overlap property*, with an overlap constant ε that depends on d and on the coboundary expansion constants of X. Gromov (2010) showed that various other families of d-dimensional complexes are coboundary expanders, hence topological expanders, e.g., spherical buildings; however, none of these examples are of *bounded degree*, i.e., for each of these complexes, the number of d-faces containing a given vertex (or even containing a given $(d-1)$-face) tends to infinity with the size of the complex.

This naturally raises the question whether there are, for instance, families of 2-dimensional topological expanders that are of *bounded degree*, either *in the weak sense* that every edge is contained in a bounded number of triangles, or *in the strong sense* that every vertex is contained in a bounded number of triangles.

Both of these questions have been answered affirmatively, the first by Lubotzky and Meshulam (2015), using a probabilistic construction based on *random Latin squares*, and the second by Kaufman, Kazhdan, and Lubotzky (2016), using a construction of *Ramanujan complexes* given by Lubotzky, Samuels, and Vishne (2005).

[2] As explained in Guth (2014), the argument can be seen as analogous to the proof of the *Waist Inequality* in Gromov (1983).

Let us state these results. For the first, let $n \in \mathbb{N}$ and let $T_n = V_1 * V_2 * V_3$ be the *complete tripartite* 2*-dimensional complex* on three pairwise disjoint sets V_1, V_2, V_3 of n vertices each. (Thus, a subset $\sigma \subseteq V_1 \sqcup V_2 \sqcup V_3$ is a face of $V_1 * V_2 * V_3$ if and only if $|\sigma \cap V_i| \leqslant 1$ for $i = 1, 2, 3$.) Thus, T_n has $3n$ vertices, $3n^2$ edges (1-simplices), and n^3 triangles (2-simplices).

For our purposes, a *Latin square* is a collection L of triangles of T_n such that every edge of T_n is contained in *exactly one* triangle in L. (Hence, for every vertex $v \in V_i$ of T_n, the *link* $L_v := \{\sigma \setminus v \mid \sigma \in L\}$ forms a *perfect matching* in the complete bipartite graph $V_j * V_k$ on the remaining two vertex sets, $j, k \neq i$.) Let $\mathscr{L}_n$ denote the set of all Latin squares. For $D \in \mathbb{N}$, define a random subcomplex $Y(n, D)$ as follows: Choose D Latin squares $L_1, \ldots, L_D \in \mathscr{L}_n$ independently uniformly at random, and let $Y(n, D)$ be the subcomplex of T_n that has the same 1-skeleton as T_n as whose triangles are exactly the triangles in $L_1 \cup \cdots \cup L_D$.

Theorem 2.4 (Lubotzky and Meshulam). *There exist constants $D \in \mathbb{N}$ and $\varepsilon > 0$ such that asymptotically almost surely (with probability tending to* 1 *as $n \to \infty$), the random complex $Y(n, D)$ has the ε-topological overlap property. Thus, there exists an infinite family of* 2*-dimensional topological expanders that are of bounded degree in the weak sense.*

More precisely, Lubotzky and Meshulam show that, asymptotically almost surely, $Y(n, D)$ has 2-dimensional coboundary expansion at least η, for some other constant $\eta > 0$. The topological overlap property then follows from Gromov's result (since the 1-skeleton of $Y(n, D)$, which is a complete tripartite graph, is a very good edge expander).

The second construction, of a family of 2-dimensional topologocial expanders that are of bounded degree in the strong sense that the number of triangles containing a given vertex is bounded by some uniform constant for all complexes is the family, is considerably more elaborate, and we will treat it mostly as a "black box", focusing on the properties used in Kaufman, Kazhdan, and Lubotzky (2016) to prove the topological overlap property.

Let q be a large but fixed prime power. For an integer $r \geqslant 2$. The *spherical building* $S(r, q)$ is defined as the complex of *flags* of nonempty proper linear subspaces of $\mathbb{F}_q^r$, i.e., the vertices of $S(d, q)$ are the nonempty proper linear subspaces $W \subset \mathbb{F}_q^r$, and a set $\{W_0, W_1, \ldots, W_k\}$ of subspaces forms a k-dimensional simplex of $S(r, q)$ if and only if $W_0 \subset W_1 \subset \cdots \subset W_k$ (possibly after reordering the W_i). Thus, $S(r, q)$ is a simplicial complex of dimension $r - 2$.

Let us say that a finite 3-dimensional complex X is *magical* if it has the following properties:

1. For every vertex v of X, the *link* X_v of v in X is isomorphic to $S(4, q)$. It follows that the 1-skeleton $X^{(1)}$ of X is a k-regular graph, where $k \sim q^4$ is the number of vertices of $S(4, q)$ (proper nonempty subspaces of $\mathbb{F}_q^4$).