Volume distortion in groups

HANNA BENNETT

Given a space Y in X, a cycle in Y may be filled with a chain in two ways: either by restricting the chain to Y or by allowing it to be anywhere in X. When the pair (G, H) acts on (X, Y), we define the k–volume distortion function of H in G to measure the large-scale difference between the volumes of such fillings. We show that these functions are quasi-isometry invariants, and thus independent of the choice of spaces, and provide several bounds in terms of other group properties, such as Dehn functions. We also compute the volume distortion in a number of examples, including characterizing the k–volume distortion of $\mathbb{Z}^k$ in $\mathbb{Z}^k \rtimes_M \mathbb{Z}$, where M is a diagonalizable matrix. We use this to prove a conjecture of Gersten.

20F65; 20F67, 57M07

1 Introduction

1.1 Overview

Consider a geodesic metric space X with geodesic subspace Y. Given a pair of points in Y, there are two ways to measure the distance between them: we can consider the minimum of the lengths of paths between them that lie entirely in Y, or we can allow paths to lie anywhere in X. Depending on how Y is embedded in X, the latter distance may be much shorter. This idea can be generalized to higher dimensions: given a $(k-1)$–cycle z, we call the smallest volume of a k–chain whose boundary is z the *filling volume* of z and denote this $\mathrm{FV}^k(z)$. When z lies in Y, there are two possibilities: we may fill z with a chain that lies anywhere in X, giving us $\mathrm{FV}^k_X(z)$, or we might require that the chain be restricted to the subspace Y, which gives us $\mathrm{FV}^k_Y(z)$. How do these two volumes compare? The *volume distortion function* provides a measurement of the difference in a large-scale sense.

We are particularly interested in the case in which X and Y are spaces on which a group and subgroup act cocompactly and properly discontinuously by isometries. We can always find such spaces by constructing a $K(G,1)$ CW–complex that contains a $K(H,1)$ complex and then considering their universal covers. Counting k–cells gives us a combinatorial definition of k–volume in these spaces. If the Eilenberg–Mac Lane

Published: 1 March 2011 DOI: 10.2140/agt.2011.11.655

spaces have a finite k–skeleton, that is, H and G are of type F_k, then we may speak of *subgroup volume distortion*, by which we mean volume distortion in the spaces on which the groups act.

Definition 1 Let H be a subgroup of G, both groups of type F_k, and let X be the universal cover of an Eilenberg–Mac Lane space of G and $Y \subset X$ the universal cover of an Eilenberg–Mac Lane spaces for H. The *k–volume distortion function* function of H in G is a function $\mathrm{VolD}^k_{(G,H)}\colon \mathbb{N} \to \mathbb{N}$ given by

$$\mathrm{VolD}^k_{G,H}(n) = \max\{\, \mathrm{FV}^k_Y(z) \mid z \text{ is a } (k-1)\text{–cycle in } Y \text{ and } \mathrm{FV}_X(z) \le n\}$$

Notice that if the filling volume is the same in the subspace as the ambient space, we get a linear volume distortion function. Thus we say that a subgroup H is *k–volume undistorted* in G if the volume distortion function is linear.

While length distortion is well-understood and area distortion has been studied to some extent (see Gersten [10]), higher-dimensional volume distortion is new.

We prove a number of foundational facts in Section 3: up to linear terms, the distortion functions of two pairs of quasi-isometric CW–complexes are equivalent (Theorem 3.1.1), and thus that volume distortion is independent of the choice of spaces. In this section we also provide bounds in terms of k–th order Dehn functions and discuss the computability of volume distortion functions. We then compute a number of examples in Section 4.

In [10], Gersten proves that the copy of $\mathbb{Z}^2$ is always area-undistorted in $\mathbb{Z}^2 \rtimes_M \mathbb{Z}$ (note that here $M \in \mathrm{GL}(2,\mathbb{Z})$). He gives the following conjecture.

Conjecture 1 (Gersten [10, page 19]) The group $\mathbb{Z}^k$, $k \ge 3$, is area undistorted in $\mathbb{Z}^k \rtimes_M \mathbb{Z}$ if and only if M is of finite order in $\mathrm{GL}(k,\mathbb{Z})$.

In Section 4.4.3 we prove a generalization of this conjecture; we allow M to be any m–by–m integer-entry matrix, and consider the group

$$\Gamma_M = \langle x_1, \dots, x_m, t \mid [x_i, x_j] = 1,\ t x_i t^{-1} = \phi(x_i) \text{ for } 1 \le i,\ j \le m \rangle,$$

where ϕ is a homomorphism taking x_i to $x_1^{a_1} x_2^{a_2} \cdots x_m^{a_m}$, where the a_j form the i–th column of M. When $\det M \ne 0$, Γ_M is an ascending HNN extension of $\mathbb{Z}^m$; in the special case that $\det M = 1$, we can write Γ_M as the semidirect product $\mathbb{Z}^m \rtimes_M \mathbb{Z}$. We then prove the following theorem.

Theorem 1.1.1 *$\mathbb{Z}^m$ is area-undistorted in Γ_M if and only if M has finite order.*

Conjecture 1 is then the special case when $\det M = 1$.

Theorem 1.1.1 is proved by identifying different cases and calculating a lower bound for the area distortion function in each case. This is illustrated in Figure 1, which charts the possible cases and the resulting area distortion in each case.

We generalize this to look at higher volume distortion of $\mathbb{Z}^k$ in Γ_M. In the case of diagonalizable matrices, we can characterize the k–volume distortion completely. In the following theorem, $W(n)$ is the Lambert W function, that is, the inverse of xe^x.

Theorem 1.1.2 *Let M be an integer-entry k–by–k diagonalizable matrix with $\det(M) = d \geq 1$, and let λ_i denote the absolute value of the i–th eigenvalue. Then the k–volume distortion of $\mathbb{Z}^k$ in Γ_M depends only on the eigenvalues of M. If M has at least two eigenvalues off the unit circle, the volume distortion is*

$$\mathrm{VolD}^{(k)}(n) \asymp n^{1+\log d/\log\alpha}, \quad \textit{where } \alpha = \frac{1}{d}\prod_{i=1}^{k}\max\{\lambda_i, d\}.$$

If M has exactly one eigenvalue off the unit circle,

$$\mathrm{VolD}^{(k)}\, n \asymp \left(\frac{n^k}{W(n)}\right)^{1/(k-1)}.$$

Otherwise, $\mathrm{VolD}^{(k)}(n) \asymp n$.

Note that here we obtain a sharp bound. When M is not diagonalizable, we provide a lower bound for volume distortion in Section 4.4.2.

We can look more generally at groups of the form $G = H \rtimes_\phi \langle t\rangle$, where H is any group and ϕ is an automorphism on H. There is a natural surjection to $\mathbb{Z}$, given by the second factor, that allows us a well-defined notion of *height* in the group and in a $\widetilde{K(G,1)}$ such that the height zero subspace is a $\widetilde{K(H,1)}$. Then we can think of ϕ as sending a k–cell at height h to its image under ϕ at height $h-1$. This corresponds to conjugating by t in the group presentation.

The dynamical properties of ϕ can thus be used to find bounds on the volume distortion of H. In Section 4.2 we take the idea of complexity from [10] and alter it, so that the *complexity* of ϕ, denoted $c_k(\phi)$, is the maximal k–volume of the image of a k–cell.

Theorem 1.1.3 *Let ϕ be an automorphism on H, a group of type F_k, and $m = \max\{c_k(\phi), c_k(\phi^{-1})\}$. Then $\mathrm{VolD}^k_{(H\rtimes\mathbb{Z},H)}(n) \leq n\cdot m^n$.*

Corollary 1.1.4 *When ϕ has complexity $m = 1$, then G is k–undistorted in $G \rtimes_\phi \mathbb{Z}$.*

In particular, if a $K(G,1)$ has only one k–cell, then ϕ must send this k–cell to itself, because it induces an automorphism on the k–skeleton. Thus the k–complexity is one, so G is k–undistorted in $G \rtimes_\phi \mathbb{Z}$.

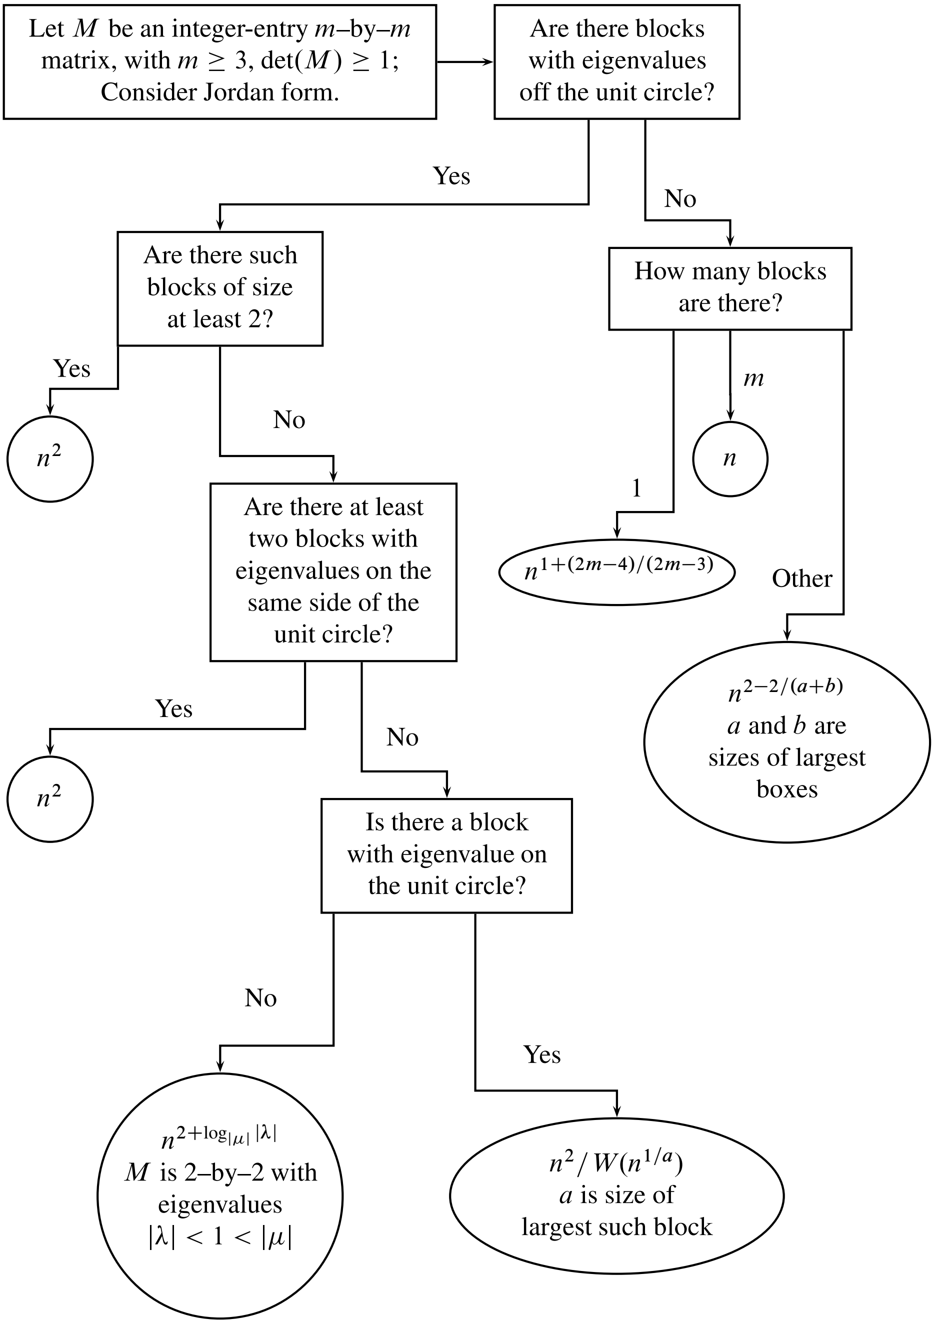

Figure 1: Area distortion flow chart. Circles indicate sharp bounds while ovals are lower bounds only; n^2 is always an upper bound. The function W is the Lambert W function.

Acknowledgements I would like to thank my advisor, Benson Farb, for all his input and support, Noel Brady and Max Forester, for providing a key step to the proof of Section 4.4, and everyone else whose conversations helped make this paper what it is, including Shmuel Weinberger, Robert Young, Nathan Broaddus, Irine Peng and Tim Riley.

2 Background

Let G be a group with presentation $\mathcal{P} = \langle S \mid R \rangle$. We say a word in $F(S)$, the free group on S, is *null-homotopic* if it represents the identity in G, that is, it can be written as a product of conjugates of relators. The *area* of a null-homotopic word is the minimal number of such relators necessary. The Dehn function for G, denoted by $\delta\colon \mathbb{N} \to \mathbb{N}$, is defined by

$$\delta(n) = \{A(w) \mid l(w) \leq n\}$$

which provides an upper bound on the area of a word in terms of its length. While this function appears to depend on the presentation, we can create a relation of functions $f \preceq g$ when there exists some $C > 0$ such that

$$f(x) \leq Cg(Cx + C) + Cx + C \tag{2-1}$$

and we say $f \asymp g$ if $f \preceq g$ and $g \preceq f$. Under this equivalence, the Dehn function is a quasi-isometry invariant, and so in particular independent of presentation.

This function can be used to answer the *word problem*, first asked by Dehn in [6]: given a word in $F(S)$, is there an algorithm for determining whether this word represents the identity? The answer is yes if and only if the Dehn function is computable. However, the algorithm provided by the Dehn function may not be very efficient. For example, if a group with exponential Dehn function can be embedded in a group with quadratic Dehn function, we can use the ambient group to more easily solve the word problem in the subgroup. In such a case, we may think of the embedding as being (area) distorted.

2.1 Definitions

2.1.1 Area distortion In [10], Gersten defines a function, similar to the Dehn function, which measures this area distortion. Precisely, let G be a group with finite presentation $\mathcal{P} = \langle S \mid R \rangle$, and let H be a subgroup with presentation $\mathcal{Q} = \langle S' \mid R' \rangle$, where $\mathcal{Q}$ is a subpresentation of $\mathcal{P}$, that is, $S' \subset S$ and $R' \subset R$. Then the *area distortion function of H in G*, AD: $\mathbb{N} \to \mathbb{N}$, is given by

$$\mathrm{AD}(n) = \max\{\mathrm{Area}_H(w) \mid \mathrm{Area}_G(w) \leq n, w \in N(R')\}.$$

It is not *a priori* clear that such a maximum must exist—perhaps we could find a sequence of words representing the identity in H with area in G bounded by n, but area in H growing arbitrarily large. This, however, cannot happen, precisely because G and H are finitely presented.

Proposition 2.1.1 (Gersten) *The area distortion function is well-defined.*

Proof Let m be the length of the longest relator in R, and let w be a word with van Kampen diagram of area at most n. Separate the diagram into a collection of topological circles, each with area $n_1, n_2, \ldots, n_k$ (note that k and the n_i are all bounded above by n). In such a topological circle of area n_i, the length of the boundary cannot be more than $m \cdot n_i$. Thus there is a finite number of possible loops for the topological circle; for each, we can fill in H with some area. Combining these for each i gives an upper bound on the area in H of the word w which depends only on n. Thus $\mathrm{AD}(n)$ is bounded above for each n, so the function is well-defined. □

Note the importance here of dividing w into pieces that contribute to the area. While we cannot bound the length of w, bounding the lengths of these pieces will often suffice for our purposes. We will continue to use this approach to bounding volume distortion, and so it will benefit us to give a name to the boundary of the "area-contributing" pieces of w. Let D be a van Kampen diagram for w. Define the *frontier* of D by $\mathrm{FR}(D) = \partial(D^\circ)$. (Note that generally frontier is used as a synonym for boundary; we are modifying the definition to a subset of the boundary that will play an important role in bounding volumes.)

We will use the same equivalence for distortion functions as used for Dehn functions, given in Equation (2-1). If AD is linear, then we say that the area of H is *undistorted* in G, as this means that there is essentially no advantage to filling in G over restricting to H. While this function is closely related to the Dehn functions of both the group and subgroup, the distortion function cannot in general be written simply as some combination of the Dehn functions of H and G. In Section 4.4, we shall see a class of examples of groups in which the Dehn functions of H and G are quadratic and exponential, respectively, but the area distortion varies between n and n^2.

Another related concept is *length distortion*, often called simply *subgroup distortion*, which compares the lengths of elements in the subgroup to the lengths in the ambient group. These two concepts are independent: groups may have distorted length but undistorted area (eg Sol groups), or undistorted length and distorted area (eg examples constructed in Baumslag et al [3]).

2.1.2 Volume distortion Just as Dehn functions have been generalized to higher dimensions, as "higher order" Dehn functions (see Alonso, Wang and Pride [2] and Chapter 10 of Epstein et al [7]), denoted $\delta^{(k)}$, we would like to generalize from area distortion to volume distortion. For both these definitions, we will need to take a more geometric approach. We will first define volume distortion on CW–complexes, and then define volume distortion for groups in terms of complexes on which they act.

In order for the distortion function to be well-defined, we need to put some conditions on the CW–complex X. The conditions needed are exactly those given in [2, Section 3] for a k–Dehn complex. X is *k–Dehn* if X is k–connected, the m–order Dehn functions are well-defined for $m \leq k$, and there is a uniform bound, say r, on the number of faces on an m–cell, for $m \leq k+1$.

In such a space X, a cellular k–chain is denoted by $z = \sum \alpha_i \sigma_i$ where the α_i are integers and σ_i are k–cells. The *volume* of z is $V^k(z) = \sum |\alpha_i|$. Given a k–cycle z, we define the *filling volume* of z, $\mathrm{FV}^{k+1}(z)$, to be minimal volume over all k–chains which extend z, that is,

$$\mathrm{FV}^{k+1}(z) = \min\{V^{k+1}(u) \mid \partial u = z\}.$$

Since X is k–connected, every cycle z is the boundary of some chain. Note that this is the definition given in [7, Chapter 10], but what we call volume they call mass, and what we call filling volume, they call volume.

Given a subcomplex Y of X, the *k–volume distortion function* $\mathrm{VolD}^k_{(X,Y)}\colon \mathbb{N} \to \mathbb{N}$ is defined by

$$\text{(2-2)} \quad \mathrm{VolD}^k(n) = \max\{\,\mathrm{FV}^k_Y(z) \mid \mathrm{FV}^k_X(z) \leq n, \text{ where } z \text{ is a } (k-1)\text{–cycle in } Y\}$$

The uniform bound on the size of the boundary of an m–cell serves the same purpose as the finite presentation in the definition of AD, that is, it ensures that a maximum exists. In particular, note that while we have no bound on the volume of z, we do obtain rn as a bound on the volume of the frontier of any filling of z.

Let $H \subset G$ be groups and $Y \subset X$ CW–complexes, where G act cellularly on X so that restricting to H gives a well-defined cellular action on Y. The choice of $Y = \widetilde{K(H,1)}$ and $X = \widetilde{K(G,1)}$ with Y in X, will satisfy this. Then the volume distortion of H in G is defined by $\mathrm{VolD}^k_{(G,H)}(n) := \mathrm{VolD}^k_{(X,Y)}(n)$. This definition only makes sense if the spaces are k–Dehn; this will happen when H and G are of type F_k, that is, their $K(\pi,1)$'s have finite k–skeleton.

Note that AD and VolD^2 are actually defined differently: in the former case, the function involves homotopy, while in the latter the function involves homology. VolD^2 is referred to as *weak distortion* in [10].

Both homology and homotopy definitions exist for higher-dimensional Dehn functions as well. If G acts by cellular automorphisms on X, the homological k–th order Dehn function $\delta^{(k)}\colon \mathbb{N} \to \mathbb{N}$ is defined as

$$\delta^{(k)}(n) = \max\{\, \mathrm{FV}^k(z) \mid z \text{ a } (k-1)\text{–cycle with } V^{k-1}(z) \le n\}.$$

In the homotopic definition, we restrict to filling spheres by balls. While these two definitions are often the same in examples in which they are easy to compute, they are not equivalent in general; an exploration of the differences can be found in Young [17]. One would expect the same to be true of volume distortion, but we do not currently have any such examples.

2.1.3 Riemannian manifolds Sometimes we will be able to determine the volume distortion by considering group actions on Riemannian manifolds. In a general sense, the function will work the same way; what changes is the manner in which we define the volume and filling volume.

Because we are using pairs of groups, we must take some care to specify that the action respects this pairing. A group G acts geometrically on a space X if it acts cocompactly and properly discontinuously by isometries. We say the pair (G, H) acts geometrically on (X, Y) if G acts geometrically on X and we can restrict this to a well-defined geometric action of H on Y. Note that in particular we may construct a $K(H, 1)$ inside of a $K(G, 1)$ so that (G, H) will act geometrically on $(\widetilde{K(G, 1)}, \widetilde{K(H, 1)})$; often this is what we will be considering.

Let M be a connected Riemannian manifold with submanifold N, such that a pair of groups (G, H) act on (M, N) properly discontinuously by isometries. We will work with lipschitz K–chains, that is, formal finite sums with coefficients in $\{+1, -1\}$ of maps $f_i\colon \Delta^k \to M$ where f_i is K–lipschitz for some universally fixed K. We choose lipschitz maps so that the functions are differentiable almost everywhere, leading to a well-defined idea of volume, and so that under quasi-isometry the composition with a lipschitz map is a bounded distance away from a lipschitz map.

We find the volume of a k–chain in the following way: for each lipschitz map f, we can consider $D_x f$ at almost every point in the domain. This map sends an orthonormal basis in $T_x\Delta^k$ to a set of vectors in $T_{f(x)}M$. These vectors give a parallelepiped; call its volume V(x). This is the k–dimensional Jacobian of f at x, and can be found by considering the matrix

$$A = \begin{pmatrix} D_x f(e_1) \\ D_x f(e_2) \\ \vdots \\ D_x f(e_k) \end{pmatrix}$$

and taking $V(x) = \sqrt{\det(A \cdot g \cdot A^T)}$, where $g = (g_{i,j})$ is the Riemannian metric.

Now we integrate over $\triangle^k$:

$$V(f) = \int_{\triangle^k} V(x)\, dx$$

More generally, the volume of a k–chain is the sum of the volumes of the component maps f_i.

With this new definition of volume, we may now define the filling volume and distortion function just as before: given a lipschitz $(k-1)$–cycle z,

$$\mathrm{FV}^{(k)}(z) = \inf\{V(u) \mid u \text{ is a lipschitz } k\text{–chain with } \partial u = z\}. \tag{2-3}$$

Before we can define a volume distortion function in this case, we need one last requirement. Again we define the *frontier* of a $(k-1)$–cycle u by $\mathrm{FR}(u) = \partial(z^\circ)$ where z is a minimal-volume filling of u. Choose some $c \in \mathbb{R}^+$. Just as the value of K does not matter so long as our maps are K–lipschitz for some K, this choice of c will not affect the distortion function up to the usual equivalence of functions, which allows us to discuss "the" distortion function without specifying c or K.

The k–volume distortion function (with respect to c) is a function $\mathrm{VolD}^{(k)}_{(M,N)}\colon \mathbb{N} \to \mathbb{N}$ with

$$\mathrm{VolD}(n) = \sup\{\, \mathrm{FV}^{(k)}_N(z) \mid \exists\ k\text{–chain } u \text{ with } V^k_M(u) \le n,\ V^{(k-1)}_N(\mathrm{FR}(u)) \le cn\,\}.$$

Note that, in the case of CW–complexes with a cocompact group action, there is a natural choice of c: the maximal boundary volume of a k–cell. Once again, the restriction on the size of the frontier of a filling gives this function an upper bound (in terms of the Dehn function of the subspace), so a supremum exists.

In the following section we will show that the volume distortion in this case is equivalent to the version obtained by taking a triangulation that is invariant under (G, H).

3 General theory

3.1 Equivalence of definitions

The definition of volume distortion appears to depend on the choice of spaces; however, we will show that up the equivalence of functions given above, it is a quasi-isometry invariant, and therefore in particular a group invariant.

A (K,C)*–quasi-isometric embedding* is a map $f\colon X \to X'$ where

$$\frac{1}{K}d(x,y) - C \leq d(f(x), f(y)) \leq Kd(x,y) + C$$

for all $x, y \in X$. This function is a quasi-isometry given the additional requirement that for all $y \in X'$ there is some $x \in X$ with $d(y, f(x)) \leq C$. This is equivalent to saying that f has a (K,C)–quasi-isometric inverse $g\colon X' \to X$ with $d(x, g(f(x))) \leq C$ and $d(x, f(g(x))) \leq C$.

As with group actions, we will be considering pairs of spaces (X,Y) where $Y \subset X$. We will say the pairs of spaces (X_1, Y_1) and (X_2, Y_2) are quasi-isometric if there is a quasi-isometry $f\colon X_1 \to X_2$ with $f(Y_1) \subseteq Y_2$ and $f|_{Y_1}\colon Y_1 \to Y_2$ also a quasi-isometry.

Theorem 3.1.1 *Suppose that (X_1, Y_1) and (X_2, Y_2) are k–Dehn spaces which are quasi-isometric as pairs. Then the distortion functions VolD_1^k of (X_1, Y_1) and VolD_2^k of (X_2, Y_2) are equivalent.*

Proof Suppose we have (K,C)–quasi-isometries $f\colon (X_1, Y_1) \to (X_2, Y_2)$ and quasi-isometric inverse $g\colon (X_2, Y_2) \to (X_1, Y_1)$.

See Figure 2 for an illustration of this proof. We start, as in Figure 2(a), with a $(k-1)$–cycle z in Y_1, which is filled in X_1 with a k–chain u, where $V(u) = n$. We want to fill z in Y_1 with volume linear in n and VolD_2^k.

In Figure 2(b) we construct a k–chain v in X_2 using $f(u^{(0)})$, that is, the image of the vertices in u. We want v to have boundary in Y_2 and volume at most $A_k n$, where A_k is independent of u. We do this by noting that for any m–cell σ in X_1, we can construct an m–chain $c(\sigma)$ in X_2 with m–volume bounded by some constant, say A_m. Further, if the m–cell was originally in Y_1, then we can construct the new m–chain in Y_2 with the same bound A on the volume. We do this construction by induction on m: when $m = 1$, we simply note that the distance between the image of the boundary points of a cell are at a distance at most $K + C$ apart, and choose a geodesic between these points; this gives a 1–chain with length at most $C + K$.

Now suppose such a construction exists for dimension $(m-1)$ and let σ be an m–cell. Carry out the construction on the boundary of σ. Since the spaces are k–Dehn, there is a universal bound r on the volume of the boundary of any m–cell in X_1 when $m \leq k+1$. Thus we have constructed a $(k-1)$–cycle of $(k-1)$–volume at most $A_{k-1}r$. Again because the spaces are k–Dehn, we know that the k–th order Dehn function of X_2 is well-defined, so the cycle can be filled in X_2 with volume at most $\delta_{X_2}^{(k)}(A_{k-1}r)$. If the original k–cell is in Y_1, then we can do the same thing, now

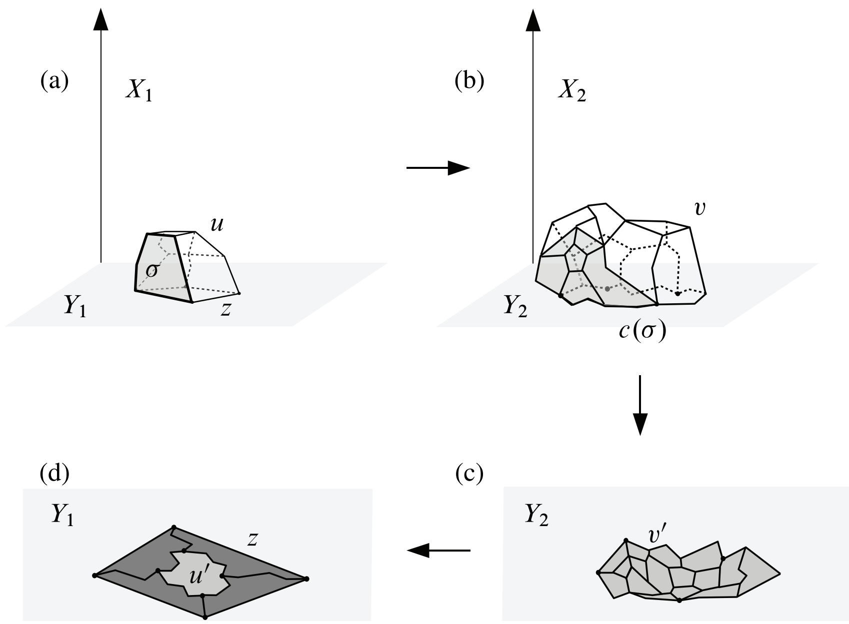

Figure 2: Quasi-isometries preserve volume distortion

filling with volume bounded by $\delta^{(k)}_{Y_2}(A_{k-1}r)$. The maximum of these two numbers is then the value of A_k. This constant depends on the spaces, quasi-isometries, and k, but is independent of u.

Using this procedure on each k–cell in u, we construct a k–chain v in X_2 with boundary in Y_2 and k–volume at most $A_k n$. By definition, the cycle ∂v can be filled in Y_2 with k–volume at most $\mathrm{VolD}_2^k(A_k n)$. Call the chain with this filling v'; note that this chain lies entirely in Y_2. This is illustrated in Figure 2(c). Just as above we can now construct a k–chain u' in Y_1 which contains the 0–skeleton of $g(v')$ and has with volume at most $B_k \cdot \mathrm{VolD}_2^k(A_k n)$, again with B_k independent of u. The boundary of this chain will probably not be z, but each vertex in z will correspond to a vertex in u' which is a distance at most C from the vertex in z; we can use this fact to construct a homotopy between the two cycles z and $\partial u'$. Again we do this by building up by dimension on each cell of z. First we construct paths of length at most C between a vertex in z and the corresponding vertex in u'. Given two vertices in z connected by an edge, their corresponding vertices are connected by a path of length at most $K(K+C)+C$. This gives us a loop of length at most $K(K+C)+3C+1$; this can be filled in with area at most $\delta_{Y_1}(K(K+C)+3C+1)$. See Figure 2(d).

Continue this construction one dimension at a time: Given an m–cell in z, we can construct a homotopy between each of the boundary cells and a corresponding $(m-1)$–chain in u' with some volume bounded by D_{m-1}, again, agreeing on their boundaries. This gives an m–cycle with volume at most $1 + rD_{m-1} + A_m B_m$, which can be filled with m–volume at most $D_k = \delta_{Y_1}^{(m)}(1 + rD_{m-1} + A_m B_m)$. Ultimately this gives us a homotopy between z and $\partial u'$ with k–volume at most $D_k V(z)$. Since $V(z) \leq rn$, combining this with u' gives us a filling of u with volume bounded above by $D_k rn + B_k \operatorname{VolD}_2^k(A_k n)$. Thus $\operatorname{VolD}_1^k(n) \leq D_k rn + B_k \operatorname{VolD}_2^k(A_k n)$. We can use the same process to reverse the roles of VolD_1^k and VolD_2^k, so the two functions are equivalent. □

We now would like to say that given an appropriate action by groups on Riemannian manifolds, the volume distortion function for the groups is equivalent to the volume distortion function of the spaces. To be precise, let M be a $(k-1)$–connected Riemannian manifold, with $(k-1)$–connected submanifold N, and let the pair (G, H) be groups of type F_k that act geometrically as a pair on (M, N). Let τ be a G–invariant triangulation of M, with τ_1 an H–invariant subtriangulation of $\tau|_N$. Then by definition, the distortion function of H in G is the distortion function $\operatorname{VolD}_\tau^k$ of the k–skeleton of τ_1 in τ.

Theorem 3.1.2 *With (G, H), (M, N), and (τ, τ_1) as above, the distortion function of H in G is equivalent to the geometric distortion function* $\operatorname{VolD}_{(M,N)}^k$.

Proof The work done by Burillo and Taback in [5] to show the isoperimetric version of this theorem provides us with all the tools we need for this proof. In particular, they prove that, given M, G, and τ as above, the following holds.

Lemma 3.1.3 (Pushing Lemma [5]) *There exists a constant C, depending only on M and τ, with the following property: Let T be a lipschitz $(k-1)$–chain in M, such that ∂T is included in $\tau^{(k-2)}$. Then there exists another lipschitz $(k-1)$–chain R, with $\partial R = \partial T$, which is included in $\tau^{(k-1)}$, and a lipschitz k–chain S, with $\partial S = T - R$, satisfying $V(R) \leq C V(T)$ and $V(S) \leq C V(T)$.*

Essentially, this is saying that $(k-1)$–chains in M are very near $(k-1)$–chains in $\tau^{(k-1)}$ of comparable area.

Now let z be a $(k-1)$–cycle in N, and let u be a k–chain in M with $\partial u = z$ and $V^k(u) \leq n$. Because z has no boundary, we can apply Lemma 3.1.3 to get a $(k-1)$–cycle z' in τ_1, with homotopy between them given by the chain S in N, where $V^k(S) \leq Cn$. Now z' is a cycle in τ_1; we can fill it with the chain $u - S$, so $\mathrm{FV}(z') \le (C+1)n$. Next we can apply Lemma 3.1.3 to $u - S$; this gives us a chain in τ with volume at most $C(C+1)n$ with boundary z'. By definition, we can fill z' with some chain in τ_1 with volume at most $\mathrm{VolD}_\tau(C(C+1)n)$. Since G acts cocompactly on τ, there is some maximal volume, say A, of any k–cell in τ; so we can now fill z', say by u', in N with volume at most $A\,\mathrm{VolD}_\tau(C(C+1)n)$. Then $u' - S$ fills z with volume at most $A\,\mathrm{VolD}_\tau(C(C+1)n) + Cn$.

On the other hand, suppose z is a cycle in τ_1, with a filling u in τ which has volume at most n. Then we can fill z in M with volume at most An. This means we can fill z in N with a chain u with volume at most $\mathrm{VolD}_{(M,N)}(An)$. By applying Lemma 3.1.3, we can find a filling u' of z in τ_1 with volume at most $C\,\mathrm{VolD}_{(M,N)}(An)$. □

Note that the above proof is independent of the c used to bound the volume of the boundary in the Riemannian manifold case; thus this constant does not affect the distortion function.

By combining Theorem 3.1.1 and Theorem 3.1.2, we obtain the following.

Theorem 3.1.4 *Given pairs of spaces (M_1, N_1) and (M_2, N_2) that are quasi-isometric, and groups (G_1, H_1) and (G_2, H_2) where (G_i, H_i) act cocompactly and properly discontinuously by isometries on (M_i, N_i), the distortion function $\mathrm{VolD}^k_{(M_1,N_1)}$ is equivalent to the function $\mathrm{VolD}^k_{(M_2,N_2)}$.*

Proof Because the group actions are geometric, each pair (G_i, H_i) is quasi-isometric to (M_i, N_i). Since the (M_i, N_i) are quasi-isometric to each other by assumption, the (G_i, H_i) must be as well. Thus by Theorem 3.1.1 their distortion functions are equivalent. By Theorem 3.1.4, the distortion functions of the pairs of Riemannian spaces are equivalent to those of the respective pair of groups. Thus the distortion functions of the spaces are equivalent as well. □

3.2 Distortion and Dehn functions

Because of the closely related definitions of Dehn functions and volume distortion functions, it can be tempting to believe that one can express the distortion function easily in terms of the Dehn functions of the group and subgroup. The reality is not so simple—for example, the subgroup may have a greater or smaller Dehn function than the ambient group. However, we can use Dehn functions to provide certain bounds for volume distortion functions.

Theorem 3.2.1 *Let $H \subset G$ be F_k groups, and let $\delta_H^{(k-1)}$ be the $(k-1)$–order Dehn function of H. Then $\mathrm{VolD}_{(G,H)} \preceq \delta_H^{(k-1)}$.*

Proof Suppose z is a $(k-1)$–cycle in H, with filling in G of k–volume n. Because these groups are F_k, there is an upper bound on the number of boundary faces in a k–cell in G, say r. Then the $(k-1)$–volume of $\mathrm{FR}(u)$ is at most rn. We can then fill $\mathrm{FR}(u)$, and therefore z, in H with volume at most $\delta_H^{(k-1)}(rn)$. □

Note that under the area distortion definition used by Gersten, the upper bound is instead $\mathrm{AD} \preceq n\delta_H$: this happens because the frontier may be disconnected, which introduces the presence of a summation. However, it is conjectured (see, for example, Sapir, Birget and Rips [16]) that Dehn functions are superadditive, in which case we would regain the simpler bound $\mathrm{AD} \preceq \delta_H$. When we choose homology, however, this complication disappears, because there is no requirement that the boundary be connected.

Theorem 3.2.2 *Let $H \subset G$ be F_k groups, and let δ_H and δ_G be their respective $(k-1)$–order Dehn functions. Suppose δ_G is an invertible function. Then $\mathrm{VolD}^k_{(G,H)} \succeq \delta_H \circ \delta_G^{-1}$.*

Proof Let z_n be a sequence of $(k-1)$–cycles with

$$V_H(z) = \delta_G^{-1}(n)$$

and

$$\mathrm{FV}_H(z) = \delta_H(\delta_G^{-1}(n)),$$

that is, a sequence of maximally "hard to fill" cycles. Then we know that

$$\mathrm{FV}_G(z) \le \delta_G^{-1}(\delta_G(n)) = n.$$

□

3.3 Subgroups

Theorem 3.3.1 *Suppose $K \subset H \subset G$ are F_k groups. Then*

(i) $\mathrm{VolD}^k_{(G,K)} \preceq \mathrm{VolD}^k_{(H,K)} \circ \mathrm{VolD}^k_{(G,H)}$.

(ii) $\mathrm{VolD}^k_{(H,K)} \preceq \mathrm{VolD}^k_{(G,K)}$.

Proof For (i), let z be a $(k-1)$–cycle in K so that $\mathrm{FV}_G(z) \le n$. Then $\mathrm{FV}^k_H(z) \le \mathrm{VolD}^k_{(G,H)}(n)$ and so $\mathrm{FV}^k_K(z) \le \mathrm{VolD}^k_{(K,H)}(\mathrm{VolD}^k_{(G,H)}(n))$.

For (ii), once again let z be a $(k-1)$–cycle in K, but now suppose that $\mathrm{FV}_H(z) \le n$ but $\mathrm{FV}_K(z) = \mathrm{VolD}^k_{(K,H)}(n)$. Since $H \subset G$, a filling in H is also a filling in G, so $\mathrm{FV}_G(z) \le n$ as well. Thus we have constructed examples of cycles in K whose filling volume in G is at most n, but whose filling volume in K is $\mathrm{VolD}^k_{(K,H)}(n)$, so this is a lower bound for the distortion of K in G. □

This theorem is of particular interest if one of the embeddings is undistorted.

Corollary 3.3.2 *Given $K \subset H \subset G$ as above,*

(i) *if* $\mathrm{VolD}^k_{(G,H)}$ *is linear, then* $\mathrm{VolD}^k_{(H,K)} \asymp \mathrm{VolD}^k_{(G,K)}$.

(ii) *if* $\mathrm{VolD}^k_{(H,K)}$ *is linear, then* $\mathrm{VolD}^k_{(G,K)} \preceq \mathrm{VolD}^k_{(G,H)}$.

(iii) *if* $\mathrm{VolD}^k_{(G,K)}$ *is linear, then* $\mathrm{VolD}^k_{(H,K)}$ *is linear.*

For an example applying these theorems, see Section 4.3.

3.4 Subcomputability of distortion functions

Papasoglu shows [15, Proposition 2.3] that if $H \subseteq G$ are both finitely presented groups, with a presentation for G that contains a subpresentation for H, then the area distortion function for H in G is computable. This contrasts significantly with length distortion and first-order Dehn functions, in which uncomputable functions can be obtained (see Farb [8]).

In higher dimensions, we can partially generalize.

Theorem 3.4.1 *Given groups $H \subset G$, where (H, G) act cocompactly by properly discontinuous cellular automorphisms as a pair on k–Dehn CW–complexes (Y, X), the function* $\mathrm{VolD}^k_{(G,H)}$ *is bounded above by a computable function for $k \geq 2$.*

Proof To compute a bound on $\mathrm{VolD}^k(n)$, note that to fill a k–cycle, it suffices to fill the connected components of its frontier. We start with a list of all k–cells in the quotient Y/G. Since there are finitely many of them, there are finitely many possible ways to combine at most n of them along boundary cells.

We can make a list of all ways of combining at most n k–cells to make chains in X_G; now pick the subset whose boundary lies entirely in X_H. Because we know these must be trivial in X_H, we can find a volume (and thus a minimal volume) of a filling in X_H. The greatest volume necessary to fill any of the chains on our list will provide an upper bound on the volume distortion. Because we can do this in the process described, this upper bound is a computable function. □

Note that we must know in advance that the groups act on CW–complexes with finite k–skeleton, which requires that we already know something fairly significant about the groups in question. This is still notable, however, as even in pairs of groups with particularly badly-behaved $(k-1)$–order Dehn functions the distortion function itself will still be subcomputable.

While this means that volume distortion is in some ways more nicely behaved than length distortion, Madlener and Otto[13] provide the following theorem:

Theorem 3.4.2 (Madlener–Otto) *Given a computable function f, there exists an example of a subgroup, group pair with area distortion bounded below by f.*

4 Examples

4.1 Motivational examples

There are many geometric examples which demonstrate the concept of area distortion; we will cover some of them here.

4.1.1 Hyperbolic space Consider a horosphere in three-dimensional hyperbolic space $\mathbb{H}^3$. Hyperbolic space has a linear isoperimetric function, so that a loop of length l can be filled in area approximately l. When restricted to the horosphere, however, we encounter Euclidean geometry, for which we have quadratic isoperimetric function. As a result, we can find loops with area l in $\mathbb{H}^3$ but l^2 in the horosphere, giving quadratic area distortion. We cannot hope for larger distortion, since the quadratic isoperimetric function provides an upper bound as well.

4.1.2 Sol geometry Consider the three-dimensional Riemannian manifold Sol that is topologically $\mathbb{R}^3$ and has the metric $ds^2 = \lambda^{-2t}dx^2 + (1/\lambda)^{-2t}dy^2 + dt^2$. If we project y to zero, we obtain a hyperbolic plane; projecting x to zero gives an upside-down hyperbolic plane, that is, length in the y–direction increases exponentially as t increases. Projecting t to zero gives a Euclidean plane. Then the xy–plane is exponentially length-distorted in Sol, since the point on the x axis of distance n from the origin can be reached via a geodesic in the hyperbolic plane in the x and z directions of length approximately $\log n$. The y–axis is similarly distorted, but the geodesics now travel down instead of up. However, when we consider area, the two factors cancel each other out: it is exactly as good in the x–direction to go up as it is bad in the y–direction. This gives us the following well-known result.

Proposition 4.1.1 *The xy–plane is area-undistorted in Sol.*

Proof Given any chain z, we know we can find the volume of z as

$$V(z) = \int \sqrt{(\lambda^{-t}\left(\frac{1}{\lambda}\right)^{-t} dxdy)^2 + (\lambda^{-t}dxdt)^2 + \left(\left(\frac{1}{\lambda}\right)^{-t} dydt\right)^2} \tag{4-1}$$

$$= \int \sqrt{(dxdy)^2 + (\lambda^{-t}dxdt)^2 + \left(\left(\frac{1}{\lambda}\right)^{-t} dydt\right)^2} \tag{4-2}$$

If we project t to zero, we lose the second and third term, which can only decrease the overall volume; the $dxdy$ term is unaffected since the scalars cancelled out. Thus projecting t to zero can only decrease area, which means that the xy–plane is area-undistorted in Sol. Note that here the distortion function is precisely the identity and not just equivalent to it. □

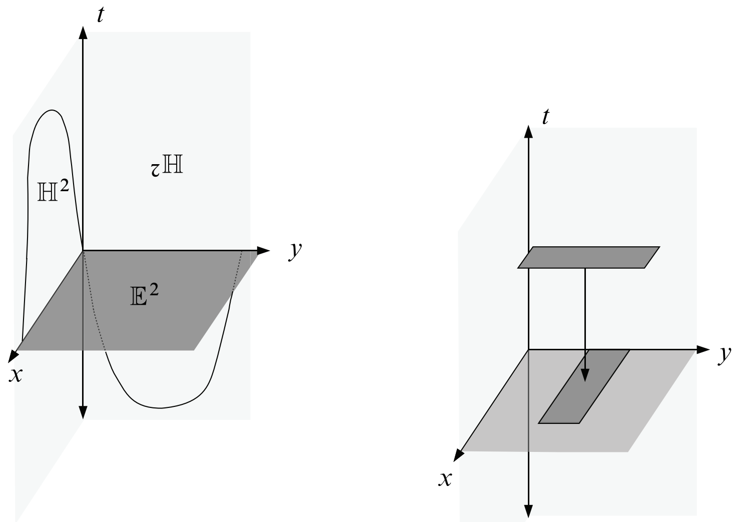

Figure 3: In Sol, length is distorted (left) but area is not.

4.1.3 Sullivan's theorem The example above has a generalization due to Sullivan. If M is a 3–manifold and $\mathcal{F}$ is a codimension 1 foliation on M which is transversely oriented, and such that there is a transverse closed curve through every leaf, then there exists a Riemannian metric on M for which every leaf of $\mathcal{F}$ is quasi-area minimizing. As a special case, $\mathbb{Z}^2$ in $\mathbb{Z}^2 \rtimes_\phi \mathbb{Z}$ has undistorted area when $\phi \in \mathrm{GL}(n, \mathbb{Z})$. Gersten proves this using a concept he calls complexity; we will adjust the definition slightly, and generalize it to higher dimensions.

4.2 Complexity

The concept of *complexity* is defined in [10] for area distortion of G in $G \rtimes_\phi \mathbb{Z}$, where ϕ is an automorphism of G and G is finitely presented.

Let G be F_k and let ϕ an be automorphism of G. Then ϕ can be thought of as a map on the edges of a CW–complex $Y = K(G, 1)$, where an edge labeled s is sent to a word representing $\phi(s)$. Then each 2–cell is sent to a closed loop, and so we can fill it in some way in Y. Choose one of minimal area for each 2–cell, and call this

$\phi(r)$. Continue with this process, inductively extending the map to the n–skeleton of Y given the image of the $(n-1)$–skeleton, until we have a map $\phi\colon Y \to Y$. We can then lift this to a map $\widetilde{\phi}\colon \widetilde{Y} \to \widetilde{Y}$.

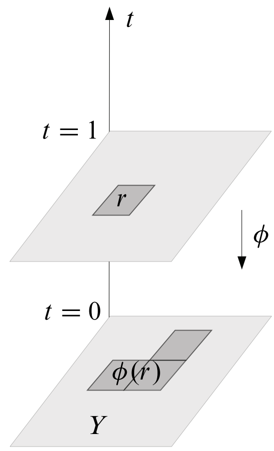

Figure 4: We can project to height zero via iterations of ϕ.

We can use this map to construct a $K(G \rtimes_\phi \mathbb{Z}, 1)$, given by

$$X = Y \times I/[(y,1) \equiv (\phi(y), 0)],$$

which has universal cover that is setwise given by $\widetilde{X} = \widetilde{Y} \times \mathbb{R}$, made by taking a copy of $\widetilde{Y} \times I$ for each integer, and identifying $(y,1)_i$ with $(\phi(y),0)_{i+1}$. This construction gives us a natural projection $\pi\colon \widetilde{X} \to \mathbb{R}$, where we call $\pi(x)$ the *height* of the point $x \in \widetilde{X}$. We can describe cells in this complex as follows: either they are cells inherited from $\widetilde{Y}$, at a height h, or they are built inductively, with two-cells having boundary $tst^{-1}\phi(s)$ and higher-dimensional cells made by constructing the cells on the boundary and then filling them. See Figure 4 for a pictorial representation.

Let $c(\phi)$ be the maximal k–volume of the image of a k–cell under $\widetilde{\phi}$. Since G is F_k, $c(\phi)$ is finite. We will call $c(\phi)$ the *k–complexity* of ϕ.

Remark 1 Gersten defines (2–)complexity in a somewhat different manner: in [10], the complexity of a map is the sum of the volumes of the images of all 2–cells, minus the number of 2–cells. His definition of complexity zero will coincide with our definition of complexity one. With his definition, one can only relate the distortion to complexity when the complexity is zero. We will be able to create a more general upper bound on the distortion, which depends on the value of $c(\phi)$.

Theorem 1.1.3 *Let G be an F_k group and let ϕ be an automorphism on G, and define $m = \max\{c_k(\phi), c_k(\phi^{-1})\}$, with $c_k(\phi)$ as defined above. Then the distortion of G in $G \rtimes_\phi \mathbb{Z}$ is bounded above by $n \cdot m^n$.*

Corollary 1.1.4 *When ϕ has complexity $m = 1$, then G is k–volume undistorted in $G \rtimes_\phi \mathbb{Z}$.*

In particular, if a $K(G,1)$ has only one k–cell, then G is k–volume undistorted in $G \rtimes_\phi \mathbb{Z}$ for any automorphism ϕ.

Proof of Corollary 1.1.4 The first statement is trivial; the second is proven by noting that ϕ must send this k–cell to itself, because it induces an automorphism on the k–homotopy of the $(k-1)$–skeleton of the space. Thus the k–complexity is 1. □

Remark 2 [10, Theorem B] proves the corollary in the case of $k = 2$ and [10, Theorem 5.1] is related to Theorem 1.1.3; however, Gersten formulates his bound in such a way that Corollary 1.1.4 does not follow from Theorem 1.1.3.

Proof of Theorem 1.1.3 Let G and ϕ be as in the theorem, and let $\widetilde{X}$ be the universal cover of a $K(G \rtimes_\phi \mathbb{Z}, 1)$ as constructed above, with the height projection $\pi\colon \widetilde{X} \to \mathbb{R}$. Let z be a $(k-1)$–cycle in $\widetilde{Y}$ with u a k–chain in $\widetilde{X}$, such that $\partial u = z$, and let $n = V^k(u)$. Assume u has no connected components which are cycles, since these could be removed to decrease the volume of u, still giving a cycle with boundary z. Denote by $\triangle_p$ the subset of u at height p, that is, $\triangle_p = \pi^{-1}(p) \cap u$.

Since u is closed, the image of u under π is a bounded subset of $\mathbb{R}$. Note that $\widetilde{Y} = \pi^{-1}(0)$, so if $\pi(u) = 0$, we are done because u is actually a chain in $\widetilde{Y}$. Otherwise, we wish to use u to construct a new chain that is in $\widetilde{Y}$, whose volume is bounded above by nm^n. To do this, we will first consider $\pi^{-1}((0,\infty))$, and then $\pi^{-1}((-\infty,0))$, which will work similarly.

If $\pi^{-1}((0,\infty))$ is empty, we proceed directly to the second set. Otherwise, we may choose some $p \in \mathbb{R}^+$ such that p is not an integer, $\Delta_p \neq \varnothing$, and $\Delta_{p+1} = \varnothing$, that is, a height near the top of u. Let U be the set of all cells in u which intersect Δ_p. This set may consist of a number of different connected components, but we need to break it up a bit more carefully: partition U into subsets $U_1, U_2, \dots, U_l$ so that in each subset, any two cells can be connected to each other in u without going below height p. Then any two cells in the same connected component of U will be in the same partition, but a partition may contain more than one connected component in U.

Each U_i then separates u into two pieces, and gives a homology between some $(k-1)$–cycle v_i at height $h = \lceil p \rceil$ and the $(k-1)$–cycle $\phi(v_i)$ at height $h-1$. The cycle v_i

must be the boundary for some subchain c_i of u; by our choice of p, c_i must be entirely at height h and thus lie in a copy of $\widetilde{Y}$. Then we can perform surgery on u, removing U_i and c_i and replacing them with $\widetilde{\phi}(c_i)$ at height $h-1$. After doing this for each i, we have created a new k–chain u' with height one less than the height of u. We may continue to do this until the maximal height of our new chain is zero.

We then do the same thing with $p < 0$ with the obvious adjustment on the choice of height, where we now use $\widetilde{\phi}^{-1}$ to move the chain upward.

We must now calculate how much this surgery has increased the area. Any k–cell that had been at height h has now been moved to height 0, each time multiplying the volume by at most m, giving a total volume of m^h. How big can h be? At most n, since there must be at least one cell at each height for a top cell to be connected to z. Thus our n k–cells have been replaced by at most nm^n cells at height zero. □

Note that this upper bound is often much larger than the actual distortion. By Theorem 3.2.1, if $m > 1$, then the k–th order Dehn function of the subgroup must be greater than exponential for the complexity bound to be greater than the one provided by the Dehn function. One problem is the height—the upper bound of n is almost certainly too large. In Section 4.4, we will find other ways to bound the height in a particular class of examples so that m^h can be made much smaller.

4.3 Heisenberg groups

One special case of a group $G = \mathbb{Z}^2 \rtimes_\phi \mathbb{Z}$ is the *Heisenberg group*, where $\phi = \left(\begin{smallmatrix} 1 & 0 \\ 1 & 1 \end{smallmatrix}\right)$.

This is also commonly written as

$$\mathcal{H}^3 = \langle x, y, z \mid [x, y] = z, [x, z] = [y, z] = 1 \rangle.$$

(Note that y and z generate the $\mathbb{Z}^2$, and conjugation by x corresponds to the given isometry.) We can create "higher-dimensional" Heisenberg groups $\mathcal{H}^{2n+1}$ with pairs of generators x_i, y_i, along with z, such that each commutator $[x_i, y_i]$ is z, and all other pairs commute. Note that any Heisenberg group is embedded in Heisenberg groups of higher order.

These groups are interesting to us in part because, while $\mathcal{H}^3$ has cubic Dehn function (see [7]), all higher dimensional Heisenberg groups have quadratic Dehn function, a theorem proven analytically by Allcock in [1] and later combinatorially by Olshanskii and Sapir in [14]. Thus, by Theorem 3.2.2, the distortion of $\mathcal{H}^3$ in $\mathcal{H}^5$ or any higher-order Heisenberg group is at least $n^{3/2}$.

Further, we can construct an upper bound as follows: consider the intermediate group $\mathcal{H}^3 \times \mathbb{Z}$, where $\mathbb{Z}$ is generated by y_2. Then we have

$$\mathcal{H}^3 \subset \mathcal{H}^3 \times \mathbb{Z} \subset \mathcal{H}^5 = (\mathcal{H}^3 \times \mathbb{Z}) \rtimes \mathbb{Z},$$

with the action given by x_2 commuting with $\mathcal{H}^3$ and sending y_2 to $y_2 z$. The first containment is undistorted, so by Corollary 3.3.2 (ii), the distortion of $\mathcal{H}^3$ in $\mathcal{H}^5$ is at most the distortion of $\mathcal{H}^3 \times \mathbb{Z}$ in $\mathcal{H}^5$. Directly applying Theorem 1.1.3 would give us an exponential upper bound, but we can modify it a bit: notice that the automorphism preserves all relators except commutators $[y_2, s]$ for $s = x_1, y_1$, each of which goes to two relators: a copy of itself, and the commutator $[z, y_2]$. Thus repeated applications of the automorphism only increase the image by one. Thus the final volume is at most nh, for height h, which is at most n. Therefore n^2 is an upper bound for the area distortion, an improvement over the bound of n^3 given by the Dehn function of $\mathcal{H}^3$.

Conjecture 2 The area distortion of $\mathcal{H}^3$ in $\mathcal{H}^5$ is $n^{3/2}$.

The reason for this is that the upper bound fails to take into account the "side area" coming from any filling; conceptually, any y_2 edge of height h ought to be creating a side with $(h-t)$ relators at height t, creating a total area of h^2.

4.4 Abelian–by–cyclic groups

Note that in this section, upper bounds are found without regard to the topology of the objects, and lower bounds are given by filling spheres with balls, so that the methods and results described work equally well if distortion is defined via homotopy rather than homology.

A group Γ is *abelian–by–cyclic* if there is an exact sequence $1 \to A \to \Gamma \to \mathbb{Z} \to 1$, where A is abelian. By a theorem of Bieri and Strebel, given a finitely presented, torsion-free abelian–by–cyclic group Γ, there is an $m \times m$ matrix M with integer entries so that Γ has the presentation

$$\Gamma_M = \langle x_1, \dots, x_m, t \mid [x_i, x_j] = 1, t x_i t^{-1} = \phi(x_i) \text{ for } 1 \leq i, j \leq m \rangle,$$

where ϕ is a homomorphism taking x_i to $x_1^{a_1} x_2^{a_2} \cdots x_m^{a_m}$, where the a_j form the i–th column of M. Several ideas from Farb and Mosher [9] will help us to find the k–volume distortion of $\mathbb{Z}^m$ in Γ_M.

We can construct a space X_M on which Γ_M acts properly discontinuously and cocompactly by isometries, so that X_M and Γ_M are quasi-isometric. Topologically, this space is $\mathbb{R}^m \times T_M$, where T_M is a directed tree with one edge entering each vertex and $\det(M)$ edges leaving each vertex. Let each edge of T_M have length 1, and fix a particular vertex v_0. This choice of v_0 gives us a height function from T_M to $\mathbb{R}$, where $h(v_0) = 0$. We can extend this to a height function $h\colon X_M \to \mathbb{R}$, where $h(v_0) = 0$, by projecting to T_M. Then $(\mathbb{Z}^m \rtimes \mathbb{Z}, \mathbb{Z}^m)$ and $(X_M, R^m \times \{v_0\})$ are quasi-isometric as pairs.

We can also consider the continuous Lie group $G_M = \mathbb{R}^m \rtimes_M \mathbb{R}$. Here multiplication is given by $(x,t)\cdot(y,s) = (x + M^t y, t+s)$. This space is a Riemannian manifold with left-invariant metric

$$g_{ij}(x,t) = \begin{pmatrix} (M^{-t})^T M^{-t} & 0 \\ 0 & 1 \end{pmatrix}.$$

While this metric involves M, choosing any power of M will give us a quasi-isometric space, so that we may replace M with M^2 (for example, if $\det(M) < 0$) or M^{-1} (if $0 < \det(M) < 1$; this amounts to flipping the space vertically) as we wish.

We will then be able to integrate using this metric to find volumes. We also have a natural height function on G_M given by the last coordinate.

The spaces X_M and G_M generally differ; if $|\det(M)| > 1$, then there is branching in X_M, and it is not a manifold. But there is a relationship between them that allows us to say that they have the same k–volume distortion. This relationship can be seen in the form of a commutative diagram, as seen in Figure 5.

$$\begin{tikzcd} & X_M \arrow[dl, "g_M"'] \arrow[dd, "h"] \arrow[dr, "\pi_M"] & \\ G_M \arrow[dr, "\pi"'] & & T_M \arrow[dl] \\ & \mathbb{R} & \end{tikzcd}$$

Figure 5: Relationship between X_M and Γ_M

Lemma 4.4.1 *The volume distortion function* VolD_X *of* $\mathbb{R}^m \times \{v_0\}$ *in* X_M *is equivalent to* VolD_G *of* $G_0 = \pi^{-1}(0)$ *in* G_M.

Proof Choose some cross-section $\hat{g}\colon G_M \to X_M$ so that the image contains v_0; this gives an isomorphic embedding of G_M into X_M. Given a chain in X_M with boundary in $\mathbb{R}^m \times \{v_0\}$, we can map it under g_M to a chain in G_M with boundary in G_0. We can then fill the chain in G_0 and use $\hat{g}$ to pull this back to a filling in $\mathbb{R}^m \times \{v_0\}$; this tells us $\mathrm{VolD}_X \preceq \mathrm{VolD}_G$. The same process can be used to show the inequality holds the other way, giving us equivalent functions. □

Note that pulling back via $\widehat{g}$ is possible only because our filling lies entirely at height zero, on which g and $\widehat{g}$ act as inverses. (This issue is important when considering Dehn functions in these groups, where the boundary is no longer restricted to a particular vertex in the tree; see Brady and Forester [4].)

We may now simplify the situation to considering the height zero subspace $\mathbb{R}^m$ inside of G_M. By [9], this group is quasi-isometric to G_N, where N is the *absolute Jordan form* for M, that is, a matrix with the absolute values of eigenvalues along the diagonal and ones and zeroes elsewhere, in accordance with the Jordan form. This quasi-isometry preserves the height-zero subspace, so we may restrict our attention to matrices in Jordan form with positive real eigenvalues.

4.4.1 Diagonalizable matrices Suppose M is in absolute Jordan form, $\det(M) \geq 1$. Call the (i,i)–th entry λ_i. If M has ones on the superdiagonal, the situation gets somewhat more complicated; we shall first restrict our attention to the case that M is diagonal. In this case, the geometry of the resulting Riemannian metric is particularly easy to understand. Topologically, we have a space of the form $\mathbb{R}^m \times \mathbb{R}$, where the last coordinate, denoted t, will be considered the height. The metric either expands (if $\lambda_i < 1$) or contracts (if $\lambda_i > 1$) the x_i direction as the height increases. We have the metric

$$ds^2 = dt^2 + \sum_{i=1}^{m} \lambda_i^{-2t} dx_i^2.$$

Further, given a map $g\colon \sigma^k \to G_M$, with image given by $(g_1, g_2, \dots, t)$, the k–volume of g in G_M is

$$\int_{\sigma^k} \left(\sum \lambda_I^{-2t} |D_I g_x|^2 \right)^{1/2} dx,$$

where I is a choice of k of the basis vectors, $|D_I g_x|$ is the determinant of Dg restricted to those k vectors, λ_I is the product of the λ_i of M for $i \in I$. Note that the choice $m+1$ gives the t direction; thus $\lambda_{m+1} = 1$.

Theorem 4.4.2 *Let M be a diagonal $k \times k$ matrix, with (i,i)–th entry $\lambda_i \in \mathbb{R}^+$ and determinant $d > 1$ and at least two eigenvalues off the unit circle. Then the k–volume distortion of the height-zero copy of $\mathbb{R}^k$ in G_M is the function*

$$\mathrm{VolD}^{(k)}(n) = n^{1+\log d/\log \alpha}, \qquad \text{where } \alpha = \left(\prod_{i=1}^{k} \max\{d, \lambda_i\} \right) \Big/ d.$$

This proof has benefitted, both in scope and simplicity, from ideas provided by Brady and Forester in [4].

The proof will be approached in this way: in order to show that this is the exact distortion function, we want to show that it is both an upper and a lower bound. For the latter we need merely exhibit an example of a cycle whose filling volume in the ambient space is bounded above by n and whose filling volume in the subspace is at least $n^{1+\log d/\log \alpha}$. By the definition of the volume distortion function, this must be a lower bound.

Proving that it is an upper bound is more complicated, and we will do this first. Ideally, we would use an argument analogous to the one that shows that the xy–plane is area-undistorted in Sol: start with a chain with boundary in $\mathbb{R}^k$ and volume n. Take the integral representing this volume, and show that we can control the increase in volume that occurs when the height t is projected to zero. This would work if we could bound the height of the filling by a certain value $h = h(M, n)$, but unfortunately we cannot. Instead, we will only project the piece that is below the height h. This projection may not completely fill the cycle in $\mathbb{R}^k$, but we will show that the chain needed to finish filling the cycle cannot be too big, which we do by containing it in a k–dimensional "box" and bounding the volumes of the sides of this box in order to bound the total volume. We do this by showing that when we project x_i to zero, the projection of the original chain must contain a cylinder whose base is one of these sides and height is h; since this cannot be greater than n, we get the desired bound on the remaining volume. This entire process is illustrated in Figure 6–Figure 10.

Proof To simplify later calculations, define $p_i = d/\lambda_i$ and $p = \prod_{i=1}^{k} \min\{p_i, 1\}$. We can then compute that α can also be written as

$$\alpha = \frac{d^{k-1}}{p};$$

this formulation requires more notation, but will better match the approach taken in the proof.

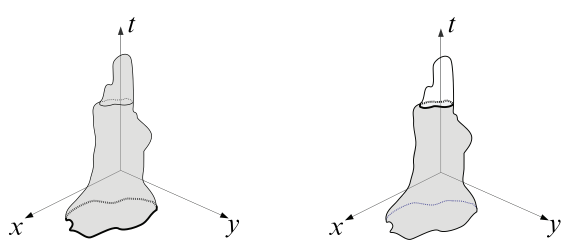

Figure 6: Start with a cycle with a filling of volume n and cut off at height h.

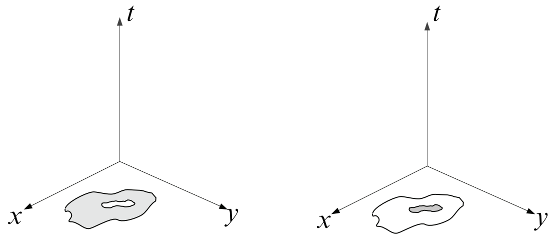

Figure 7: Project low piece down; we want to bound the volume of the pieces not filled.

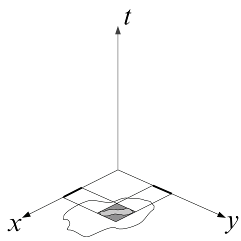

Figure 8: Do this by projecting each coordinate to zero; bound the volumes of the projections.

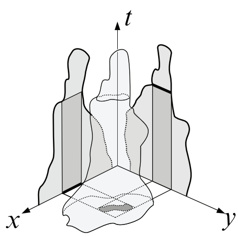

Figure 9: The projection must contain a cylinder; this bounds the volume at height zero.

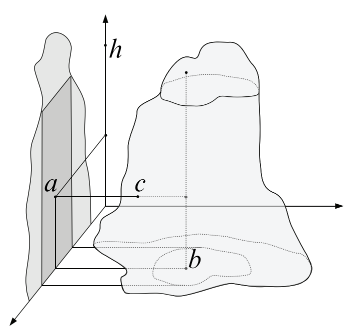

Figure 10: A cylinder with base w_i (and thus base volume v_i) and height h must appear in the projection.

We begin by showing this function is an upper bound. Suppose u is a k–chain in G_M with volume n, and boundary contained in $\mathbb{R}^k$. Then it will suffice to prove that we can fill ∂u in $\mathbb{R}^k$ with volume equivalent to $n^{1+\log d/\log\alpha}$.

First, break u into two pieces: the "low" piece u_L consisting of all of u with height less than $h = \log n/\log\alpha$, and the "high" piece u_H.

The projection π_t sending t to 0 sends u to a filling of ∂u, but increases the volume. We know that the volume of $\pi_t(u_L)$ is an increase of the volume of u_L by a factor of at most $d^h = d^{\log n/\log\alpha} = n^{\log d/\log\alpha}$.

Now we need only bound the total volume of w, the regions interior to ∂u not covered by $\pi(u_L)$. First note that the map π_i projecting x_i to 0 is volume nonincreasing. We will find bounds V_i on the $(k-1)$–volume of $w_i = \pi_i(w)$; this allows us to bound the final volume

$$V(w) \le \Big(\prod V_i\Big)^{1/(k-1)} \tag{4-3}$$

In particular, we will show that if $p_i > 1$, then $V_i \le C_i n$, where C_i is independent of n, and if $p_i < 1$, then $V_i \le C_i n p_i^h$.

Given these bounds on the V_i, notice that when we multiply them together, we get a factor of p_i^h for each $p_i < 1$. This is exactly the definition of p^h. Thus

$$\prod V_i \le C n^k p^h. \tag{4-4}$$

We can use the relationships between p, α, h, and n and basic properties of logarithms to find that

$$p^h = n^{(k-1)\log d/\log\alpha - 1}. \tag{4-5}$$

We can now substitute Equation (4-4) and Equation (4-5) into Equation (4-3) to find

$$V(w) \le (C' n^{(k-1)(1+\log d/\log\alpha)})^{1/(k-1)} = C' n^{1+\log d/\log\alpha}. \tag{4-6}$$

This allows us to fill all of ∂u at height zero with volume at most $(1+C)n^{1+\log d/\log\alpha}$, proving our upper bound.

To finish the proof of the upper bound, it remains to show that these bounds on the V_i are valid. The process we describe is illustrated in Figure 10.

Without loss of generality, assume $i = k$. Let a be a point in the cylinder, say with (x, x_k, t)–coordinates given by $(q, 0, h_0)$. Then $(q, 0, 0)$ is interior to w_k, so there is some q_k such that $b = (q, q_k, 0) \in u$. Now consider the line (q, q_k, t). This will intersect u for the first time at some height $H > h$. In particular, the point (q, q_i, h_0)

is interior to u. Then the line (q, x_i, h) must intersect u at some point c. Thus $\pi_k(c) = a$.

The volume of the cylinder is given by $\int_0^h p_i^{-t} V_i \, dt$. We will restrict the height further: if $p_i > 1$, we will consider only the cylinder from height 0 to 1; if $p_i < 1$ then we will consider the cylinder from height $h-1$ to h. In the former case we have

$$\begin{aligned} \int_0^1 p_i^{-t} V_i \, dt &= -\frac{1}{\log p_i} (p_i^{-t} V_i)|_0^1 \\ &= \frac{V_i}{\log \pi_i}(1 - p_i^{-1}). \end{aligned}$$

Since this volume is less than n, we must have

$$V_i \leq n(\log p_i)(1 - p_i^{-t})^{-1} = C_i n.$$

Similarly, in the latter case we have

$$\begin{aligned} \int_{h-1}^h p_i^{-t} V_i \, dt &= -\frac{1}{\log p_i} (p_i^{-t} V_i)|_{h-1}^h \\ &= \frac{V_i}{|\log pi|} p_i^{-h}(1 - p_i). \end{aligned}$$

Thus $V_i \leq |\log p_i| n p_i^h (1 - p_i)^{-1} = Cnp^h$, as desired.

To show this function is also a lower bound, we need to construct an example exhibiting this amount of distortion. We do this by constructing a $(k-1)$–dimensional box with side lengths chosen so that each projection π_i gives an object with $(k-1)$–volume equivalent to the upper bound V_i found above, that is, $V_i = n$ when $p > 1$, and $V_i = np_i^h$ when $p_i < 1$.

In order to do this, set $l_i = (\prod V_i)^{1/(k-1)} / V_j$. Now build a $(k-1)$–hyper-rectangle with the length in the x_i direction equal to l_i.

This box can be filled in the ambient space by flowing the box up to the height h, where again $h = \log n / \log \alpha$ and then filling the resulting box at height h. This height was chosen so that the volume obtained by flowing each side of the box is some Cn, and the volume at height h is also Cn. Thus the overall volume of this filling is bounded above by $C'n$.

The subspace is Euclidean, so we know the best filling for the box, which is exactly $\prod V_i = Cn^{1+\log d/\log \alpha}$, as calculated in Equation (4-3) through Equation (4-6), which is exactly the value of the volume distortion function. $\square$

In the case M has one eigenvalue off the unit circle, we must amend the bound somewhat. In particular, it is necessary to use the *Lambert W function*, that is, the inverse of the function $e^n n$.

Theorem 4.4.3 *If M is as above, but has exactly one eigenvalue off the unit circle, say $\lambda > 1$, then the k–volume distortion function is $(n^k/W(n))^{1/(k-1)}$, where W is the Lambert W function.*

Proof The reason we cannot use exactly the proof of Theorem 4.4.2 is that when we project x_1 to zero, we obtain a space with a Euclidean metric. Thus the volume of the cylinder will be $h \cdot V_1$.

Because of this, we must change h to $W(n)$ and set $V_1 = n/h$. Notice that all other V_i will still be n, since these projections behave as before. With these changes, the argument for the proof of Section 4.4 works exactly, giving a k–volume distortion function of $\lambda^h n$, which is equivalent to $(n^k/W(n))^{1/(k-1)}$. □

By Lemma 4.4.1 above, the distortion of $\mathbb{Z}^k$ in Γ_M is the same as that of $\mathbb{R}^k$ in G_N, where N is the absolute Jordan form for M, so we immediately have:

Theorem 1.1.2 *Let M be an integer-entry k–by–k diagonalizable matrix with $\det(M) = d \geq 1$, and let λ_i denote the absolute value of the i–th eigenvalue. Then the k–volume distortion of $\mathbb{Z}^k$ in Γ_M depends only on the eigenvalues of M. If M has at least two eigenvalues off the unit circle, the volume distortion is*

$$\mathrm{VolD}^{(k)}(n) \asymp n^{1+\log d/\log \alpha}, \quad \textit{where } \alpha = \Big(\prod_{i=1}^{k} \max\{\lambda_i, d\}\Big)/d.$$

If M has exactly one eigenvalue off the unit circle,

$$\mathrm{VolD}^{(k)}\, n \asymp \left(\frac{n^k}{W(n)}\right)^{1/(k-1)}.$$

Otherwise, $\mathrm{VolD}^{(k)}(n) \asymp n$.

The case in which M has determinant $d = 1$ is covered by Corollary 1.1.4, since in this case M gives an automorphism, and $\mathbb{Z}^k$ has a unique k–cell. This presents one extreme, the case in which volume is undistorted.

At the other extreme, when all eigenvalues are at least 1, and when at least two eigenvectors are greater than one, then the k–volume distortion is maximal, ie $n^{k/(k-1)}$, which is the $(k-1)$–order Dehn function for $\mathbb{Z}^k$.

We may wish to consider the k–volume distortion of $\mathbb{Z}^m$ in Γ_M with $k < m$, that is, the distortion of a smaller-dimensional volume. We can easily bound this below: consider projecting $m-k$ of the x_i to zero; this gives us some $\mathbb{Z}^k \subset G_{M'}$, where M' is a k–by–k matrix. Each projection is volume nonincreasing, so these bound the distortion from below. Thus the largest such distortion out of all choices of $m-k$ dimensions provides a lower bound.

Corollary 4.4.4 *For any integers $m > 1$ and $1 < k < m$, there exists a pair (G, H) with distorted k–volume but undistorted m–volume.*

Proof Simply choose a group $G = \Gamma_M$ with M an m–by–m matrix with $\det(M) = 1$ and at least one eigenvector off the unit circle, and let $H = \mathbb{Z}^m$. □

Ideally, we would like groups that exhibit stronger behavior: for example, a pair in which only the 3–volume is distorted. Examples exist for the area case, (for example, in [10]) but have proven more difficult to construct in general.

4.4.2 Other matrices When the matrix M is not diagonalizable, the situation becomes more complicated. In this case, the automorphism no longer preserves the eigendirections, but also changes lengths along other directions at a rate proportional to a polynomial in the height. This means that techniques involving projection become more difficult to use.

Conjecture 3 Given a matrix M with at least one eigenvalue off the unit circle, the distortion of $\mathbb{Z}^k$ in $\mathbb{Z}^k \rtimes_M \mathbb{Z}$ is the same as that of $\mathbb{R}^k$ in $\mathbb{R}^k \rtimes_N \mathbb{R}$, where N is a diagonal matrix with diagonal entries given by the norms of the eigenvalues of M.

The idea behind this conjecture is that the exponential change in length created by the eigenvalues dominates the polynomial change given by the ones on the superdiagonal, and so, in a large-scale sense, we should be able to ignore the polynomial contribution.

This leaves one more case: when all of the eigenvalues are on the unit circle. In this case, there is no exponential growth coming from the eigenvalues, and so it is the polynomial effect that comes into play.

Such a matrix will have absolute Jordan form in which each block has ones on the diagonal and superdiagonal, and zeroes everywhere else. Thus it can be described completely by the number of blocks, say c, and the size of each block, which we will denote by a_i for $i \in \{1, 2, \dots, c\}$.

Theorem 4.4.5 *Let M be a nondiagonalizable integer-entry m–by–m matrix whose eigenvalues are all on the unit circle. Suppose the absolute Jordan form for M has b blocks, and denote by a_i the size of the i–th block. Then for $k < m$, $\mathbb{Z}^m$ is k–volume distorted in Γ_M.*

In fact, by choosing the last k_i dimensions from the i–th block, so that $\sum k_i = k$, we can bound the distortion below by the function

$$n^{1+\beta/\alpha}$$

where
$$\alpha = (k-1)\sum_i k_i(a_i - k_i) + \frac{1}{2}\sum_i k_i^2 - \frac{k}{2}$$

and
$$\beta = \sum_i k_i(a_i - k_i).$$

We will prove this by providing a family of examples with k–volume of n in G_M and $n^{1+\alpha/\beta}$ in $\mathbb{R}^m$. Essentially, we do this by constructing the top and sides of a box in G_M, with the lengths and directions of the edges chosen in a way that depends on M and the k_i. Then this is a chain whose boundary is a cycle in $\mathbb{R}^m$; we can either fill this with the box given above, or the bottom of the box, which lies entirely in $\mathbb{R}^m$. These will give us the distortion function in question.

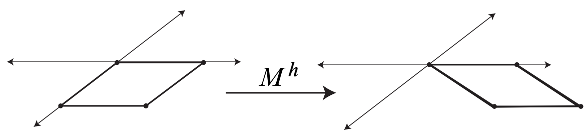

Figure 11: Multiply by M^h.

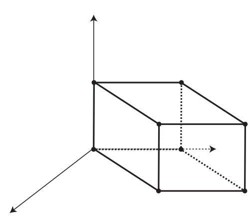

Figure 12: The chain is the top and sides of a box.

Proof As in the diagonalizable case, we need only to find the k–volume distortion of $\mathbb{R}^m$ in $G_M = \mathbb{R}^m \rtimes_M \mathbb{R}$, where M is in absolute Jordan form. We will show that distortion exists by providing a family of examples of k–cycles with a particular distortion. Note that the k–order Dehn function for Euclidean space, $n^{k/(k-1)}$, gives an upper bound on the volume distortion [12]. However, this will generally not match the lower bound we will be constructing.

To construct examples, choose a set $J = \{j_1, j_2, \dots, j_k\} \subset \{1, 2, \dots, m\}$ where the corresponding set of rows of M are subject to the following conditions:

- There is at least one block of size $a_{j_i} > 1$ such that at least one, but not all, rows from that block are included.
- If one row from a block is included, than all later rows in the block are also included.

The former condition is necessary to ensure that distortion exists; the latter will help us to calculate the exact distortion, and ensures that the largest possible distortion is attained in the process that follows.

We will also need k positive real numbers, say $l_1, \dots, l_k$, which will represent the lengths of the sides of the k–box we will be making.

Let $f\colon I^k \to R^m$ map the unit box to the box with edges parallel to the x_{j_i} with length l_i, that is, $f_{j_i}(u) = l_i u_i$ and f_i is zero otherwise. Let $p_i\colon I^{k-1} \to I^k$ be the map embedding I^{k-1} in I^k with $u_i = 0$ and q_i the corresponding map with $u_i = l_i$. From these maps, we construct the final chain: let $F\colon I^k \to G_m$ be the function $(M^h f, h)$, and $F_i = (M^h f \circ p_i, hu_k)$, and $H_i = (M^h f \circ q_i, hu_k)$. The chain will be

$$F \cup \bigcup_{i=1}^{k} F_i \cup \bigcup_{i=1}^{k} H_i.$$

The function F will correspond to the top of the box, and the F_i and H_i to the sides. In the case of two dimensions, we get the picture in Figure 12.

Next we need to compute the volumes for this example. By Section 2.1.3 we know that the volume is

$$\int \sqrt{\det\left((DF_i)^T (M^{-t})^T M^{-t} DF_i\right)}$$

for F_i, and similarly for H_i and F.

Case 1 *M has one block of size m.*

In this case, we must choose $J = \{m-k+1, m-k, \dots, m\}$. Let M_k be the matrix containing just the last k columns of M. Then, for F, $t = h$ and $DF = M^h Df$. So

$$\int \sqrt{\det\left((DF^T (M^{-t})^T M^{-t} DF\right)} = \int \sqrt{\det(DF^T DF)}$$

which is the Euclidean volume, or $l_1 l_2 \cdots l_k$.

Next we note that $DF_i = M^h Df$, and Df is a matrix with l_i in the (j_i, i) position and 0 elsewhere, but t is now allowed to vary. So we have

$$\begin{aligned}\int \sqrt{\det\big((DF_i)^T (M^{-t})^T M^{-t} DF_i\big)} &= \int \sqrt{\det (Df)^T (M^{h-t})^T M^{h-t} Df} \\ &= \int \sqrt{\det (Df)^T (M^v)^T M^v Df} \\ &= l_1 l_2 \cdots l_k \int \sqrt{\det (M_i^v)^T M_i^v}\end{aligned}$$

where M_i is the matrix made of the columns of M corresponding to the elements of J, except j_i. By the Cauchy–Binet theorem, this equals

$$l_1 \cdots l_k \int \sqrt{\sum_\sigma \det M_{i,\sigma}}$$

where $M_{i,\sigma}$ is the $(k-1)\times(k-1)$ matrix consisting of a choice of $k-1$ rows of M_i (with order preserved).

The largest such determinant is found by choosing the last $k-1$ rows of M_i; this gives a matrix of the form

$$\begin{pmatrix} v^{m-k+1}/(m-k+1)! & v^{m-k}/(m-k)! & \cdots & v^{m-1}/(m-1)! \\ v^{m-k}/(m-k)! & v^{m-k+1}/(m-k+1)! & \cdots & v^{m-2}/(m-2)! \\ \vdots & \cdots & \ddots & \vdots \\ v^{m-2k+1}/(m-2k)! & \cdots & \cdots & v^{m-k}/(m-k)! \end{pmatrix}.$$

Note that, from left to right and bottom to top, the power of v increases by one, but will jump and increase by 2 at the i–th column, because the j_i column was left out.

The determinant of this matrix is $v^{(k-1)(m-k)+(k-i))}$. Thus the volume we are seeking is the integral of a polynomial of degree $(k-1)(m-k)+(k-i)$, which gives a final volume of

$$\frac{l_1 \cdots l_k}{l_i} h^{(k-1)(m-k)+k-i+1}.$$

The volume for H_i is identical. Next we want to choose the values for l_i and h so that the distortion is as large as possible. Do this by setting the volumes of the top and sides all equal to each other. In particular, this will happen when

$$l_i = h^{(k-1)(m-k)+k-i+1}.$$

With this choice of l_i, we find that the total volume in G_M is a polynomial of h of degree

$$\sum_{i=1}^{k}((k-1)(m-k)+k-i+1)=k(k-1)(m-k)+k^2+\frac{k^2-k}{2}.$$

Note that this is α when there is only one block.

Next we find the volume in $\mathbb{R}^m$ of the boundary of this chain: this is

$$l_i\cdots l_k\cdots\det(M^h)=h^\alpha\cdot h^{k(m-k)}.$$

If we choose h so that the area in the ambient space is n, then we have $h=n^{1/\alpha}$, and the area in the subspace will be $h^{\alpha+\beta}=n^{1+\beta/\alpha}$.

Case 2 *M has multiple blocks.*

In this case, the calculations proceed much as above, but now the different blocks contribute powers of different sizes. In particular, the volume of F_j is now equivalent to

$$\frac{l_1\cdots l_k}{l_j}h^{\left(\sum_{i=1}^{b}k_i(a_i-k_i)\right)-a_j+k_j+\delta_j}$$

where the sum is taken over the blocks and δ_j is a number between 1 and the number of dimensions chosen for the block containing the j–th choice. The volume of F is still the product of the l_j, so we choose l_i, as before, to be power the of h given in the volume of F_i. Then the volume of the chosen chain is a polynomial in h of degree

$$\sum_{j=1}^{k}\sum_{i=1}^{b}k_i(a_i-k_i)-a_j+k_j+\delta_j=\sum_{i=1}^{b}(k-1)k_i(a_i-k_i)+\frac{k_i^2+k_i}{2}=\alpha.$$

Further, the volume in the subgroup is $l_i\cdots l_k\det(M_J^h)$, where M_J is the $k\times k$ matrix containing the j–th rows of M, with $j\in J$. This is a polynomial in h of degree $\alpha+\beta$, with

$$\beta=\sum_{i=1}^{b}k_I(a_i-k_i).$$

As the k_i have been chosen so that there is some $(a_i-k_i)>1$, this shows that volume distortion exists. $\square$

4.4.3 Area distortion We can combine these results to answer a question of Gersten [10, page 19]:

Theorem 1.1.1 *The group $\mathbb{Z}^m$, $m \geq 3$, is area undistorted in Γ_M if and only if M has finite order.*

See Figure 1 for a flow chart for the various possible area distortion functions.

Proof The latter condition is equivalent to saying that M is diagonalizable and all eigenvalues of M are roots of unity, which, by a theorem of Kronecker, is true if and only if all eigenvalues of M are on the unit circle (see for example [11]). Then by Theorem 4.4.2, $\mathbb{Z}^m$ is undistorted.

Otherwise, let us consider the possible cases.

Case 1 *All eigenvalues are on the unit circle.*

Since no power of M is the identity, it must be the case that some Jordan block of M has ones along the superdiagonal. Then Theorem 4.4.5 gives us a lower bound on area distortion.

Case 2 *There is a block of size more than one with eigenvalue off the unit circle.*

It will suffice in this case to show that area is distorted in the case $M = \left(\begin{smallmatrix} \lambda & 1 \\ 0 & \lambda \end{smallmatrix}\right)$, as this will always be a subgroup in G_M, giving a lower bound on the volume distortion.

Consider the square of side length $n\lambda^{-h}$ at height $h = \log n / \log \lambda$. Projecting this to height zero gives a parallelogram in $\mathbb{R}^2$ with area at least n^2; however, we can fill it in G_M with five parallelograms each of area linear in n. Thus area is quadratically distorted.

Case 3 *There are at least two eigenvalues off the unit circle.*

Then by Theorem 4.4.2 it must be the case that the distortion is nonlinear.

Notice that if there are at least three such eigenvalues, then two must lie on the same side of the unit circle, which means that the area distortion is quadratic, the maximum possible.

Case 4 *There is exactly one eigenvalue off the unit circle.*

In this case, Theorem 4.4.3 tells us the distortion is bounded below by $n^2/W(n)$. The distortion may in fact be higher if we have large blocks associated to unit-length eigenvectors. □

In the case $m = 2$, we can also classify area distortion, though the conditions are different: here area is undistorted if and only if $\det(M) = 1$. Otherwise, M has two eigenvalues, say with absolute values λ and μ, and by Theorem 4.4.2, area distortion is quadratic (maximal) if λ and μ are both greater than one, and $n^{2+\log_\lambda(\mu)}$ if $\lambda > 1 > \mu$ and $\lambda\mu > 1$. If $\lambda > 1$ and $\mu = 1$, then the distortion is $n^2/W(n)$. The examples from Section 4.1.2 and Section 4.1.3 (Sol and Nil geometry) are cases of this sort where $\det(M) = 1$.

While the cases become more complicated with higher dimensions, it should be possible to generalize Gersten's conjecture as follows.

Conjecture 4 Let M be a square matrix of size at least m with nonzero determinant, and let $k < m$. Then the group $\mathbb{Z}^m$ is k–volume undistorted in Γ_M if and only if M has finite order.

References

[1] **D Allcock**, *An isoperimetric inequality for the Heisenberg groups*, Geom. Funct. Anal. 8 (1998) 219–233 MR1616147

[2] **J M Alonso**, **X Wang**, **S J Pride**, *Higher-dimensional isoperimetric (or Dehn) functions of groups*, J. Group Theory 2 (1999) 81–112 MR1670329

[3] **G Baumslag**, **M R Bridson**, **C F Miller, III**, **H Short**, *Finitely presented subgroups of automatic groups and their isoperimetric functions*, J. London Math. Soc. (2) 56 (1997) 292–304 MR1489138

[4] **N Brady**, **M Forester**, *Density of isoperimetric spectra*, Geom. Topol. 14 (2010) 435–472 MR2578308

[5] **J Burillo**, **J Taback**, *Equivalence of geometric and combinatorial Dehn functions*, New York J. Math. 8 (2002) 169–179 MR1934388

[6] **M Dehn**, *Über unendliche diskontinuierliche Gruppen*, Math. Ann. 71 (1911) 116–144 MR1511645

[7] **D B A Epstein**, **J W Cannon**, **D F Holt**, **S V F Levy**, **M S Paterson**, **W P Thurston**, *Word processing in groups*, Jones and Bartlett, Boston (1992) MR1161694

[8] **B Farb**, *The extrinsic geometry of subgroups and the generalized word problem*, Proc. London Math. Soc. (3) 68 (1994) 577–593 MR1262309

[9] **B Farb**, **L Mosher**, *On the asymptotic geometry of abelian-by-cyclic groups*, Acta Math. 184 (2000) 145–202 MR1768110

[10] **S M Gersten**, *Preservation and distortion of area in finitely presented groups*, Geom. Funct. Anal. 6 (1996) 301–345 MR1384614

[11] **G Greiter**, *A simple proof for a theorem of Kronecker*, Amer. Math. Monthly 85 (1978) 756–757 MR514044

[12] **M Gromov**, *Filling Riemannian manifolds*, J. Differential Geom. 18 (1983) 1–147 MR697984

[13] **K Madlener**, **F Otto**, *Pseudonatural algorithms for the word problem for finitely presented monoids and groups*, J. Symbolic Comput. 1 (1985) 383–418 MR849044

[14] **A Y Olshanskii**, **M V Sapir**, *Quadratic isometric functions of the Heisenberg groups. A combinatorial proof*, J. Math. Sci. 93 (1999) 921–927 MR1698761

[15] **P Papasoglu**, *Isodiametric and isoperimetric inequalities for complexes and groups*, J. London Math. Soc. (2) 62 (2000) 97–106 MR1771853

[16] **M V Sapir**, **J-C Birget**, **E Rips**, *Isoperimetric and isodiametric functions of groups*, Ann. of Math. (2) 156 (2002) 345–466 MR1933723

[17] **R Young**, *Homological and homotopical higher-order filling functions*, to appear in Groups, Geom. Dyn. arXiv:0805.0584

Department of Mathematics, University of Michigan
2074 East Hall, 530 Church St, Ann Arbor MI 48109, USA

hbennett@umich.edu

http://www.math.lsa.umich.edu/~hbennett/

Received: 12 February 2010 Revised: 21 December 2010