Geom Dedicata (2016) 180:323–338
DOI 10.1007/s10711-015-0104-6

ORIGINAL PAPER

Using simplicial volume to count multi-tangent trajectories of traversing vector fields

Hannah Alpert[1] · **Gabriel Katz**[1]

Received: 31 March 2015 / Accepted: 27 July 2015 / Published online: 4 August 2015

Abstract For a non-vanishing gradient-like vector field on a compact manifold X^{n+1} with boundary, a discrete set of trajectories may be tangent to the boundary with reduced multiplicity n, which is the maximum possible. (Among them are trajectories that are tangent to ∂X exactly n times.) We prove a lower bound on the number of such trajectories in terms of the simplicial volume of X by adapting methods of Gromov, in particular his "amenable reduction lemma". We apply these bounds to vector fields on hyperbolic manifolds.

Keywords Traversing vector field · Simplicial volume · Simplicial norm · Amenable group

Mathematics Subject Classification (2010) 53C23 · 57N80 · 58K45

1 Introduction

In this paper we consider a smooth vector field v on a space X, which is a compact smooth manifold with boundary, with dim $X = n + 1$. For any such vector field, we may form the space of trajectories, denoted $\mathcal{T}(v)$, of the flow along v, and the quotient map is denoted $\Gamma : X \to \mathcal{T}(v)$. In general $\mathcal{T}(v)$ may not be a nice space, but it is nicer if v is a ***traversing vector field***: a non-vanishing vector field such that every trajectory is either a singleton in ∂X or a closed segment. Figure 1 depicts a traversing vector field on a 2-dimensional space, and the associated trajectory space. One of the authors has explored this general setup in multiple papers beginning with [7], and in the paper [8] he introduces the class of ***traversally generic*** vector fields, which have certain nice properties. In Theorem 3.5 of that paper, he proves that the traversally generic vector fields form an open and dense subset of the traversing

✉ Hannah Alpert
hcalpert@math.mit.edu

Gabriel Katz
gabkatz@gmail.com

1 MIT, Cambridge, MA 02139, USA

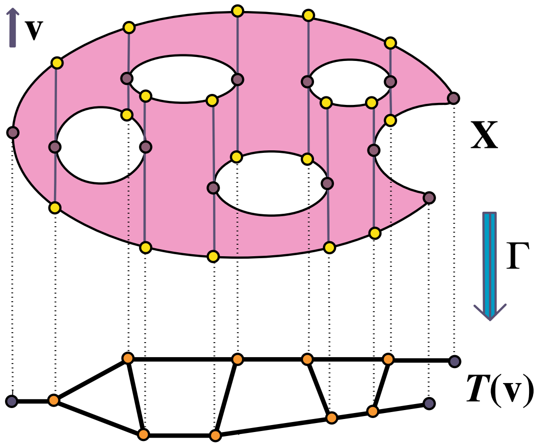

Fig. 1 The typical trajectory of a traversally generic vector field v is a path meeting ∂X twice, each time with multiplicity 1. When X is 2-dimensional, the possible sequences of multiplicities along a trajectory are (11), (2), and (121); in the trajectory space $\mathcal{T}(v)$, these correspond to edge points, vertices of degree 1, and vertices of degree 3

vector fields. Therefore, we study only the traversally generic vector fields; the definition and relevant properties appear in our Sect. 2.

Every traversally generic vector field v has a well-defined ***multiplicity*** $m(a)$ with which v meets ∂X at a point a, and every trajectory γ has a ***reduced multiplicity*** $m'(\gamma)$, which is the sum over all $a \in \gamma \cap \partial X$ of $(m(a) - 1)$. (The full definition of multiplicity appears in Sect. 2.) Every trajectory of a traversally generic vector field v on a manifold X^{n+1} has reduced multiplicity at most n, and so we denote by $\mathsf{max\text{-}mult}(v)$ the set of ***maximum-multiplicity trajectories***; that is, those trajectories γ with $m'(\gamma) = n$.

Theorem 1 *Let M be a closed, oriented hyperbolic manifold of dimension $n + 1 \geq 2$, and let X be the space obtained by removing from M an open set U satisfying the following properties:*

- *The boundary $\partial U = \partial X$ is a closed submanifold of M, possibly with multiple connected components; and*
- *The closure $\overline{U}$ is contained in a topological open ball of M, possibly very far from round.*

Let v be a traversally generic vector field on X. Then we have

$$\# \mathsf{max\text{-}mult}(v) \geq \mathrm{const}(n) \cdot \mathrm{Vol}\, M.$$

In particular, because Vol M is nonzero, there must be at least one maximum-multiplicity trajectory. This theorem generalizes Theorem 7.5 of [7], which addresses the case where $n + 1 = 3$ and U is any finite disjoint union of balls, with constant $1/\mathrm{Vol}(\Delta^3)$, where Δ^3 denotes the regular ideal simplex in hyperbolic 3-space.

Theorem 1 is a special case of the main theorem of this paper. The main theorem is a variant of the theorem "Δ-Inequality for Generic Maps" in Section 3.3 of Gromov's paper [5]. It requires the notion of simplicial volume, which was introduced in [4] and is defined as follows. For every singular chain c with real coefficients, the norm of c, denoted $\|c\|_\Delta$, is the sum of absolute values of the coefficients. For every real homology class h, the ***simplicial norm*** (really a semi-norm) of h, denoted $\|h\|_\Delta$, is the infimum of $\|c\|_\Delta$ over all cycles c representing h. The simplicial norm is often called the ***simplicial volume*** because it generalizes hyperbolic volume: if M is any closed, oriented hyperbolic manifold of dimension n, then $\mathrm{Vol}\, M = \mathrm{const}(n) \cdot \|[M]\|_\Delta$ (Proportionality Theorem, p. 11 of [4]).

Our main theorem is stated as follows. If X is an oriented manifold with boundary, then let $D(X)$ denote the double of X, which is the oriented manifold obtained by gluing two copies of X along their boundary ∂X.

Theorem 2 *Let X be a compact, oriented manifold with boundary, with* $\dim X = n+1$. *Let Z be a space with contractible universal cover, and let $\alpha : D(X) \to Z$ be a continuous map. Assume that for each connected component of the boundary ∂X, the corresponding subgroup of $\pi_1(Z)$ is an amenable group. Then for every traversally generic vector field v, we have*

$$\#\,\mathsf{max\text{-}mult}(v) \geq \mathrm{const}(n) \cdot \|\alpha_*[D(X)]\|_\Delta.$$

That is, the topological quantity $\|\alpha_*[D(X)]\|_\Delta$ is an obstruction to the existence of a traversally generic vector field without maximum-multiplicity trajectories. In Sect. 2 we summarize which properties of traversally generic vector fields are needed in order to apply the methods of Gromov from [5]. In Sect. 3 we present full details for the Amenable Reduction Lemma and Localization Lemma of [5], which are used to prove the "Δ-Inequality for Generic Maps" there—Gromov's presentation is rough, so for the reader's convenience we include a full development of the proofs—and then we prove Theorem 2 and Theorem 1. The new insight of this paper is to bring Gromov's methods to the setting of traversally generic vector fields.

2 Traversally generic vector fields

The purpose of this section is to prove Lemma 1, which describes the nice properties of a traversally generic vector field and which is a consequence of the work of one of the authors in the paper [8]. That paper introduces the definition of traversally generic (Definition 3.2) and proves that the traversally generic vector fields form an open dense subset of the traversing vector fields (Theorem 3.5). The machinery behind the proof of density comes from the theory of singularities of generic maps, in particular from Thom-Boardman theory (see Theorem 5.2 from Chapter VI of [3]). Below, before stating Lemma 1 we give the definitions of traversally generic vector fields and the reduced multiplicity of a trajectory.

The definition of traversally generic includes the notion of boundary generic (Definition 2.1 in [8]), which is defined as follows. Given a traversing vector field v on X, we let $\partial_2 X$ denote the set of points where v is tangent to ∂X. Alternatively, we view $v|_{\partial X}$ as a section of the normal bundle $TX/T\partial X$ of ∂X in X, and let $\partial_2 X$ be the zero locus. If the section corresponding to v is transverse to the zero section, then $\partial_2 X$ is a submanifold of ∂X with codimension 1. Then we repeat the process using the following iterative construction. Let $\partial_0 X = X$ and $\partial_1 X = \partial X$. Once the submanifolds $\partial_j X$ have been defined for all $j \leq k$, we view $v|_{\partial_k X}$ as a section of the normal bundle of $\partial_k X$ in $\partial_{k-1} X$, and if it is transverse to the zero section, then the zero locus $\partial_{k+1} X$ is a submanifold of $\partial_k X$ with codimension 1. We say v is ***boundary generic*** if for all k, when we view $v|_{\partial_k X}$ as a section of the normal bundle of $\partial_k X$ in $\partial_{k-1} X$, this section is transverse to the zero section.

If v is boundary generic, then the ***multiplicity*** $m(a)$ of any point $a \in X$ is defined to be the greatest j such that $a \in \partial_j X$. By definition, if $m(a) = j > 0$, this means that v is tangent to $\partial_{j-1} X$ at a but not tangent to $\partial_j X$ there. Because each $\partial_j X$ has dimension $n+1-j$, the greatest possible multiplicity is $m(a) = n+1$.

Being traversally generic is a property of each trajectory γ of v. Using the v-flow along γ, we may identify all fibers of the normal bundle of γ in X; we denote the resulting quotient by T_*, so that T_* is an n-dimensional vector space. For each point $a_i \in \gamma \cap \partial X$, the tangent space $T\partial_{m(a_i)} X$ is transverse to γ, so it can be viewed as a subspace $\mathsf{T}_i \subseteq \mathsf{T}_*$. We say that a traversing vector field v is ***traversally generic*** if v is boundary generic and if for every trajectory γ of v, the collection of subspaces $\{\mathsf{T}_i(\gamma)\}_i$ is generic in T_*; that is, the quotient map

$$\mathsf{T}_* \to \bigoplus_{a_i \in \gamma \cap \partial X} \mathsf{T}_*/\mathsf{T}_i$$

is surjective. Equivalently, for every subcollection of the subspaces, the sum of their codimensions is equal to the codimension of their intersection (and is in particular nonnegative). Recall that the ***reduced multiplicity*** of every trajectory γ is the sum over all $a_i \in \gamma \cap \partial X$ of $m(a_i) - 1$. Thus, because $\dim \mathsf{T}_* = n$ and $\dim T_i = n - (m(a_i) - 1)$, the property of being traversally generic implies $m'(\gamma) \leq n$.

Lemma 1 describes how every traversally generic vector field v gives rise to stratifications of $\mathcal{T}(v)$ and of X; following [5] we define a ***stratification*** of a space to be any partition with the following property: if a stratum S intersects the closure $\overline{S'}$ of another stratum S', then $S \subseteq \overline{S'}$.

Lemma 1 *Let X be a compact manifold with boundary, with* $\dim X = n+1$. *The traversally generic vector fields v on X satisfy the following properties:*

1. *For* $k = 0, \ldots, n$, *define* $Y_k \subseteq \mathcal{T}(v)$ *by*

$$Y_k := \{\gamma \in \mathcal{T}(v) : m'(\gamma) = n - k\}.$$

   Then every Y_k *is a k-dimensional manifold, the connected components of all* Y_k *constitute a stratification of* $\mathcal{T}(v)$, *and the boundary of each stratum is a union of smaller-dimensional strata.*
2. *Let S be any stratum of* $\mathcal{T}(v)$. *Then the restriction* $\Gamma| : \Gamma^{-1}(S) \to S$ *has the structure of a trivial bundle with fiber equal to either an interval or a point; and the restriction* $\Gamma| : \partial X \cap \Gamma^{-1}(S) \to S$ *has the structure of a finite covering space.*
3. *For* $k = 0, \ldots, n$, *define* $X_k^\partial \subseteq \partial X$ *by*

$$X_k^\partial := \Gamma^{-1}(Y_k) \cap \partial X,$$

   and for $k = 1, \ldots, n+1$, *define* $X_k^\circ \subseteq X \setminus \partial X$ *by*

$$X_k^\circ := \Gamma^{-1}(Y_{k-1}) \setminus \partial X.$$

   Then every X_k^∂ *and* X_k° *is a k-dimensional submanifold, the connected components of all* X_k^∂ *and* X_k° *constitute a stratification of X, and the boundary of each stratum is a union of smaller-dimensional strata.*
4. *There is a finite collection, depending only on the dimension* $n+1$ *and not on v or X, of stratified local models covering X. That is, each local model is an* $(n+1)$*-dimensional stratified space with finitely many strata, and every point in X has a neighborhood diffeomorphic to one of the local models in a way that preserves the stratification.*

The paper [8] proves (Theorem 3.1) an equivalent characterization of traversally generic vector fields, called versal vector fields (Definition 3.5); the main ingredient in the proof is the Malgrange preparation theorem (see, for instance, Theorem 2.1 from Chapter IV of [3]). When we use the description of versal vector fields, our Lemma 1 becomes straightforward. In the remainder of the section, we define versal vector fields and explain why they satisfy the properties in Lemma 1.

For a vector field to be versal means that in a neighborhood of each trajectory there are local coordinates of a certain form. Figure 2 depicts the local geometry of these neighborhoods in the case $\dim X = 3$. In preparation for the definition of a versal vector field, we first

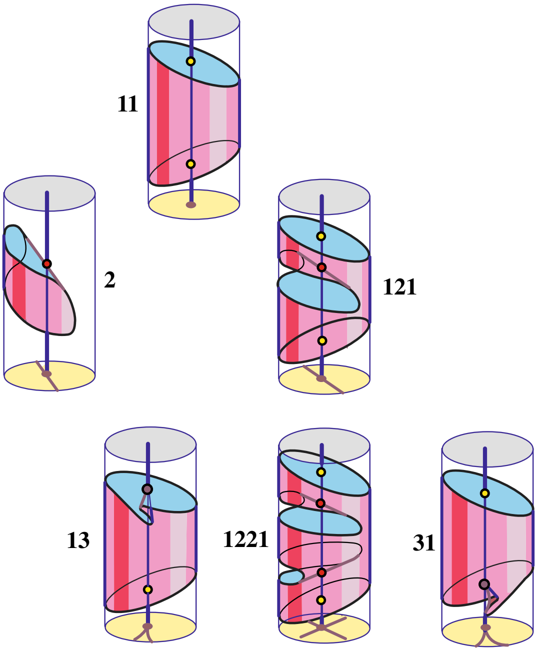

Fig. 2 When X is 3-dimensional, there are six possibilities for the local geometry of the v-flow in the vicinity of a trajectory, corresponding to the sequence of multiplicities with which that trajectory meets ∂X

define local coordinates near one point. For each m with $1 \leq m \leq n+1$, we use variables $u \in \mathbb{R}$, $\mathbf{x} = (x_0, \ldots, x_{m-2}) \in \mathbb{R}^{m-1}$, and $\mathbf{y} \in \mathbb{R}^{n-(m-1)}$, and define

$$P_m(u, \mathbf{x}) = u^m + \sum_{\ell=0}^{m-2} x_\ell u^\ell = u^m + x_{m-2}u^{m-2} + \cdots + x_1 u + x_0.$$

We consider the vector field $\frac{\partial}{\partial u}$ on the space

$$X_+ = \{(u, \mathbf{x}, \mathbf{y}) : P_m(u, \mathbf{x}) \geq 0\}$$

or on the space

$$X_- = \{(u, \mathbf{x}, \mathbf{y}) : P_m(u, \mathbf{x}) \leq 0\}.$$

The trajectories above each fixed $(\mathbf{x}, \mathbf{y})$ stretch between the roots of $P_m(u, \mathbf{x})$ as a function of u. If m is odd, then X_+ has unbounded trajectories in the positive direction (that is, $u \to +\infty$), and X_- has unbounded trajectories in the negative direction ($u \to -\infty$). If m is even, then X_+ has unbounded trajectories in both directions, and X_- has only bounded trajectories. In particular, if m is even, then the trajectory in X_+ through the point $(u, \mathbf{x}, \mathbf{y}) = 0$ is only that one point. The vector fields in these local models are boundary generic, and the multiplicity of each point $(u, \mathbf{x}, \mathbf{y})$ in the sense defined earlier is equal to the multiplicity of vanishing of $P_m(u, \mathbf{x})$ as a function of u.

A ***versal vector field*** is described by local coordinates in a neighborhood of each trajectory γ, as follows. Suppose γ enters X at $a_1 \in \partial X$, and then meets ∂X at $a_2, \ldots, a_p \in \partial X$, in order, exiting at a_p. For each i with $1 \leq i \leq p$, let $\mathbf{x}_i$ denote a variable in $\mathbb{R}^{m(a_i)-1}$, and let $\mathbf{y}$ denote a variable in $\mathbb{R}^{n-m'(\gamma)}$. Then the coordinates are

$$(u, \mathbf{x}_1, \ldots, \mathbf{x}_p, \mathbf{y}) \in \mathbb{R}^{n+1},$$

and X corresponds to the subset

$$\left\{(u, \mathbf{x}_1, \ldots, \mathbf{x}_p, \mathbf{y}) : P_{m(a_i)}(u - i, \mathbf{x}_i) \geq 0 \, \forall i < p, \; P_{m(a_p)}(u - p, \mathbf{x}_p) \leq 0\right\},$$

and v corresponds to the vector field $\frac{\partial}{\partial u}$. The trajectory γ corresponds to the line $(\mathbf{x}_1, \ldots, \mathbf{x}_p, \mathbf{y}) = 0$, and the points $a_1, \ldots, a_p$ correspond to the points $u = 1, \ldots, p$. Note that for X to be nonempty, we must have either $p = 1$ and $m(a_p)$ is even, or $p > 1$ and both $m(a_1)$ and $m(a_p)$ are odd while all other $m(a_i)$ are even.

Proof (Proof of Lemma 1) Because every traversally generic vector field is versal (Theorem 3.1 of [8]), it suffices to check Lemma 1 for the versal vector fields. Part 4 is immediate: there is one local model for each sequence of multiplicities corresponding to reduced multiplicity at most n. Parts 1 and 3 follow from examining the local models: near each trajectory γ, the only trajectories with reduced multiplicity equal to $m'(\gamma)$ are those with all coordinates $\mathbf{x}_i$ equal to 0 (with $\mathbf{y}$ and u varying). To prove Part 2, we see from the local models that Γ is a locally trivial bundle map over each stratum of $\mathcal{T}(v)$, with fiber equal to either an interval or a point. Then, because the interval is oriented, and every bundle of oriented intervals is trivial, the bundle must be trivial. □

3 Main theorem

Let Z be a topological space, and $c \in C_*(Z)$ be a singular cycle. We use the following general strategy to find an upper bound for the simplicial norm $\|[c]\|_\Delta$: first we generate a large number of cycles homologous to c, all with the same norm and with many simplices in common (but with different signs). Then we take the average of these cycles; the result is homologous to c, and because of the cancellation it has small norm.

In order for this cancellation to be possible, we use the simplex-straightening technique from the case where Z is a complete hyperbolic manifold, but we apply a slight generalization for the case where Z is any space with contractible universal cover—that is, $Z = K(\pi_1(Z), 1)$.

Lemma 2 ([4], p. 48) *Let Z be a space with contractible universal cover $\widetilde{Z}$. Then there is a "straightening" operator*

$$\mathsf{straight} : C_*(Z) \to C_*(Z)$$

with the following properties:

- *For each simplex $\sigma \in C_*(Z)$, the straightened version* straight(σ) *is a simplex of the same dimension with the same sequence of vertices.*
- *If two simplices $\sigma_1, \sigma_2 \in C_*(Z)$ have the same sequence of vertices, and their lifts $\widetilde{\sigma}_1, \widetilde{\sigma}_2 \in C_*(\widetilde{Z})$ to the universal cover also have the same sequence of vertices, then* straight$(\sigma_1) =$ straight(σ_2).
- straight *commutes with the boundary map ∂; that is,* straight *is a chain-complex endomorphism.*
- straight *commutes with the standard action of the symmetric group S_{j+1} on each $C_j(Z)$.*
- straight *is chain homotopic to the identity.*

Proof We construct the straightening operator and the chain homotopy simultaneously, one dimension at a time. The chain homotopy will be, for each simplex $\sigma : \Delta \to Z$, a map

$H(\sigma) : \Delta \times I \to Z$ such that the restriction of $H(\sigma)$ to $\Delta \times 0$ is σ, the restriction to $\Delta \times 1$ is $\mathsf{straight}(\sigma)$, and for each face $F \in \partial\Delta$, the restriction of $H(\sigma)$ to $F \times I$ is $H(\sigma|_F)$.

For every 0-dimensional simplex $\sigma^0 \in C_0(Z)$, we have $\mathsf{straight}(\sigma^0) = \sigma^0$ and a constant homotopy $H(\sigma^0)$. For a simplex σ with $\dim \sigma = j > 0$, suppose the straightening and chain homotopy are already defined for every dimension less than j. In particular, $\mathsf{straight}(\partial\sigma)$ depends only on the sequence of vertices of the lift $\widetilde{\sigma}$ of σ to the universal cover $\widetilde{Z}$. We lift $\mathsf{straight}(\partial\sigma)$ to $\widetilde{Z}$; because $\widetilde{Z}$ is contractible, there is some simplex filling in the lift of $\mathsf{straight}(\partial\sigma)$, and we can choose $\mathsf{straight}(\sigma)$ to be the corresponding simplex in Z. We make this choice only once per orbit of $\pi_1(Z) \times S_{j+1}$ on the set $(\widetilde{Z})^{j+1}$ of sequences of vertices in $\widetilde{Z}$.

Having chosen $\mathsf{straight}(\sigma)$, we lift σ, $\mathsf{straight}(\sigma)$, and $H(\partial\sigma)$ to $\widetilde{Z}$ to form a sphere of dimension j. Because $\widetilde{Z}$ is contractible, we can fill in this sphere by a map $\widetilde{H(\sigma)}$ on $\Delta \times I$ that has the prescribed boundary, and let $H(\sigma)$ be the corresponding map into Z. □

We also use the anti-symmetrization operator,

$$\mathsf{symm} : C_*(Z) \to C_*(Z)$$

given by

$$\mathsf{symm}(\sigma^j) = \frac{1}{(j+1)!} \sum_{q \in S_{j+1}} \mathrm{sign}(q) \cdot q(\sigma^j)$$

for every simplex $\sigma^j \in C_j(Z)$. Gromov states that this operator is chain homotopic to the identity ([4], p. 29), and Fujiwara and Manning give the proof in [2].

Lemma 3 ([2], Appendix B) *Let Z be any topological space. The anti-symmetrization operator $\mathsf{symm} : C_*(Z) \to C_*(Z)$ is chain homotopic to the identity.*

Thus, the composition of these two operators

$$\mathsf{symm} \circ \mathsf{straight} : C_*(Z) \to C_*(Z),$$

satisfies the following properties:

1. $\mathsf{symm} \circ \mathsf{straight}(\sigma)$ depends only on the list of vertices of the lift $\widetilde{\sigma}$ to the universal cover.
2. For every $q \in S_{j+1}$ and every $\sigma^j \in C_j(Z)$, we have
$$\mathsf{symm} \circ \mathsf{straight} \circ q(\sigma^j) = \mathrm{sign}(q) \cdot \mathsf{symm} \circ \mathsf{straight}(\sigma^j).$$
3. $\mathsf{symm} \circ \mathsf{straight}$ is a chain map, chain homotopic to the identity.
4. $\mathsf{symm} \circ \mathsf{straight}$ does not increase the norm; that is, for every chain $c \in C_*(Z)$, we have
$$\| \mathsf{symm} \circ \mathsf{straight}(c) \|_\Delta \le \|c\|_\Delta.$$

Property 2 is our reason for introducing symm at all: it allows homotopic simplices with opposite orientations to cancel in a sum or average.

Below we state the setup for the next lemma. Let X be a topological space, and let $c \in C_j(X)$ be a singular cycle of dimension j. We can construct a space Σ from c as follows. For each singular simplex $\sigma_i : \Delta^j \to X$ appearing in c, where Δ^j denotes the abstract j-simplex, there are $j+1$ face maps from Δ^{j-1} to X obtained by restricting σ_i. We form Σ by taking one copy of Δ^j for each σ_i and identifying the faces that have the same face map. (A similar construction appears on pages 108–109 of Hatcher's textbook [6].) Note that every face must be glued to at least one other face, because otherwise it would appear with nonzero coefficient in the linear combination ∂c of face maps, contradicting the cycle hypothesis $\partial c = 0$. Then we can view c as a triple (Σ, c_Σ, f), where c_Σ is a simplicial cycle on Σ, and $f : \Sigma \to X$ is a continuous map such that $c = f_* c_\Sigma$. Note that the space Σ is not necessarily an honest simplicial complex, linearly embeddable in Euclidean space, but is what Hatcher calls a Δ-complex, which may have (for instance) edges in its 1-skeleton that are self-loops (i.e., both endpoints are the same vertex).

By a ***partial coloring*** of c we mean a list $V_1, V_2, \ldots, V_\ell, \ldots$ of disjoint subsets of the set of vertices of Σ. According to the partial coloring we classify each simplex as either essential or non-essential; the non-essential simplices are the ones that can be made to disappear in a certain sense. A simplex σ of Σ is a non-essential simplex of c if either of the following conditions holds:

– σ has two distinct vertices in the same V_ℓ (the vertices are permitted to have the same image in X as long as they are distinct in Σ); or
– σ has two vertices that are the same point of Σ, and the edge between them is a null-homotopic loop in X.

An ***essential simplex*** of c is any simplex σ of Σ that is not non-essential; that is, σ is essential if in the 1-skeleton of σ in Σ, every vertex is in a different V_ℓ, and any edges that are self-loops map to non-contractible loops in X. (In particular, any simplex σ with vertices in $\dim(\sigma) + 1$ different sets V_ℓ is essential.)

Let Z be a space with contractible universal cover and let $\alpha : X \to Z$ be a continuous map that sends all vertices of c in X to the same point of Z. Let Γ_ℓ denote the subgroup of $\pi_1(Z)$ generated by the α-images of the edges of c for which both endpoints are in V_ℓ.

Lemma 4 (Amenable Reduction Lemma, p. 25 of [5]) *Let c be a cycle on X with a partial coloring $\{V_\ell\}$, let Z and $\alpha : X \to Z$ be as above, and suppose that Γ_ℓ is an amenable group for every ℓ. Then the simplicial norm of the α-image of the homology class $[c] \in H_*(X)$ represented by the cycle $c = \sum r_i \sigma_i$ (where $r_i \in \mathbb{R}$ are coefficients and σ_i are simplices) satisfies*

$$\|\alpha_*[c]\|_\Delta \le \sum_{\sigma_i \text{ essential}} |r_i| \,.$$

Proof Given a singular simplex $\sigma \in C_*(X)$ and a path $\gamma : [0, 1] \to X$ beginning at one vertex of σ, there is a homotopy σ_t pushing the vertex along γ; the image of each σ_t is the union of the image of σ with the partial path $\gamma|_{[0,t]}$. Given a singular cycle $c \in C_*(X)$ and a path γ beginning at one vertex of c, we may apply this process to every simplex of c containing that vertex, to obtain a homotopic (and thus homologous) cycle $\gamma * c$. More precisely, if c is (Σ, c_Σ, f) then we modify f by a homotopy supported in a neighborhood of one vertex of Σ. Likewise, we may take a path γ in Z rather than in X, and obtain a cycle $\gamma * \alpha_* c$, for which the straightened cycle $\mathsf{straight}(\gamma * \alpha_* c)$ depends on γ only up to homotopy. That is, if c is (Σ, c_Σ, f) then $\alpha_* c$ is $(\Sigma, c_\Sigma, \alpha \circ f)$ and we homotope $\alpha \circ f$.

Applying this process to every vertex of Σ in $\bigcup_\ell V_\ell$ simultaneously, we obtain an action of the product group $\times_\ell (\Gamma_\ell)^{V_\ell}$ on c, given by

$$g \mapsto \mathsf{symm} \circ \mathsf{straight}(g * \alpha_* c), \;\; g \in \times_\ell (\Gamma_\ell)^{V_\ell}.$$

That is, suppose that g is an element of the product group $\times_\ell (\Gamma_\ell)^{V_\ell}$, and for each vertex v in the union $\bigcup_\ell V_\ell$, let γ_v denote the corresponding coordinate of g; we can think of γ_v as a loop in Z. Then to find $g * \alpha_* c$ we choose disjoint neighborhoods in Σ of all vertices $v \in \bigcup_\ell V_\ell$

and for each v we homotope $\alpha \circ f$ in the chosen neighborhood of v to push $\alpha \circ f(v)$ along γ_v in Z. We will take the average of cycles $\mathsf{symm} \circ \mathsf{straight}(g * \alpha_* c)$ as g ranges over a large finite subset of $\times_\ell (\Gamma_\ell)^{V_\ell}$. To choose this subset, we use the definition of amenable group.

One characterization of (discrete) amenable groups is the Følner criterion: for every amenable group Γ, every finite subset $S \subset \Gamma$, and every $\varepsilon > 0$, there is a finite subset $A \subset \Gamma$ satisfying the inequality

$$\frac{|xA \,\Delta\, A|}{|A|} \leq \varepsilon \;\; \forall x \in S,$$

where Δ denotes the symmetric difference. In our setting, we choose S_ℓ to be the set of α-images of edges in c with both endpoints in V_ℓ, and then apply the Følner criterion to find $A_\ell \subset \Gamma_\ell$. We take the average of $\mathsf{symm} \circ \mathsf{straight}(g * \alpha_* c)$ for $g \in \times_\ell (A_\ell)^{V_\ell} \subset \times_\ell (\Gamma_\ell)^{V_\ell}$; the result is some cycle homologous to $\alpha_* c$ which we show has small norm.

First we show that if σ is *not* an essential simplex of c, then the average of $\mathsf{symm} \circ \mathsf{straight}(g * \alpha_* \sigma)$ has norm at most ε. If one edge of σ is a contractible loop in X, then every $\mathsf{symm} \circ \mathsf{straight}(g * \alpha_* \sigma)$ is equal to 0 (using properties 1 and 2 of $\mathsf{symm} \circ \mathsf{straight}$), so the average is 0. Thus, we address the case where σ has two distinct vertices v_1 and v_2 in some V_ℓ. When averaging over all $g \in \times_\ell (A_\ell)^{V_\ell}$, we average separately over each slice where only the v_1 and v_2 components of g vary and all other components are fixed. It suffices to show that the average of $\mathsf{symm} \circ \mathsf{straight}(g * \alpha_* \sigma)$ over each such slice $A_\ell \times A_\ell$ has norm at most ε.

Having fixed all components of g other than the v_1 and v_2 components, we let $g_{(\gamma_1, \gamma_2)}$ denote the element of $\times_\ell (A_\ell)^{V_\ell}$ for which the v_1 and v_2 components are γ_1 and γ_2 in A_ℓ and the other components have the specified fixed values. Let $x \in \Gamma_\ell$ denote the α-image of the edge in c between v_1 and v_2. Then the edge x in $\alpha_* \sigma$ becomes an edge $\gamma_1^{-1} x \gamma_2$ in $g_{(\gamma_1, \gamma_2)} * \alpha_* \sigma$. Consider the involution

$$(\gamma_1, \gamma_2) \mapsto (x\gamma_2, x^{-1}\gamma_1)$$

on the square subset

$$(\gamma_1, \gamma_2) \in (xA_\ell \cap A_\ell) \times (x^{-1} A_\ell \cap A_\ell) \subset A_\ell \times A_\ell.$$

The path resulting from $(x\gamma_2, x^{-1}\gamma_1)$ is the inverse of the path resulting from (γ_1, γ_2), and thus (using property 2 of $\mathsf{symm} \circ \mathsf{straight}$) we have

$$\sum_{(\gamma_1, \gamma_2) \in (xA_\ell \cap A_\ell) \times (x^{-1} A_\ell \cap A_\ell)} \mathsf{symm} \circ \mathsf{straight}(g_{(\gamma_1, \gamma_2)} * \alpha_* \sigma) = 0.$$

In other words, only those (γ_1, γ_2) outside the square subset contribute to the average. By the Følner criterion we have

$$|xA_\ell \cap A_\ell| \geq \left(1 - \frac{\varepsilon}{2}\right) |A_\ell|,$$

and so

$$\left|(xA_\ell \cap A_\ell) \times (x^{-1} A_\ell \cap A_\ell)\right| \geq \left(1 - \frac{\varepsilon}{2}\right)^2 |A_\ell \times A_\ell| > (1 - \varepsilon) |A_\ell \times A_\ell|.$$

Thus the average over $A_\ell \times A_\ell$ satisfies

$$\frac{1}{|A_\ell \times A_\ell|} \sum_{(\gamma_1,\gamma_2)\in A_\ell\times A_\ell} \mathsf{symm}\circ\mathsf{straight}(g_{(\gamma_1,\gamma_2)} * \alpha_*\sigma)$$
$$\leq \frac{1}{|A_\ell \times A_\ell|} \sum_{(\gamma_1,\gamma_2)\notin(xA_\ell\cap A_\ell)\times(x^{-1}A_\ell\cap A_\ell)} 1 < \varepsilon.$$

Taking the sum over all simplices σ_i of c, we obtain the inequality

$$\|\alpha_*[c]\|_\Delta \leq \sum_{\sigma_i \text{ essential}} |r_i| + \sum_{\sigma_i \text{ not essential}} \varepsilon\, |r_i|\,,$$

and taking the limit as $\varepsilon \to 0$ we obtain the inequality of the lemma statement.

Recall the definition of stratification: if a stratum S intersects the closure $\overline{S'}$ of another stratum S', then $S \subseteq \overline{S'}$. In this case we write $S \preceq S'$. If neither $S \preceq S'$ nor $S' \preceq S$, then we say the two strata are ***incomparable***.

The next lemma involves the notion of ***stratified simplicial norm*** (as with simplicial norm, it is really a semi-norm), which for a homology class h on a stratified space X is the infimum of norms of all cycles c representing h that are consistent with the stratification, in the following sense illustrated in Fig. 3. Gromov gives two conditions: *ord*(er) and *int*(ernality) ([5], p. 27). We use these two conditions plus two more:

- We require that for each simplex of c, the image of the interior of each face (of any dimension) must be contained in one stratum. (This condition may be implicit in Gromov's paper.) We call this the *cellular* condition.
- The (*ord*) condition states that the image of each simplex of c must be contained in a totally ordered chain of strata; that is, the simplex does not intersect any two incomparable strata.
- The (*int*) condition states that for each simplex of c, if the boundary of a face (of any dimension) maps into a stratum S, then the whole face maps into S.
- For technical reasons involving the amenable reduction lemma (Lemma 4), we require that if two vertices of a simplex of c map to the same point $v \in X$, then the edge between them must be constant at v. We call this the *loop* condition.

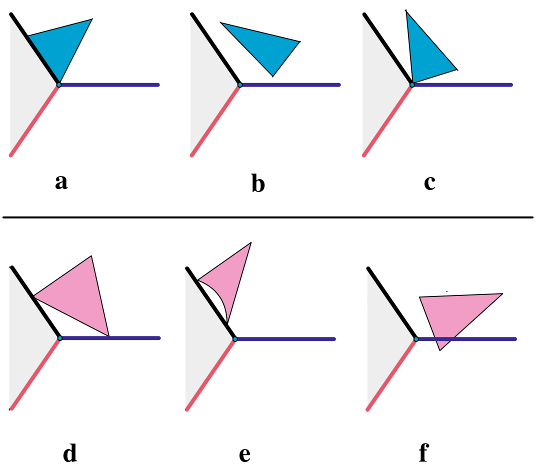

Fig. 3 An example stratification of the plane: one 0-dimensional stratum, three 1-dimensional strata, and three 2-dimensional strata. In diagrams **a**–**c**, the *shaded triangle* satisfies all four criteria for simplices used to compute the stratified simplicial norm. Diagram **d** violates the (*ord*) condition, diagram **e** violates the (*int*) condition, and diagram **f** violates the *cellular* condition

The stratified simplicial norm of a homology class h on a space X with stratification $\mathcal{S}$ is denoted by $\|h\|_\Delta^{\mathcal{S}}$.

Lemma 5 (Localization Lemma, p. 27 of [5]) *Let X be a closed manifold with stratification $\mathcal{S}$ consisting of finitely many connected submanifolds, and let j be an integer between* 0 *and* dim X. *Let Z be a space with contractible universal cover, and let $\alpha : X \to Z$ be a continuous map such that the α-image of the fundamental group of each stratum of codimension less than j is an amenable subgroup of $\pi_1(Z)$. Let $X_{-j} \subseteq X$ denote the union of strata with codimension at least j, and let U be a neighborhood of X_{-j} in X. Then the α-image of every j-dimensional homology class $h \in H_j(X)$ satisfies the bound*

$$\|\alpha_* h\|_\Delta \le \|h_U\|_\Delta^{\mathcal{S}},$$

where $h_U \in H_j(U, \partial U)$ denotes the restriction of h to U, obtained from the composite homomorphism

$$H_j(X) \to H_j(X, X \setminus U) \to H_j(U, \partial U),$$

where the last map is the excision isomorphism.

To prove the localization lemma, we construct a partition of X as follows.

Lemma 6 *Let X be a closed manifold, with a metric space structure and stratified by finitely many submanifolds. For every $\varepsilon > 0$, there is a partition of X consisting of one subset P_S for each stratum S, and some $\delta > 0$, with the following properties:*

- *If S and S' are incomparable strata, then* $\operatorname{dist}(P_S, P_{S'}) > \delta$, $\operatorname{dist}(P_S, S') > \delta$, *and* $\operatorname{dist}(S, P_{S'}) > \delta$.
- *If $x \in P_S$, then* $\operatorname{dist}(x, S) < \varepsilon$.
- *Let $N_\delta(P_S)$ denote the δ-neighborhood of P_S. There is a homotopy beginning with the inclusion $N_\delta(P_S) \hookrightarrow X$ and ending with a map with image in S.*

Proof Figure 4 depicts the relationship between the strata S and the subsets P_S. The sets P_S are determined by a choice of small numbers

$$0 < \varepsilon_{\dim X} < \cdots < \varepsilon_0 < \varepsilon,$$

which we choose inductively.

- Step 0: We choose $\varepsilon_0 < \varepsilon$ such that for every 0-dimensional stratum S_0, the ball $N_{3\varepsilon_0}(S_0)$ is Euclidean and has the following property: if S_0 is disjoint from the closure $\overline{S'}$ of another stratum S' (i.e., if $S_0 \not\le S'$), then $N_{3\varepsilon_0}(S_0)$ is also disjoint from $\overline{S'}$. We put $P_{S_0} = N_{\varepsilon_0}(S_0)$.

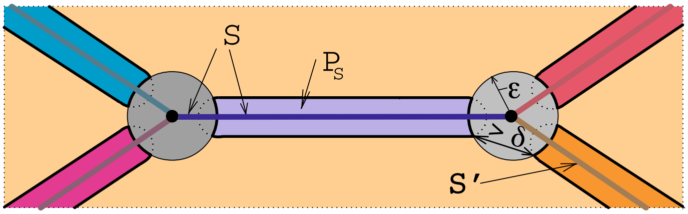

Fig. 4 For each stratum S, its approximation P_S lies within the ε-neighborhood of S. For every two incomparable strata S and S', the distance between the sets P_S and $P_{S'}$ exceeds $\delta = \delta(\varepsilon)$

– Step 1: First we find a tubular neighborhood U_{S_1} of every 1-dimensional stratum S_1, such that if S_1 is disjoint from some $\overline{S'}$, then U_{S_1} is also disjoint from $\overline{S'}$. The portion of S_1 that is outside the union of all P_{S_0} is a compact set. Therefore, we can choose $\varepsilon_1 < \varepsilon_0$ such that for every 1-dimensional stratum S_1, we have

$$N_{3\varepsilon_1}\left(S_1 \setminus \bigcup_{S_0} P_{S_0}\right) \subseteq U_{S_1}.$$

We put

$$P_{S_1} = N_{\varepsilon_1}\left(S_1 \setminus \bigcup_{S_0} P_{S_0}\right) \setminus \bigcup_{S_0} P_{S_0}.$$

– Step k: Having chosen $\varepsilon_{k-1} < \cdots < \varepsilon_0 < \varepsilon$, we choose $\varepsilon_k < \varepsilon_{k-1}$ much as in Step 1. We find a tubular neighborhood U_{S_k} of every k-dimensional stratum S_k, such that if S_k is disjoint from some $\overline{S'}$, then U_{S_k} is also disjoint from $\overline{S'}$. We choose $\varepsilon_k < \varepsilon_{k-1}$ such that for every k-dimensional stratum S_k, we have

$$N_{3\varepsilon_k}\left(S_k \setminus \bigcup_{i=0}^{k-1} \bigcup_{S_i} P_{S_i}\right) \subseteq U_{S_k}.$$

We put

$$P_{S_k} = N_{\varepsilon_k}\left(S_k \setminus \bigcup_{i=0}^{k-1} \bigcup_{S_i} P_{S_i}\right) \setminus \bigcup_{i=0}^{k-1} \bigcup_{S_i} P_{S_i}.$$

– Step $\dim X$: Formally, the procedure from Step k applies. However, we note that each tubular neighborhood $U_{S_{\dim X}}$ is equal to all of $S_{\dim X}$, and so we have

$$P_{S_{\dim X}} = S_{\dim X} \setminus \bigcup_{i=0}^{\dim X-1} \bigcup_{S_i} P_{S_i}.$$

We do choose $\varepsilon_{\dim X}$ just as in Step k.

We choose $\delta < \varepsilon_{\dim X}$. The second and third properties in the lemma statement are immediate. For the first property, suppose S and S' are incomparable. In particular $S \not\preceq S'$. Then all of S' lies at least $3\varepsilon_{\dim S}$ away from $S \cap P_S$, whereas all of P_S lies within $\varepsilon_{\dim S}$ of $S \cap P_S$, and so we have

$$\operatorname{dist}(P_S, S') \geq 2\varepsilon_{\dim S} > \delta,$$

and likewise with S and S' reversed. To check $\operatorname{dist}(P_S, P_{S'})$, assume without loss of generality $\dim S \leq \dim S'$. Then we have

$$\operatorname{dist}(P_S, P_{S'}) \geq 2\varepsilon_{\dim S} - \varepsilon_{\dim S'} \geq \varepsilon_{\dim S} > \delta.$$

Proof (Proof of Lemma 5) Here is the rough idea of the proof: we extend each relative cycle c_U representing h_U to a cycle c representing h. We construct a partial coloring on the vertices of c with one subset V_ℓ for each stratum S_ℓ of codimension less than j. If c is chosen carefully, then every simplex of $c - c_U$ is not essential, so the amenable reduction lemma (Lemma 4) implies that these new simplices do not contribute to the simplicial norm of $\alpha_* h$.

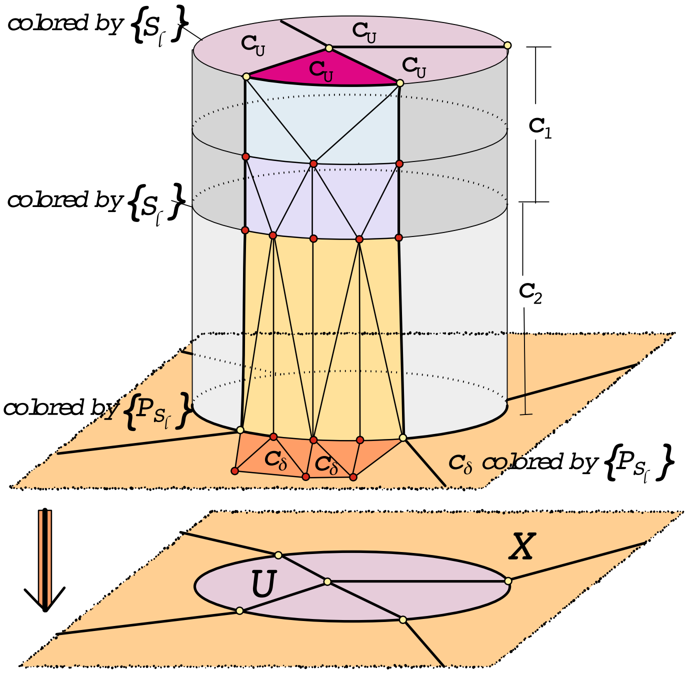

Fig. 5 The singular cycle c is constructed as the sum $c = c_U + c_1 + c_2 + c_\delta$, where c_U is the restriction to U, and c_δ is the restriction to the complement of U, and the cylinders c_1 and c_2 connect c_U to c_δ. The vertices of c_U and c_1 are *colored* according to which stratum S_ℓ they are in, the vertices of c_δ are *colored* according to which subset P_{S_ℓ} they are in, and the cylinder c_2 connects the two coloring methods. (Color figure online)

In fact, the proof gets more complicated because we need to guarantee that the α-images of the edges of c with both endpoints in V_ℓ generate a subgroup of $\alpha_*\pi_1(S_\ell)$, and thus an amenable subgroup of $\pi_1(Z)$. (Every subgroup of an amenable group is amenable.) We construct the partition $\{P_S\}$ with ε chosen to be smaller than the distance from X_{-j} to the complement of U, and use the δ arising from the construction of $\{P_S\}$. Then we construct chains c_1, c_2, and c_δ so that $c = c_U + c_1 + c_2 + c_\delta$ is a cycle homologous to h, by the following method depicted in Fig. 5.

- To construct c_δ, we first take a cycle c' representing h such that c_U is its restriction to U, and then obtain c_δ by iterated barycentric subdivision of $c' - c_U$ so that the diameter of each simplex is less than δ. For each vertex v of c_δ, if $v \in P_{S_\ell}$, then $v \in V_\ell$.
- c_2 is the cylinder $-\partial c_\delta \times [0, 1]$, triangulated in such a way that no new vertices are created, and mapped to X by the projection $-\partial c_\delta \times [0, 1] \to -\partial c_\delta$. The vertices of $-\partial c_\delta \times 1$ are identified with the vertices of c_δ, so their partial coloring is determined by their membership in P_{S_ℓ}. The vertices of $-\partial c_\delta \times 0$ have a different partial coloring: if $v \in S_\ell$, then $v \in V_\ell$.
- c_1 is a subdivision of the cylinder $\partial c_U \times [0, 1]$, mapped to X by the projection $\partial c_U \times [0, 1] \to \partial c_U$. The end $\partial c_U \times 0$ is identified with ∂c_U and is not subdivided. The end $\partial c_U \times 1$ is divided by barycentric subdivision so that it may be identified with the 0 end of c_2, which is equal to $-\partial c_\delta$. The middle of the cylinder is subdivided by concatenating the chain homotopies corresponding to barycentric subdivision, one for each iteration. For each vertex v of c_1, if $v \in S_\ell$, then $v \in V_\ell$.

First we verify that every simplex in c_1, c_2, and c_δ is not essential. The (*ord*) condition on ∂c_U and the choice of δ imply that the labels of the $(j+1)$ vertices of each simplex correspond to S_ℓ in a totally ordered chain of strata. Because each S_ℓ has codimension between 0 and $j-1$, and every two strata of the same dimension are incomparable, two of the vertices must have the same label. If these two vertices are identical in c, then the edge between them must be constant; this results from the *loop* property on c_U and the fact that barycentric subdivision destroys loops.

Next we need to check that the α-images of the edges with both endpoints in each subset V_ℓ generate an amenable subgroup of $\pi_1(Z)$. In the current setup, the α-images of these edges are not even loops in Z. We correct this problem by modifying c by a homotopy that adds a path to each vertex, as in the proof of Lemma 4; the path is chosen as follows. For each stratum S_ℓ, we choose one special point $x_\ell \in S_\ell$. We homotope c so that every vertex $v \in V_\ell$ travels along some path ending at x_ℓ, chosen as follows:

- If $v \in c_\delta$ and $v \in V_\ell$, then $v \in P_{S_\ell}$, so we choose the path to be the trajectory of v under the homotopy sending $N_\delta(P_{S_\ell})$ into S_ℓ. Then we concatenate this path with any path contained in S_ℓ and ending at x_ℓ.
- If v is in the 0 end of c_2, and $v \in V_\ell$, then $v \in S_\ell$. If v has an edge to some vertex in the 1 end of c_2 (and thus in c_δ) that is in P_{S_ℓ}, then $v \in N_\delta(P_{S_\ell})$, so as above we take the trajectory of v under the homotopy sending $N_\delta(P_{S_\ell})$ into S_ℓ. If v does not have such an edge, then we take a constant path at v instead. Then we concatenate this first path (either of the two options) with any path contained in S_ℓ and ending at x_ℓ.
- If v is any other vertex in c_1—that is, not in the 1 end—and $v \in V_\ell$, then we take a path contained in S_ℓ from v to x_ℓ.

Then we homotope α (or c again) so that the image of every x_ℓ is the same point in Z. Now the α-images induced by a given V_ℓ do generate a subgroup of $\pi_1(Z)$. In order to show that this subgroup is amenable, it suffices to show that it is contained in the amenable subgroup $\alpha_* i_* \pi_1(S_\ell)$, where $i : S_\ell \hookrightarrow X$ denotes the inclusion. Every edge γ with both endpoints in V_ℓ is a loop at x_ℓ; we show that its homotopy class $[\gamma]$ in $\pi_1(X, x_\ell)$ is in $i_*\pi_1(S_\ell)$—that is, γ is homotopic through loops at x_ℓ to a loop entirely contained in S_ℓ. Then $\alpha_*[\gamma] \in \alpha_* i_* \pi_1(S_\ell)$.

We construct the homotopy on γ as follows. If γ is in c_2 or c_δ, and at least one endpoint is in P_{S_ℓ}, then all of γ is in $N_\delta(P_{S_\ell})$, so we homotope γ by the homotopy sending $N_\delta(P_{S_\ell})$ into S_ℓ. If γ is in c_2 or c_1, and both endpoints are in S_ℓ, then we use the fact that the *cellular* property and the (*int*) property together are preserved by barycentric subdivision. Thus γ is already contained in S_ℓ.

By this method, we produce a cycle c, extending c_U and homotopic to h, with a partial coloring such that every new simplex is not essential and such that, after a homotopy of α, every edge of c with both endpoints in V_ℓ is a loop representing an element of $\alpha_*\pi_1(S_\ell)$. Applying the amenable reduction lemma (Lemma 4), we obtain

$$\|\alpha_* h\|_\Delta \leq \|c_U\|_\Delta,$$

and taking the infimum over all such c_U,

$$\|\alpha_* h\|_\Delta \leq \|h_U\|_\Delta^{\mathcal{S}}.$$

Proof (Proof of Theorem 2) From Part 3 of Lemma 1, the vector field v gives rise to a stratification of X; doubling this stratification produces a stratification of the closed manifold $D(X)$ by submanifolds. From Part 4 of Lemma 1, there are only finitely many strata, because the compact set X can be covered by finitely many neighborhoods each matching one of the local models.

In order to apply the localization lemma (Lemma 5), we need to check that for each stratum S of X, the subgroup $\alpha_* \pi_1(S)$ of $\pi_1(Z)$ is an amenable group. We have assumed that this is true if $S \subseteq \partial X$. (Every subgroup of an amenable group is amenable.) Otherwise, we apply Parts 2 and 3 of Lemma 1: S is one connected component of $\Gamma^{-1}(\sigma) \setminus \partial X$ for some stratum σ of $\mathcal{T}(v)$, and the entire preimage $\Gamma^{-1}(\sigma)$ is a trivial bundle $\sigma \times F$, for some fiber F. Under the stratification of X, the fiber F is an interval subdivided by finitely many points from ∂X. There is a homotopy on the 1-dimensional part of F that pushes each open subinterval to the next point of ∂X, which gives a homotopy on $\Gamma^{-1}(\sigma) \setminus \partial X$ that starts with the inclusion into X and ends with a map into ∂X. Applying this homotopy to loops in S we see that $\pi_1(S)$ is contained in $\pi_1(\partial X)$, so its α-image is an amenable group.

Now we apply the localization lemma (Lemma 5), with $j = \dim X = n + 1$. Then X_{-j} consists of the 0-dimensional strata, which are the intersections of the maximum-multiplicity trajectories with ∂X. Let $x_1, \ldots, x_r$ denote these 0-dimensional strata; then we have

$$r \leq (n + 2) \cdot \textsf{max-mult}(v),$$

because each trajectory has at most n intermediate points of ∂X, and $n + 2$ points in total. Applying Part 4 of Lemma 1, around each point x_i we choose a neighborhood $U_i \subseteq X$ matching one of the local models, small enough that the various U_i are disjoint, and let $D(U_i) \subseteq D(X)$ denote the double of U_i. We take $U = \bigcup_{i=1}^{r} D(U_i)$. If $\mathcal{S}$ denotes the stratification on $D(X)$, then there exists some constant M_n depending only on n, satisfying

$$\|[D(U_i), \partial D(U_i)]\|_{\Delta}^{\mathcal{S}} \leq M_n$$

for all i. Thus, the conclusion of the localization lemma gives

$$\|\alpha_*[D(X)]\|_{\Delta} \leq \|[U, \partial U]\|_{\Delta}^{\mathcal{S}} \leq \sum_{i=1}^{r} M_n \leq M_n \cdot (n + 2) \cdot \textsf{max-mult}(v).$$

Proof (Proof of Theorem 1) We construct a degree-1 map $\alpha : D(X) \to M$ that sends all of the boundary ∂X to a single point. Given such a map, we have $\alpha_* \pi_1(\partial X) = 0$ (and 0 is an amenable group), so Theorem 2 gives

$$\textsf{max-mult}(v) \geq \text{const}(n) \|\alpha_*[D(X)]\|_{\Delta} = \text{const}(n) \|[M]\|_{\Delta} \geq \text{const}(n) \cdot \text{Vol}\, M,$$

where the value of const(n) is not fixed and may change between inequalities, but is always positive. To construct α, let B be an open ball containing $\overline{U}$, and let B' be a slightly smaller ball with $\overline{U} \subset B' \subset \overline{B'} \subset B$. There is a degree-1 map $M \to M$ obtained by collapsing $\overline{B'}$ to a single point $*$, and stretching the cylinder $B \setminus B'$ to fill $B \setminus *$. We define α on X as the restriction of this map on M, and define α on the second copy of X as the constant map at $*$. Then α on all of $D(X)$ has degree 1, and we have $\alpha(\partial X) = *$.

4 Future directions

One immediate follow-up question is how large the constant should be in Theorem 2. The 3-dimensional case of Theorem 1 (Theorem 7.5 of [7]) suggests that we might hope for a constant of 1 for every n. However, the constant obtained in our proof is much weaker and is a little confusing to compute. It would be nice to compute an explicit upper bound for the stratified simplicial volume of the coordinate neighborhood of each trajectory of a versal vector field. Also useful would be to check whether any examples might refute a possible constant of 1 in Theorem 2.

A second question for further study comes from a special case of Theorem 1. Let $f : M \to \mathbb{R}$ be a Morse function, and let X be the space obtained by deleting a small open ball around each critical point of f. If the negative gradient field $v = -\nabla f$ is traversally generic, then there are finitely many maximum-multiplicity trajectories in X, and if f satisfies Morse-Smale transversality, then there are finitely many n-times-broken trajectories in M. We might hope that these two sets of trajectories correspond in a fixed ratio depending on n. Thus, Theorem 1 suggests the following conjecture.

Conjecture 1 Let M be a closed, oriented hyperbolic manifold of dimension $n + 1 \geq 2$, and let $f : M \to \mathbb{R}$ be a Morse-Smale function. Then we have

$$\#(n\text{-times-broken trajectories of } -\nabla f) \geq \text{const}(n) \cdot \text{Vol}\, M.$$

This conjecture is proven in a forthcoming paper [1]. Actually the theorem in [1] is stronger than the conjecture: it holds for the specific constant $\text{const}(n) = (\text{Vol}\, \Delta^{n+1})^{-1}$, where $\text{Vol}\, \Delta^{n+1}$ denotes the supremal volume of a straight simplex in $(n+1)$-dimensional hyperbolic space. The proof avoids using traversally generic vector fields but does use a modified version of the localization lemma, along with some known theorems about Morse-Smale gradient vector fields.

Acknowledgments The authors would like to thank Larry Guth (Hannah's advisor) for initiating the collaboration, actively supervising most of our meetings, and improving the exposition in the paper. The authors also thank the referee for clarifying wording and suggesting future directions.

References

1. Alpert, H.: Using simplicial volume to count maximally broken Morse trajectories (2015). http://arxiv.org/pdf/1506.04789v1
2. Fujiwara, K., Manning, J.: Simplicial volume and fillings of hyperbolic manifolds. Algebr. Geom. Topol. **11**(4), 2237–2264 (2011)
3. Golubitsky, M., Guillemin, V.: Stable mappings and their singularities. In: Halmos, P.R., Moore, C.C. (eds.) Graduate Texts in Mathematics, vol. 14. Springer, New York (1974)
4. Gromov, M.: Volume and bounded cohomology. Inst. Hautes Études Sci. Publ. Math. **56**, 5–99 (1982)
5. Gromov, M.: Singularities, expanders and topology of maps. I. Homology versus volume in the spaces of cycles. Geom. Funct. Anal. **19**(3), 743–841 (2009)
6. Hatcher, A.: Algebraic Topology. Cambridge University Press, Cambridge (2002)
7. Katz, G.: Convexity of Morse stratifications and gradient spines of 3-manifolds. JP J. Geom. Topol. **9**(1), 1–119 (2009)
8. Katz, G.: Traversally generic & versal vector flows: semi-algebraic models of tangency to the boundary. Asian J. Math. (2014). http://arxiv.org/pdf/1407.1345v1