

On expansion and topological overlap

Dominic Dotterrer¹ · Tali Kaufman² · Uli Wagner³

Received: 23 October 2015 / Accepted: 19 September 2016
© The Author(s) 2017. This article is an open access publication

Abstract We give a detailed and easily accessible proof of Gromov's *Topological Overlap Theorem*. Let X be a finite simplicial complex or, more generally, a finite polyhedral cell complex of dimension d . Informally, the theorem states that if X has sufficiently strong *higher-dimensional expansion properties* (which generalize edge expansion of graphs and are defined in terms of cellular cochains of X) then X has the following *topological overlap property*: for every continuous map $X \rightarrow \mathbb{R}^d$ there exists a point $p \in \mathbb{R}^d$ that is contained in the images of a positive fraction $\mu > 0$ of the d -cells of X . More generally, the conclusion holds if \mathbb{R}^d is replaced by any d -dimensional piecewise-linear manifold M , with a constant μ that depends only on d and on the expansion properties of X , but not on M .

Keywords Expansion · Cell complexes · Topological overlapping · High dimensional expansion

Mathematics Subject Classification 05E45 · 58K15 · 53C23

Research supported by the Swiss National Science Foundation (Project SNSF-PP00P2-138948).
An extended abstract of this paper [5] appeared in the Proceedings of the 32nd International Symposium on Computational Geometry (SoCG 2016).

✉ Uli Wagner
uli@ist.ac.at

Dominic Dotterrer
d.dotterrer@math.uchicago.edu

Tali Kaufman
kaufmant@macs.biu.ac.il

¹ Department of Mathematics, University of Chicago, 5734 S. University Avenue, Chicago, IL 60637, USA

² Department of Computer Science, Bar-Ilan University, 5290002 Ramat Gan, Israel

³ IST Austria, Am Campus 1, 3400 Klosterneuburg, Austria

1 Introduction

Let X be a finite polyhedral cell complex¹ of dimension $\dim X = d$. Gromov [8] recently showed that if X has sufficiently strong *higher-dimensional expansion properties* (which generalize edge expansion of graphs, see below for the precise definition) then X has the following *topological overlap property*: For every every continuous map $f: X \rightarrow \mathbb{R}^d$, there exists a point $p \in \mathbb{R}^d$ that is contained in the images of some positive fraction of the d -cells of X , i.e.,

$$|\{\sigma \in \Sigma_d(X): p \in f(\sigma)\}| \geq \mu \cdot |\Sigma_d(X)|, \quad (1)$$

where $\Sigma_k(X)$ denotes the set of k -dimensional cells of X , $0 \leq k \leq d$, and $\mu > 0$. More generally, the same conclusion holds if the target space \mathbb{R}^d is replaced by a d -dimensional manifold M , and the *overlap constant* $\mu > 0$ depends only on the dimension d and on the constants quantifying the expansion properties of X , but not on M . For technical reasons, we will assume that the manifold M admits a *piecewise-linear (PL) triangulation*, so that we can apply standard tools to perturb a given map to *general position*. We refer to the book by Rourke and Sanderson [15] or to the lecture notes by Zeeman [16] for background and standard facts about piecewise-linear topology.

In the special case where X is the n -dimensional simplex Δ^n (or its d -dimensional skeleton), determining the optimal overlap constant for maps $\Delta^n \rightarrow \mathbb{R}^d$ is a classical problem in discrete geometry, also known as the *point selection problem* [1, 2] and originally only considered for affine maps. Apart from the generalization from affine to arbitrary continuous maps, Gromov's proof also led to improved estimates for the point selection problem, and a number of papers have appeared with expositions and simplified proofs of Gromov's result in this special case $X = \Delta^n$, see [9, 13] and [4, Sec. 7.8].

The goal of the present paper is to provide a detailed and easily accessible proof of Gromov's result for general complexes X , see Theorem 8 below. This is a crucial ingredient for obtaining examples of simplicial complexes X of *bounded degree* (i.e., such that every vertex is incident to a bounded number of simplices) that have the topological overlap property [6, 7]. The basic idea of the proof is the same as Gromov's, but we present a simplified and streamlined version of the proof that uses only elementary topological notions (general position for piecewise-linear maps, algebraic intersection numbers, cellular chains and cochains, and chain homotopies) and avoids much of the machinery used in Gromov's original paper (in particular, the simplicial set of cocycles).

For stating the result formally, we need to discuss higher-dimensional expansion properties of cell complexes. The relevant notion of expansion originated in the work of Linial and Meshulam [10] and of Gromov [8] and generalizes edge expansion of graphs (which corresponds to 1-dimensional expansion). To define k -dimensional expansion, we need two ingredients: first, information about incidences between cells of dimensions k and $k - 1$ and, second, a notion of *discrete volumes* in X . To define these, it is convenient to use the language of *cellular cochains* of X .

1.1 Cellular cochains

Let X be a polyhedral cell complex, let $\Sigma_k(X)$ denote the set of k -dimensional cells of X , and let $C^k(X) := C^k(X; \mathbb{F}_2) := \mathbb{F}_2^{\Sigma_k(X)}$ be the space of k -dimensional cellular cochains with coefficients in the field \mathbb{F}_2 ; in other words $C^k(X)$ is the space of functions $a: \Sigma_k(X) \rightarrow$

¹ See [3, Sec. 12] or [16, Ch. I] for more background on polyhedral cell complexes (in [16], they are called convex linear cell complexes).

$\mathbb{F}_2 = \{0, 1\}$. For a pair $(\sigma, \tau) \in \Sigma_k(X) \times \Sigma_{k-1}(X)$, let $[\sigma : \tau]$ be 1 or 0 depending on whether τ is incident to σ (i.e., whether τ is contained in the boundary $\partial\sigma$) or not. This incidence information is recorded in the *coboundary operator*, which is a linear map $\delta: C^{k-1}(X) \rightarrow C^k(X)$ given by $\delta a(\sigma) := \sum_{\tau \in \Sigma_{k-1}(X)} [\sigma : \tau] a(\tau)$.

The elements of the subspaces $Z^k(X) := \ker(\delta: C^k(X) \rightarrow C^{k+1}(X))$ and $B^k(X) := \text{im}(\delta: C^{k-1}(X) \rightarrow C^k(X))$ are called *k-dimensional cocycles* and *coboundaries*, respectively. The composition of consecutive coboundary operators is zero, i.e., $B^k(X) \subseteq Z^k(X)$, and $H^k(X) = Z^k(X)/B^k(X)$ is the *k-dimensional homology group* (with \mathbb{F}_2 -coefficients) of X . This information is customarily recorded in the *cellular cochain complex*² of X :

$$0 \rightarrow \mathbb{F}_2 = C^{-1}(X) \xrightarrow{\delta} C^0(X) \xrightarrow{\delta} C^1(X) \xrightarrow{\delta} \dots \xrightarrow{\delta} C^{d-1}(X) \xrightarrow{\delta} C^d(X) \rightarrow 0 \quad (2)$$

1.2 Norm, cofilling, expansion, and systoles

For $\alpha \in C^k(X)$, let $|\alpha|$ denote the *Hamming norm* of α , i.e., the cardinality of the *support* $\text{supp}(\alpha) := \{\sigma \in \Sigma_k(X) : \alpha(\sigma) \neq 0\}$, which we think of as a measure of “discrete k -dimensional volume.” In fact, it will be convenient to allow more general norms on cochains; the following definition summarizes the properties that we will need.

Definition 1 (*Norm on cochains*) A *norm* on the group $C^*(X) = \bigoplus_{k=0}^d C^k(X)$ of cellular cochains of X with \mathbb{F}_2 -coefficients is a function $\|\cdot\|: C^*(X; \mathbb{F}_2) \rightarrow \mathbb{R}_{\geq 0}$ that satisfies the following properties for all cochains $\alpha, \beta \in C^k(X)$, $0 \leq k \leq d$:

1. $\|0\| = 0$.
2. *Triangle inequality*: $\|\alpha + \beta\| \leq \|\alpha\| + \|\beta\|$.

Furthermore, we will assume throughout that the norm satisfies the following:

3. *Monotonicity*: $\|\alpha\| \leq \|\beta\|$ whenever $\text{supp}(\alpha) \subseteq \text{supp}(\beta)$.

From now on, we work with a fixed norm on the cochains of X . We assume that the norm is normalized such that $\|\mathbb{1}_X^k\| = 1$ for $0 \leq k \leq d$, where $\mathbb{1}_X^k \in C^k(X)$ assigns 1 to every k -cell of X . In particular, when working with the Hamming norm, we will consider its normalized version

$$\|\alpha\|_H := \frac{|\alpha|}{|\Sigma_k(X)|}.$$

Given $\beta \in B^k(X)$, we say that $\alpha \in C^{k-1}(X)$ *cofills* b if $\beta = \delta\alpha$. Once we have a notion of discrete volumes, we can consider the following (*co*)*isoperimetric question*: Can we bound the minimum norm of a cofilling for a coboundary β in terms of the norm of β ?

Definition 2 (*Cofilling/coisoperimetric inequality*) Let $L > 0$. We say that X satisfies a *L-cofilling inequality* (or *coisoperimetric inequality*) in dimension k if, for every $\beta \in B^k(X)$, there exists some $\alpha \in C^{k-1}(X)$ such that $\delta\alpha = \beta$ and $\|\alpha\| \leq L\|\beta\|$.

Any two cofillings of a given coboundary differ by a cocycle. Thus, X satisfies an *L-cofilling inequality* in dimension k if and only if

$$\|\delta\alpha\| \geq \frac{1}{L} \cdot \min\{\|\alpha + \zeta\| : \zeta \in Z^{k-1}(X)\} \quad \text{for all } \alpha \in C^{k-1}(X). \quad (3)$$

² More precisely, we work with the *augmented* cellular cochain complex of X , unless stated otherwise, i.e., we consider X to have a unique (-1) -dimensional cell, the empty cell \emptyset , which is incident to every vertex of X .

We can strengthen (3) by replacing cocycles with coboundaries and obtain a condition that also allows us to draw conclusions about the cohomology of X . For $\alpha \in C^{k-1}(X)$, let

$$\|[\alpha]\| := \min\{\|\alpha + \beta\| : \beta \in B^{k-1}(X)\} \tag{4}$$

denote the distance (with respect to the norm $\|\cdot\|$) of α to the space $B^{k-1}(X)$ of coboundaries.

Definition 3 (*Coboundary expansion*) Let $\eta > 0$. We say that X is η -expanding in dimension k , if for every $(k - 1)$ -cochain $\alpha \in C^{k-1}(X)$,

$$\|\delta\alpha\| \geq \eta \cdot \|[\alpha]\|. \tag{5}$$

Lemma 4 Let $\eta > 0$. A complex X is η -expanding in dimension k if and only if $H^{k-1}(X) = 0$ and X satisfies a $1/\eta$ -coisoperimetric inequality in dimension k .

Proof Suppose that X is η -expanding in dimension k . Clearly, (5) implies (3), i.e., X satisfies a $1/\eta$ -cofilling inequality. Moreover, if $\alpha \in C^{k-1}(X) \setminus B^{k-1}(X)$ then $\|[\alpha]\| > 0$, hence $\|\delta\alpha\| > 0$, hence $\alpha \notin Z^{k-1}(X)$. Thus, $Z^{k-1}(X) = B^{k-1}(X)$, i.e., $H^{k-1}(X) = 0$.

Conversely, assume that $H^{k-1}(X) = 0$. Then $Z^{k-1}(X) = B^{k-1}(X)$, so (5) and (3) are equivalent. \square

In some cases, however, vanishing of $H^{k-1}(X)$ turns out to be too stringent a requirement, and we can replace it by the condition that every nontrivial cocycle has large norm:

Definition 5 (*Large cosystoles*) Let $\vartheta > 0$. We say that X has ϑ -large cosystoles in dimension j if $\|\alpha\| \geq \vartheta$ for every $\alpha \in Z^j(X) \setminus B^j(X)$.

Example 6 Consider the case $k = 1$, with the normalized Hamming norm. In this case, η -expansion in dimension 1 corresponds to η -edge expansion of a graph (the 1-skeleton of the complex). An L -cofilling inequality in dimension 1 means that every connected component of the graph is $1/L$ -edge expanding. Having ϑ -large cosystoles in dimension 0 means that every connected component contains at least a ϑ -fraction of the vertices.

1.3 Local sparsity of X

For the formal statement of the overlap theorem, we need one more technical condition on X . For a cell τ of X , let ι_τ^k be the k -dimensional cochain that assigns 1 to k -cells of X that have nonempty intersection with τ and 0 otherwise.

Definition 7 (*Local sparsity*) Let $\varepsilon > 0$. We say that X is locally ε -sparse (with respect to a given norm $\|\cdot\|$) if $\|\iota_\tau^k\| \leq \varepsilon$ for every nonempty cell τ of X and every $k, 0 \leq k \leq d$.

For example, in the case of the normalized Hamming norm $\|\cdot\|_H$, local sparsity means that

$$|\{\sigma \in \Sigma_k(X) : \tau \cap \sigma \neq \emptyset\}| \leq \varepsilon |\Sigma_k(X)|,$$

for every nonempty cell τ of X .

1.4 Formal statement of the theorem

We are now ready to state Gromov’s theorem.

Theorem 8 (Gromov’s Topological Overlap Theorem [8]) *For every $d \geq 1$ and $L, \vartheta > 0$ there exists $\varepsilon_0 = \varepsilon_0(d, L, \vartheta) > 0$ such that the following holds:*

Let X be a finite cell complex of dimension d , and let $\|\cdot\|$ be a norm on the cochains of X . Suppose that

1. X satisfies a L -cofilling inequality in dimensions $1, \dots, d$;
2. X has ϑ -large cosystoles in dimensions $0, \dots, d - 1$; and
3. X is locally ε -sparse for some $\varepsilon \leq \varepsilon_0$.

Then for every continuous map $f: X \rightarrow M$ into a compact connected d -dimensional piecewise-linear (PL) manifold M , there exists a point $p \in M$ such that³

$$\|\{\sigma \in \Sigma_d(X) \mid p \in f(\sigma)\}\| \geq \mu, \tag{6}$$

where $\mu = \mu(d, \varepsilon, L, \vartheta) > 0$.

Remark 9 The assumption that the manifold M is compact is not essential; moreover, we may assume without loss of generality that M has no boundary. Indeed, since X is compact, the image $f(X)$ is compact and hence contained in a compact submanifold N of M with boundary ∂N ; we can turn N into a compact manifold without boundary by doubling, i.e., by glueing two copies of N along their boundary.

If a complex X satisfies the conclusion of the theorem, we also say that X is *topologically μ -overlapping* for maps into d -dimensional PL manifolds. If the conclusion holds true just for *affine maps* and $M = \mathbb{R}^d$, we say that X is *geometrically μ -overlapping*.

2 Preliminaries from piecewise-linear topology

2.1 Assumptions on M

We assume that M is a compact connected piecewise-linear (PL) d -dimensional manifold, without boundary. That is, we assume that M admits a triangulation⁴ T with the property that the link of every nonempty simplex τ of T is a PL sphere of dimension $d - 1 - \dim(\tau)$; throughout this paper, we only consider triangulations of M that have this property.

2.2 Approximation by PL maps

We can fix a metric on M , e.g., by fixing a triangulation T of M and by considering each simplex of T as a regular simplex with edge length 1. By subdividing a given triangulation T sufficiently often, we can pass to a new triangulation T' in which each simplex has diameter at most $\rho > 0$, for a given ρ (see, e.g., [12, Sec. 1.7]).

By the standard *simplicial approximation theorem* [14], given the triangulation T' of M and a continuous map $f: X \rightarrow M$, there is a simplicial approximation of f , i.e., there is a subdivision X' of X and a simplicial map $g: X' \rightarrow T'$ such that, for each point $x \in X$, the image $g(x)$ belongs to the (uniquely defined) simplex of T' whose relative interior contains $f(x)$. (In fact, g is even homotopic to f , but we will not need that.) This map g is a PL map

³ Here, we use that a subset of $\Sigma_k(X)$ can be identified with a k -dimensional cellular cochain, its indicator function.

⁴ The triangulation is necessarily finite, since M is compact.

$X \rightarrow M$ and the distance between $g(x)$ and $f(x)$ is at most the maximum diameter of any simplex in T' , hence at most ρ , for every $x \in X$.

Thus, by the preceding discussion and the following lemma, it suffices to prove Theorem 8 for PL maps.

Lemma 10 *Let $f: X \rightarrow M$ be a continuous map, and let $g_n: X \rightarrow M$ be a sequence of continuous maps that converges to f uniformly, i.e., $\max_{x \in X} \text{dist}(g_n(x), f(x)) \rightarrow 0$ as $n \rightarrow \infty$. Suppose that for every g_n there exists a point $p_n \in M$ such that $\|\{\sigma \in \Sigma_d(X) \mid p_n \in g_n(\sigma)\}\| \geq \mu$. Then there exists a point $p \in M$ such that (6) holds.*

Proof By compactness, there is a subsequence of the points p_n that converges to a point p . We claim that p is the desired point. Since there are only finitely many cells in X , there is some $\rho > 0$ such that for every d -cell σ of X with $p \notin f(\sigma)$, the distance between p and $f(\sigma)$ is at least ρ . Choose n sufficiently large so that the distance between p_n and p is less than $\rho/2$, and the distance between $f(x)$ and $g_n(x)$ is at most $\rho/2$, for every $x \in X$. If $p_n \in g_n(\sigma)$, then the distance between p and $f(\sigma)$ is less than ρ , so by the choice of ρ , we have $p \in f(\sigma)$. Therefore, $\{\sigma \in \Sigma_d(X) \mid p \in f(\sigma)\} \subseteq \{\sigma \in \Sigma_d(X) \mid p_n \in g_n(\sigma)\}$, and the desired conclusion follows by the monotonicity property of the norm. \square

2.3 General position

We refer to [16, Ch. VI] for a comprehensive treatment of general position for PL maps. The following definition summarizes the properties that we will need.

Definition 11 Let X be a finite polyhedral cell complex, M a PL manifold, and let $f: X \rightarrow M$ be a PL map.

1. We say that f is in *strongly general position* (with respect to the given decomposition of X into polyhedral cells) if, for every $r \geq 1$ and pairwise disjoint cells $\sigma_1, \dots, \sigma_r$ of X ,

$$\dim \left(\bigcap_{i=1}^r f(\sigma_i) \right) \leq \max \left\{ -1, \left(\sum_{i=1}^r \dim \sigma_i \right) - d(r-1) \right\}. \tag{7}$$

In particular, if the number of the right-hand side is -1 , then the intersection is empty.

2. Given a triangulation T of M , we say that f is in *general position with respect to T* if, for every simplex σ of X and every simplex τ of T , $\dim(f(\sigma) \cap \tau) \leq \max\{-1, \dim \sigma + \dim \tau - d\}$; moreover, if $\dim \sigma + \dim \tau = d$ then we require that $f(\sigma)$ and τ intersect transversely (either the intersection is empty, or they intersect locally like complementary linear subspaces).

The main fact that we will need is that any map $f: X \rightarrow M$ can be approximated arbitrarily closely by a PL map that is in general position:

Lemma 12 [16, Ch. VI] *Let $f: X \rightarrow M$ be a PL map and let T be a triangulation of M . Then, up to a small perturbation, we may assume that f is in general position with respect to T and in strongly general position.*

Furthermore, we will need the following notion of *sufficiently fine triangulations*:

Definition 13 Let T be a triangulation of M and let $f: X \rightarrow M$ be a PL map in general position with respect to T . We say that T is *sufficiently fine* with respect to f if, for every $k > 0$ and every k -simplex τ of T ,

$$\|\{\sigma \in \Sigma_{d-k}(X) : f(\sigma) \cap \tau \neq \emptyset\}\| \leq \frac{d}{k} \max\{\|\iota_{\sigma'}^{d-k}\| : \sigma' \in \Sigma_{d-k}(X)\}.$$

Lemma 14 *Suppose that $f: X \rightarrow M$ be a PL map in strongly general position and in general position with respect to a triangulation T of M . Then (by refining T , if necessary), we may assume furthermore that T is sufficiently fine with respect to f .*

Proof If f is in general position with respect to T , then by choosing points at which we subdivide T in a sufficiently generic way, we can assume that f is also in general position with respect to the subdivision T' . Thus, we may assume that T already has the property that every simplex of T has diameter smaller than some specified parameter $\rho > 0$.

Now suppose that $\sigma_1, \dots, \sigma_r$ are pairwise distinct simplices of X with $f(\sigma_1) \cap \dots \cap f(\sigma_r) = \emptyset$. By compactness, there exists $\rho = \rho(\sigma_1, \dots, \sigma_r) > 0$ such that no matter how we select $x_i \in f(\sigma_i)$, some pair x_i, x_j has distance at least ρ . Since X is finite, there is some $\rho > 0$ that works for all finite collections of simplices whose images do not have a common point of intersection. Suppose now that we have chosen T such that all simplices in T have diameter at most $\rho/2$.

Given $\tau \in T$ of dimension $k > 0$ consider $S(\tau) := \{\sigma \in \Sigma_{d-k}(X) : f(\sigma) \cap \tau \neq \emptyset\}$. We claim that $\bigcap_{\sigma \in S(\tau)} f(\sigma) \neq \emptyset$. Otherwise, for every choice of points $x_\sigma \in f(\sigma)$, $\sigma \in S(\tau)$, there would be some pair σ, σ' such that x_σ and $x_{\sigma'}$ have distance at least ρ . However, by the definition of $S(\tau)$, we can choose each x_σ to lie in the intersection $f(\sigma) \cap \tau$, from which it follows that for every pair $\sigma, \sigma' \in S(\tau)$, the distance between x_σ and $x_{\sigma'}$ is at most the diameter of τ , i.e., at most $\rho/2$.

Let $\{\sigma_1, \dots, \sigma_r\} \subseteq S(\tau)$ be an inclusion-maximal subset with $\sigma_i \cap \sigma_j = \emptyset$ (i.e., the σ_i are pairwise vertex-disjoint; we can pick this subset greedily). Since f is in strongly general position and $\bigcap_{\sigma \in S(\tau)} f(\sigma) \neq \emptyset$, it follows that $\sum_{i=1}^r (d - k) - d(r - 1) \geq 0$; this implies $r \leq d/k$. Now, every other simplex $\sigma \in S(\tau)$ intersects one of the σ_i . Thus, by monotonicity of the norm and by the triangle inequality, $\|S(\tau)\| \leq \frac{d}{k} \max_{1 \leq i \leq r} \|\sigma_i\|^{d-k}$. □

2.4 Intersection numbers

Definition 15 (*Intersection numbers*) If T is a PL triangulation of M and if $f: X \rightarrow M$ is a PL map in general position with respect to T , then for every pair of chains $a \in C_{d-k}(X; \mathbb{F}_2)$ and $b \in C_k(T; \mathbb{F}_2)$, we can define their (algebraic) *intersection number*

$$f(a) \cdot b \in \mathbb{F}_2$$

as follows: If σ is a $(d - k)$ -dimensional cell of X and if τ is a k -dimensional simplex of T , then by general position, the intersection $f(\sigma) \cap \tau$ consists of a finite number of points, and the intersection number $f(\sigma) \cdot \tau$ is defined as the number of intersections⁵ modulo 2. This definition is extended by linearity (over \mathbb{F}_2) to arbitrary chains.

This yields, for $0 \leq k \leq d$, an *intersection number homomorphism*

$$f^{\frown}: C_k(T) \rightarrow C^{d-k}(X), \tag{8}$$

defined by $f^{\frown}(b)(a) = f(a) \cdot b$ for each $a \in C_{d-k}(X)$.

It is well-known that the intersection number homomorphism is a *chain-cochain map*, i.e., it commutes with the boundary and coboundary operators in the following sense (see, e.g., [11, Sec. 2.2] for a detailed review of this and other properties of intersection numbers).

⁵ There is a small caveat: In the case $k = 0$, an intersection point in $f(\sigma) \cdot \tau$ may have several preimages in σ and should be counted with the corresponding multiplicity; equivalently, the intersection number is defined as the number of points in $\sigma \cap f^{-1}(\tau)$ modulo 2.

Lemma 16

$$f^{\#}(\partial a) = \delta f^{\#}(a).$$

For the proof of the main theorem, we need the following definition:

Definition 17 (*Chain-cochain homotopy*) Consider two chain-cochain maps $\varphi, \psi : C_k(M) \rightarrow C^{d-k}(X)$ from the (non-augmented) chain complex of M to the cochain complex of X . A *chain-cochain homotopy* between φ and ψ is a family of linear maps $h : C_k(M) \rightarrow C^{d-k-1}(X)$ such that $\varphi - \psi = h\partial + \delta h$. To keep track of the various maps, it is convenient to keep in mind the following diagram:

$$\begin{array}{ccccccccccc}
 0 & \longrightarrow & C_d(M) & \xrightarrow{\partial} & C_{d-1}(M) & \xrightarrow{\partial} & \cdots & \xrightarrow{\partial} & C_1(M) & \xrightarrow{\partial} & C_0(M) & \longrightarrow & 0 \\
 & & \searrow h & & \searrow h & & & & \searrow h & & \searrow h & & \\
 & & \varphi \downarrow & & \varphi \downarrow & & & & \varphi \downarrow & & \varphi \downarrow & & \\
 & & \psi & & \psi & & & & \psi & & \psi & & \\
 & & \downarrow & & \downarrow & & & & \downarrow & & \downarrow & & \\
 0 & \longrightarrow & C^0(X) & \xrightarrow{\delta} & C^1(X) & \xrightarrow{\delta} & \cdots & \xrightarrow{\delta} & C^{d-1}(X) & \xrightarrow{\delta} & C^d(X) & \longrightarrow & 0
 \end{array} \tag{9}$$

3 Proof of the overlap theorem

Proof of Theorem 8 Let μ and ε_0 be parameters that we will determine in the course of the proof. We assume that X satisfies the assumptions of the theorem, in particular that it is locally ε -sparse for some $\varepsilon \leq \varepsilon_0$.

Let $f : X \rightarrow M$ be a map. By the discussion in Sect. 2.2 and by Lemmas 12 and 14, we may assume that f is PL and in general position with respect to a sufficiently fine PL triangulation T of M .

We wish to show that there is a vertex v of T such that the intersection number cochain $f^{\#}(v) \in C^d(X)$ satisfies $\|f^{\#}(v)\| \geq \mu$. We assume that this is not the case and we proceed to derive a contradiction.

Let v_0 be a fixed vertex of T ; by assumption, $\|f^{\#}(v_0)\| < \mu$. (Note that if f is not surjective then we can choose the triangulation T and v_0 so that $\|f^{\#}(v_0)\| = 0$.)

We define a chain-cochain map⁶

$$G : C_*(T) \rightarrow C^{d-*}(X)$$

by setting $G(v) := f^{\#}(v_0)$ for every vertex v of T and $G(c) = 0$ for every $c \in C_k(T; \mathbb{F}_2)$, $k > 0$.

We will construct a *chain-cochain homotopy* $H : C_*(T) \rightarrow C^{d-1-*}(X)$ between $f^{\#}$ and G ; that is, for every k , we construct a homomorphism

$$H : C_k(T) \rightarrow C^{d-1-k}(X)$$

such that

$$f^{\#}(c) - G(c) = H(\partial c) + \delta H(c) \tag{10}$$

for $c \in C_k(T)$. We stress that for this proof, we work with *non-augmented* chain and cochain complexes as in (9), i.e., we use the convention that $C^{-1}(X) = 0$. It follows that $G(c) = 0$ for $k > 0$ and that $H(c) = 0$ for $c \in C_d(M)$.

⁶ That is, a homomorphism $G : C_k(T) \rightarrow C^{d-k}(X)$ for every k such that $G(\partial c) = \delta G(c)$ for $c \in C_k(T)$.

The chain-cochain homotopy H will yield the desired contradiction: Given the triangulation T of M , the formal sum of all d -dimensional simplices of T is a d -dimensional cycle ζ_M (here we use that M has no boundary). Note that $f^{\frown}(\zeta_M) = \mathbb{1}_X^0$ (every vertex v of X is mapped into the interior of a unique d -simplex of M) but $G(\zeta_M) = 0$. This is a contradiction, since

$$0 \neq \mathbb{1}_X^0 = f^{\frown}(\zeta_M) - G(\zeta_M) = \underbrace{H(\partial\zeta_M)}_{=0 \text{ since } \partial\zeta_M=0} + \delta \underbrace{H(\zeta_M)}_{=0} = 0.$$

To complete the proof, it remains construct H , which we will do by induction on k .

For $k = 0$, we observe that for every vertex v of T , the cochains $f^{\frown}(v)$ and $G(v) = f^{\frown}(v_0)$ are cohomologous, i.e., their difference is a coboundary: We assume that M is connected, hence there is a 1-chain (indeed, a path) c in T with $\partial c = v - v_0$, and so $f^{\frown}(v) - G(v) = f^{\frown}(v - v_0) = \delta f^{\frown}(c)$. For every vertex v of T , we set $H(v)$ to be a cofilling of $f^{\frown}(v) - G(v)$ of minimal norm (if there is more than one minimal cofilling, we choose one arbitrarily). Thus, the homotopy condition (10) is satisfied for 0-chains (since chains and cochains of dimension less than zero or larger than d are, by convention, zero).

By choice of $H(v)$ and the coisoperimetric assumption on X , we have

$$\|H(v)\| \leq L \underbrace{\|f^{\frown}(v) - f^{\frown}(v_0)\|}_{<2\mu} < s_0 := 2L\mu.$$

Inductively, assume that we have already defined H on chains of dimension less than k and that $\|H(\rho)\| < s_i$ for every i -simplex of T , $i < k$, where s_i is a parameter that we will determine inductively. Thus, if τ is a k -simplex of T , then $H(\partial\tau)$ is already defined and has norm less than $(k + 1)s_{k-1}$.

Moreover, we have $\|f^{\frown}(\tau)\| \leq \frac{d}{k}\varepsilon \leq d\varepsilon$, by the sparsity assumption on X and since the triangulation T is sufficiently fine.

By construction, $z := f^{\frown}(\tau) - H(\partial\tau)$ is a $(d - k)$ -dimensional cocycle, and

$$\|z\| \leq \|f^{\frown}(\tau) - H(\partial\tau)\| < d\varepsilon + (k + 1)s_{k-1}. \tag{11}$$

If z is cohomologically trivial, i.e., $z \in B^{d-k}(X)$, then we define $H(\tau)$ to be a minimal cofilling of z and extend H to $C_k(T)$ by linearity. By assumption on X , we get

$$\|H(\tau)\| < s_k := L(d\varepsilon + (k + 1)s_{k-1}).$$

Note that this recursion yields $s_k = d\varepsilon(L + \dots + L^k) + (k + 1)!L^{k+1}2\mu$.

If z is nontrivial,⁷ then by the assumption on large cosystoles and (11),

$$\vartheta \leq \|z\| < d\varepsilon + (k + 1)s_{k-1},$$

which is a contradiction if we choose μ and ε_0 (and hence ε) sufficiently small with respect to d, L and ϑ . □

Remarks 18 1. In many interesting cases, X belongs to an infinite family of complexes for which the local sparsity parameter ε tends to zero as the size of the complex increases.

For instance, if X is the d -skeleton of the n -simplex, $n \rightarrow \infty$, then we have $\varepsilon = O(1/n)$.

For complexes with local sparsity $\varepsilon = o(1)$, the above proof yields $\mu \geq \frac{\vartheta}{2(k+1)!L^k} + o(1)$.

If M is unbounded, then, as remarked in the proof, we can take the vertex v_0 to satisfy

⁷ Note that in the special case that X is connected and $k = d$, the only nontrivial 0-cocycle is $z = \mathbb{1}_X^0$, hence $\|z\| = 1$.

$f^{\#}(v_0) = 0$, which improves the estimate by a factor of 2.

More quantitative information and better bounds on the overlap constant (which are of interest for specific families of complexes, e.g., skeleta of simplices) can be gleaned from the proof by a more refined analysis through the *cofilling profiles* of X [8], which estimate the size of a minimal cofilling of a cocycle b as a possibly nonlinear function of $\|b\|$. Further improvements in the estimates are possible through the notion of *pagodas* [13].

2. The proof of the overlap theorem is very robust and easily generalizes to other settings, in particular to other coefficient rings and other norms. Suppose that R is a fixed ring of coefficients (commutative, with 1), and consider (co)chains and (co)homology with R -coefficients. If R is not of characteristic 2, we need to add some minor assumptions to deal with orientations. First, we need to assume that the target manifold M is R -orientable, i.e., that $H_d(M; R) \cong R$, generated by a fundamental homology class $[M]$. The definition of the intersection number changes slightly: if two oriented linear simplices σ, τ of complementary dimensions in M intersect transversely in a single point, then their orientations determine a local orientation of M , and we set the intersection number $\sigma \cdot \tau$ to be $+1$ or -1 depending on whether this orientation agrees with the chosen global orientation of M or not.

Second, we need to assume that the norm of a cochain is invariant under sign changes in the values of the cochain, i.e., if two k -cochains $c, c' \in C^k(X; R)$ satisfy $c(\sigma) = \pm c'(\sigma)$ for every orientated k -cell σ of X (the signs may be different for different σ), then $\|c\| = \|c'\|$.

With these additional assumptions, the proof of Theorem 8 goes through also for R -coefficients and yields that for every $f: X \rightarrow M$, there exists $p \in M$ such (6) holds.

3. For norms other than the normalized Hamming norm, $\|f^{\#}(p)\| \geq \mu$ does not necessarily imply that (1) holds. For instance, suppose that $R = \mathbb{R}$ and that we work with the ℓ_2 -norm. In this case, large norm $\|f^{\#}(p)\|$ might be caused by a single d -simplex σ such that $f^{\#}(p)(\sigma)$ is a large integer, i.e., $f(\sigma)$ intersects p with large multiplicity. However, this problem does not occur if we impose additional assumptions on the map f , e.g., that $f^{\#}(p)(\sigma)$ is bounded by some constant K in absolute value (e.g., if f is linear, then we can take $K = 1$).
4. We used the assumption that M is piecewise-linear in order to apply standard general position arguments from piecewise-linear topology. We believe that the result holds more generally if M is a homology manifold. General position arguments for homology manifolds are much more subtle, but for the proof we do not really need to perturb the map f to general position (which may not be possible), we only need a general position chain map that is close to the chain map induced by f . We plan to investigate this in more detail in a future paper.

Acknowledgements Open access funding provided by Institute of Science and Technology (IST Austria). We would like to thank the anonymous referees for many helpful remarks concerning the presentation.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Bárány, I.: A generalization of Carathéodory's theorem. *Discret. Math.* **40**(2–3), 141–152 (1982)

2. Boros, E., Füredi, Z.: The number of triangles covering the center of an n -set. *Geom. Dedicata* **17**, 69–77 (1984)
3. Björner, A.: Topological methods. In: Graham, R.L., Grötschel, M., Lovász, L. (eds.) *Handbook of Combinatorics*, vol. 2, pp. 1819–1872. Elsevier, Amsterdam (1995)
4. Burago, D., Eliashberg, Y., Bestvina, M., Forstnerič, F., Guth, L., Nabutovsky, A., Phillips, A., Roe, J.: A few snapshots from the work of Mikhail Gromov. In: Holden, H., Piene, R. (eds.) *The Abel Prize, 2008–2012*, pp. 139–234. Springer, Berlin (2014)
5. Dotterer, D., Kaufman, T., Wagner, U.: On expansion and topological overlap. In: Fekete, S., Lubiw, A. (eds.) *32nd International Symposium on Computational Geometry (SoCG 2016)*, volume 51 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pp. 35:1–35:10. Dagstuhl, Germany, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik (2016)
6. Evra, S., Kaufman, T.: Systolic Expanders of Every Dimension. Preprint, <http://arxiv.org/abs/1510.00839v2>
7. Kaufman, T., Kazhdan, D., Lubotzky, A.: Isoperimetric inequalities for Ramanujan complexes and topological expanders. *Geom. Funct. Anal.* **26**(1), 250–284 (2016)
8. Gromov, M.: Singularities, expanders and topology of maps. Part 2: from combinatorics to topology via algebraic isoperimetry. *Geom. Funct. Anal.* **20**(2), 416–526 (2010)
9. Karasev, R.: A simpler proof of the Boros–Füredi–Bárány–Pach–Gromov theorem. *Discret. Comput. Geom.* **47**(3), 492–495 (2012)
10. Linial, N., Meshulam, R.: Homological connectivity of random 2-complexes. *Combinatorica* **26**(4), 475–487 (2006)
11. Mabillard, I., Wagner, U.: Eliminating higher-multiplicity intersections, I. A Whitney trick for Tverberg-type problems. Preprint, <http://arxiv.org/abs/1508.02349v1>
12. Matoušek, J.: *Using the Borsuk-Ulam theorem*. Springer, Berlin (2003)
13. Matoušek, J., Wagner, U.: On Gromov’s method of selecting heavily covered points. *Discret. Comput. Geom.* **52**(1), 1–33 (2014)
14. Prasolov, V.: *Elements of Combinatorial and Differential Topology*. American Mathematical Society, Providence (2006)
15. Rourke, C.P., Sanderson, B.J.: *Introduction to Piecewise-Linear Topology*. Springer, New York (1972)
16. Zeeman, E.C.: *Seminar on Combinatorial Topology*. Lecture Notes, Institut des Hautes Études Scientifiques. <http://www.maths.ed.ac.uk/~aar/papers/zeemanpl> (1963)