




Contents lists available at ScienceDirect

Advances in Mathematics

journal homepage: www.elsevier.com/locate/aim

On the Čech cohomology of Morse boundaries

Elia Fioravanti^{a,*,1}, Annette Karrer^{b,2}, Alessandro Sisto^c,
Stefanie Zbinden^c^a Max Planck Institute for Mathematics, Bonn, Germany^b McGill University, Montreal, Canada^c Maxwell Institute and Department of Mathematics, Heriot-Watt University, Edinburgh, UK

ARTICLE INFO

Article history:

Received 2 May 2023

Received in revised form 13

February 2024

Accepted 25 February 2024

Available online xxxx

Communicated by C. Leininger

MSC:

primary 20F65, 20F67

secondary 22E40, 57M07

Keywords:

Morse boundary

Čech cohomology

Discrete homology

Cusped hyperbolic manifold

Relatively hyperbolic

Sphere boundary

ABSTRACT

We consider cusped hyperbolic n -manifolds, and compute Čech cohomology groups of the Morse boundaries of their fundamental groups. In particular, we show that the reduced Čech cohomology with real coefficients vanishes in dimension at most $n-3$ and does not vanish in dimension $n-2$. A similar result holds for relatively hyperbolic groups with virtually nilpotent peripherals and Bowditch boundary homeomorphic to a sphere; these include all non-uniform lattices in rank-1 simple Lie groups.

© 2024 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY-NC license (<http://creativecommons.org/licenses/by-nc/4.0/>).

* Corresponding author.

E-mail addresses: elia.fioravanti@kit.edu (E. Fioravanti), annette.karrer@mcgill.ca (A. Karrer), a.sisto@hw.ac.uk (A. Sisto), sz2020@hw.ac.uk (S. Zbinden).¹ Fioravanti thanks the Max Planck Institute for Mathematics in Bonn for their hospitality and financial support while this work was being completed. He was also partially supported by Emmy Noether grant 515507199 of the Deutsche Forschungsgemeinschaft (DFG).² Karrer was partially supported by CIRGET and CRM and the Israel Science Foundation (grant no. 1562/19).

Contents

0.	Introduction	2
1.	Discrete chains and super-refinements	7
1.1.	Čech cohomology and discrete chains	7
1.2.	Super-refinements	9
1.3.	Direct limits of compact spaces	14
1.4.	Discrete chains in metric spaces	17
1.5.	Discretisation of singular simplices	17
2.	Filling discrete cycles in the Bowditch boundary	18
2.1.	Cusped space and Bowditch boundary	18
2.2.	Setup and statement of main proposition	18
2.3.	Geometry of nilpotent groups	19
2.4.	Preliminary lemmas	21
2.5.	The main argument	24
3.	Filling discrete cycles in the Morse boundary	31
3.1.	Notation	31
3.2.	Assumptions	31
3.3.	Detouring	33
4.	Vanishing of Čech cohomology	40
4.1.	Subdivision results	40
4.2.	Main result	49
5.	Non-vanishing	50
	References	52

0. Introduction

The Morse boundary ∂_*X of a metric space X , as introduced in [12,11], is a topological space that encodes the “hyperbolic-like” directions of X . It extends to all metric spaces Gromov’s classical concept of boundary for hyperbolic spaces [14]: when X is hyperbolic, ∂_*X is naturally identified with the Gromov boundary of X . The Morse boundary also enjoys another fundamental property of Gromov boundaries, namely quasi-isometric invariance: quasi-isometric metric spaces have homeomorphic Morse boundaries. As a consequence, the Morse boundary ∂_*G of a finitely generated group G is well-defined up to homeomorphism. Quasi-isometry invariance of ∂_*G was in fact used in [12] to distinguish certain right-angled Coxeter groups up to quasi-isometry.

In view of this, it is interesting to study topological features of Morse boundaries. Since ∂_*X is not even metrisable when X is non-hyperbolic [8], it was thought until rather recently that “fully understanding” the topology of Morse boundaries of non-hyperbolic spaces would be an impossible feat. In contrast to this prediction, we show that in some important cases one can compute a rather sophisticated and powerful topological invariant, namely Čech cohomology:

Theorem A. *Let G be the fundamental group of a cusped hyperbolic n -manifold for $n \geq 3$. Then*

$$(1) \quad \check{H}^i(\partial_*G, \mathbb{R}) = \{0\} \text{ for } 1 \leq i \leq n - 3;$$

(2) $\check{H}^{n-2}(\partial_* G, \mathbb{R})$ is infinite dimensional,

where \check{H}^* denotes reduced Čech cohomology.

The theorem follows from Theorem B below, which is more general in two ways. First, it covers all relatively hyperbolic groups with virtually nilpotent peripherals and Bowditch boundary homeomorphic to a sphere; these include all non-uniform lattices in rank-1 simple Lie groups. Secondly, it allows for more general coefficients.

We note that, for $i = 0$, our result would be equivalent to $\partial_* G$ being connected, but one can even prove that $\partial_* G$ is path connected in the setting of Theorem A. This follows from results in [19], as we argue in Theorem 3.2. We will actually further develop various tools from that paper and use them extensively in our arguments.

Our hope is that the techniques we develop to prove the main theorem will be useful to study Morse boundaries of other groups and to distinguish quasi-isometry classes of groups via their Morse boundary; see below for some concrete candidates.

Theorem A also provides the first examples of relatively hyperbolic groups with high-dimensional Morse boundary not witnessed by an “obvious” high-dimensional hyperbolic subgroup. For instance, for $n \geq 4$, it is still an open problem whether the fundamental group of a (cusped) hyperbolic n -manifold always admits codimension-1 subgroups. In relation to this, we mention that Cordes previously showed the existence of high-dimensional spheres in Morse boundaries of mapping class groups [11, Theorem 4.2], using quasi-convex copies of \mathbb{H}^n in the thick part of Teichmüller space constructed by Leininger and Schleimer [18]. It is open whether any of these corresponds to a subgroup of the mapping class group, and some certainly do not.

Our main theorem is in a similar line of research as the papers [7,27,26], where Morse boundaries of certain groups, including 3-manifold groups, are fully described. A similar goal seems out of reach in our case (except when $n = 3$, as shown in [7]), but we also think of our theorem as a proof of concept: even when the Morse boundary is not “fully describable”, powerful topological invariants can sometimes still be computed.

As a final note, there is a different topology on Morse boundaries [10] which has the advantage of being metrisable. However, we do not know whether analogues of our main theorem (or of the other results hinted at above) hold for that topology.

0.1. Distinguishing groups up to quasi-isometry

To provide further motivation for the study of Čech cohomology of Morse boundaries, we now construct a collection of groups that can conjecturally be distinguished up to quasi-isometry using this invariant, but where other usual quasi-isometry invariants do not seem to be helpful. We believe that the results and techniques in this paper go a long way towards proving the conjecture, and we point out below what is missing.

Fix for every integer $n \geq 1$ the fundamental group G_n of an $(n+2)$ -dimensional cusped hyperbolic manifold. Let \bar{n} be a finite collection of integers at least 1. Following [15,

Section 6], one can construct a certain finitely generated, infinite group $G(\bar{n})$ with several properties, including that $G(\bar{n})$ contains copies of all G_n with $n \in \bar{n}$ as hyperbolically embedded subgroups, but $G(\bar{n})$ is not (non-trivially) relatively hyperbolic.

The following conjecture would allow one to distinguish the quasi-isometry classes of the various $G(\bar{n})$:

Conjecture 1. $\check{H}^i(\partial_*G(\bar{n}), \mathbb{R})$ is non-trivial for $i \in \bar{n}$ and trivial for $i \notin \bar{n}$, with $1 < i \leq \max(\bar{n})$.

The intuition behind the conjecture is that $\partial_*G(\bar{n})$ should contain copies of the ∂_*G_n for $n \in \bar{n}$, and those should contribute cohomology in the appropriate degrees, but $\partial_*G(\bar{n})$ should otherwise be 1-dimensional. Proving the conjecture requires understanding Morse boundaries of small cancellation free products (as $\bar{G}(n)$ is of this form). We believe this task to be of independent interest, but also not within the scope of this paper, so we leave it for future work. In fact, more work is still required in the context of classical small cancellation groups (which have non-trivial Morse boundary by [15] and [23], except in trivial cases).

We could not think of methods that would distinguish the quasi-isometry classes of the $G(\bar{n})$ as \bar{n} varies, other than (variations of) the conjecture. For example, the groups are not relatively hyperbolic, which rules out a lot of tools. Similarly, they are not CAT(0), hierarchically hyperbolic, etc, as they are not finitely presented. Also, their asymptotic dimension presumably only depends on $\max(\bar{n})$, as should the “stable asymptotic dimension” introduced in [9].

0.2. Questions

There are various natural questions that arise at this point. Let us first consider those related to the same groups as in Theorem A.

Question 2. For G as in Theorem A, does $\check{H}^i(\partial_*G, \mathbb{R})$ vanish for all $i > n - 2$?

We believe that the answer is yes, and we plan future work on covering dimension of Morse boundaries that should imply this. Related to this, we note that it is easy to show that $\check{H}^i(K, \mathbb{Z}) = \{0\}$ for all $i > n - 2$ and all compact subsets $K \subseteq \partial_*G$. Indeed, any such K is homeomorphic to a closed nowhere-dense subset of S^{n-1} , and the latter all have covering dimension at most $n - 2$ (for instance, by [21, Theorem 19]).

As mentioned above, Theorem A holds for more general coefficients. However, certain natural coefficients are not covered, especially for the non-vanishing part, so in particular we can ask:

Question 3. For G as in Theorem A, is $\check{H}^{n-2}(\partial_*G, \mathbb{Z})$ non-zero?

We believe that answering this question requires deeper understanding of discrete cycles (see discussion on this notion below), and hence working towards it could lead to interesting discoveries.

Moving on, it is natural to ask what the Čech cohomology is for Morse boundaries of other groups of interest. For example:

Problem 4. *Compute Čech cohomology groups $\check{H}^i(\partial_*G, \mathbb{R})$ for G a mapping class group of a finite-type surface.*

As mentioned above our hope is that our strategy will lead to the use of Morse boundaries as effective quasi-isometry invariants. More generally than in Conjecture 1, we then ask:

Problem 5. *Use topological invariants of Morse boundaries such as Čech cohomology to distinguish quasi-isometry classes of groups.*

Another natural class of groups to look into regarding this problem is that of right-angled Coxeter groups.

Finally, looking at Čech cohomology of Morse boundaries was also partly inspired by the fact that, for hyperbolic groups, the Čech cohomology of the Gromov boundary has been fruitfully studied and contains important information about the group itself, see e.g. [4]. Therefore, in a different direction than the other questions, it is natural to ask:

Problem 6. *What kind of information about a group can be gleaned from the Čech cohomology of its Morse boundary?*

0.3. Outline

We deduce Theorem A as a consequence of the following more general result.

Theorem B. *Let (G, \mathcal{P}) be a relatively hyperbolic pair. Suppose that \mathcal{P} is a nonempty finite collection of finitely generated, virtually nilpotent subgroups, and that the Bowditch boundary is homeomorphic to the $(n - 1)$ -sphere, for some $n \geq 3$. Then*

- (1) $\check{H}^i(\partial_*G, \mathcal{R}) = \{0\}$ for $1 \leq i \leq n - 3$ and any principal ideal domain \mathcal{R} ;
- (2) $\check{H}^{n-2}(\partial_*G, \mathcal{F})$ is infinite dimensional for any field \mathcal{F} of characteristic 0.

The assumptions of Theorem B are satisfied by all fundamental groups of non-compact, finite-volume, complete Riemannian manifolds with pinched negative sectional curvature (by the Margulis lemma). These include non-uniform lattices in rank-1 simple Lie groups.

Theorem B is proved by combining Theorem 4.1 and Theorem 5.2 in the main body of the text, which deal with the vanishing and non-vanishing part, respectively.

Our approach to these results relies heavily on the notion of *discrete chains*, as described in Subsection 1.1. Indeed, we will translate vanishing and non-vanishing of Čech cohomology for “sufficiently nice” spaces into statements about discrete cycles being, or not being, boundaries of discrete chains. This is one of the main goals of Section 1, see in particular Corollary 1.10.

The topological spaces to which this strategy can be applied are those that we call *super-refinable*, see Definition 1.6. Roughly, these are spaces where each open cover admits a “very small” refinement. Another main goal of Section 1 is to show that (countable) direct limits of compact metric spaces are super-refinable, see Proposition 1.12. This is relevant because Morse boundaries are, by definition, direct limits of compact metric spaces (possibly over an uncountable index set, but not in our case).

In Section 2 we prove Proposition 2.1, which allows us to fill discrete cycles in the Bowditch boundary of the relatively hyperbolic groups that we are interested in. This result is near-obvious in the setting of Theorem A, so Section 2 is only required to reach the generality of Theorem B.

In Section 3 we relate Morse boundaries to Bowditch boundaries. Roughly, the strata of the Morse boundary are obtained by removing from the Bowditch boundary a collection of balls of specific radii around parabolic points. Hence, in order to fill a discrete cycle in the Morse boundary, the idea is to start with a filling in the Bowditch boundary (as constructed in Section 2) and “detour” it away from parabolic points. We implement this strategy to prove Proposition 3.6, which is one of the main results of the section. We note that a similar strategy was adopted in [19] to construct suitable arcs in Bowditch boundaries, and in fact we further develop the techniques of that paper for our purposes.

In Section 3 we also show Proposition 3.8 in support of our non-vanishing result, Theorem 5.2. Roughly, considering the Bowditch boundary minus a finite collection of parabolic points F , this proposition provides, for each homology class, a representative that is supported in a fixed stratum of the Morse boundary (depending on F). In other words, these representatives stay away not just from F , but from *all* parabolic points, in a suitable sense.

In Section 4 we conclude the proof of our vanishing result, Theorem 4.1. The main point here is that the fillings constructed in Section 3 need to be improved in order to apply Corollary 1.10, and to do so we have to “make certain discrete chains finer”. In fact, we will use fillings to do so.

Finally, in Section 5 we prove our non-vanishing result, Theorem 5.2. Besides the above-mentioned Proposition 3.8, this boils down to Lemma 5.1, which reduces non-vanishing of Čech cohomology (over a field of characteristic 0) to finding discrete cycles that are not boundaries of specific chains. This can be seen as a counterpart to Corollary 1.10.

Acknowledgements. We thank the referee for their several useful comments.

1. Discrete chains and super-refinements

In this section we define and study discrete chains and related notions. We could not find any papers in the literature studying the exact versions of these notions that we need. However, there are a few places where similar concepts have been considered. The closest we could find is [2], where the authors define discrete homology groups at a given scale, and then consider a direct limit as the scale goes to infinity. Instead, here we are interested in the behaviour at small scales.

1.1. Čech cohomology and discrete chains

Let X be a topological space and let \mathcal{O} be an open cover of X .

We denote by $N(\mathcal{O})$ the nerve of \mathcal{O} , and by $C_*(N(\mathcal{O}))$ its chain complex of simplicial chains with \mathbb{Z} -coefficients. In other words, for each $n \geq 0$, the group $C_n(N(\mathcal{O}))$ is the free \mathbb{Z} -module generated by ordered $(n + 1)$ -tuples $[O_0, \dots, O_n]$, where $O_0, \dots, O_n \in \mathcal{O}$ and $O_0 \cap \dots \cap O_n \neq \emptyset$.

If $\mathcal{O}' < \mathcal{O}$ is a refinement, there are well-defined maps in homology and cohomology (see e.g. part (1) of Lemma 1.8 below):

$$H_*(N(\mathcal{O}'), \mathbb{Z}) \rightarrow H_*(N(\mathcal{O}), \mathbb{Z}), \quad H^*(N(\mathcal{O}), \mathbb{Z}) \rightarrow H^*(N(\mathcal{O}'), \mathbb{Z}).$$

The Čech homology³ and cohomology of X are then defined as:

$$\check{H}_*(X, \mathbb{Z}) = \varinjlim_{\mathcal{O}} H_*(N(\mathcal{O}), \mathbb{Z}), \quad \check{H}^*(X, \mathbb{Z}) = \varprojlim_{\mathcal{O}} H^*(N(\mathcal{O}), \mathbb{Z}),$$

where the two limits are taken over the directed set of open covers \mathcal{O} of X with the order given by refinements. Čech (co)homology with other coefficients is defined analogously.

Rather than working with open covers and their nerves, it is often simpler to study *discrete chains*, which we now introduce. We will relate them to Čech (co)homology in Subsection 1.2.

We define a chain complex $\mathcal{C}_*(X)$ as follows. For each $n \geq 0$, the \mathbb{Z} -module $\mathcal{C}_n(X)$ is freely generated by ordered $(n + 1)$ -tuples $[x_0, \dots, x_n]$, where $x_0, \dots, x_n \in X$. As customary, the boundary operator is the \mathbb{Z} -linear extension of:

$$\partial[x_0, \dots, x_n] = \sum_{i=0}^n (-1)^i \cdot [x_0, \dots, \hat{x}_i, \dots, x_n].$$

We refer to tuples $[x_0, \dots, x_n]$ as *discrete n -simplices*, and to general elements of $\mathcal{C}_n(X)$ as *discrete n -chains*. A discrete chain c is a *cycle* if $\partial c = 0$, and a *boundary* if $c = \partial d$ for

³ We caution the reader that Čech homology fails to satisfy the Eilenberg–Steenrod exactness axiom and, thus, it is not a homology theory as such. Instead, Čech cohomology does constitute an honest cohomology theory.

some discrete chain d . It is convenient to also introduce the module of *reduced* n -chains $\overline{\mathcal{C}}_n(X)$: for $n \geq 1$ we have $\overline{\mathcal{C}}_n(X) = \mathcal{C}_n(X)$, while $\overline{\mathcal{C}}_0(X)$ is the kernel of the coefficient-sum map $\mathcal{C}_0(X) \rightarrow \mathbb{Z}$.

The *support* of a discrete simplex $[x_0, \dots, x_n]$ is the set $\{x_0, \dots, x_n\} \subseteq X$. If $c = \sum_i \sigma_i$ is a discrete chain, written as a sum of discrete simplices without any cancellations, we define its *support* as $\text{supp}(c) := \bigcup_i \text{supp}(\sigma_i)$.

With a slight abuse, we will often confuse a discrete chain c with its support. For instance, for a subset $A \subseteq X$, we will write “ $c \subseteq A$ ” meaning “ $\text{supp}(c) \subseteq A$ ”. Similarly, we will speak of “vertices $x \in c$ ”, rather than “points $x \in \text{supp}(c)$ ”.

We will often be interested in chains that are “small” relative to a fixed cover, in one of two possible different ways, as described in the following definition.

Definition 1.1. Given an open cover \mathcal{O} of X , a subset $A \subseteq X$ is \mathcal{O} -*tiny* if there exists $O \in \mathcal{O}$ such that $A \subseteq O$. A discrete chain is \mathcal{O} -*tiny* if its support is. A discrete chain is \mathcal{O} -*fine*⁴ if all of its simplices are \mathcal{O} -tiny. We denote by $\mathcal{C}_*(X, \mathcal{O}) \subseteq \mathcal{C}_*(X)$ the chain subcomplex of \mathcal{O} -fine discrete chains. For a refinement $\mathcal{O}' < \mathcal{O}$, we have $\mathcal{C}_*(X, \mathcal{O}') \subseteq \mathcal{C}_*(X, \mathcal{O})$.

Although not strictly necessary for our purposes, the following definitions are helpful to illustrate our viewpoint on Čech (co)homology. Let $\mathcal{H}_*(X, \mathcal{O})$ denote the homology of the chain complex $\mathcal{C}_*(X, \mathcal{O})$. For each refinement $\mathcal{O}' < \mathcal{O}$, the inclusion $\mathcal{C}_*(X, \mathcal{O}') \subseteq \mathcal{C}_*(X, \mathcal{O})$ gives rise to a homomorphism $\mathcal{H}_*(X, \mathcal{O}') \rightarrow \mathcal{H}_*(X, \mathcal{O})$. We can then define the *discrete homology* of X as:

$$\mathcal{H}_*(X) := \varprojlim_{\mathcal{O}} \mathcal{H}_*(X, \mathcal{O}),$$

where the inverse limit is again taken over the directed set of open covers \mathcal{O} of X with the order given by refinements. For nice spaces, we will show in Corollary 1.9 that $\mathcal{H}_*(X)$ is naturally isomorphic to $\check{H}_*(X, \mathbb{Z})$ (so discrete homology also does not satisfy the Eilenberg–Steenrod axioms).

We conclude this subsection with a couple of definitions and an observation, which will be used multiple times throughout the paper.

Definition 1.2 (Cones). For a discrete simplex $\sigma = [x_0, \dots, x_n]$ and a point $x \in X$, the *cone* over σ from x is the simplex $\text{cone}(x, \sigma) := [x, x_0, \dots, x_n]$. In particular, if σ is the empty -1 -simplex, we define $\text{cone}(x, \sigma) := [x]$.

For a discrete chain $c = \sum_j \sigma_j$, where the σ_j are discrete simplices, the *cone* over c from x is the chain $\text{cone}(x, c) := \sum_j \text{cone}(x, \sigma_j)$.

⁴ In the literature, it is common to refer to both \mathcal{O} -tiny subsets and \mathcal{O} -fine chains simply as “ \mathcal{O} -small”. We prefer to use distinct terms for the two concepts, also to highlight the analogy between \mathcal{O} -fine chains and “ δ -fine” chains introduced in Subsection 1.4.

Lemma 1.3. *Let c be a discrete chain in a topological space X . For every point $x \in X$, we have $\partial \operatorname{cone}(x, c) = c - \operatorname{cone}(x, \partial c)$. In particular, if $c \in \overline{\mathcal{C}}_n(X)$ is a reduced cycle, then $\partial \operatorname{cone}(x, c) = c$.*

Proof. For a discrete simplex σ , a simple computation shows that $\partial \operatorname{cone}(x, \sigma) = \sigma - \operatorname{cone}(x, \partial \sigma)$. The statement about general chains then follows from linearity of the coning operator. \square

Definition 1.4 (Face complex). Let $c = \sum_i \sigma_i$ be a discrete chain, written as a sum of discrete simplices without cancellations. The *face complex* of c is the chain subcomplex $\mathcal{F}_*(c) \subseteq \mathcal{C}_*(X)$ generated by the σ_i and all their lower-dimensional faces.

1.2. Super-refinements

Let X be a topological space.

The following strengthening of the notion of refinement will be crucial for us to relate properties of nerves of covers to properties of discrete chains, and hence to reduce Čech cohomology statements to statements about discrete chains.

Definition 1.5. Let \mathcal{O} be an open cover of X . An open cover \mathcal{O}' is a *super-refinement* of \mathcal{O} (denoted $\mathcal{O}' \ll \mathcal{O}$) if, for each $x \in X$, there exists $O \in \mathcal{O}$ with:

$$\bigcup \{O' \mid O' \in \mathcal{O}', x \in O'\} \subseteq O.$$

An open set O satisfying the above inclusion is called a *parent* of x .

Definition 1.6. A topological space is *super-refinable* if every open cover admits a super-refinement.

It is readily seen that compact metric spaces are super-refinable due to the existence of Lebesgue numbers for covers (that is, the fact that, for every open cover, there exists $\epsilon > 0$ such that every ϵ -ball is contained in an open set from the cover). We will show in Subsection 1.3 that countable direct limits of compact metrisable spaces are super-refinable as well.

In order to connect discrete chains in X to simplicial chains in nerves of open covers of X , we introduce the following terminology.

Definition 1.7. Let \mathcal{O} be an open cover of X .

- A map $f: \mathcal{O} \rightarrow X$ is a *child map* if $f(O) \in O$ for all $O \in \mathcal{O}$.

Let $\mathcal{O}' < \mathcal{O}$ be a refinement.

- A map $f: \mathcal{O}' \rightarrow \mathcal{O}$ is a *spouse map* if $\mathcal{O}' \subseteq f(\mathcal{O}')$ for all $\mathcal{O}' \in \mathcal{O}'$.

Suppose now that $\mathcal{O}' \ll \mathcal{O}$ is a super-refinement.

- A map $f: X \rightarrow \mathcal{O}$ is an $(\mathcal{O}', \mathcal{O})$ -parent map if $\mathcal{O}' \subseteq f(x)$ for all $\mathcal{O}' \in \mathcal{O}'$ with $x \in \mathcal{O}'$ (that is, $f(x)$ is a parent of x according to Definition 1.5).
- A map $f: X \rightarrow X$ is an $(\mathcal{O}', \mathcal{O})$ -sibling map if it is the composition of an $(\mathcal{O}', \mathcal{O})$ -parent map $X \rightarrow \mathcal{O}$ and a child map $\mathcal{O} \rightarrow X$.

Each of the maps appearing in Definition 1.7 naturally induces a chain map between the corresponding chain complexes. This is the content of the next lemma.

Lemma 1.8. *Let \mathcal{O} be an open cover of X with a refinement $\mathcal{O}' < \mathcal{O}$.*

- (1) *Every spouse map $f: \mathcal{O}' \rightarrow \mathcal{O}$ naturally induces a chain map $f_*: C_*(N(\mathcal{O}')) \rightarrow C_*(N(\mathcal{O}))$. If $g: \mathcal{O}' \rightarrow \mathcal{O}$ is another spouse map, then f_* and g_* are chain-homotopic.*

Suppose now that $\mathcal{O}' \ll \mathcal{O}$ is a super-refinement.

- (2) *Every child map $f: \mathcal{O}' \rightarrow X$ naturally induces a chain map $f_*: C_*(N(\mathcal{O}')) \rightarrow \mathcal{C}_*(X, \mathcal{O})$. If $g: \mathcal{O}' \rightarrow X$ is another child map, then f_* and g_* are chain-homotopic.*
- (3) *Every $(\mathcal{O}', \mathcal{O})$ -parent map $f: X \rightarrow \mathcal{O}$ naturally induces a chain map $f_*: \mathcal{C}_*(X, \mathcal{O}') \rightarrow C_*(N(\mathcal{O}))$. If $g: X \rightarrow \mathcal{O}$ is another $(\mathcal{O}', \mathcal{O})$ -parent map, then f_* and g_* are chain-homotopic.*

Finally, consider a chain of super-refinements $\mathcal{O}'' \ll \mathcal{O}' \ll \mathcal{O}$.

- (4) *Every $(\mathcal{O}'', \mathcal{O}')$ -sibling map $f: X \rightarrow X$ naturally induces a chain map $f_*: \mathcal{C}_*(X, \mathcal{O}'') \rightarrow \mathcal{C}_*(X, \mathcal{O})$ and this is chain-homotopic to the standard inclusion.*

Proof. Part (1) is classical (see e.g. [6, Lemma 10.4.2]) and the proof of the other three parts is very similar, so we will omit most of the computations.

Recall that $\mathcal{C}_n(X, \mathcal{O})$ is the free \mathbb{Z} -module generated by \mathcal{O} -tiny discrete simplices $[x_0, \dots, x_n]$, while $C_n(N(\mathcal{O}))$ is the free \mathbb{Z} -module generated by the elements $[O_0, \dots, O_n]$, where $O_0, \dots, O_n \in \mathcal{O}$ and $O_0 \cap \dots \cap O_n \neq \emptyset$.

Part (1). The chain map $f_n: C_n(N(\mathcal{O}')) \rightarrow C_n(N(\mathcal{O}))$ takes $[O'_0, \dots, O'_n]$ to $[f(O'_0), \dots, f(O'_n)]$.
 If $g: \mathcal{O}' \rightarrow \mathcal{O}$ is another spouse map, a chain homotopy $H_n: C_n(N(\mathcal{O}')) \rightarrow C_{n+1}(N(\mathcal{O}))$ between f_* and g_* is given by the formula:

$$H_n[O'_0, \dots, O'_n] := \sum_{i=0}^n (-1)^i \cdot [f(O'_0), \dots, f(O'_i), g(O'_i), \dots, g(O'_n)].$$

Part (2). The chain map $f_n: C_n(N(\mathcal{O}')) \rightarrow \mathcal{C}_n(X, \mathcal{O})$ takes $[O'_0, \dots, O'_n]$ to $[f(O'_0), \dots, f(O'_n)]$. We prove that f_n is well-defined. Since \mathcal{O}' is a super-refinement of \mathcal{O} , there exists a parent $O \in \mathcal{O}$ of some point $x \in O'_0 \cap \dots \cap O'_n$, and we have:

$$\{f(O'_0), \dots, f(O'_n)\} \subseteq O'_0 \cup \dots \cup O'_n \subseteq O.$$

Thus, the discrete simplex $[f(O'_0), \dots, f(O'_n)]$ is indeed \mathcal{O} -tiny.

Now, if $g: \mathcal{O}' \rightarrow X$ is another child map, a chain homotopy $H_n: C_n(N(\mathcal{O}')) \rightarrow \mathcal{C}_{n+1}(X, \mathcal{O})$ between f_* and g_* is given by the formula:

$$H_n[O'_0, \dots, O'_n] := \sum_{i=0}^n (-1)^i \cdot [f(O'_0), \dots, f(O'_i), g(O'_i), \dots, g(O'_n)].$$

This map is again well-defined because all $f(O'_i)$ and $g(O'_i)$ lie in any parent $O \in \mathcal{O}$ of any point of $O'_0 \cap \dots \cap O'_n$. We leave to the reader the straightforward check that $\partial H_n + H_{n-1} \partial = g_n - f_n$ for all $n \geq 0$.

Part (3). The chain map $f_n: \mathcal{C}_n(X, \mathcal{O}') \rightarrow C_n(N(\mathcal{O}))$ takes $[x_0, \dots, x_n]$ to $[f(x_0), \dots, f(x_n)]$. If $\{x_0, \dots, x_n\}$ is contained in $O' \in \mathcal{O}'$, then

$$f(x_0) \cap \dots \cap f(x_n) \supseteq O' \neq \emptyset,$$

by definition of parent map. This proves that f_n is well-defined.

If $g: X \rightarrow \mathcal{O}$ is another parent map, a chain homotopy $H_n: \mathcal{C}_n(X, \mathcal{O}') \rightarrow C_{n+1}(N(\mathcal{O}))$ between f_* and g_* is given by the formula:

$$H_n[x_0, \dots, x_n] := \sum_{i=0}^n (-1)^i \cdot [f(x_0), \dots, f(x_i), g(x_i), \dots, g(x_n)].$$

Again, this is a well-defined map because all $f(x_i)$ and $g(x_i)$ contain any $O' \in \mathcal{O}'$ containing all x_i . As above, we have $\partial H_n + H_{n-1} \partial = g_n - f_n$ for all $n \geq 0$.

Part (4). The chain map $f_n: \mathcal{C}_n(X, \mathcal{O}'') \rightarrow \mathcal{C}_n(X, \mathcal{O})$ takes $[x_0, \dots, x_n]$ to $[f(x_0), \dots, f(x_n)]$. Let us show that this is well-defined. First, by definition of sibling map, there exists an $(\mathcal{O}'', \mathcal{O}')$ -parent map $F: X \rightarrow \mathcal{O}'$ such that $f(x) \in F(x)$ for all $x \in X$. If $\{x_0, \dots, x_n\}$ is \mathcal{O}'' -tiny, we have $F(x_0) \cap \dots \cap F(x_n) \neq \emptyset$ as in part (3). It follows that $F(x_0) \cup \dots \cup F(x_n)$ is \mathcal{O} -tiny as in part (2). Since this union contains all $f(x_i)$, we conclude that the discrete simplex $[f(x_0), \dots, f(x_n)]$ is indeed \mathcal{O} -tiny.

Now, if $g: X \rightarrow X$ is another sibling map, a chain homotopy $H_n: \mathcal{C}_n(X, \mathcal{O}') \rightarrow \mathcal{C}_{n+1}(X, \mathcal{O})$ between f_* and g_* is given by the formula:

$$H_n[x_0, \dots, x_n] := \sum_{i=0}^n (-1)^i \cdot [f(x_0), \dots, f(x_i), g(x_i), \dots, g(x_n)].$$

The map H_n is well-defined by the argument in the previous paragraph, and the usual computation shows that $\partial H_n + H_{n-1}\partial = g_n - f_n$ for all $n \geq 0$.

This concludes the proof of the lemma. \square

Part (1) of Lemma 1.8 shows that each refinement $\mathcal{O}' < \mathcal{O}$ gives rise to well-defined maps between the (co)homology groups of the nerves of \mathcal{O} and \mathcal{O}' , which we used to define the Čech (co)homology groups of X in Subsection 1.1.

The rest of Lemma 1.8 allows us to relate the Čech homology of X to its discrete homology, provided that X is super-refinable. While we will not need the next result (preferring instead its reformulation given by Corollary 1.10), it motivates our approach to Čech (co)homology.

Corollary 1.9. *If X is super-refinable, its Čech homology $\check{H}_*(X, \mathbb{Z})$ is naturally isomorphic to its discrete homology $\mathcal{H}_*(X)$.*

Proof. Recall that we have:

$$\check{H}_*(X, \mathbb{Z}) = \varprojlim_{\mathcal{O}} H_*(N(\mathcal{O}), \mathbb{Z}), \quad \mathcal{H}_*(X) = \varprojlim_{\mathcal{O}} \mathcal{H}_*(X, \mathcal{O}).$$

For each super-refinement $\mathcal{O}' \ll \mathcal{O}$, parts (2) and (3) of Lemma 1.8 give maps:

$$H_*(N(\mathcal{O}'), \mathbb{Z}) \rightarrow \mathcal{H}_*(X, \mathcal{O}), \quad \mathcal{H}_*(X, \mathcal{O}') \rightarrow H_*(N(\mathcal{O}), \mathbb{Z}).$$

In addition, for a chain of super-refinements $\mathcal{O}'' \ll \mathcal{O}' \ll \mathcal{O}$, parts (1) and (4) of Lemma 1.8 guarantee that the compositions

$$\begin{aligned} H_*(N(\mathcal{O}''), \mathbb{Z}) &\rightarrow \mathcal{H}_*(X, \mathcal{O}') \rightarrow H_*(N(\mathcal{O}), \mathbb{Z}), \\ \mathcal{H}_*(X, \mathcal{O}'') &\rightarrow H_*(N(\mathcal{O}'), \mathbb{Z}) \rightarrow \mathcal{H}_*(X, \mathcal{O}). \end{aligned}$$

coincide with the standard homomorphisms.

Since X is super-refinable, the above maps give rise to homomorphisms

$$\check{H}_*(X, \mathbb{Z}) \rightarrow \mathcal{H}_*(X), \quad \mathcal{H}_*(X) \rightarrow \check{H}_*(X, \mathbb{Z}),$$

and both their compositions are the identity. This proves the corollary. \square

In order to show vanishing of $\check{H}_n(X, \mathbb{Z})$ and $\mathcal{H}_n(X)$ for $n \geq 1$, it is sometimes more practical to check that X satisfies the following (a priori stronger) condition⁵:

(DA_n) for each open cover \mathcal{O} there exists a refinement \mathcal{O}' such that every discrete \mathcal{O}' -fine n -cycle is the boundary of a discrete \mathcal{O} -fine $(n + 1)$ -chain.

⁵ The letters D and A stand for Discrete and Acyclic.

We only define (DA_n) as above for $n \geq 1$, declaring instead (DA_0) to be vacuous.

Corollary 1.10. *Let X be super-refinable and let $n \geq 1$. Let \mathcal{F} be a field and \mathcal{R} a principal ideal domain.*

- (1) *If X satisfies (DA_n) , then $\check{H}_n(X, \mathbb{Z}) = \{0\}$ and $\check{H}^n(X, \mathcal{F}) = \{0\}$.*
- (2) *If X satisfies (DA_n) and (DA_{n-1}) , then $\check{H}^n(X, \mathcal{R}) = \{0\}$.*

Proof. Suppose that X satisfies (DA_n) and let \mathcal{O} be an open cover of X .

For part (1), we want to show that there exists a refinement $\mathcal{O}' < \mathcal{O}$ and a spouse map $s: \mathcal{O}' \rightarrow \mathcal{O}$ such that the induced maps

$$s_n: H_n(N(\mathcal{O}'), \mathbb{Z}) \rightarrow H_n(N(\mathcal{O}), \mathbb{Z}), \quad s^n: H^n(N(\mathcal{O}), \mathcal{F}) \rightarrow H^n(N(\mathcal{O}'), \mathcal{F})$$

are trivial. It then follows by definition of Čech (co)homology that $\check{H}_n(X, \mathbb{Z}) = 0$ and $\check{H}^n(X, \mathcal{F}) = 0$.

Consider a chain of refinements $\mathcal{O}' \ll \tilde{\mathcal{O}}' < \tilde{\mathcal{O}} \ll \mathcal{O}$, where $\tilde{\mathcal{O}}' < \tilde{\mathcal{O}}$ is a refinement provided by (DA_n) . Let $f: \mathcal{O}' \rightarrow X$ be a child map and let $g: X \rightarrow \mathcal{O}$ be an $(\tilde{\mathcal{O}}, \mathcal{O})$ -parent map. Let f_*, g_* be the chain maps provided by Lemma 1.8. Observe that $s := g \circ f: \mathcal{O}' \rightarrow \mathcal{O}$ is a spouse map.

Now, let $c \in C_n(N(\mathcal{O}'))$ be a cycle. Then $f_*(c) \in \mathcal{C}_n(X, \tilde{\mathcal{O}}')$ is an $\tilde{\mathcal{O}}'$ -fine discrete cycle and hence there exists an $\tilde{\mathcal{O}}$ -fine discrete chain d with $f_*(c) = \partial d$. Since f_* and g_* are chain maps, we have that:

$$\partial g_*(d) = g_*(\partial d) = g_*f_*(c) = s_*(c).$$

This shows that the induced map in homology $s_n: H_n(N(\mathcal{O}'), \mathbb{Z}) \rightarrow H_n(N(\mathcal{O}), \mathbb{Z})$ is trivial. The induced map in cohomology $s^n: H^n(N(\mathcal{O}), \mathcal{F}) \rightarrow H^n(N(\mathcal{O}'), \mathcal{F})$ is then also trivial, by the universal coefficients theorem. This proves part (1).

In order to prove part (2), let the refinement \mathcal{O}' and the spouse map $s: \mathcal{O}' \rightarrow \mathcal{O}$ be those constructed in part (1). Then consider a refinement $\mathcal{O}'' < \mathcal{O}'$ and a spouse map $t: \mathcal{O}'' \rightarrow \mathcal{O}$ obtained by applying the same procedure to the cover \mathcal{O}' . We will show that, when X satisfies (DA_{n-1}) , the spouse map $s \circ t$ induces the trivial map $H^n(N(\mathcal{O}), \mathcal{R}) \rightarrow H^n(N(\mathcal{O}''), \mathcal{R})$.

The universal coefficients theorem (in the form stated e.g. in [24, Section 5.5, Theorem 3]) gives the following commutative diagram with exact rows:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & \text{Ext}(H_{n-1}(N(\mathcal{O}), \mathbb{Z}), \mathcal{R}) & \longrightarrow & H^n(N(\mathcal{O}), \mathcal{R}) & \longrightarrow & \text{Hom}(H_n(N(\mathcal{O}), \mathbb{Z}), \mathcal{R}) & \longrightarrow & 0 \\ & & \downarrow (s_{n-1})^* & & \downarrow s^n & & \downarrow (s_n)^* & & \\ 0 & \longrightarrow & \text{Ext}(H_{n-1}(N(\mathcal{O}'), \mathbb{Z}), \mathcal{R}) & \longrightarrow & H^n(N(\mathcal{O}'), \mathcal{R}) & \longrightarrow & \text{Hom}(H_n(N(\mathcal{O}'), \mathbb{Z}), \mathcal{R}) & \longrightarrow & 0 \\ & & \downarrow (t_{n-1})^* & & \downarrow t^n & & \downarrow (t_n)^* & & \\ 0 & \longrightarrow & \text{Ext}(H_{n-1}(N(\mathcal{O}''), \mathbb{Z}), \mathcal{R}) & \longrightarrow & H^n(N(\mathcal{O}''), \mathcal{R}) & \longrightarrow & \text{Hom}(H_n(N(\mathcal{O}''), \mathbb{Z}), \mathcal{R}) & \longrightarrow & 0 \end{array}$$

Since X satisfies (DA_n) and (DA_{n-1}) , we have shown that the maps in homology s_n and t_{n-1} are trivial, provided that $n \geq 2$. Hence, in the above diagram, the left-hand map $(t_{n-1})^*$ and the right-hand map $(s_n)^*$ are trivial. This latter statement also holds for $n = 1$, since homology is free in degree 0, which leads to vanishing of the Ext terms.

A quick diagram chase then shows that the composition $t^n \circ s^n$ is trivial, as required. \square

We conclude this subsection with the following observation needed in Section 4.

Remark 1.11. Consider an open cover \mathcal{O} and a double super-refinement $\mathcal{O}'' \ll \mathcal{O}' \ll \mathcal{O}$. Then, for any pairwise-intersecting collection $\{O''_i\}_{i \in I} \subseteq \mathcal{O}''$, there exists some $O \in \mathcal{O}$ containing all O''_i .

Indeed, fixing some $i_0 \in I$, there exist elements $O'_i \in \mathcal{O}'$ with $O''_{i_0} \cup O''_i \subseteq O'_i$, since $\mathcal{O}'' \ll \mathcal{O}'$ and $O''_{i_0} \cap O''_i \neq \emptyset$. Now, all O'_i contain any given point of O''_{i_0} , hence their union is contained in an element of \mathcal{O} , again because $\mathcal{O}' \ll \mathcal{O}$.

1.3. Direct limits of compact spaces

We will prove the vanishing part of Theorems A and B by checking that our Morse boundaries satisfy condition (DA_n) in the required range and applying Corollary 1.10. Thus, we first need to show that these spaces are super-refinable, which is the goal of this subsection (Proposition 1.12).

Recall that, given a sequence of topological spaces $X_n \subseteq X_{n+1}$, the direct limit $X = \varinjlim X_n$ is the union of the X_n , endowed with the topology where $U \subseteq X$ is open if and only if $U \cap X_n$ is open for all n . All our Morse boundaries will be spaces of this form.

Proposition 1.12. *Any countable direct limit of compact metrisable spaces is super-refinable.*

Before proving the proposition, we need to make a couple of observations. First, we can and will assume in the following discussion that the union of the X_n is equipped with a single metric $d(\cdot, \cdot)$ that induces the topology of each X_n . (Of course, we stress that $d(\cdot, \cdot)$ does not induce the topology of the direct limit X , which is strictly finer.)

In general, the existence of the metric $d(\cdot, \cdot)$ is guaranteed by the following result; see [16, Lemma 1.4] or [25] for short proofs. For our Morse boundaries, we will actually simply consider the restriction of a visual metric on the Bowditch boundary.

Lemma 1.13. *Let A be a closed subspace of a compact metrisable space B . Every metric $d_A: A \times A \rightarrow \mathbb{R}$ inducing the topology of A can be extended to a metric $d_B: B \times B \rightarrow \mathbb{R}$ inducing the topology of B .*

The following will be our main recipe to construct open covers of countable direct limits of compact metrisable topological spaces.

Construction 1.14. Consider a countable direct limit $X = \varinjlim X_n$ of metric spaces and a sequence of positive reals $\underline{\epsilon} = (\epsilon_n)_{n \geq 0}$. We describe a family $\mathcal{U}(\underline{\epsilon})$ of open sets of X that will come in handy multiple times in the future.

There is an element $U_p \in \mathcal{U}(\underline{\epsilon})$ for each point $p \in X$, constructed as follows. If $p \in X_k - X_{k-1}$ for some $k \geq 0$, set $U_p(n) = \emptyset$ for $n < k$, and define $U_p(k)$ as the open ϵ_k -ball around p within X_k . For $n > k$, we define iteratively $U_p(n)$ as the open ϵ_n -neighbourhood of $U_p(n-1)$ within X_n . Finally, set $U_p := \bigcup_n U_p(n)$.

Note that, since $U_p(n)$ is an open subset of X_n , we have that $U_p \cap X_n$ is the union of the open sets $U_p(m) \cap X_n$ for $m \geq n$, and hence it is open. This shows that U_p is open in X , by definition of the direct limit topology.

Note that it is possible for $U_p(n)$ to be strictly contained in the intersection $U_p \cap X_n$. For instance, if $p \in X_0$, then $U_p \cap X_0$ is a ball of radius $\sum_{n \geq 0} \epsilon_n$, whereas $U_p(0)$ is just an ϵ_0 -ball. It is also useful to introduce the notation $\mathcal{U}_k(\underline{\epsilon}) := \{U_p \mid p \in X_k - X_{k-1}\}$ and $\mathcal{U}_k^n(\underline{\epsilon}) = \{U_p(n) \mid p \in X_k - X_{k-1}\}$.

If \mathcal{U} is a family of subsets of a topological space X and $A \subseteq X$ is another subset, it is convenient to write:

$$\mathcal{U}[A] := \bigcup_{U \in \mathcal{U}, U \cap A \neq \emptyset} \overline{U};$$

here \overline{U} denotes the closure of U in X . We are finally ready to prove Proposition 1.12.

Proof of Proposition 1.12. As observed above, it suffices to consider a sequence of isometric embeddings of compact metric spaces $\dots \hookrightarrow X_n \hookrightarrow X_{n+1} \hookrightarrow \dots$ with $n \geq 0$. Let $X = \varinjlim X_n$ be their direct limit. Let \mathcal{O} be an open cover of X . Our aim is to construct a super-refinement of \mathcal{O} .

We rely on Construction 1.14, for a sequence $\underline{\epsilon} = (\epsilon_n)_{n \geq 0}$ to be determined later. For every $k \geq 0$, choose a finite subset $I_k \subseteq X_k - X_{k-1}$ such that the ϵ_k -balls $\{U_i(k) \in \mathcal{U}_k^k(\underline{\epsilon}) \mid i \in I_k\}$ cover $X_k - X_{k-1}$. (For instance, cover X_k by finitely many $\epsilon_k/2$ -balls, select only those that intersect $X_k - X_{k-1}$, for each of them pick a point in its intersection with $X_k - X_{k-1}$, then take the finitely many ϵ_k -balls with these centres.)

Define $I := \bigsqcup_{n \geq 0} I_n$ and $\mathcal{U}_k^n := \{U_i(n) \mid i \in I_k\}$ (that is, the elements of \mathcal{U}_k^n are open sets of X_n constructed as in Construction 1.14 starting from the points of I_k). We say that a finite subset $J \subseteq I_0 \cup \dots \cup I_k$ is k -good if $J \cap I_k \neq \emptyset$ and the union of closures $\bigcup_{j \in J} \overline{U_j(k)}$ is contained in some $O \in \mathcal{O}$; we fix one such open set O_J for every good set J .

Claim. *There exist numbers $\epsilon_n, \eta_n > 0$ such that the following hold for $n \geq 0$.*

- (A) *If B is an η_n -ball in X_n , then the set $(\mathcal{U}_0^n \cup \dots \cup \mathcal{U}_n^n)[B]$ is contained in an element of \mathcal{O} .*
- (B) *For every $0 \leq k < n$ and every k -good set J , we have $\bigcup_{j \in J} \overline{U_j(n)} \subseteq O_J$.*

First, assuming the Claim, we construct a super-refinement of \mathcal{O} . For each $i \in I_k$, set $V_i := U_i - X_{k-1}$ and define $\mathcal{V} = \{V_i \mid i \in I\}$. It is clear that this is an open cover of X . Let us show that it super-refines \mathcal{O} .

Consider a point $x \in X_n$ and let $J \subseteq I$ be the subset of indices j such that $x \in V_j$. Note that $J \subseteq I_0 \cup \dots \cup I_n$, so J is finite. It follows that, for some $m > n$, the point x lies in $U_j(m)$ for all $j \in J$. By property (A), there exists $O \in \mathcal{O}$ containing $\bigcup_{j \in J} \overline{U_j(m)}$. Thus, J is k -good for some k and property (B) ensures that $O_J \in \mathcal{O}$ contains $\bigcup_{j \in J} \overline{U_j(m')}$ for all $m' > k$. This implies that O_J contains all U_j with $j \in J$, hence all V_j with $j \in J$, as we wanted.

Note that we have only used property (A) with the ball B replaced by a single point, and the need for closures of the $U_j(n)$ is also not apparent yet. The full strength of property (A) is needed in the proof of the Claim.

We conclude by proving the Claim, that is, by showing that ϵ_n, η_n can be chosen in a way that (A) and (B) hold. Let λ_n be a Lebesgue number for the open cover $\{O \cap X_n \mid O \in \mathcal{O}\}$ of X_n . We proceed by induction on n .

Consider the base step $n = 0$. Property (B) holds vacuously in this case (note the “ $k < n$ ”). Regarding property (A), consider an η_0 -ball $B \subseteq X_0$ and recall that \mathcal{U}_0^0 is a finite cover of X_0 by ϵ_0 -balls. The diameter of $\mathcal{U}_0^0[B]$ is at most $4\epsilon_0 + 2\eta_0$. Choosing ϵ_0, η_0 with $4\epsilon_0 + 2\eta_0 < \lambda_0$, we ensure that $\mathcal{U}_0^0[B]$ is contained in an element of \mathcal{O} .

For the inductive step, suppose that $\epsilon_0, \dots, \epsilon_n$ and η_0, \dots, η_n have been chosen to satisfy (A) and (B), and let us look for suitable values of ϵ_{n+1} and η_{n+1} .

Observe that there are only finitely many k -good sets J with $k \leq n$ and, for each of them, we have

$$d\left(\bigcup_{j \in J} \overline{U_j(n)}, X_{n+1} - O_J\right) > 0,$$

as this is the distance between two disjoint compact sets. Thus, we can choose a number $\delta_{n+1} > 0$ that is smaller than all these distances. We then choose ϵ_{n+1} and η_{n+1} so that the following inequalities hold:

$$4\epsilon_{n+1} + 2\eta_{n+1} < \lambda_{n+1}, \quad 2\epsilon_{n+1} + 2\eta_{n+1} < \eta_n, \quad 3\epsilon_{n+1} + 2\eta_{n+1} < \delta_{n+1}.$$

This immediately guarantees that property (B) is satisfied, as $\epsilon_{n+1} < \delta_{n+1}$ and each $U_j(n+1)$ is the ϵ_{n+1} -neighbourhood of $U_j(n)$ in X_{n+1} .

Regarding property (A), consider an η_{n+1} -ball $B \subseteq X_{n+1}$. Let $J \subseteq I_0 \cup \dots \cup I_n$ be the set of indices j such that $U_j(n+1) \cap B \neq \emptyset$.

If $J = \emptyset$, then $(\mathcal{U}_0^{n+1} \cup \dots \cup \mathcal{U}_{n+1}^{n+1})[B] = \mathcal{U}_{n+1}^{n+1}[B]$. As in the base step, this is a set of diameter at most $4\epsilon_{n+1} + 2\eta_{n+1} < \lambda_{n+1}$, which ensures that it is contained in an element of \mathcal{O} .

If instead $J \neq \emptyset$, we can choose for every $j \in J$ a point $x_j \in U_j(n)$ with $d(x_j, B) < \epsilon_{n+1}$. The set $\{x_j \mid j \in J\}$ is contained in X_n and it has diameter at most $2\epsilon_{n+1} + 2\eta_{n+1} <$

η_n , so it is contained in an η_n -ball $B' \subseteq X_n$. For every $j \in J$, the set $U_j(n)$ meets B' , so property (A) for n implies that J is k -good for some $k \leq n$. Let $\mathcal{O}_J \in \mathcal{O}$ be the corresponding open set.

Now, the set $(\mathcal{U}_0^{n+1} \cup \dots \cup \mathcal{U}_{n+1}^{n+1})[B]$ is contained in the $(3\epsilon_{n+1} + 2\eta_{n+1})$ -neighbourhood of $\bigcup_{j \in J} \overline{U_j(n)}$ within X_{n+1} . Since $3\epsilon_{n+1} + 2\eta_{n+1} < \delta_{n+1}$, this set is contained in $\mathcal{O}_J \in \mathcal{O}$, as required. \square

1.4. Discrete chains in metric spaces

Let (X, ρ) be a metric space. In this short subsection, we introduce some important notation relating to chains in the presence of a metric.

If c is a discrete or singular chain in X , we write $\text{diam}(c) := \text{diam}(\text{supp}(c))$ for short. With a slight abuse, we will sometimes also write things like “ $\rho(x, c) \geq 1$ ” for a point $x \in X$, rather than the more correct “ $\rho(x, \text{supp}(c)) \geq 1$ ”.

If $c = \sum_i \sigma_i$, written as a sum of discrete (resp. singular) simplices without cancellations, we define

$$\text{fine}(c) := \sup_i \text{diam}(\sigma_i).$$

If $\text{fine}(c) \leq \delta$ for some $\delta > 0$, we say that c is δ -fine.

1.5. Discretisation of singular simplices

Given a singular chain d , one can associate to it a discrete chain that we will denote by $\text{disc}(d)$, informally by only keeping track of the vertices. More precisely, we have the following.

Definition 1.15. If $\sigma: \Delta^i \rightarrow X$ is a singular simplex, then $\text{disc}(\sigma_i)$ is the discrete simplex $[x_0, \dots, x_i]$, where x_i is the image under σ of the i -th vertex of the standard simplex Δ^i . This definition is then extended linearly; namely, for $d = \sum c_i \sigma_i$ we set $\text{disc}(d) := \sum c_i \text{disc}(\sigma_i)$.

Lemma 1.16. Let X be a metric space and let $\hat{d} \subseteq X$ be a singular chain. Then there exists a discrete chain $d \subseteq X$ with $\partial d = \text{disc}(\partial \hat{d})$ and $\text{fine}(d) \leq 2 \text{fine}(\partial \hat{d})$.

Proof. It suffices to perform an iterated relative barycentric subdivision of \hat{d} without ever altering the cell structure on $\partial \hat{d}$. More precisely, recall that, in the usual barycentric subdivision, the new vertices are the barycentres of all simplices. Instead, in this relative subdivision, the new vertices are the barycentres of simplices that are not part of $\partial \hat{d}$, as well as the vertices of simplices that are part of $\partial \hat{d}$.

Iterating this relative subdivision procedure, the fineness of \hat{d} converges to the fineness of $\partial \hat{d}$. Thus, we can just take $d = \text{disc}(\hat{d}')$ for a sufficiently high subdivision \hat{d}' of the original \hat{d} . \square

2. Filling discrete cycles in the Bowditch boundary

The goal of this section is to prove Proposition 2.1 below, which gives a way of filling cycles in the Bowditch boundary of certain relatively hyperbolic groups. We now recall some background on relatively hyperbolic groups, introduce the setup and state the proposition. We note that no other results from this section are used in the rest of the paper.

2.1. Cusped space and Bowditch boundary

We now briefly discuss cusped spaces and their boundaries, and we refer the reader to [19, Subsection 2.1] for more details, as we will mostly use techniques from that paper. These notions were first considered in [5].

Given a connected graph Γ , one can construct a hyperbolic space, called the *combinatorial horoball* $\mathcal{H}(\Gamma)$, by gluing strips of the hyperbolic plane onto the edges of the graph, see [19, Definition 2.1]. We refer to the copy of the graph contained in the combinatorial horoball as the *horosphere*, and we recall that, for $x, y \in \Gamma$, we have that $d_{\mathcal{H}(\Gamma)}(x, y)$ and $2 \log(d_{\Gamma}(x, y))$ differ by a uniformly bounded amount, see [19, Lemma 2.2].

A relatively hyperbolic group pair (G, \mathcal{P}) is a finitely generated group G together with a collection of finitely generated subgroups \mathcal{P} , called *peripheral subgroups*, such that one obtains a hyperbolic space by gluing combinatorial horoballs onto the cosets of the peripheral subgroups, see [19, Definition 2.3]. This glued-up space is the *cusped space* $\mathcal{X}(G, \mathcal{P})$, and its Gromov boundary is the *Bowditch boundary* $\partial_B(G, \mathcal{P})$. The combinatorial horoballs inside the cusped space lie at bounded Hausdorff distance from horoballs as defined in terms of Busemann functions, see [5, Lemma 2.5].

2.2. Setup and statement of main proposition

Let (G, \mathcal{P}) be a relatively hyperbolic pair. We assume that G is finitely generated and that \mathcal{P} consists of finitely many finitely generated, virtually nilpotent subgroups. In addition, we assume that the Bowditch boundary $\mathcal{S} := \partial_B(G, \mathcal{P})$ is homeomorphic to a sphere S^{k+1} for some $k \geq 1$.

For $a, b > 0$ and $K \geq 1$, we write $a \sim_K b$ as shorthand for $1/K \leq a/b \leq K$. As customary, we also write $a \vee b := \max\{a, b\}$ and $a \wedge b := \min\{a, b\}$.

Let $(\mathcal{X}, d_{\mathcal{X}})$ be the cusped space for (G, \mathcal{P}) and let δ_0 be its hyperbolicity constant. For simplicity we choose, for all points $x \in \mathcal{X}$ and $y \in \mathcal{X} \cup \partial\mathcal{X}$, a geodesic from x to y , which we denote by $[x, y]$.

Let ϵ_0 be a parameter for visual metrics, which can be chosen depending only on δ_0 (e.g. as in [3, Proposition III.H.3.21]). Denote by ρ_x the visual metric on $\mathcal{S} = \partial\mathcal{X}$ based at a point $x \in \mathcal{X}$. There exists a constant C (depending only on δ_0, ϵ_0) such that $\rho_x(\xi, \xi') \sim_C e^{-\epsilon_0(\xi|\xi')_x}$ for all $x \in \mathcal{X}$ and $\xi, \xi' \in \partial\mathcal{X}$. Here $(\cdot|\cdot)_x$ denotes Gromov products based at x . Choose a basepoint $o \in \mathcal{X}$ and set $\rho := \rho_o$.

For each parabolic point $p \in \partial\mathcal{X}$, denote by $\mathcal{H}_p \subseteq \mathcal{X}$ the corresponding horoball and by $\mathcal{K}_p \subseteq \mathcal{H}_p$ its boundary horosphere. Choose a point $e_p \in \mathcal{K}_p$ nearest to o and set $r_p := e^{-\epsilon_0 d_{\mathcal{X}}(o, \mathcal{H}_p)}$.

If Z is a metric space, $x \in Z$ and $0 < r_1 < r_2$, we denote by $A_Z(x; r_1, r_2)$ the closed metric annulus around x with inner radius r_1 and outer radius r_2 . When the metric space under consideration is clear, we simply write $A(x; r_1, r_2)$. Similarly, $B_Z(x, r)$ or $B(x, r)$ is the closed r -ball around x .

We are ready to state the main result of Section 2:

Proposition 2.1. *Given the setup above, there exist a function $g: \mathbb{R}^+ \rightarrow \mathbb{R}^+$ and a constant $K \geq 1$ such that the following statements hold for every $\delta > 0$. Let $c \in \overline{\mathcal{C}}_i(\mathcal{S})$ be a reduced discrete cycle with $\text{fine}(c) \leq g(\delta)$.*

- (0) *If $i \leq k$, then $c = \partial d$ for a δ -fine discrete chain $d \subseteq \mathcal{S}$ with $\text{diam}(d) \leq K\sqrt{\text{diam}(c)}$.*
- (1) *If $i < k$ and $c \subseteq A(p; r_1, r_2)$ for a parabolic point $p \in \mathcal{S}$ and radii $\delta \leq r_1 < r_2 \leq r_p/K$, then $c = \partial d$ for a δ -fine discrete chain $d \subseteq \mathcal{S}$ with $\text{diam}(d) \leq K\sqrt{\text{diam}(c)}$ and $d \subseteq A(p; r_1/K, Kr_2)$.*
- (2) *If $i \leq k$ and $c \subseteq B(p, r_2)$ for a parabolic point $p \in \mathcal{S}$ and a radius $0 < r_2 \leq r_p/K$, then $c = \partial d$ for a δ -fine discrete chain $d \subseteq \mathcal{S}$ with $\text{diam}(d) \leq K\sqrt{\text{diam}(c)}$ and $d \subseteq B(p, Kr_2)$.*

We note that, if we are only interested in fundamental groups of finite-volume real hyperbolic manifolds, then we can use \mathbb{H}^n instead of the cusped space \mathcal{X} . In this case, the boundary is isometric to a round sphere and the proposition is almost immediate.

In particular, the reader only interested in the proof of Theorem A can safely skip the rest of Section 2, which is more technical than the rest of the paper.

2.3. Geometry of nilpotent groups

The following lemma can be easily deduced from a result of Karidi on balls in nilpotent Lie groups [17]. Roughly speaking it says that, in an n -dimensional, connected, nilpotent Lie group, discrete i -cycles can be filled within annuli by chains of roughly the same fineness for $i < n - 1$, while $(n - 1)$ -cycles can be “discretely homotoped” far away, again with the same fineness.

Lemma 2.2. *Let \mathcal{N} be an n -dimensional, connected, nilpotent Lie group, equipped with a left-invariant Riemannian metric. There exists $L \geq 1$ such that the following holds for all $r_2 > r_1 \geq L$, all $x \in \mathcal{N}$ and all discrete cycles $c \in \overline{\mathcal{C}}_i(A_{\mathcal{N}}(x; r_1, r_2))$.*

- (1) *If $0 \leq i < n - 1$ and $\text{fine}(c) \leq r_1/L$, then $c = \partial d$ for some discrete chain $d \subseteq A_{\mathcal{N}}(x; r_1/L, Lr_2)$ with $\text{fine}(d) \leq L(\text{fine}(c) \vee 1)$.*

- (2) If $i = n - 1$ and $\text{fine}(c) \leq r_1/L$, then $c - \partial d \subseteq A_{\mathcal{N}}(x; r_2/L, Lr_2)$ for some discrete chain $d \subseteq A_{\mathcal{N}}(x; r_1/L, Lr_2)$ with $\text{fine}(d) \leq L(\text{fine}(c) \vee 1)$.

Proof. For simplicity, we write $B(r) := B_{\mathcal{N}}(1, r)$ and $A(r_1, r_2) := A_{\mathcal{N}}(1; r_1, r_2)$, where $1 \in \mathcal{N}$ is the identity. By the main result of [17], there exist a constant $a > 1$ and a collection of topological n -discs $1 \in \mathfrak{B}(r) \subseteq \mathcal{N}$ such that, for every $r > 1$, we have $\mathfrak{B}(r/a) \subseteq B(r) \subseteq \mathfrak{B}(ar)$. In addition, the sets $\mathfrak{A}(r_1, r_2) := \mathfrak{B}(r_2) - \mathfrak{B}(r_1)$ are homeomorphic to $S^{n-1} \times (0, 1]$.

First of all, observe that for every (reduced) singular cycle $\hat{c} \subseteq \mathcal{N}$, we have $\hat{c} = \partial \hat{d}$ for a singular chain \hat{d} with $\text{diam}(\hat{d}) \leq 2a^2 \cdot (\text{diam}(\hat{c}) \vee 2)$. Indeed, setting $D := \text{diam}(\hat{c}) \vee 2$, we have $\hat{c} \subseteq B(D) \subseteq \mathfrak{B}(aD)$ up to left multiplication. Since $\mathfrak{B}(aD)$ is a topological disc, there exists a singular chain $\hat{d} \subseteq \mathfrak{B}(aD)$ with $\partial \hat{d} = \hat{c}$. Finally, since $\mathfrak{B}(aD) \subseteq B(a^2D)$, we have $\text{diam}(\hat{d}) \leq 2a^2D$.

Consider now a discrete chain c . We want to argue that there exists a singular chain \hat{c} with $\text{disc}(\hat{c}) = c$ and fineness controlled in terms of the fineness of c . We can obtain this by “filling in” discrete 1-simplices with singular 1-simplices, at which point discrete 2-simplices yield singular 1-cycles, and we can fill in those as well, and proceed inductively. It follows that, for every discrete i -chain $c \subseteq \mathcal{N}$, there exists a singular i -chain $\hat{c} \subseteq \mathcal{N}$ with $c = \text{disc}(\hat{c})$ and

$$\text{fine}(\hat{c}) \leq \frac{1}{2}(2a)^{2i} \cdot (\text{fine}(c) \vee 2).$$

Now, consider a discrete cycle $c \in \overline{\mathcal{C}}_i(A_{\mathcal{N}}(x; r_1, r_2))$ with $\text{fine}(c) \vee 2 \leq (2a)^{-2i}r_1$ and $i \geq 0$. Up to left multiplication, we can assume that $c \subseteq A(r_1, r_2)$. By the previous paragraph, we have $c = \text{disc}(\hat{c})$ for a (reduced) singular cycle \hat{c} with $\text{fine}(\hat{c}) \leq r_1/2$. In particular, we have

$$\hat{c} \subseteq A(r_1/2, 2r_2) \subseteq \mathfrak{A}(r_1/2a, 2ar_2) := \mathfrak{A}.$$

Note that $\mathfrak{A} \subseteq A(r_1/2a^2, 2a^2r_2)$.

Recall that \mathfrak{A} is homeomorphic to $S^{n-1} \times (0, 1]$. Thus, if $i \neq n - 1$, there exists a singular chain $\hat{d} \subseteq \mathfrak{A}$ with $\hat{c} = \partial \hat{d}$. If $i = n - 1$, there exists nevertheless $\hat{d} \subseteq \mathfrak{A}$ with $\partial \hat{d} = \hat{c} - \hat{c}'$ where $\hat{c}' \subseteq \mathfrak{A}(r_2/2a, 2ar_2)$ is a singular chain homologous to \hat{c} , which we can and will arrange to be finer than \hat{c} (e.g. by barycentric subdivision). Note that \hat{c}' exists since $\mathfrak{A}(r_2/2a, 2ar_2)$ is a deformation retract of \mathfrak{A} .

Finally, since $\text{fine}(\partial \hat{d}) = \text{fine}(\hat{c}) \leq \frac{1}{2}(2a)^{2i}(\text{fine}(c) \vee 2)$, Lemma 1.16 yields a discrete chain $d \subseteq \mathfrak{A}$ with $\partial d = \text{disc}(\partial \hat{d})$ and $\text{fine}(d) \leq (2a)^{2i}(\text{fine}(c) \vee 2)$. For $i \neq n - 1$, we have $\partial d = c$ and, for $i = n - 1$, we have $c - \partial d = \text{disc}(\hat{c}') \subseteq \mathfrak{A}(r_2/2a, 2ar_2)$, as required. \square

2.4. Preliminary lemmas

In the rest of Section 2, we consider the setup described in Subsection 2.2. Also recall that $\overline{\mathcal{C}}_*(\mathcal{S})$ denotes the chain complex of reduced discrete chains in the Bowditch boundary \mathcal{S} , where 0-chains are required to have zero coefficient sum.

The following lemma says that we can fill discrete cycles while staying away from a specified point in the boundary, and that each point has neighbourhoods where cycles can be filled.

For short, we write $f = o(1)$ for a function $f : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ with $\lim_{t \rightarrow 0} f(t) = 0$.

Lemma 2.3. *There exist functions $f, h = o(1)$ such that the following statements hold for every point $\xi \in \mathcal{S}$ and every $\delta, \epsilon > 0$.*

- (1) *If $c \in \overline{\mathcal{C}}_*(\mathcal{S})$ is a discrete δ -fine cycle with $\rho(\xi, c) > \epsilon$, then there exists a discrete chain $d \subseteq \mathcal{S}$ with $\partial d = c$ and $\text{fine}(d) \leq f(\delta)$ and $\rho(\xi, d) > h(\epsilon)$.*
- (2) *There exists an open subset $\xi \in W \subseteq \mathcal{S}$ such that $\text{diam}(W) \leq \epsilon$ and such that, for every discrete δ -fine cycle $c \in \overline{\mathcal{C}}_*(W)$, there exists a discrete chain $d \subseteq W$ with $\partial d = c$ and $\text{fine}(d) \leq f(\delta)$.*

The reason why, in the above statement, we introduce δ and conclude that $\text{fine}(d) \leq f(\delta)$, rather than simply writing $\text{fine}(d) \leq f(\text{fine}(c))$, is to deal with 0-cycles. These are 0-fine but cannot be filled with a 0-fine 1-chain.

Proof. The lemma follows from the fact that \mathcal{S} is homeomorphic to a sphere, and the fact that the statement holds for a sphere with its round metric. More precisely, let us consider two metrics on \mathcal{S} , namely ρ and a metric that makes it isometric to a round sphere.

For item (2) we can take W to be a ball in the round metric. In said ball, any $c \in \overline{\mathcal{C}}_*(W)$ can be filled by a chain with the same fineness (with respect to the round metric). Comparing ρ to the round metric, using the fact that they induce the same topology on \mathcal{S} , yields item (2).

Item (1) is similar, with the small difference that, in the complement of a ball in the round metric, any $c \in \overline{\mathcal{C}}_*(W)$ can be filled by a chain whose fineness is at most a fixed multiple of $\text{fine}(c)$. \square

For each parabolic point $p \in \mathcal{S} = \partial\mathcal{X}$, we can define two maps projecting the Bowditch boundary to the horosphere around p and vice versa:

$$\Psi_p : \mathcal{S} - \{p\} \rightarrow \mathcal{K}_p, \quad \Phi_p : \mathcal{K}_p \rightarrow \mathcal{S} - \{p\}.$$

More precisely, we define these maps as follows. For every $\xi \in \mathcal{S} - \{p\}$, the point $\Psi_p(\xi)$ is within distance 1 of a bi-infinite geodesic from p to ξ . This would not be strictly




Fig. 1. A point in the boundary and its “projection” to the horosphere corresponding to the parabolic point p , as in the proof of item (1) of Lemma 2.4.

necessary, but for convenience, we will require Ψ_p to be injective. This can always be arranged, since we allow Ψ_p to take values outside the vertex set of \mathcal{K}_p and any set of cardinality 2^{\aleph_0} (a suitable neighbourhood of ξ in \mathcal{S}) can be injected into any set of the same cardinality (a suitable neighbourhood of a point in \mathcal{K}_p). Conversely, for every $x \in \mathcal{K}_p$, the point $\Phi_p(x)$ is the endpoint at infinity of a bi-infinite geodesic from p that passes within distance 1 of x . Again for convenience, we will require that $\Phi_p \Psi_p(\xi) = \xi$ for all $\xi \in \mathcal{S} - \{p\}$.

The next three lemmas are general facts about cusped spaces, which do not require the Bowditch boundary to be a sphere, nor the peripherals to be virtually nilpotent. They are similar to [19, Lemma 5.6] and [20, Lemma 6.1], and can be proven with the same method. Namely, one can use that finite configurations of points in a hyperbolic space can be approximated by trees, and compute all relevant Gromov products in said trees (see [19, Figures 1 and 2]). We give more details below only for the first lemma.

We denote by $d_{\mathcal{K}_p}(\cdot, \cdot)$ the intrinsic path metric of a horosphere \mathcal{K}_p .

Lemma 2.4. *There exists a constant $C \geq 1$, depending only on δ_0, ϵ_0 , such that the following statements hold for all $x, y \in \mathcal{S} - \{p\}$:*

- (1) $d_{\mathcal{K}_p}(e_p, \Psi_p(x)) \sim_C 1 \vee \left(\frac{r_p}{\rho(p,x)}\right)^{1/\epsilon_0}$;
- (2) if $\rho(x, y) \leq \rho(p, \{x, y\})$, then $d_{\mathcal{K}_p}(\Psi_p(x), \Psi_p(y)) \sim_C 1 \vee \left(\frac{r_p \cdot \rho(x, y)}{\rho(p, x)^2}\right)^{1/\epsilon_0}$.

Proof sketch. The outline for item (1) is as follows, in the case that $\Psi_p(x)$ is sufficiently far from e_p . Fig. 1 shows an approximating tree for the relevant points. The two highlighted segments have the same length ℓ because their endpoints lie on a horosphere. We have that $(p|x)_o$ is equal to $d(o, e_p) + \ell$ up to bounded additive error, and we also have that $d_{\mathcal{K}_p}(e_p, \Psi_p(x)) \sim_D e^\ell$ for some D depending only on the cusped space \mathcal{X} . Recalling




Fig. 2. Illustration of the proof idea of item (2) of Lemma 2.4.

the definition of the visual metric, we have that $\rho(p, x) \sim_{D'} r_p e^{-\epsilon_0 \ell}$ for some D' . Hence, $(r_p/\rho(p, x))^{1/\epsilon_0} \sim_{D''} e^\ell$ for some D'' and, as mentioned above, the right-hand side is $\sim_D d_{K_p}(e_p, \Psi_p(x))$, as required.

In item (2), in principle there would be two approximating trees to consider, up to swapping x and y . These are pictured in Fig. 2. However, in the right-hand configuration, $(p|x)_o - (x|y)_o$ is approximately t as marked in the Fig. 2. Thus, if t is large we have $\rho(x, y) \gg \rho(p, x, y)$, while for small t the right-hand configuration collapses into the left-hand one. Estimates analogous to those used for item (1) conclude the proof. \square

Lemma 2.5. *There exists a constant $C \geq 1$, depending only on δ_0, ϵ_0 , such that the following statements hold for all points $u, v \in K_p$ outside the C -ball around e_p :*

- (1) $\rho(p, \Phi_p(u)) \sim_C \frac{r_p}{d_{K_p}(e_p, u)^{\epsilon_0}}$;
- (2) if $d_{K_p}(u, v) \leq \frac{1}{2} \cdot d_{K_p}(e_p, u)$, then $\rho(\Phi_p(u), \Phi_p(v)) \leq C \cdot \frac{r_p \cdot (1 \vee d_{K_p}(u, v))^{\epsilon_0}}{d_{K_p}(e_p, u)^{2\epsilon_0}}$.

For a discrete chain $c \subseteq \mathcal{S}$, we define the *barycentre* $\mathfrak{b}(c) \in \mathcal{X} \cup \mathcal{S}$ as follows: choose a vertex $\xi \in c$ and let $\mathfrak{b}(c) \in [o, \xi]$ be the point with $d_{\mathcal{X}}(o, \mathfrak{b}(c)) = \inf_{\eta, \eta' \in c} (\eta|\eta')_o$. Observe that $\text{diam}(c) \sim_C e^{-\epsilon_0 d_{\mathcal{X}}(o, \mathfrak{b}(c))}$, for a constant C only depending on δ_0, ϵ_0 .

Lemma 2.6. *For every $K \geq 1$, there exists $D \geq 0$ such that the following holds for all parabolic points $p \in \mathcal{S}$, all points $\xi \in B(p, r_p) \subseteq \mathcal{S}$ and all points $z \in [0, \xi]$.*

- (1) If $\frac{1}{K} \cdot \frac{\rho(p, \xi)^2}{r_p} \leq e^{-\epsilon_0 d_{\mathcal{X}}(o, z)} \leq r_p$, then $d_{\mathcal{X}}(z, \mathcal{H}_p) \leq D$.
- (2) If $e^{-\epsilon_0 d_{\mathcal{X}}(o, z)} \leq K \cdot \frac{\rho(p, \xi)^2}{r_p}$, then $d_{\mathcal{X}}(z, \mathcal{X} - \mathcal{H}_p) \leq D$.

In addition, there exists K_0 , depending only on δ_0, ϵ_0 , with the following property.

(3) If $z \in \mathcal{H}_p$, then $\frac{1}{K_0} \cdot \frac{\rho(p, \xi)^2}{r_p} \leq e^{-\epsilon_0 d_{\mathcal{X}}(o, z)}$.

Finally, the next lemma is a consequence of Lemma 2.2 and we provide a complete proof.

Lemma 2.7. *There exists a constant $L \geq 1$ such that the following holds for every horosphere $\mathcal{K} = \mathcal{K}_p \subseteq \mathcal{X}$ with $p \in \mathcal{S}$ a parabolic point, all $r_2 > r_1 \geq L$ and all discrete cycles $c \in \overline{\mathcal{C}}_i(A_{\mathcal{K}}(e_p; r_1, r_2))$.*

- (1) *If $i < k$ and $\text{fine}(c) \leq r_1/L$, then $c = \partial d$ for some discrete chain $d \subseteq A_{\mathcal{K}}(e_p; r_1/L, Lr_2)$ with $\text{fine}(d) \leq L(1 \vee \text{fine}(c))$.*
- (2) *If $i = k$ and $\text{fine}(c) \leq r_1/L$, then $c - \partial d \subseteq A_{\mathcal{K}}(e_p; r_2/L, Lr_2)$ for some discrete chain $d \subseteq A_{\mathcal{K}}(e_p; r_1/L, Lr_2)$ with $\text{fine}(d) \leq L(1 \vee \text{fine}(c))$.*

Proof. Since there are only finitely many G -orbits of horospheres in \mathcal{X} , and peripherals are virtually nilpotent, it suffices to prove the statement in every finitely generated nilpotent group N (equipped with a word metric). Passing to a finite-index subgroup, we can assume that N is torsion-free; see e.g. [1, Theorem 2.1]. Set $n := \sum \text{rk}(N^{(i)}/N^{(i+1)})$.

The group N is a uniform lattice in its Mal'cev completion \mathcal{N} , which is an n -dimensional, connected, nilpotent Lie group; see e.g. [13, Theorem 13.40].

Observe that $n = k + 1$. Indeed, N acts freely and cocompactly both on \mathcal{N} and $\mathcal{S} - \{*\}$ (since the parabolic point corresponding to N is a *bounded parabolic point*, see e.g. [5, Section 6]). These spaces are homeomorphic, respectively, to \mathbb{R}^n and \mathbb{R}^{k+1} , hence the cohomological dimension of N equals both n and $k + 1$.

Now, since the inclusion $N \hookrightarrow \mathcal{N}$ is a quasi-isometry, the statement follows from Lemma 2.2. \square

2.5. The main argument

In order to simplify notation throughout this subsection, we will write $a \sim b$, for two quantities $a, b > 0$, when there exists a constant $C \geq 1$, depending only on δ_0, ϵ_0 , such that $a \sim_C b$. Similarly, we write $a \lesssim b$ when $a \leq Cb$ for such a constant C .

For a constant $D \geq 0$, we say that a point $x \in \mathcal{X}$ is D -deep in a horoball \mathcal{H}_p if $x \in \mathcal{H}_p$ and $d_{\mathcal{X}}(x, \mathcal{K}_p) \geq D$. Recall that $\overline{\mathcal{C}}_*(\mathcal{S})$ denotes *reduced* discrete chains in \mathcal{S} ; this means that 0-chains are required to have zero coefficient sum.

We now embark in the proof of Proposition 2.1. The following result gives us some initial control on fineness and diameter of fillings.

Lemma 2.8. *For every $D \geq 0$, there exist a function $f = o(1)$ and a constant $K \geq 1$ with the following property. Let $c \in \overline{\mathcal{C}}_*(\mathcal{S})$ be a discrete cycle such that some point $z \in [o, \mathbf{b}(c)]$ is not D -deep in any horoball. If c is δ -fine for some $\delta > 0$, then $c = \partial d$ for a discrete chain $d \subseteq \mathcal{S}$ with $\text{fine}(d) \leq f(e^{\epsilon_0 d_{\mathcal{X}}(o, z)} \cdot \delta)$ and $\text{diam}(d) \leq Ke^{-\epsilon_0 d_{\mathcal{X}}(o, z)}$.*




Fig. 3. The points relevant for case (1) in the proof of Lemma 2.8.

Proof. Throughout the proof, balls and diameters will usually refer to the metric $\rho = \rho_o$; in case of ambiguity, we will indicate the relevant metric as a subscript.

The reason for introducing δ in the statement is the same as for Lemma 2.3. In the argument below, we will instead use $\text{fine}(c)$ with a slight abuse, even though for 0-cycles we should use some $\delta > 0$ instead.

We distinguish two cases in the proof.

Case (1): the geodesic $[o, z]$ goes at least $(D + K' + 10\delta_0)$ -deep into a horoball \mathcal{H}_p , for a constant K' only depending on D to be determined below.

Consider the metrics $\rho := \rho_o$ and $\rho' := \rho_z$. Since $[0, z] \subseteq [0, \mathbf{b}(c)]$ goes $(D + 10\delta_0)$ -deep into \mathcal{H}_p , while z is not D -deep in \mathcal{H}_p , all geodesic lines from p to the vertices of c intersect the $10\delta_0$ -neighbourhood of z . Hence $\rho'(p, c)$ is bounded away from zero purely in terms of δ_0, ϵ_0 , see Fig. 3.

Since z is not D -deep in any horoball, there exists $g \in G$ such that the metrics ρ' and ρ_{go} are K -bi-Lipschitz equivalent, for a constant K only depending on D . In addition, g gives an isometry between (\mathcal{S}, ρ_o) and (\mathcal{S}, ρ_{go}) , and Lemma 2.3(1) holds in the former space. We conclude that there exists $d \subseteq \mathcal{S}$ with $\partial d = c$ and

$$\text{fine}_{\rho'}(d) \leq f_D(\text{fine}_{\rho'}(c)), \quad \rho'(p, d) \geq \eta_D,$$

for a function $f_D = o(1)$ and a constant $\eta_D > 0$ depending on D . In particular, we have $\sup_{\xi \in d}(\xi|p)_z \leq K'$, for a constant K' only depending on D and δ_0 .

Let $x, y \in [o, z]$ be the points with $d_{\mathcal{X}}(o, x) = (p|z)_o$ and $d_{\mathcal{X}}(y, z) = \sup_{\xi \in d}(\xi|p)_z$. Assuming now that $[o, z]$ goes at least $(D + K' + 10\delta_0)$ -deep into \mathcal{H}_p , we have $d_{\mathcal{X}}(o, y) \geq d_{\mathcal{X}}(o, x) + 10\delta_0$. Thus, all rays from o to the vertices of d track the geodesic $[o, z]$ at least up to y , hence $\inf_{\xi, \xi' \in d}(\xi|\xi')_o \geq d_{\mathcal{X}}(o, y) - 10\delta_0$ and $\text{diam}_{\rho}(d) \lesssim e^{-\epsilon_0 d_{\mathcal{X}}(o, y)} \leq e^{\epsilon_0 K'} e^{-\epsilon_0 d_{\mathcal{X}}(o, z)}$, as required.

We are left to control the fineness of d . Let $E(d)$ be the set of pairs $(\xi, \xi') \in \mathcal{S}^2$ corresponding to the edges of d . Recall that $\text{fine}_{\rho'}(d) \leq f_D(\text{fine}_{\rho'}(c))$. In addition, ρ and ρ' are $e^{\epsilon_0 d_{\mathcal{X}}(o,z)}$ -bi-Lipschitz equivalent. It follows that, for some constant C depending only on δ_0, ϵ_0 , we have:

$$\inf_{(\xi, \xi') \in E(d)} (\xi|\xi')_z \geq -\frac{\log f_D(\text{fine}_{\rho'}(c))}{\epsilon_0} - C \geq -\frac{\log f_D(e^{\epsilon_0 d_{\mathcal{X}}(o,z)} \text{fine}_{\rho}(c))}{\epsilon_0} - C.$$

Suppose for a moment that, for some $t \geq 1$, the right-hand side is at least $K' + (\log t)/\epsilon_0$. Then, recalling that $K' \geq \sup_{\xi \in d} (\xi|p)_z = d_{\mathcal{X}}(y, z)$, we obtain:

$$\inf_{(\xi, \xi') \in E(d)} (\xi|\xi')_o \geq d_{\mathcal{X}}(o, y) + (\log t)/\epsilon_0 - 10\delta_0.$$

We then have $\text{fine}_{\rho}(d) \lesssim e^{-\epsilon_0 d_{\mathcal{X}}(o,y)}/t \lesssim 1/t$. In conclusion, if $e^{\epsilon_0 d_{\mathcal{X}}(o,z)} \text{fine}_{\rho}(c)$ is smaller than a certain constant depending on D and t , then $\text{fine}_{\rho}(d) \leq 1/t$.

This settles case (1).

Case (2): *the geodesic $[o, z]$ does not go $(D + K' + 10\delta_0)$ -deep into any horoball.*

In this case we will use a form of (partial) self-similarity of the boundary made precise in [19, Lemma 4.6].

Let $D' \geq 0$ be a constant such that the geodesic $[o, z]$ is contained in the D' -neighbourhood of the orbit $G_o \subseteq \mathcal{X}$. Let $L_0 = L_0(D')$ be the constant provided by [19, Lemma 4.6]. Let f be the function provided by Lemma 2.3, and choose an open cover \mathscr{W} of \mathcal{S} by the sets provided by Lemma 2.3(2) with $\epsilon = 1/L_0$. Let α be a Lebesgue number for \mathscr{W} .

Let $y \in [o, z]$ be the point with $d_{\mathcal{X}}(y, z) = D''$, where D'' is a sufficiently large constant to be determined below. (If $d_{\mathcal{X}}(o, z) < D''$, we can directly apply Lemma 2.3(1) because, in this case, the condition on the diameter is vacuous if K is large enough.) Choose an element $g \in G$ with $d_{\mathcal{X}}(g^{-1}o, y) \leq D'$. Choose a vertex $\xi \in c$.

By [19, Lemma 4.6], g induces an L_0 -bi-Lipschitz map from the rescaled ball $(B(\xi, r), \rho_o/r)$ to an open subset $U \subseteq \mathcal{S}$ containing $B(g\xi, 1/L_0)$, for some radius $r \sim e^{-\epsilon_0 d_{\mathcal{X}}(o,y)}$. Recall that $\text{diam}(c) \sim e^{-\epsilon_0 d_{\mathcal{X}}(o,b(c))} \leq e^{-\epsilon_0 D''} e^{-\epsilon_0 d_{\mathcal{X}}(o,y)}$. Thus, if D'' is sufficiently large, we have $\text{diam}(c) \leq r$, hence $c \subseteq B(\xi, r)$ and:

$$\text{diam}(gc) \leq L_0 \text{diam}(c)/r \lesssim L_0 e^{-\epsilon_0 d_{\mathcal{X}}(y,b(c))} \leq L_0 e^{-\epsilon_0 D''}.$$

Now, if D'' is sufficiently large (depending only on D', δ_0, ϵ_0), we also have $\text{diam}(gc) \leq \alpha$. Since α is a Lebesgue number for \mathscr{W} , the cycle gc is then contained in some $W \in \mathscr{W}$. Since $\text{diam}(W) \leq 1/L_0$ and $g\xi \in gc \subseteq W$, we have $W \subseteq B(g\xi, 1/L_0) \subseteq U$. In addition, by our choice of \mathscr{W} , there exists a discrete chain $d' \subseteq W$ with $\partial d' = gc$ and $\text{fine}(d') \leq f(\text{fine}(gc))$.

In conclusion, setting $d := g^{-1}d'$, we have $c = \partial d$ and:

$$\text{diam}(d) \leq rL_0 \text{diam}(d') \leq rL_0 \text{diam}(\mathcal{S}) \lesssim L_0 e^{-\epsilon_0 d_{\mathcal{X}}(o,y)} = L_0 e^{\epsilon_0 D''} e^{-\epsilon_0 d_{\mathcal{X}}(o,z)},$$

$$\text{fine}(d) \leq rL_0 \text{fine}(d') \leq rL_0 f(\text{fine}(gc)) \lesssim L_0 f(L_0 \text{fine}(c)/r),$$

where $r \sim e^{\epsilon_0 D''} e^{-\epsilon_0 d_{\mathcal{X}}(o,z)}$. Since L_0 and D'' only depend on D' (and δ_0, ϵ_0), they only depend on D . This settles case (2), proving the lemma. \square

As a consequence of the previous lemma, we can already fill discrete cycles c in the Bowditch boundary with good fineness and diameter control, provided that the barycentre $\mathfrak{b}(c)$ is not deep in any horoball. This is the content of the following corollary.

Corollary 2.9. *For every $D \geq 0$, there exist a function $g : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ and $K \geq 1$ such that the following hold for all $\delta > 0$ and all discrete cycles $c \in \overline{\mathcal{C}}_*(\mathcal{S})$ with $\text{fine}(c) \leq g(\delta)$.*

- (1) *If some point $z \in [o, \mathfrak{b}(c)]$ is not D -deep in any horoball, then $c = \partial d$ for a discrete chain $d \subseteq \mathcal{S}$ with $\text{fine}(d) \leq \delta$ and $\text{diam}(d) \leq Ke^{-\epsilon_0 d_{\mathcal{X}}(o,z)}$.*
- (2) *If $\mathfrak{b}(c)$ is not D -deep in any horoball, then $c = \partial d$ for a discrete chain $d \subseteq \mathcal{S}$ with $\text{fine}(d) \leq \delta$ and $\text{diam}(d) \leq K \text{diam}(c)$.*

Proof. As in the previous proof, we will use $\text{fine}(c)$ with a slight abuse, rather than introducing a positive fineness constant.

Observe that part (2) follows from part (1) simply taking $z = \mathfrak{b}(c)$ and recalling that $\text{diam}(c) \sim e^{-\epsilon_0 d_{\mathcal{X}}(o, \mathfrak{b}(c))}$. (This is assuming that $\mathfrak{b}(c) \in \mathcal{X}$. If instead $\mathfrak{b}(c) \in \mathcal{S}$, the support of c is a single point, and we can take d supported on the same point.) We now prove part (1).

Let f be as in Lemma 2.8. Define $g(\delta) := \delta \cdot \tau$, for any $\tau > 0$ such that $f(\tau') \leq \delta$ for all $\tau' \leq \tau$. By the lemma, we have $c = \partial d$ for a discrete chain with $\text{diam}(d) \leq Ke^{-\epsilon_0 d_{\mathcal{X}}(o,z)}$ and $\text{fine}(d) \leq f(e^{\epsilon_0 d_{\mathcal{X}}(o,z)} \text{fine}(c))$. If $e^{\epsilon_0 d_{\mathcal{X}}(o,z)} \text{fine}(c) \leq \tau$, we have $\text{fine}(d) \leq \delta$ as required.

Otherwise $e^{-\epsilon_0 d_{\mathcal{X}}(o,z)} \leq \text{fine}(c)/\tau \leq \delta$ (since $\text{fine}(c) \leq g(\delta)$ by assumption), so we can just define d' as the cone over c from any of its vertices (Definition 1.2). Then $\partial d' = c$ and $\text{fine}(d') = \text{diam}(d') = \text{diam}(c)$, where $\text{diam}(c) \sim e^{-\epsilon_0 d_{\mathcal{X}}(o, \mathfrak{b}(c))} \leq e^{-\epsilon_0 d_{\mathcal{X}}(o,z)} \leq \delta$. \square

Just like Proposition 2.1, the next two lemmas roughly say that discrete cycles in annuli around parabolic points can be filled in slightly larger annuli. However, these two lemmas require an additional technical condition (cycle barycentres should lie in horoballs) that we can only drop later on by a bootstrapping argument.

Recall that the face complex $\mathcal{F}_*(d)$ of a discrete chain d was introduced in Definition 1.4.

Lemma 2.10. *There exists a constant $K \geq 1$ such that the following holds for every parabolic point $p \in \mathcal{S}$ and all $0 < r_1 < r_2 \leq r_p/K$.*

If c is a discrete cycle in $\overline{\mathcal{C}}_i(A(p; r_1, r_2))$ with $\mathfrak{b}(c) \in \mathcal{H}_p$ and $\text{fine}(c) \leq r_1^2/r_p$, then there exists a discrete chain $d \subseteq A(p; r_1/K, Kr_2)$ with $\text{diam}(d) \leq K\sqrt{\text{diam}(c)}$ and the following properties.

- (*) For every simplex $\sigma \in \mathcal{F}_*(d)$, we have $\text{diam}(\sigma) \leq K \cdot \frac{\rho(p,\sigma)^2}{r_p}$.
- (**) If $i < k$, we have $\partial d = c$. If $i = k$, we have $c - \partial d \subseteq B(p, Kr_1)$.

Proof. Choose a vertex $\xi \in c$ maximising the distance $\rho(p, \xi)$. Without loss of generality, we can assume that $r_2 = \rho(p, \xi)$. Since $\mathfrak{b}(c) \in \mathcal{H}_p$, Lemma 2.6(3) implies that $\text{diam}(c) \gtrsim r_2^2/r_p$.

By Lemma 2.4(1), we have $\Psi_p(c) \subseteq A_{\mathcal{K}_p}(e_p; R_1, R_2)$ with $R_1 \sim (r_p/r_2)^{1/\epsilon_0}$ and $R_2 \sim (r_p/r_1)^{1/\epsilon_0}$. Since $\text{fine}(c) \leq r_1^2/r_p$ by assumption, Lemma 2.4(2) also yields $\text{fine}(\Psi_p(c)) \lesssim 1$.

In view of Lemma 2.7 and the previous paragraph, there exists a constant L such that the following holds. Provided that $r_2 \leq r_p/L$, there exists a chain $\mathcal{D} \subseteq A_{\mathcal{K}_p}(e_p; R_1/L, LR_2)$ such that $\text{fine}(\mathcal{D}) \leq L$ and:

- if $i < k$, then $\partial \mathcal{D} = \Psi_p(c)$;
- if $i = k$, then $\Psi_p(c) - \partial \mathcal{D}$ is disjoint from $B_{\mathcal{K}_p}(e_p, R_2/L)$.

Now, set $d := \Phi_p(\mathcal{D})$. Lemma 2.5 guarantees that there exists a constant K depending on L, δ_0, ϵ_0 such that, if $r_2 \leq r_p/K$, the following holds: $d \subseteq A_{\mathcal{S}}(p; r_1/K, Kr_2)$ and, for every simplex $\sigma \subseteq d$, we have $\text{diam}(\sigma) \leq K \cdot \frac{\rho(p,\sigma)^2}{r_p}$. In particular, we have:

$$\text{diam}(d) \leq 2Kr_2 \lesssim 2K\sqrt{r_p \text{diam}(c)} \leq 2K\sqrt{\text{diam}(c)}.$$

Finally, if $i < k$, we have $\partial d = \Phi_p(\partial \mathcal{D}) = \Phi_p \Psi_p(c) = c$. If instead $i = k$, we obtain $c - \partial d \subseteq B(p, Kr_1)$ by another application of Lemma 2.5(1). \square

The fineness of the filling d provided by the previous lemma is not small enough for our purposes. The next result will solve this problem with a refining procedure.

Lemma 2.11. *There exist g and $K \geq 1$ such that the following holds for all parabolic points $p \in \mathcal{S}$ and all $\delta > 0$.*

Let $c \in \overline{\mathcal{C}}_i(\mathcal{S})$ be a cycle with $\text{fine}(c) \leq g(\delta)$ and $\mathfrak{b}(c) \in \mathcal{H}_p$. Then:

- (1) *if $i < k$ and $c \subseteq A(p; r_1, r_2)$ for radii $\delta \leq r_1 < r_2 \leq r_p/K$, we have $c = \partial d$ for some $d \subseteq A(p; r_1/K, Kr_2)$ with $\text{fine}(d) \leq \delta$ and $\text{diam}(d) \leq K\sqrt{\text{diam}(c)}$;*
- (2) *if $i \leq k$ and $c \subseteq B(p, r_2)$ for a radius $0 < r_2 \leq r_p/K$, we have $c = \partial d$ for some $d \subseteq B(p, Kr_2)$ with $\text{fine}(d) \leq \delta$ and $\text{diam}(d) \leq K\sqrt{\text{diam}(c)}$.*

Proof. We prove in detail only part (1), since the argument for part (2) is identical.

Part (1). Assume that $c \subseteq A(p; r_1, r_2)$ and $i < k$.

Ensuring that $g(\delta) \leq \delta^2$, we have $\text{fine}(c) \leq \delta^2 \leq r_1^2 \leq r_1^2/r_p$. Thus, Lemma 2.10 gives $c = \partial d$ for some $d \subseteq A(p; r_1/K, Kr_2)$ with $\text{diam}(d) \leq K\sqrt{\text{diam}(c)}$. In addition, for every simplex $\sigma \in \mathcal{F}_*(d)$, we have $\text{diam}(\sigma) \leq K \cdot \frac{\rho(p,\sigma)^2}{r_p}$.

Our only task is now to replace d by a δ -fine chain with the same properties. We will achieve this by constructing a chain map $R_*: \mathcal{F}_*(d) \rightarrow \mathcal{C}_*(\mathcal{S})$ such that:

- $R_*(\sigma) = \sigma$ for every η_0 -fine simplex $\sigma \in \mathcal{F}_*(d)$ (in particular, $R_0 = \text{id}$);
- for every simplex $\sigma \in \mathcal{F}_j(d)$ with $j \geq 1$, we have $\text{fine}(R_j(\sigma)) \leq \eta_j$ and $\text{diam}(R_j(\sigma) \cup \sigma) \leq M_j \cdot \frac{\rho(p, \sigma)^2}{r_p}$.

The constants η_j and M_j are defined inductively as follows. Set $M_0 := K$ and $\eta_k := \delta$. For each $j \geq 0$, let D_j be the constant provided by Lemma 2.6 in relation to the constant $2^5 M_j$. We then let K_{j+1} and g_j be the constant and function provided by Corollary 2.9 applied to the constant D_j . Finally, we set $\eta_j := g_j(\eta_{j+1})$ and $M_{j+1} := 3^5 K_{j+1} M_j + K$.

We now construct the chain map R_* . We proceed by induction on j , setting $R_0 = \text{id}$. For the inductive step, suppose that the degree- j map R_j has been defined with the above properties and let us construct R_{j+1} .

Without loss of generality, assume that $r_2 \leq r_p/2M_j$. This implies that, for every simplex $\tau \in \mathcal{F}_j(d)$, we have $\text{diam}(R_j(\tau) \cup \tau) \leq \frac{M_j r_2}{r_p} \cdot \rho(p, \tau) \leq \frac{1}{2} \rho(p, \tau)$. In particular, this guarantees that $\rho(p, R_j(\tau)) \geq \frac{1}{2} \rho(p, \tau)$. Similarly, assuming that $r_2 \leq r_p/2K$, we have $\text{diam}(\sigma) \leq \frac{1}{2} \rho(p, \sigma)$ for every simplex $\sigma \in \mathcal{F}_*(d)$.

As a consequence, for each simplex $\sigma \in \mathcal{F}_{j+1}(d)$ we have:

$$\begin{aligned} \text{diam}(R_j(\partial\sigma)) &\leq 2M_j \cdot \max_{\tau \subseteq \partial\sigma} \frac{\rho(p, \tau)^2}{r_p} \leq 2^3 M_j \cdot \frac{\rho(p, \partial\sigma)^2}{r_p} \leq 2^5 M_j \cdot \frac{\rho(p, R_j(\partial\sigma))^2}{r_p}, \\ \rho(p, R_j(\partial\sigma)) &= \min_{\tau \subseteq \partial\sigma} \rho(p, R_j(\tau)) \leq \min_{\tau \subseteq \partial\sigma} \frac{3}{2} \rho(p, \tau) \leq \frac{9}{4} \rho(p, \sigma). \end{aligned}$$

Now, for a simplex $\sigma \in \mathcal{F}_{j+1}(d)$, we need to define $R_{j+1}(\sigma)$ so that $\partial R_{j+1}(\sigma) = R_j(\partial\sigma)$. If σ is η_0 -fine, then so is $\partial\sigma$, hence $R_j(\partial\sigma) = \partial\sigma$ and we set $R_{j+1}(\sigma) := \sigma$.

Suppose instead that σ is not η_0 -fine. By the inductive hypothesis, we have $\text{fine}(R_j(\partial\sigma)) \leq \eta_j = g_j(\eta_{j+1})$. In addition, by the above inequalities, there exists a point $z \in [o, \mathfrak{b}(R_j(\partial\sigma))]$ with $e^{-\epsilon_0 d_{\mathcal{X}}(o, z)} = 2^5 M_j \cdot \frac{\rho(p, R_j(\partial\sigma))^2}{r_p}$. By Lemma 2.6 and our choice of D_j , we then have $d_{\mathcal{X}}(z, \mathcal{K}_p) \leq D_j$. In particular, z is not D_j -deep in any horoball, so Corollary 2.9(1) implies that $R_j(\partial\sigma) = \partial\sigma'$ for a discrete $(j+1)$ -chain $\sigma' \subseteq \mathcal{S}$ with $\text{fine}(\sigma') \leq \eta_{j+1}$ and $\text{diam}(\sigma') \leq K_{j+1} e^{-\epsilon_0 d_{\mathcal{X}}(o, z)}$. It follows that:

$$\text{diam}(\sigma \cup \sigma') \leq K \cdot \frac{\rho(p, \sigma)^2}{r_p} + 2^5 K_{j+1} M_j \cdot \frac{\rho(p, R_j(\partial\sigma))^2}{r_p} \leq (K + 3^5 K_{j+1} M_j) \cdot \frac{\rho(p, \sigma)^2}{r_p}.$$

Thus, we can set $R_{j+1}(\sigma) := \sigma'$.

This completes the construction of R_* . We now bound the diameter of $R_{i+1}(d)$ and show that this chain is contained in a suitable annulus around p .

Recall that $d \subseteq A(p; r_1/K, Kr_2)$ and that, for every simplex $\sigma \in \mathcal{F}_{i+1}(d)$, we have $\text{diam}(R_{i+1}(\sigma) \cup \sigma) \leq \frac{1}{2} \rho(p, \sigma)$ and $\rho(p, R_{i+1}(\sigma)) \geq \frac{1}{2} \rho(p, \sigma)$. This immediately implies that $R_{i+1}(d) \subseteq A(p; r_1/2K, 2Kr_2)$.

Without loss of generality, assume that $r_2 = \max_{\xi \in c} \rho(p, \xi)$. Since $\mathfrak{b}(c) \in \mathcal{H}_p$, Lemma 2.6(3) guarantees that $\text{diam}(c) \gtrsim r_2^2/r_p$. Set $M := \max_{j \leq k} M_j$. For every simplex $\sigma \in \mathcal{F}_*(d)$, we have:

$$\text{diam}(R_*(\sigma) \cup \sigma) \leq M \cdot \frac{\rho(p, \sigma)^2}{r_p} \leq MK^2 \cdot \frac{r_2^2}{r_p} \lesssim MK^2 \text{diam}(c) \lesssim MK^2 \sqrt{\text{diam}(c)}.$$

Since $\text{diam}(d) \leq K \sqrt{\text{diam}(c)}$, we obtain $\text{diam}(R_{i+1}(d)) \lesssim (MK^2 + K) \sqrt{\text{diam}(c)}$.

Finally, define $g(\delta) = \delta^2 \wedge \eta_0$, recalling that $\eta_0 = g_0 g_1 \dots g_{k-1}(\delta)$. Also note that $c \in \mathcal{F}_*(d)$ since $c = \partial d$. Thus, if c is $g(\delta)$ -fine, we have $R_i(c) = c$ and, in particular, $c = \partial R_{i+1}(d)$. As shown above, $R_{i+1}(d)$ is δ -fine, it is contained in a suitable annulus around p and it admits the required diameter bound. This proves part (1).

Part (2). Assume that $c \subseteq B(p, r_2)$ and $i \leq k$. Set $r_1 := \delta$.

Exactly as in part (1), the combination of Lemma 2.10 and the construction of a refining chain map R_* (this time defined on the chain complex $\mathcal{F}_*(d) \oplus \mathcal{F}_*(c)$) yields a discrete chain d' and a constant $K' \geq 1$ such that $\text{fine}(d') \leq \delta$ and $\text{diam}(d') \leq K' \sqrt{\text{diam}(c)}$ and $c - \partial d' \subseteq B(p, K'\delta)$. If $r_2 \leq \delta$, we simply take $d' = 0$.

Now, define d'' as the cone over $c - \partial d'$ from p (Definition 1.2). We have $\partial d'' = c - \partial d'$ and $d'' \subseteq B(p, K'\delta)$, hence $\text{fine}(d'') \leq \text{diam}(d'') \leq 2K'\delta$.

In conclusion, setting $d := d' + d'' \subseteq B(p, K'r_2)$, we have $\partial d = c$ and $\text{fine}(d) \leq 2K'\delta$. In addition, note that we can assume that $\text{diam}(c) \geq \delta^2$ (otherwise, we could have just taken d to be instead the cone over c from any of its vertices, since $\delta^2 \leq \delta$). Thus, we have $\text{diam}(d) \leq 2K'\delta + K' \sqrt{\text{diam}(c)} \leq 3K' \sqrt{\text{diam}(c)}$.

This proves part (2). \square

Finally, we can use Corollary 2.9 to remove from Lemma 2.11 the requirement that $\mathfrak{b}(c) \in \mathcal{H}_p$. That is, we can prove Proposition 2.1.

Proof of Proposition 2.1. Let g_1, K_1 be the function and constant provided by Lemma 2.11. Choose $D \geq 0$ such that, if $\mathfrak{b}(c)$ is D -deep in some horoball \mathcal{H}_q , then $c \subseteq B(q, r_q/K_1)$. Now, let g_2, K_2 be the function and constant provided by Corollary 2.9 in relation to the constant D . Finally, set $g := g_1 \wedge g_2$ and $K' := K_1 \vee K_2$.

Suppose that $c \in \overline{\mathcal{E}}_i(\mathcal{S})$ is $g(\delta)$ -fine and $i \leq k$. If $\mathfrak{b}(c)$ is not D -deep in any horoball, then $c = \partial d$ with $\text{fine}(d) \leq \delta$ and $\text{diam}(d) \leq K' \text{diam}(c)$, by Corollary 2.9. If $\mathfrak{b}(c)$ is D -deep in a horoball \mathcal{H}_q , then $c = \partial d$ with $\text{fine}(d) \leq \delta$ and $\text{diam}(d) \leq 2r_q \wedge K' \sqrt{\text{diam}(c)}$, by Lemma 2.11(2).

This proves part (0). Regarding parts (1) and (2), we are left to show that the filling d constructed in the previous paragraph lies in the required annulus/ball around the parabolic point p . For this, we can assume that $\mathfrak{b}(c) \notin \mathcal{H}_p$, otherwise a suitable filling of c is provided by Lemma 2.11.

Since $\mathfrak{b}(c) \notin \mathcal{H}_p$, there exists a constant C , depending only on δ_0, ϵ_0 , such that $\text{diam}(c) \leq C \rho(p, c)^2/r_p$. In addition, if $\mathfrak{b}(c) \in \mathcal{H}_q$ for some parabolic point $q \in \mathcal{S}$,

then we similarly have $r_q \leq C\rho(p, c)^2/r_p$. Note that $\rho(p, c) \leq r_2$ in part (2), while we can assume without loss of generality that $\rho(p, c) = r_1$ in part (1).

Now, recall that the filling d satisfies $\text{diam}(d) \leq 2r_q$ if $\mathfrak{b}(c)$ is deep in a horoball \mathcal{H}_q , and $\text{diam}(d) \leq K' \text{diam}(c)$ otherwise. Thus, provided that $r_2 \leq \frac{r_p}{2K'C} \wedge \frac{r_p}{4C}$, we have $\text{diam}(d) \leq r_2/2$ in part (2), and $\text{diam}(d) \leq r_1/2$ in part (1). This implies that d is contained in the ball $B(p, 2r_2)$ in part (2), and in the annulus $A(p; r_1/2, 2r_2)$ in part (1). \square

3. Filling discrete cycles in the Morse boundary

This section has two main goals, namely Proposition 3.6 and Proposition 3.8, which both hold for relatively hyperbolic groups with sphere boundary (see below for the precise assumptions). Proposition 3.6 roughly says that sufficiently low-dimensional cycles in a given stratum of the Morse boundary can be filled in a controlled stratum and with controlled fineness and diameter. While Proposition 3.6 will be used for our vanishing results for Čech cohomology, Proposition 3.8 will be used for our non-vanishing results. This proposition says that, for a given finite collection of parabolic points F in the Bowditch boundary \mathcal{S} , we can represent the entire homology of $\mathcal{S} - F$ by cycles of arbitrarily small fineness that are contained in some fixed stratum of the Morse boundary.

3.1. Notation

If B is a metric ball in a metric space, we denote by $\mathfrak{r}(B)$ its radius, and by λB the ball with the same centre as B and $\mathfrak{r}(\lambda B) = \lambda\mathfrak{r}(B)$. If \mathcal{B} is a family of metric balls, we write $\lambda\mathcal{B} = \{\lambda B \mid B \in \mathcal{B}\}$, and denote by $\mathfrak{r}_{\max}(\mathcal{B})$ and $\mathfrak{r}_{\min}(\mathcal{B})$, respectively, the supremum and the infimum of radii of balls in \mathcal{B} .

Recall that we use the notation $a \vee b := \max\{a, b\}$ and $a \wedge b := \min\{a, b\}$.

3.2. Assumptions

Let (G, \mathcal{P}) be a relatively hyperbolic pair where G is finitely generated and \mathcal{P} consists of finitely many finitely generated subgroups. Throughout most of Section 3, we will further assume as in Section 2 that the elements of \mathcal{P} are virtually nilpotent and that the Bowditch boundary \mathcal{S} is homeomorphic to a sphere S^{k+1} for some $k \geq 1$. The only exceptions will be Proposition 3.1 and Theorem 3.2, which we state under more general hypotheses.

Choose a visual metric ρ on \mathcal{S} . Let g and K be the corresponding function and constant provided by Proposition 2.1. Without loss of generality, we can assume that $g(\delta) \leq \delta$ for all $\delta > 0$. To simplify the proof of Lemma 3.5 below, we also assume that $K \geq 9$.

Each parabolic point $p \in \mathcal{S}$ comes with a characteristic radius r_p , as defined in Subsection 2.2. We fix a constant $M \geq 2K^2$ and consider the following collection of balls in \mathcal{S} :

$$\mathfrak{B} := \{B(p, r_p/M) \mid p \in \mathcal{S} \text{ parabolic}\}.$$

Up to enlarging M , we can and will assume (see e.g. [19, Lemma 3.2]) that, whenever two balls $B, B' \in \mathfrak{B}$ with $\mathfrak{r}(B) \leq \mathfrak{r}(B')$ have nonempty intersection, we have:

$$\mathfrak{r}(B) \leq \frac{1}{K^4} \cdot \mathfrak{r}(B'). \quad (*)$$

For every $n \geq 1$, we define the sets:

$$X_n := \mathcal{S} - \bigcup_{B \in \mathfrak{B}} \frac{1}{n}B.$$

We endow each X_n with the subspace topology induced by \mathcal{S} .

Under weak assumptions on the peripherals, the Morse boundary of a relatively hyperbolic group can be described as the direct limit of the strata X_n (we do not yet assume that the Bowditch boundary is a sphere or that the peripherals are virtually nilpotent).

Proposition 3.1. *Let (G, \mathcal{P}) be a relatively hyperbolic pair where G is finitely generated and \mathcal{P} consists of finitely many finitely generated groups with empty Morse boundary. Then the Morse boundary $\partial_* G$ is homeomorphic to $\varinjlim X_n$.*

Proof. This is proven for fundamental groups of finite volume hyperbolic manifolds in [7, Proposition 6.4], but the argument works more generally replacing \mathbb{H}^3 with the cusped space for G , up to minor modifications that we now discuss. All references below are from [7].

Lemma 6.1 uses CAT(0)-ness of the neutered space under consideration, but only to deduce that Morse geodesics are strongly contracting (quantitatively). This still holds since the neutered space is relatively hyperbolic, and by [22] a geodesic in a relatively hyperbolic space is strongly contracting if for every $C > 0$ there exists $B > 0$ such that the intersections of the geodesic with the C -neighbourhoods of the peripheral subsets have diameter bounded by B . Lemma 6.1 also uses that horospheres are flats, but only to the extent that they do not contain Morse rays. In Lemma 6.2 the visual metric used is $e^{-(\cdot)}$, while in general it is $e^{-\epsilon(\cdot)}$ for some small positive ϵ ; this does not affect the argument. The proof of part (2) of the same lemma uses convexity of horoballs to conclude that the intersection of the ray γ under consideration and the horoball is a subgeodesic. In a general cusped space, this intersection lies within uniformly bounded Hausdorff distance of a subgeodesic, so this does not cause issues. In the same part of the proof, the argument involving the midpoint of a geodesic in a horoball with endpoints on the horosphere can be performed equally well in a combinatorial horoball rather than in the half-space model for \mathbb{H}^3 . Finally, the argument in the proof of Proposition 6.4 again uses that horospheres are flats, but again to the same extent as the proof of Lemma 6.1 \square

As an aside, in view of the description of Morse boundaries given above, we point out that results of [19] yield path-connectedness (under weaker assumptions than the boundary being a sphere):

Theorem 3.2. *Let (G, \mathcal{P}) be a relatively hyperbolic pair where G is finitely generated and \mathcal{P} consists of finitely many finitely generated, virtually nilpotent subgroups. Suppose that G is one-ended and does not split over a subgroup conjugate into some peripheral group. Then ∂_*G is path-connected.*

Proof. The hypotheses of the theorem are the same as [19, Theorem 1.2], which is a special case of [19, Theorem 1.3]. In the proof of said theorem it is argued that [19, Corollary 7.4] applies, which is a result that allows one to connect pairs of points in the Bowditch boundary with arcs that avoid balls around parabolic points. By the description above, each of these arcs is contained in (a stratum of) the Morse boundary. \square

The next lemma allows us to approximate each stratum X_n by the complement of a finite collection of balls in \mathcal{S} .

Lemma 3.3. *For every $n \geq 1$ and $\epsilon > 0$ there exists $r = r(n, \epsilon) > 0$ such that X_n is ϵ -dense in the set:*

$$\mathcal{S} - \bigcup \left\{ \frac{1}{n}B \mid B \in \mathfrak{B}, \tau(B) \geq r \right\}.$$

Proof. Fix $n \geq 1$ and $\epsilon > 0$. Consider, for $m \geq 1$, the compact sets:

$$K_m := \mathcal{S} - \bigcup \left\{ \frac{1}{n}B \mid B \in \mathfrak{B}, \tau(B) \geq \frac{1}{m} \right\}.$$

Note that we have $\bigcap_{m \geq 1} K_m = X_n$. If there exists \overline{m} such that X_n is ϵ -dense in $K_{\overline{m}}$, we can set $r(n, \epsilon) := 1/\overline{m}$. Otherwise, for each $m \geq 1$ there exists a point $x_m \in K_m$ with $\rho(x_m, X_n) \geq \epsilon$. Passing to a subsequence, we have $x_m \rightarrow x$ with both $x \in \bigcap_{m \geq 1} K_m = X_n$ and $\rho(x, X_n) \geq \epsilon > 0$, a contradiction. \square

3.3. Detouring

Recall that, from now on, we assume that \mathcal{P} consists of virtually nilpotent subgroups and the Bowditch boundary \mathcal{S} is homeomorphic to a sphere S^{k+1} with $k \geq 1$.

In order to construct fillings in the Morse boundary $X = \varinjlim X_n$, our strategy will be to construct fillings in the Bowditch boundary \mathcal{S} using Proposition 2.1, and then to detour these around suitable collections of balls in \mathcal{S} , thus pushing the fillings into a stratum of X . The final result of this type will be Proposition 3.6.

To begin with, the following lemma explains how to detour chains in \mathcal{S} around collections of pairwise disjoint balls. Recall that g and K were introduced in Proposition 2.1.




Fig. 4. The decomposition of the discrete chain d into parts close to parabolic points plus the rest d_{out} .

Lemma 3.4. Consider a finite subset $\mathcal{B} \subseteq \frac{1}{N}\mathfrak{B}$, for some $N \geq 1$. Suppose that $2K\mathcal{B}$ is pairwise disjoint. Let $d \in \mathcal{C}_i(\mathcal{S})$ be a $g(\delta)$ -fine discrete chain, where $1 \leq i \leq k$ and $0 < \delta \leq \mathfrak{r}_{\min}(\mathcal{B})$. Suppose that $\text{supp}(\partial d)$ is disjoint from all balls in \mathcal{B} .

Then there exists a discrete chain d' in \mathcal{S} such that:

- (1) d' is δ -fine and $\partial d' = \partial d$;
- (2) $\text{supp}(d')$ is disjoint from all balls in $\frac{1}{K}\mathcal{B}$;
- (3) d and d' coincide outside the union of the balls in $2K\mathcal{B}$;
- (4) $\text{diam}(d') \leq \text{diam}(d) + 4K\sqrt{\mathfrak{r}_{\max}(\mathcal{B})}$;
- (5) if d is $g(g(\delta))$ -fine, then $d - d' = \partial e$ for a δ -fine discrete $(i + 1)$ -chain e contained in the union of the balls in $2K^2\mathcal{B}$.

Proof. We begin by noting that, for every ball $B \in \mathcal{B}$ centred at a parabolic point $p \in \mathcal{S}$, we have $2K\mathfrak{r}(B) \leq r_p/K$, since we have chosen $M \geq 2K^2$ in the definition of \mathfrak{B} . This will allow us to freely apply Proposition 2.1 later in the proof.

Write $d = d_{\text{out}} + \sum_{B \in \mathcal{B}} d_B$, where d_B is obtained by grouping all simplices in a reduced expression for d whose support is entirely contained in $2B$; this is illustrated in Fig. 4. Note that the balls in $2\mathcal{B}$ are pairwise disjoint, since $K \geq 1$.

Observe that $\text{supp}(d_{\text{out}})$ is disjoint from all balls $B \in \mathcal{B}$. Indeed, since d is $g(\delta)$ -fine and $g(\delta) \leq \delta \leq \mathfrak{r}_{\min}(\mathcal{B})$, every simplex σ appearing in d_{out} has diameter at most $\mathfrak{r}(B)$ and intersects the complement of $2B$, for each $B \in \mathcal{B}$.

Similarly, each $\text{supp}(\partial d_B)$ is contained in the annulus $2B - B$. To see this, consider a codimension-1 face τ of a simplex σ appearing in d , and suppose that $\text{supp}(\tau) \cap B \neq \emptyset$. Then $\text{supp}(\sigma) \subseteq 2B$, since d is $\mathfrak{r}(B)$ -fine, as observed above. It follows that all simplices of d containing τ appear in d_B . Since $\text{supp}(\partial d) \cap B = \emptyset$, the simplex τ does not appear in ∂d , hence it cannot appear in ∂d_B either.

Now, ∂d_B is a $g(\delta)$ -fine discrete, reduced $(i - 1)$ -chain in the annulus $2B - B$ and we have $i - 1 \leq k - 1$. Proposition 2.1(1) yields a δ -fine discrete i -chain d'_B in the annulus $2KB - \frac{1}{K}B$ with $\partial d'_B = \partial d_B$ and $\text{diam}(d'_B) \leq K\sqrt{\text{diam}(\partial d_B)}$. If $\partial d_B = 0$, we simply take $d'_B = 0$. Define $d' := d_{\text{out}} + \sum_{B \in \mathcal{B}} d'_B$.

Checking that d' satisfies (1) and (3) is straightforward. Moreover, recall that $\text{supp}(d_{\text{out}})$ is disjoint from all balls $B \in \mathcal{B}$. Since $\text{supp}(d'_B) \subseteq 2KB - \frac{1}{K}B$ and the balls in $2K\mathcal{B}$ are pairwise disjoint, we see that d' also satisfies (2).

Regarding property (4), note that every point of $\text{supp}(d'_B)$ is at distance at most $K\sqrt{\text{diam}(\partial d_B)}$ from a point of $\text{supp}(\partial d_B) \subseteq \text{supp}(d)$. Since we have both $\text{diam}(\partial d_B) \leq \text{diam}(d)$ and $\text{diam}(\partial d_B) \leq \text{diam}(2B) \leq 4\mathfrak{r}_{\max}(\mathcal{B})$, this yields the required inequality.

Finally, we prove property (5). If d is $g(g(\delta))$ -fine, Proposition 2.1(2) ensures that $d_B - d'_B = \partial e_B$ for a δ -fine discrete chain e_B contained in the ball $2K^2B$. Thus, it suffices to set $e := \sum_{B \in \mathcal{B}} e_B$. \square

We now discuss detouring around general collections of balls, which are allowed to intersect each other. For a function f and an integer $n \geq 0$, we denote by $f^{(n)}$ the n -fold composition of f (in particular, $f^{(0)} = \text{id}$ and $f^{(1)} = f$).

Lemma 3.5. *Consider a finite subset $\mathcal{B} \subseteq \frac{1}{N}\mathfrak{B}$ for some $N \geq 2K$. Define*

$$m := 1 + \lfloor \log_K(\mathfrak{r}_{\max}(\mathcal{B})/\mathfrak{r}_{\min}(\mathcal{B})) \rfloor.$$

Let $d \in \mathcal{C}_i(\mathcal{S})$ be a $g^{(m)}(\delta)$ -fine discrete chain with $1 \leq i \leq k$ and $0 < \delta \leq \mathfrak{r}_{\min}(\mathcal{B})$. Suppose that $\text{supp}(\partial d)$ is disjoint from all balls in \mathcal{B} .

Then there exists a discrete chain d' in \mathcal{S} such that:

- (1) d' is δ -fine and $\partial d' = \partial d$;
- (2) $\text{supp}(d')$ is disjoint from all balls in $\frac{1}{2K}\mathcal{B}$;
- (3) $\text{supp}(d')$ is contained in the $8K\mathfrak{r}_{\max}(\mathcal{B})$ -neighbourhood of $\text{supp}(d)$;
- (4) if $\partial d \neq 0$, then $\text{diam}(d') \leq \text{diam}(d) + 8K\sqrt{2K \text{diam}(d)}$;
- (5) if d is $g^{(m+1)}(\delta)$ -fine, then $d - d' = \partial e$ for a δ -fine discrete chain e contained in the $(4K^2 + 8K) \cdot \mathfrak{r}_{\max}(\mathcal{B})$ -neighbourhood of $\text{supp}(d)$.

Proof. We begin by partitioning $\mathcal{B} = \mathcal{B}_1 \sqcup \dots \sqcup \mathcal{B}_m$ so that a ball B lies in \mathcal{B}_i exactly when $\mathfrak{r}(B)$ lies in the half-open interval $[\mathfrak{r}_{\max}(\mathcal{B})/K^{i-1}, \mathfrak{r}_{\max}(\mathcal{B})/K^i)$. Note that, in view of Equation (*) and the fact that $N \geq 2K$, each set $2K\mathcal{B}_i$ consists of pairwise disjoint balls.

We then set $d_0 := d$ and iteratively define d_i as the chain obtained by applying Lemma 3.4 to d_{i-1} and the family of balls \mathcal{B}_i . Finally, we set $d' := d_m$.

Note that each application of Lemma 3.4 is allowed. Indeed, the chain d_i is $g^{(m-i)}(\delta)$ -fine and $g^{(m-i)}(\delta) \leq \delta \leq \mathfrak{r}_{\min}(\mathcal{B}) \leq \mathfrak{r}_{\min}(\mathcal{B}_{i+1})$. In addition, we have $\partial d_i = \partial d$, which is disjoint from all balls in \mathcal{B} .

We are left to check that the final chain $d' = d_m$ satisfies all required properties. The previous paragraph shows that d_m is δ -fine and $\partial d_m = \partial d$, proving (1).

Regarding (3), note that d_{i+1} and d_i only differ within a union of balls that intersect $\text{supp}(d_i)$ and lie in $2K\mathcal{B}_{i+1}$. Hence $\text{supp}(d_{i+1})$ is contained in a neighbourhood of

$\text{supp}(d_i)$ of radius at most $2 \cdot \tau_{\max}(2K \mathcal{B}_{i+1}) \leq 4K \cdot \tau_{\max}(\mathcal{B})/K^i$. Recalling that $K \geq 2$, we deduce that $\text{supp}(d_m)$ is contained in a neighbourhood of $\text{supp}(d)$ of radius at most:

$$4K \cdot \tau_{\max}(\mathcal{B}) \cdot \sum_{i \geq 0} \frac{1}{K^i} \leq 8K \cdot \tau_{\max}(\mathcal{B}).$$

Now, we prove (2). Consider a ball $B \in \mathcal{B}$ and say it lies in \mathcal{B}_i . By Lemma 3.4, the chain d_i is supported outside $\frac{1}{K}B$. We need to check that this property is not affected too much by subsequent detouring.

First, observe that the chain d_{i+3} is also supported outside $\frac{1}{K}B$. Indeed, setting for a moment $\mathcal{B}' := \mathcal{B}_{i+1} \cup \mathcal{B}_{i+2} \cup \mathcal{B}_{i+3}$, the chains d_i and d_{i+3} only differ within the balls in $2K\mathcal{B}'$. And, for every $B' \in \mathcal{B}'$, the ball $2KB'$ is disjoint from $\frac{1}{K}B$. Indeed, if this were not the case, the balls $2KB'$ and $2KB$ would also intersect. These two balls are rescalings of elements of \mathfrak{B} by the same factor ≥ 1 , because $N \geq 2K$ and $B, B' \in \frac{1}{N}\mathfrak{B}$. Thus, Equation (*) would imply that $\tau(B') \leq \tau(B)/K^4$, contradicting the fact that, since $B' \in \mathcal{B}'$, we have $\tau(B') > \tau(B)/K^4$.

Now, recalling that $K \geq 9$, we obtain that $\text{supp}(d_m)$ is contained in a neighbourhood of $\text{supp}(d_{i+3})$ of radius at most:

$$\begin{aligned} 4K \cdot \sum_{j \geq i+4} \tau_{\max}(\mathcal{B}_j) &\leq 4K \cdot \tau_{\max}(\mathcal{B}_{i+1}) \cdot \sum_{j \geq 3} \frac{1}{K^j} \\ &\leq 4K \cdot \tau_{\max}(\mathcal{B}_{i+1}) \cdot \frac{1}{8K^2} < \frac{1}{2K} \cdot \tau(B). \end{aligned}$$

In conclusion, $\text{supp}(d_m)$ is disjoint from $\frac{1}{2K}B$, for every $B \in \mathcal{B}$.

Let us prove (4). Let $1 \leq k \leq m$ be the smallest integer such that $d_k \neq d_{k-1}$. In particular, we have $d = d_{k-1}$ and there exists a ball $B \in \mathcal{B}_k$ such that $\text{supp}(d) \cap \frac{1}{K}B \neq \emptyset$ (otherwise we could have simply taken $d_k = d_{k-1}$, rather than applying Lemma 3.4). Since $\text{supp}(\partial d)$ is nonempty and disjoint from B , it follows that $\text{diam}(d) \geq \frac{1}{2}\tau(B) \geq \frac{1}{2K}\tau_{\max}(\mathcal{B}_k)$.

Now, by part (4) of Lemma 3.4 and recalling that $\sqrt{K} \geq 3$, we obtain:

$$\begin{aligned} \text{diam}(d_m) &\leq \text{diam}(d) + 4K \sum_{j \geq k} \sqrt{\tau_{\max}(\mathcal{B}_j)} \\ &\leq \text{diam}(d) + 4K \sqrt{\tau_{\max}(\mathcal{B}_k)} \sum_{j \geq 0} K^{-j/2} \leq \text{diam}(d) + 8K \sqrt{2K \text{diam}(d)}. \end{aligned}$$

Finally, we prove (5). Part (5) of Lemma 3.4 gives δ -fine chains e_1, \dots, e_k with $d_i - d_{i-1} = \partial e_i$ and $\text{supp}(e_i)$ contained in the union of the balls in $2K^2\mathcal{B}_i$ that intersect $\text{supp}(d_{i-1})$. In particular, $\text{supp}(e_i)$ is contained in a neighbourhood of $\text{supp}(d_{i-1})$ of radius at most $4K^2 \cdot \tau_{\max}(\mathcal{B})$, and hence in a neighbourhood of $\text{supp}(d)$ of radius at most $(4K^2 + 8K) \cdot \tau_{\max}(\mathcal{B})$, by the proof of property (3). We conclude by setting $e := e_1 + \dots + e_k$. \square

The detouring construction developed in Lemma 3.5 allows us to fill discrete cycles in a given stratum of X while remaining in a controlled larger stratum of X . This is the content of the next proposition, which is the main ingredient for the vanishing part of Theorems A and B.

Recall that $\overline{\mathcal{C}}_*(\mathcal{S})$ denotes the chain complex of reduced discrete chains in \mathcal{S} , where 0-chains are required to have zero coefficient sum.

Proposition 3.6. *There exist functions $f_1(n) \geq n$, $0 < g_1(\delta, n) \leq \delta$ and $\mathfrak{h}(r)$ (going to 0 for $r \rightarrow 0$) such that the following holds for all $\delta > 0$, $n \geq 1$, $0 \leq i < k$.*

If $c \in \overline{\mathcal{C}}_i(\mathcal{S})$ is a $g_1(\delta, n)$ -fine discrete cycle supported in X_n , then there exists a δ -fine chain d supported in $X_{f_1(n)}$ with $\partial d = c$ and $\text{diam}(d) \leq \mathfrak{h}(\text{diam}(c))$.

Proof. Define $f_1(n) := 2K(n \vee 2K)$ and set $N := f_1(n)$ for simplicity. Define $\bar{r} := \frac{1}{N} \cdot r(N, \delta/3)$, where $r(\cdot, \cdot)$ is the function in Lemma 3.3, and consider:

$$\mathcal{B} := \left\{ B \in \frac{1}{N}\mathfrak{B} \mid \mathfrak{r}(B) \geq \bar{r} \right\}.$$

Set $m = m(\delta, n) := 1 + \lfloor \log_K(\mathfrak{r}_{\max}(\mathcal{B})/\mathfrak{r}_{\min}(\mathcal{B})) \rfloor$ and $g_1(\delta, n) := g^{(m+1)}(\bar{\delta})$, where $\bar{\delta} := \delta/3 \wedge 2K\bar{r}$. Finally, define $\mathfrak{h}(r) := r + Kr^{1/2} + 16K^2r^{1/4}$.

Now, let c be a $g_1(\delta, n)$ -fine, reduced, discrete i -cycle with $i < k$ and $\text{supp}(c) \subseteq X_n$. Without loss of generality, we can assume that $\delta \leq \text{diam}(c)$. Otherwise, we could simply define d as the cone over c from any of its vertices (Definition 1.2), which would satisfy $\text{supp}(d) \subseteq X_n$, $\text{fine}(d) = \text{diam}(d) = \text{diam}(c)$ and $c = \partial d$.

By Proposition 2.1(0), there exists a $g^{(m)}(\bar{\delta})$ -fine discrete chain d' in \mathcal{S} with $\partial d' = c$ and $\text{diam}(d') \leq K\sqrt{\text{diam}(c)}$. Note that $c = \partial d'$ is disjoint from all balls in $2K\mathcal{B}$, since $\text{supp}(c) \subseteq X_n$ and every ball of $2K\mathcal{B}$ is contained in one of $\frac{1}{n}\mathfrak{B}$.

Applying Lemma 3.5 to d' and the family of balls $2K\mathcal{B}$, we obtain a $\bar{\delta}$ -fine (hence $\delta/3$ -fine) discrete chain d'' in \mathcal{S} such that $\partial d'' = \partial d'$ and $\text{supp}(d'')$ is disjoint from all balls in \mathcal{B} . In addition:

$$\text{diam}(d'') \leq \text{diam}(d') + 8K\sqrt{2K\text{diam}(d')} \leq K\sqrt{\text{diam}(c)} + 16K^2(\text{diam}(c))^{1/4}.$$

Finally, by Lemma 3.3 and our choice of \mathcal{B} , every point of $\text{supp}(d'')$ is $\delta/3$ -close to a point of X_N . Replacing every point of $\text{supp}(d'')$ with a closest point in X_N , we obtain a δ -fine chain d with $\partial d = \partial d'' = c$ and $\text{supp}(d) \subseteq X_N$. Finally, since $\delta \leq \text{diam}(c)$, we have $\text{diam}(d) \leq 2\delta/3 + \text{diam}(d'') \leq \mathfrak{h}(\text{diam}(c))$, concluding the proof. \square

The next proposition is required for the non-vanishing part of Theorems A and B. This is the only place in Section 3 where we explicitly use the fact that \mathcal{S} is homeomorphic to a sphere. In the rest of the section, we have only used this assumption indirectly, in the form of Proposition 2.1.

Before proving Proposition 3.8, we need the following observation.

Remark 3.7 (*Straight chains and straightening*). Let (S, d_\circ) be a constant-curvature sphere. If $D \subseteq S$ is a finite subset with $\text{diam}(D) \leq \text{diam}(S)/2$, then there is a unique singular simplex σ_D that has vertex set D and maps every affine line in the standard simplex to a geodesic in S parametrised proportionally to arc-length.

We refer to simplices obtained in this way as *straight simplices*. A singular chain in (S, d_\circ) is *straight* if all its simplices are straight.

If c is a $\text{diam}(S)/2$ -fine discrete chain in (S, d_\circ) , we define its *straightening* \widehat{c} . This is the singular chain obtained by replacing every discrete simplex of c with the straight singular simplex with the same (ordered) vertex set. Note that $\text{disc}(\widehat{c}) = c$.

If c is δ -fine in (S, d_\circ) , then $\text{supp}(\widehat{c})$ is contained in the δ -neighbourhood of $\text{supp}(c)$. Indeed, all metric balls in (S, d_\circ) with radius $\leq \text{diam}(S)/2$ are convex.

Proposition 3.8. *For every finite (nonempty) set of parabolic points $F \subseteq S$, there exists a stratum X_N with the following property.*

For every $\delta > 0$, every class in $H_k(S - F, \mathbb{Z})$ is represented by a singular cycle c such that $\text{disc}(c)$ is δ -fine and contained in X_N .

Proof. Let d_\circ be a constant-curvature (i.e. round) metric on S . We continue denoting by ρ our chosen visual metric on S . When speaking of neighbourhoods, radii and fineness, we will always refer to ρ unless explicitly stated otherwise.

Note that ρ and d_\circ induce the same topology on S , which is compact. Thus, there exists a weakly increasing function ξ such that $\xi(t) \rightarrow 0$ for $t \rightarrow 0$ and, for all $x, y \in S$, we have:

$$d_\circ(x, y) \leq \xi(\rho(x, y)).$$

Let $3D > 0$ be the minimum distance $d_\circ(x, y)$ for distinct points $x, y \in F$. Let Σ be the union of the D -spheres around the points of F in the metric d_\circ . Note that, since F is nonempty, the homology of $S - F$ is entirely supported on Σ .

Up to shrinking δ , we can assume that $\xi(\delta) < D/3$. We then choose the integer N so that:

$$\xi\left(\frac{\delta}{3} + (4K^2 + 8K) \cdot \frac{\text{diam}(S, \rho)}{N/2K}\right) \leq \frac{D}{3}.$$

Enlarging N if necessary, we further assume that $N \geq 4K^2$.

As in the proof of Proposition 3.6, set $\bar{r} := \frac{1}{N} \cdot r(N, \delta/3)$, where $r(\cdot, \cdot)$ is the function in Lemma 3.3, and define:

$$\mathcal{B} := \left\{ B \in \frac{1}{N}\mathfrak{B} \mid \mathfrak{r}(B) \geq \bar{r} \right\}.$$

Set $m := 1 + \lfloor \log_K(\mathfrak{r}_{\max}(\mathcal{B})/\mathfrak{r}_{\min}(\mathcal{B})) \rfloor$ and $\delta' := g^{(m+1)}(\delta/3 \wedge 2K\bar{r})$.

Now consider any nonzero class $\alpha \in H_k(\mathcal{S} - F, \mathbb{Z})$. Represent α by a straight, singular k -cycle \widehat{c} such that $c := \text{disc}(\widehat{c})$ is supported on Σ and δ' -fine (with respect to the visual metric ρ). This is possible in view of Remark 3.7.

Applying Lemma 3.5 to c and the collection of balls $2K\mathcal{B}$, we obtain a $\delta/3$ -fine discrete cycle c' supported outside all balls in \mathcal{B} . In addition, using part (5) of Lemma 3.5, there exists a $\delta/3$ -fine discrete chain e with $\partial e = c - c'$ and $\text{supp}(e)$ contained in the $(4K^2 + 8K) \cdot r_{\max}(2K\mathcal{B})$ -neighbourhood of $\text{supp}(c) \subseteq \Sigma$. Note that:

$$r_{\max}(2K\mathcal{B}) \leq \frac{\text{diam}(\mathcal{S}, \rho)}{N/2K}.$$

By Lemma 3.3 and our choice of \mathcal{B} , every point of $\text{supp}(c')$ is $\delta/3$ -close to a point of X_N . Thus, we can modify c' to a δ -fine discrete cycle c'' , and e to a δ -fine discrete chain e' , so that $\partial e' = c - c''$ and that $\text{supp}(c'') \subseteq X_N$. Note that $\text{supp}(e')$ is contained in the $\delta/3$ -neighbourhood of $\text{supp}(e)$.

Since $\xi(\delta) < D/3$, the chain e' is $D/3$ -fine with respect to the metric d_o . In addition, our choice of N guarantees that $\text{supp}(e')$ is contained in the $D/3$ -neighbourhood of $\text{supp}(c) \subseteq \Sigma$ with respect to the metric d_o . Remark 3.7 then shows that the straightening \widehat{e}' of e' is supported in the $2D/3$ -neighbourhood of Σ , which is disjoint from F .

Finally, denoting by \widetilde{c}'' the straightening of c'' , and recalling that \widehat{c} was straight to begin with, we obtain $\partial \widehat{e}' = \widehat{c} - \widetilde{c}''$. Since $\text{supp}(\widehat{e}') \cap F = \emptyset$, this shows that $\alpha = [\widehat{c}] = [\widetilde{c}'']$ in $H_k(\mathcal{S} - F, \mathbb{Z})$. Since $\text{disc}(\widetilde{c}'') = c''$, which is δ -fine and supported on X_N , this concludes the proof. \square

We will also need the following relation between singular and discrete cycles, for which it is convenient to use the constant-curvature metric d_o .

Lemma 3.9. *Let F be a finite subset of (\mathcal{S}, d_o) . Every open cover \mathcal{U} of $\mathcal{S} - F$ containing only convex open sets of diameter $\leq \text{diam}(\mathcal{S})/2$ satisfies the following property.*

If c is a straight \mathcal{U} -fine singular cycle in $\mathcal{S} - F$ and $\text{disc}(c)$ is the boundary of a \mathcal{U} -fine discrete chain, then $[c] = 0$ in the singular homology of $\mathcal{S} - F$.

Proof. Suppose that $\text{disc}(c) = \partial d$, for a \mathcal{U} -fine discrete chain d . Let \widehat{c} and \widehat{d} be the straightenings of $\text{disc}(c)$ and d , as defined in Remark 3.7. By convexity of the elements of \mathcal{U} , the singular chains \widehat{c} and \widehat{d} are again \mathcal{U} -fine (and contained in $\mathcal{S} - F$). We have $\partial \widehat{d} = \widehat{c}$ by construction and, moreover, the singular cycle $c - \widehat{c}$ is also the boundary of a singular chain that is \mathcal{U} -fine (hence contained in $\mathcal{S} - F$), again by convexity of \mathcal{U} . In conclusion, c is the boundary of singular chain contained in $\mathcal{S} - F$. \square

Remark 3.10. It is easy to see that, for every finite subset $F \subseteq \mathcal{S}$, the complement $\mathcal{S} - F$ admits a convex cover as in the lemma.

4. Vanishing of Čech cohomology

Throughout Section 4, we again consider a relatively hyperbolic group G , its Morse boundary $X := \partial_*G$ and the filtration $X = \varinjlim X_n$. Notation and assumptions are exactly the same as in Section 3 (in particular, see Subsection 3.2 for the precise assumptions on G).

We fix a visual metric ρ on the Bowditch boundary and equip X and all strata X_n with its restriction. We stress once more that ρ does not induce the topology of X , though it does induce that of each X_n . When speaking e.g. of fineness of chains below, this will always be meant with respect to this metric.

Our only tool in Section 4 is Proposition 3.6, so all results in this section hold more generally for any countable direct limit $X = \varinjlim X_n$ of compact metric spaces X_n satisfying Proposition 3.6.

The main result for this section is the vanishing of Čech cohomology in low degrees, proving the first part of Theorem B and, consequently, also the first part of Theorem A.

Theorem 4.1. *Let \mathcal{R} be a principal ideal domain. Then, for $X = \partial_*G$, we have $\check{H}^i(X, \mathcal{R}) = \{0\}$ for $1 \leq i < k$.*

All chains and simplices in the following discussion are assumed to be discrete without explicit mention.

4.1. Subdivision results

Let $\mathcal{C}_i(X_n)$ be the free \mathbb{Z} -module generated by (discrete) i -simplices with support contained in X_n . For $\delta > 0$, let $\mathcal{C}_i(X_n; \delta)$ be the submodule of $\mathcal{C}_i(X_n)$ generated by δ -fine simplices. For an open cover \mathcal{O} of X , we denote by $\mathcal{C}_i(X_n, \mathcal{O})$ the submodule of $\mathcal{C}_i(X_n)$ formed by \mathcal{O} -fine chains, and by $\mathcal{C}_i(X_n, \mathcal{O}; \delta)$ its intersection with $\mathcal{C}_i(X_n; \delta)$.

The main aim of this subsection is the following result, proved at the end.

Proposition 4.2 (*Arbitrarily fine subdivisions*). *For every open cover \mathcal{O} of X , there exist a refinement \mathcal{O}' and a function $f_2(n) \geq n$ such that the following hold.*

- (1) *For every $\delta > 0$, there is a chain map $\phi^\delta: \mathcal{C}_*(X_n, \mathcal{O}') \rightarrow \mathcal{C}_*(X_{f_2(n)}, \mathcal{O}; \delta)$ defined in degrees $\leq k$.*
- (2) *The composition $\iota\phi^\delta: \mathcal{C}_*(X_n, \mathcal{O}') \rightarrow \mathcal{C}_*(X_{f_2(n)}, \mathcal{O})$ is chain-homotopic to ι , where ι denotes all standard inclusions.*

In particular, part (2) of Proposition 4.2 implies that, for every cycle $c \in \mathcal{C}_i(X_n, \mathcal{O}')$, the cycles c and $\phi^\delta(c)$ are \mathcal{O} -homologous, in the sense that there exists an \mathcal{O} -fine chain d such that $\partial d = c - \phi^\delta(c)$.

As a first step towards the proof of Proposition 4.2, we need Proposition 4.5, which is an analogue disregarding covers. The proof of Proposition 4.5 is almost exactly the same, but it relies on Proposition 3.6 rather than Proposition 4.7 below.

Recall that the face complex $\mathcal{F}_*(d)$ of a discrete chain d was introduced in Definition 1.4.

Definition 4.3. Consider $r > 0$ and a subset $Y \subseteq X$. A chain c is (r, Y) -related to a chain c' if there exists a chain map $T : \mathcal{F}_*(c) \rightarrow \mathcal{C}_*(Y)$ such that $T(c) = c'$ and, for each simplex $\sigma \in \mathcal{F}_*(c)$, the support of the chain $T(\sigma)$ is contained in the r -neighbourhood of $\text{supp}(\sigma)$.

Remark 4.4 (*Projecting discrete chains*). Let c be an i -chain in X . If $\text{supp}(c)$ is contained in the ϵ -neighbourhood of a subspace $Y \subseteq X$, for some $\epsilon > 0$, we can project c to a chain in Y as follows.

For every point $x \in \text{supp}(c)$, choose a point $x' \in Y$ with $\rho(x, x') < \epsilon$. Write $c = \sum_j \sigma_j$, where each σ_j is an i -simplex. For each $\sigma_j = [x_0, \dots, x_i]$, define the new simplex $\sigma'_j = [x'_0, \dots, x'_i]$ and set $c' := \sum_j \sigma'_j$. We will refer to c' as an ϵ -projection of c to Y . Note that, by construction, c is (ϵ, Y) -related to c' and $\text{fine}(c') \leq \text{fine}(c) + 2\epsilon$.

Proposition 4.5. *There exist functions $f_3(n) \geq n$ and $H(r), H'(r)$ (both going to 0 for $r \rightarrow 0$) such that the following holds.*

- (1) For all $\delta > 0$ and $n \geq 1$, there is a chain map $\phi^{\delta, n} : \mathcal{C}_*(X_n) \rightarrow \mathcal{C}_*(X_{f_3(n)}; \delta)$ defined in degrees $\leq k$.
- (2) If $c \in \mathcal{C}_i(X_n; r)$ for some $r > 0$ and $0 \leq i \leq k$, then $\phi_i^{\delta, n} c$ is contained in the $H(r)$ -neighbourhood of $\text{supp}(c)$.
- (3) If a cycle $c \in \mathcal{C}_i(X_{n+1}; r)$ with $0 \leq i < k$ is (r, X_n) -related to a cycle $c' \in \mathcal{C}_i(X_n; r)$, then $\phi_i^{\delta, n+1} c - \phi_i^{\delta, n} c' = \partial B$ for a δ -fine chain B contained in the $H'(r)$ -neighbourhood of $\text{supp}(c)$ within $X_{f_3(n)}$.

Proof. Fix n and δ for the duration of the proof. We will write ϕ_i for $\phi_i^{\delta, n}$, that is, for the restriction of $\phi^{\delta, n}$ to i -chains.

Let f_1, g_1, \mathfrak{h} be the functions appearing in Proposition 3.6. Without loss of generality, we can assume that g_1 is weakly decreasing in its second argument and that f_1 is weakly increasing.

Define $n_i := f_1^{(i)}(n)$ and $m_i := f_1^{(i)}(n + 1)$, recalling that $f_1^{(i)}$ denotes the i -fold composition of f_1 . Define $\delta_{k+1} := \delta$ and inductively $\delta_{i-1} := g_1(\delta_i, m_{i-1})$; the fact that we are using m_{i-1} rather than n_{i-1} will only matter in part (3). Define $r_0 = 0$ and inductively $r_{i+1} := \mathfrak{h}(r + 2r_i) + r_i$; each r_i is a function of r . Note that n_i, δ_i and r_i all increase with i .

We now proceed by induction on $0 \leq i \leq k$ and construct maps

$$\phi_i : \mathcal{C}_i(X_n) \longrightarrow \mathcal{C}_i(X_{n_i}; \delta_i)$$

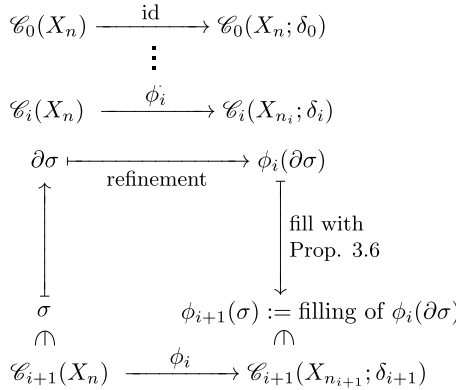


Fig. 5. Illustrative diagram of the proof of part (1) of Proposition 4.5, showing the inductive construction of the maps ϕ_i .

so that $\partial\phi_i = \phi_{i-1}\partial$. In addition, if $\sigma \in \mathcal{C}_i(X_n; r)$ is a simplex, then $\text{supp}(\phi_i\sigma)$ will be contained in the r_i -neighbourhood of $\text{supp}(\sigma)$. See Fig. 5.

In the base step, we simply take ϕ_0 to be the identity. In the inductive step, suppose that, for some $0 \leq i \leq k - 1$, we have defined all maps up to ϕ_i so that they satisfy the required properties. We want to define ϕ_{i+1} .

If $\sigma \in \mathcal{C}_{i+1}(X_n)$ is a simplex, then $\phi_i\partial\sigma$ is a cycle in $\mathcal{C}_i(X_{n_i}; \delta_i)$, by the inductive hypothesis. Note that $\delta_i \leq g_1(\delta_{i+1}, n_i)$, since $m_i \geq n_i$. In addition, if $\sigma \in \mathcal{C}_{i+1}(X_n; r)$, then $\text{supp}(\phi_i\partial\sigma)$ is contained in the r_i -neighbourhood of $\text{supp}(\sigma)$. In this case, the diameter of $\text{supp}(\phi_i\partial\sigma)$ is at most $r + 2r_i$.

By Proposition 3.6, the cycle $\phi_i\partial\sigma$ is the boundary of a chain in $\mathcal{C}_{i+1}(X_{n_{i+1}}; \delta_{i+1})$ with diameter at most $\mathfrak{h}(r + 2r_i)$. We define $\phi_{i+1}\sigma$ as this chain.

Note that $\text{supp}(\phi_{i+1}\sigma) \supseteq \text{supp}(\partial\phi_{i+1}\sigma) = \text{supp}(\phi_i\partial\sigma)$. We have already observed that the latter is contained in the r_i -neighbourhood of $\text{supp}(\sigma)$. It follows that $\text{supp}(\phi_{i+1}\sigma)$ is contained in the neighbourhood of $\text{supp}(\sigma)$ of radius $r_i + \mathfrak{h}(r + 2r_i) = r_{i+1}$.

Defining $f_3(\cdot)$ as the function $n \mapsto n_k$ and $H(\cdot)$ as the function $r \mapsto r_k$, this proves parts (1) and (2). In fact, in view of the proof of part (3) below, we should define $f_3(\cdot)$ as the (larger) function $n \mapsto m_k$.

We now prove part (3). Suppose that we have chain subcomplexes (each generated by a set of simplices)

$$E_* \subseteq \mathcal{C}_*(X_{n_{i+1}}; r), \quad F_* \subseteq \mathcal{C}_*(X_n; r),$$

and a chain map $T_*: E_* \rightarrow F_*$ such that, for every simplex $\sigma \in E_*$, the support of $T(\sigma)$ is contained in the r -neighbourhood of $\text{supp}(\sigma)$.

Define $R_{-1} := r + H(r)$, then inductively $R_{i+1} := \mathfrak{h}(r + 2R_i) + R_i$. Proceeding by induction on $0 \leq i < k$, we will construct a chain homotopy

$$H_i: E_i \rightarrow \mathcal{C}_{i+1}(X_{m_{i+1}}; \delta_{i+1})$$

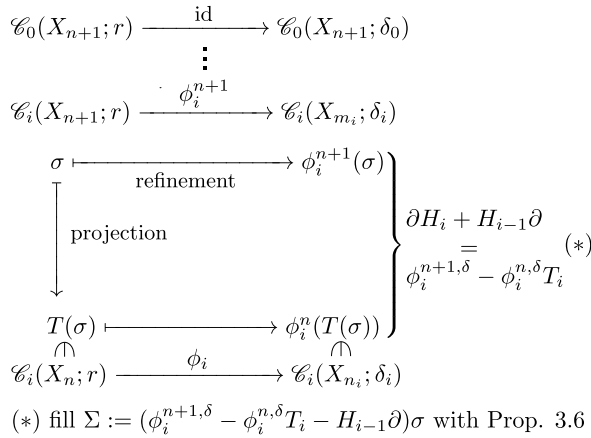


Fig. 6. Illustrative diagram of the proof of part (3) of Proposition 4.5.

so that $\partial H_i + H_{i-1} \partial = \phi_i^{n+1, \delta} - \phi_i^{n, \delta} T_i$. In addition, for every simplex $\sigma \in E_i$, the support of $H_i \sigma$ will be contained in the R_i -neighbourhood of $\text{supp}(\sigma)$. See Fig. 6.

The base step “ $i = -1$ ” is clear, taking $H_{-1} = 0$. In the inductive step, suppose that, for some $0 \leq i < k$, we have defined all maps up to H_{i-1} so that they satisfy the required properties.

In order to define H_i , consider a simplex $\sigma \in E_i$. Note that

$$\Sigma := (\phi_i^{n+1, \delta} - \phi_i^{n, \delta} T_i - H_{i-1} \partial) \sigma$$

is a cycle in $\mathcal{C}_i(X_{m_i}; \delta_i)$. (Since the function $g_1(\cdot, \cdot)$ is decreasing in its second argument, the “ δ_i ” appearing in the definition of $\phi_i^{\delta, n+1}$ is smaller than the one appearing in the definition of $\phi_i^{\delta, n}$, which is the one used here. Thus, Σ is indeed δ_i -fine.)

In addition, $\text{supp}(\Sigma)$ is contained in the neighbourhood of $\text{supp}(\sigma)$ of radius R_{i-1} , since $R_{i-1} \geq r + H(r)$ by construction. In particular, the diameter of $\text{supp}(\Sigma)$ is at most $r + 2R_{i-1}$.

By Proposition 3.6, we can fill Σ by a chain in $\mathcal{C}_{i+1}(X_{m_{i+1}}; \delta_{i+1})$ with diameter at most $\mathfrak{h}(r + 2R_{i-1})$. (Here it finally matters that we used m_{i-1} in the definition of δ_i at the beginning of the whole proof.) We define $H_i \sigma$ as this chain which fills Σ .

It is clear that $\text{supp}(H_i \sigma)$ is contained in the neighbourhood of $\text{supp}(\sigma)$ of radius:

$$\mathfrak{h}(r + 2R_{i-1}) + R_{i-1} = R_i.$$

This completes the construction of the chain homotopy H_* . If c, c' are as in the statement of part (3), we have $\phi_i^{\delta, n+1} c - \phi_i^{\delta, n} c' = \partial H_i c$. It is clear that $H_i c$ is δ -fine and that its support is contained in the neighbourhood of $\text{supp}(c)$ of radius R_i . We can thus define $H'(\cdot)$ as the function $r \mapsto R_k$.

This concludes the proof of the proposition. \square

Proposition 4.5 allows us to arbitrarily “refine” chains, but it completely disregards open covers of X . Thus, as a second step towards Proposition 4.2, we will need to be able to fill cycles while remaining within a given open set. This is the content of Proposition 4.7 below, after which we will finally prove Proposition 4.2.

Before stating Proposition 4.7, we record the following observation.

Lemma 4.6. *Let a cycle $c \in \mathcal{C}_*(X_n; \delta)$ be (δ, X_n) -related to a cycle $c' \in \mathcal{C}_*(X_n)$, for some $\delta > 0$ and $n \geq 1$. Then there exists a chain $d \in \mathcal{C}_*(X_n; 3\delta)$ such that $\partial d = c - c'$ and $\text{supp}(d)$ is contained in the δ -neighbourhood of $\text{supp}(c)$.*

Proof. The proof is similar to that of part (3) of Proposition 4.5, but simpler.

Let $T: \mathcal{F}_*(c) \rightarrow \mathcal{C}_*(X_n)$ be the chain map witnessing the fact that c and c' are (δ, X_n) -related. For each simplex $\sigma \in \mathcal{F}_*(c)$, denote by $\mathcal{T}(\sigma)$ the union of the sets $\text{supp}(T\tau)$ where τ ranges over σ and all its lower-dimensional faces. Thus, $\mathcal{T}(\sigma)$ contains $\text{supp}(T\sigma)$, but it can be larger in general. Nevertheless, $\mathcal{T}(\sigma)$ is contained in the δ -neighbourhood of $\text{supp}(\sigma)$ in X_n . In addition, recalling that c is δ -fine, we have $\text{diam}(\text{supp}(\sigma) \cup \mathcal{T}(\sigma)) \leq 3\delta$ for every simplex $\sigma \in \mathcal{F}_*(c)$.

Now, for each $i \geq 0$, we inductively construct a chain homotopy

$$H_i: \mathcal{F}_i(c) \longrightarrow \mathcal{C}_{i+1}(X_n; 3\delta)$$

satisfying $\partial H_i + H_{i-1}\partial = \text{id} - T_i$ and $\text{supp}(H_i\sigma) \subseteq \text{supp}(\sigma) \cup \mathcal{T}(\sigma)$ for every simplex $\sigma \in \mathcal{F}_i(c)$.

In the base step “ $i = -1$ ”, we simply set $H_{-1} = 0$. In the inductive step, suppose that, for some $i \geq 0$, we have defined all maps up to H_{i-1} so that they satisfy the required properties.

In order to define H_i , consider a simplex $\sigma \in \mathcal{F}_i(c)$ and observe that $\Sigma := (\text{id} - T_i - H_{i-1}\partial)\sigma$ is a cycle with support contained in $\text{supp}(\sigma) \cup \mathcal{T}(\sigma) \subseteq X_n$. Then we simply define $H_i\sigma$ as the cone over Σ from any of its vertices (Definition 1.2). We obtain $\partial H_i\sigma = \Sigma$ and $\text{supp}(H_i\sigma) \subseteq \text{supp}(\Sigma) \subseteq \text{supp}(\sigma) \cup \mathcal{T}(\sigma)$. As observed above, this guarantees that $H_i\sigma$ has diameter at most 3δ and that it is supported in the δ -neighbourhood of $\text{supp}(\sigma)$.

Finally, set $d := Hc$. Since c is a cycle, we have $\partial d = c - Tc = c - c'$. In addition, d is 3δ -fine and it is supported in the δ -neighbourhood of $\text{supp}(c)$. \square

Recall that $\overline{\mathcal{C}}_*(\mathcal{S})$ denotes the chain complex of reduced discrete chains in \mathcal{S} , where 0-chains are required to have zero coefficient sum. Also recall the notion of \mathcal{O} -tiny chain from Definition 1.1.

Proposition 4.7. *For every open cover \mathcal{O} of $X = \varinjlim X_n$, there exist a refinement \mathcal{O}' and functions $f_4(n) \geq n$ and $0 < g_2(\delta, n) \leq \delta$ such that the following holds.*

For $0 \leq i < k$, let $c \in \overline{\mathcal{C}}_i(X_n)$ be an \mathcal{O}' -tiny, $g_2(\delta, n)$ -fine cycle. Then $c = \partial d$ for an \mathcal{O} -tiny, δ -fine chain d supported within $X_{f_4(n)}$.

The proposition is proved right after the following lemma.

Lemma 4.8. *For every open cover \mathcal{O} of X , there exist functions $f_4(n) \geq n$ and $0 < g_3(x, \delta, n) \leq \delta$ such that the following holds for all $x \in X$, $\delta > 0$ and $n \geq 1$.*

There exists an \mathcal{O} -tiny open set $U_x \ni x$ with the following property. If c is a $g_3(x, \delta, n)$ -fine, reduced i -cycle with $0 \leq i < k$ and $c \subseteq X_n \cap U_x$, then $c = \partial d$ for an \mathcal{O} -tiny δ -fine chain d with $d \subseteq X_{f_4(n)}$.

Proof. We will construct the open sets U_x using the inductive Construction 1.14 (where the parameters are also determined inductively). Fix the open cover \mathcal{O} and the point $x \in X$. Let $\ell \geq 1$ be the integer such that $x \in X_\ell - X_{\ell-1}$. Choose an element $O_x \in \mathcal{O}$ with $x \in O_x$, which will be fixed throughout the proof.

Let f_1, \mathfrak{h} be the functions provided by Proposition 3.6, and f_3, H, H' those provided by Proposition 4.5. Recall that these functions are all bounded below by the identity and, without loss of generality, weakly increasing. Define $f_4 := f_1 \circ f_3$. Also, recall that we are fixing the visual metric ρ on X , although this only induces the topology of compact strata in X .

For a sequence of real numbers $\underline{\epsilon} = (\epsilon_j)_{j \geq \ell}$ and an integer $m \geq \ell$, we introduce the notation $\sigma_m(\underline{\epsilon}) := \sum_{j \geq m} \epsilon_j$. We also define the sets $U_x^\underline{\epsilon}(m)$ and $U_x^\underline{\epsilon}$ as in Construction 1.14. Namely, $U_x^\underline{\epsilon}(\ell)$ is the open ϵ_ℓ -ball around x within X_ℓ and, for $m > \ell$, the set $U_x^\underline{\epsilon}(m)$ is obtained inductively as the open ϵ_m -neighbourhood of $U_x^\underline{\epsilon}(m-1)$ within X_m . Finally, we have $U_x^\underline{\epsilon} := \bigcup_{m \geq \ell} U_x^\underline{\epsilon}(m)$, which is open in X .

In the rest of the proof, we fix a sequence of positive numbers $\underline{\epsilon} = (\epsilon_j)_{j \geq \ell}$ such that the following conditions hold:

- (1) the closure of $U_x^\underline{\epsilon}(j)$ is contained in O_x , for each $j \geq \ell$;
- (2) $H(3\sigma_{\ell+1}(\underline{\epsilon})) + \mathfrak{h}(2\epsilon_\ell + 2H(3\sigma_{\ell+1}(\underline{\epsilon}))) < \rho(U_x^\underline{\epsilon}(\ell), X_{f_4(\ell)} - O_x)$;
- (3) $\sigma_{j+1}(\underline{\epsilon}) + H'(3\sigma_j(\underline{\epsilon})) < \rho(U_x^\underline{\epsilon}(j), X_{f_3(j)} - O_x)$, for each $j \geq \ell$.

This is possible because the functions $\mathfrak{h}(r), H(r), H'(r)$ all go to zero for $r \rightarrow 0$. Thus, there exists $\epsilon_\ell > 0$ such that the sequence $(\epsilon_\ell, 0, 0, \dots)$ satisfies the above conditions. And, if a sequence $(\epsilon_\ell, \dots, \epsilon_m, 0, 0, \dots)$ satisfies these conditions, then there exists $\epsilon_{m+1} > 0$ such that the sequence $(\epsilon_\ell, \dots, \epsilon_m, \epsilon_{m+1}, 0, \dots)$ also does.

Now, set $U_x := U_x^\underline{\epsilon}$, for our fixed sequence $\underline{\epsilon}$. For simplicity, we will simply write $U_x(m)$ and σ_m in the rest of the proof, rather than $U_x^\underline{\epsilon}(m)$ and $\sigma_m(\underline{\epsilon})$.

By condition (1), we have $U_x \subseteq O_x$. We now show that this open set satisfies the property claimed by the lemma.

Fix $n \geq 1$ and $\delta > 0$. Choose $\eta > 0$ small enough that:

- (1') $\eta \vee H(\eta) \leq \delta/3$;
- (2') $H(\eta + 2\sigma_{\ell+1}) + \mathfrak{h}(2\epsilon_\ell + 2H(\eta + 2\sigma_{\ell+1})) < \rho(U_x(\ell), X_{f_4(\ell)} - O_x)$;
- (3') $\sigma_{j+1} + 3H(\eta) + H'(\eta + 2\sigma_j) < \rho(U_x(j), X_{f_3(j)} - O_x)$, for all $\ell \leq j \leq n$.

Conditions (2') and (3') can be achieved because $\underline{\epsilon}$ was chosen to satisfy conditions (2) and (3) above. The resulting value of η only depends on x, n, δ , so we can define the function in the statement of the lemma as $g_3(x, \delta, n) := \eta$.

Now, consider an η -fine reduced i -cycle c with $0 \leq i < \ell$ and $\text{supp}(c) \subseteq X_n \cap U_x$. Let us construct a δ -fine chain d with $\partial d = c$ and $\text{supp}(d) \subseteq O_x \cap X_{f_4(n)}$. It is convenient to set $\delta' := g_1(\delta, f_3(n)) \leq \delta$, where g_1 is the function provided by Proposition 3.6.

We begin by constructing a sequence of cycles $(c_j)_{\ell \leq j \leq n-1}$ with $\text{supp}(c_j) \subseteq U_x(j)$. First, we define c_{n-1} as a σ_n -projection of c to $U_x(n-1)$, in the sense of Remark 4.4; this exists because $\text{supp}(c) \subseteq X_n \cap U_x$ is contained in the σ_{n+1} -neighbourhood of $U_x(n)$, which is contained in the ϵ_n -neighbourhood of $U_x(n-1)$. Then, for each $\ell \leq j \leq n-2$, we define inductively c_j as an ϵ_{j+1} -projection of c_{j+1} to $U_x(j)$.

Note that $\text{fine}(c_j) \leq \eta + 2\sigma_{j+1}$ for each $\ell \leq j \leq n-1$. In addition, c_j is (ϵ_j, X_j) -related to c_{j-1} , and c is (σ_n, X_{n-1}) -related to c_{n-1} .

We now refine these cycles using the chain maps provided by Proposition 4.5. Specifically, recall that $\delta' = g_1(\delta, f_3(n))$ and define $c'_j := \phi^{\delta', j} c_j$ for each $\ell \leq j \leq n-1$ and $c' := \phi^{\delta', n} c$. Part (3) of the proposition guarantees the existence of δ' -fine chains B_j with $\partial B_j = c'_{j+1} - c'_j$ and $\text{supp}(B_j)$ contained in the $H'(\eta + 2\sigma_{j+1})$ -neighbourhood of $\text{supp}(c_{j+1})$ within $X_{f_3(j+1)}$, for each $\ell \leq j \leq n-2$. Since $\text{supp}(c_{j+1}) \subseteq U_x(j+1)$, this implies that $\text{supp}(B_j) \subseteq O_x$, by applying condition (3') above.

Similarly, we have $\partial B_{n-1} = c' - c'_{n-1}$ and $\text{supp}(B_{n-1})$ is contained in the neighbourhood of $\text{supp}(c) \subseteq X_n \cap U_x$ of radius $H'(\eta + 2\sigma_n)$ within $X_{f_3(n)}$. Thus, $\text{supp}(B_{n-1})$ is contained in the $(\sigma_{n+1} + H'(\eta + 2\sigma_n))$ -neighbourhood of $U_x(n)$ within $X_{f_3(n)}$. Hence $\text{supp}(B_{n-1}) \subseteq O_x$, again by condition (3').

In conclusion, the chain $(B_\ell + \dots + B_{n-1})$ is δ' -fine, supported in $O_x \cap X_{f_3(n)}$ and it satisfies $\partial(B_\ell + \dots + B_{n-1}) = c' - c'_\ell$. Thus, we are only left to similarly fill the cycles c'_ℓ and $c - c'$.

First, we deal with c'_ℓ . Observe that c'_ℓ is δ' -fine and, by part (2) of Proposition 4.5, it is supported in the $H(\eta + 2\sigma_{\ell+1})$ -neighbourhood of c_ℓ within $X_{f_3(\ell)}$. Since $\delta' = g_1(\delta, f_3(n))$ and $n \geq \ell$, Proposition 3.6 allows us to fill c'_ℓ by a δ -fine chain $B_{\ell-1}$ supported in $X_{f_1(f_3(\ell))} = X_{f_4(\ell)}$ and with $\text{diam}(B_{\ell-1}) \leq \mathfrak{h}(\text{diam}(c'_\ell))$. In particular, $\text{supp}(B_{\ell-1})$ is contained in the neighbourhood of $\text{supp}(c_\ell) \subseteq U_x(\ell)$ in $X_{f_4(\ell)}$ of radius:

$$H(\eta + 2\sigma_{\ell+1}) + \mathfrak{h}(2\epsilon_\ell + 2H(\eta + 2\sigma_{\ell+1})),$$

where we have used the fact that $\text{diam}(c_\ell) \leq 2\epsilon_\ell$, since c_ℓ is contained in an ϵ_ℓ -ball. Now, condition (2') implies that $\text{supp}(B_{\ell-1})$ is contained in O_x .

Finally, we fill $c - c'$. Note that c is η -fine and $(H(\eta), X_{f_3(n)})$ -related to c' , by part (2) of Proposition 4.5.

Thus, Lemma 4.6 yields $c - c' = \partial B_n$ for a chain B_n with $\text{fine}(B_n) \leq 3H(\eta) \leq \delta$ (by condition (1')) and $\text{supp}(B_n)$ contained in the $H(\eta)$ -neighbourhood of $\text{supp}(c)$ in $X_{f_3(n)}$. Since $\text{supp}(c)$ is contained in the σ_{n+1} -neighbourhood of $U_x(n)$, condition (3') again shows that $\text{supp}(B_n) \subseteq O_x$.

Summing up, we have $c = \partial(B_{\ell-1} + B_\ell + \dots + B_{n-1} + B_n)$, where each B_i is δ -fine and supported within $X_{f_4(n)} \cap O_x$. This concludes the proof. \square

Proof of Proposition 4.7. For each $\ell \geq 1$, choose a finite subset $A_\ell \subseteq X_\ell$ such that $X_\ell \subseteq \bigcup_{x \in A_\ell} U_x$, where the sets U_x are provided by Lemma 4.8. Define $\mathcal{O}' := \bigcup_{\ell \geq 1} \{U_x - X_{\ell-1} \mid x \in A_\ell\}$. It is clear that all elements of \mathcal{O}' are open in X , that they cover X , and that \mathcal{O}' refines \mathcal{O} .

Now, define $g_2(\delta, n) := \inf_{x \in A_0 \cup \dots \cup A_n} g_3(x, \delta, n)$; since the infimum is taken over a finite set, we have $g_2(\delta, n) > 0$. Let c be a reduced $g_2(\delta, n)$ -fine i -cycle supported within $X_n \cap \mathcal{O}'$, for some $\mathcal{O}' \in \mathcal{O}'$. We necessarily have $\mathcal{O}' = U_y - X_{\ell-1}$ for some $y \in A_\ell$ with $\ell \leq n$. It follows that c is, in particular, $g_3(y, \delta, n)$ -fine, hence Lemma 4.8 yields an \mathcal{O} -tiny δ -fine chain d with $\text{supp}(d) \subseteq X_{f_4(n)}$ and $\partial d = c$. \square

We are finally ready to prove Proposition 4.2. The argument is almost identical to the one in the proof of Proposition 4.5, except that we will use Proposition 4.7 in place of Proposition 3.6.

Proof of Proposition 4.2. Fixing $\delta > 0$, we will simply write ϕ for the chain map ϕ^δ . We will prove the proposition by induction on the degree i of chains. The key ingredient in the construction of ϕ is Proposition 4.7, while the chain homotopy will be obtained filling chains trivially via Definition 1.2.

Set $\overline{\mathcal{O}}_0 := \mathcal{O}$. Define inductively refinements $\overline{\mathcal{O}}_{i+1} < \mathcal{O}'_i < \mathcal{O}_i < \overline{\mathcal{O}}_i$ as follows. First, \mathcal{O}_i is a super-refinement of $\overline{\mathcal{O}}_i$, then \mathcal{O}'_i is the refinement of \mathcal{O}_i provided by Proposition 4.7, and finally $\overline{\mathcal{O}}_{i+1}$ is a double super-refinement of \mathcal{O}'_i (as in Remark 1.11).

Let f_i and g_i be the functions appearing in Proposition 4.7 applied to the cover \mathcal{O}_i . Define $n_0 := n$ and inductively $n_i = f_{k+1-i}(n_{i-1})$. Define $\delta_k := \delta$ and inductively $\delta_{i-1} = g_{k+1-i}(\delta_i, n_{i-1})$. Note that both n_i and δ_i increase with i .

We begin by inductively constructing maps

$$\phi_i : \mathcal{C}_i(X_n, \mathcal{O}_{k+1}) \longrightarrow \mathcal{C}_i(X_{n_i}, \mathcal{O}_{k+1-i}; \delta_i)$$

for $0 \leq i \leq k$ so that $\partial\phi_i = \phi_{i-1}\partial$. For the induction process to work, we additionally require that, for every simplex $\sigma \in \mathcal{C}_i(X_n, \mathcal{O}_{k+1})$, the set $\text{supp}(\phi_i\sigma) \cup \text{supp}(\sigma)$ be $\overline{\mathcal{O}}_{k+1-i}$ -tiny. See Fig. 7

In the base step “ $i = 0$ ”, we can simply take ϕ_0 to be the identity. In the inductive step, suppose that, for some $0 \leq i \leq k - 1$, we are given all maps up to ϕ_i so that they satisfy the required properties. In order to define ϕ_{i+1} , consider a simplex $\sigma \in \mathcal{C}_{i+1}(X_n, \mathcal{O}_{k+1})$. Note that $\phi_i\partial\sigma$ is a cycle and lies in $\mathcal{C}_i(X_{n_i}, \mathcal{O}_{k+1-i}; \delta_i)$.

We claim that $\text{supp}(\phi_i\partial\sigma) \cup \text{supp}(\sigma)$ is \mathcal{O}'_{k-i} -tiny. This is clear for $i = 0$ since ϕ_0 is the identity. For $i \geq 1$, let $\tau_0, \dots, \tau_{i+1}$ be the i -simplices in $\partial\sigma$. By our assumptions on ϕ_i , each set $\text{supp}(\phi_i\tau_j) \cup \text{supp}(\tau_j)$ is contained in some $O_j \in \overline{\mathcal{O}}_{k+1-i}$. Thus, O_0, \dots, O_{i+1} pairwise intersect, hence they are all contained in some $O' \in \mathcal{O}'_{k-i}$ by Remark 1.11, since

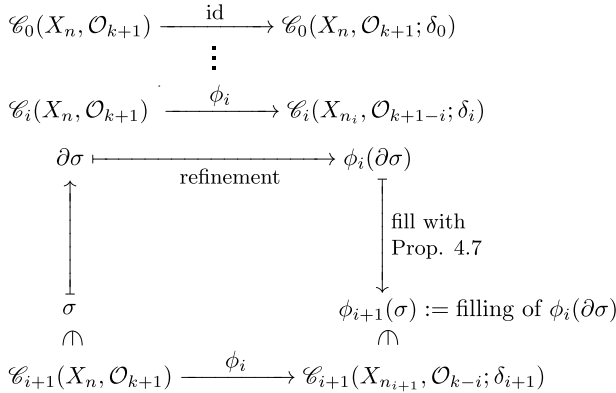


Fig. 7. Proof of part 1 of Proposition 4.2.

$\overline{\mathcal{O}}_{k+1-i}$ was defined as a double super-refinement of \mathcal{O}'_{k-i} . In conclusion, O' contains $\text{supp}(\phi_i \partial\sigma) \cup \text{supp}(\sigma)$.

Now, we have shown that $\phi_i \partial\sigma$ is \mathcal{O}'_{k-i} -tiny, δ_i -fine, and contained in X_{n_i} . Recalling that $n_{i+1} = f_{k-i}(n_i)$ and $\delta_i = g_{k-i}(\delta_{i+1}, n_i)$, Proposition 4.7 implies that $\phi_i \partial\sigma$ (which is reduced as any boundary is reduced) is the boundary of an \mathcal{O}_{k-i} -tiny, δ_{i+1} -fine chain contained in $X_{n_{i+1}}$. We define $\phi_{i+1}\sigma$ as this chain.

It is clear that we have $\partial\phi_{i+1}\sigma = \phi_i \partial\sigma$. We are left to show that $\text{supp}(\phi_{i+1}\sigma) \cup \text{supp}(\sigma)$ is $\overline{\mathcal{O}}_{k-i}$ -tiny. Recall that $\text{supp}(\phi_i \partial\sigma) \cup \text{supp}(\sigma)$ is \mathcal{O}'_{k-i} -tiny, say contained in $O' \in \mathcal{O}'_{k-i}$. In addition, $\text{supp}(\phi_{i+1}\sigma)$ is \mathcal{O}_{k-i} -tiny, say contained in $O \in \mathcal{O}_{k-i}$. Note that $O \cap O' \neq \emptyset$, since

$$\text{supp}(\phi_{i+1}\sigma) \supseteq \text{supp}(\partial\phi_{i+1}\sigma) = \text{supp}(\phi_i \partial\sigma).$$

Recalling that \mathcal{O}'_{k-i} is a refinement of \mathcal{O}_{k-i} , which is a super-refinement of $\overline{\mathcal{O}}_{k-i}$, we conclude that $O \cup O'$ is contained in an element of $\overline{\mathcal{O}}_{k-i}$. This element contains $\text{supp}(\phi_{i+1}\sigma) \cup \text{supp}(\sigma)$, as required.

Finally, we construct a chain-homotopy H . We will inductively define maps

$$H_i: \mathcal{C}_i(X_n, \mathcal{O}_{k+1}) \longrightarrow \mathcal{C}_{i+1}(X_{n_i}, \overline{\mathcal{O}}_{k-i})$$

for $0 \leq i \leq k$ so that $\phi_i - \text{id} = \partial H_i + H_{i-1} \partial$. Again, for the induction to work, we also require that, for every $\sigma \in \mathcal{C}_i(X_n, \mathcal{O}_{k+1})$, the set $\text{supp}(H_i \sigma) \cup \text{supp}(\sigma)$ be $\overline{\mathcal{O}}_{k-i}$ -tiny. See Fig. 8.

In the base step “ $i = 0$ ”, we can just take H_0 to be the zero map, since ϕ_0 is the identity. In the inductive step, suppose that we are given all maps up to H_i so that they satisfy the required properties. In order to define H_{i+1} , consider a simplex $\sigma \in \mathcal{C}_{i+1}(X_n, \mathcal{O}_{k+1})$. Note that:

$$\partial(\phi_{i+1}\sigma - \sigma - H_i \partial\sigma) = \phi_i \partial\sigma - \partial\sigma - (\phi_i - \text{id} - H_{i-1} \partial)\partial\sigma = 0.$$

$$\begin{array}{ccc}
 \mathcal{C}_0(X_n, \mathcal{O}_{k+1}) & \xrightarrow{\text{id}} & \mathcal{C}_0(X_n, \mathcal{O}_{k+1}; \delta_0) \\
 \vdots & & \\
 \mathcal{C}_i(X_n, \mathcal{O}_{k+1}) & \xrightarrow{\phi_i} & \mathcal{C}_i(X_{n_i}, \mathcal{O}_{k+1-i}; \delta_i) \\
 \sigma & \xrightarrow{\text{refinement}} & \phi_i(\sigma) \\
 & \searrow \text{id} & \downarrow \\
 & & \sigma
 \end{array}
 \left. \vphantom{\begin{array}{ccc} \mathcal{C}_0 & \xrightarrow{\text{id}} & \mathcal{C}_0 \\ \vdots & & \\ \mathcal{C}_i & \xrightarrow{\phi_i} & \mathcal{C}_i \\ \sigma & \xrightarrow{\text{refinement}} & \phi_i(\sigma) \\ & \searrow \text{id} & \downarrow \\ & & \sigma \end{array}} \right\} \phi_i - \text{id} = \partial H_i + H_{i-1} \partial \quad (*)$$

(*) fill($\phi_i \sigma - \sigma - H_i \partial \sigma$) trivially via Def. 1.2

Fig. 8. Proof of part 2 of Proposition 4.2.

We claim that the cycle $(\phi_{i+1}\sigma - \sigma - H_i\partial\sigma)$ is $\overline{\mathcal{O}}_{k-i-1}$ -tiny. Recall that $\text{supp}(\phi_{i+1}\sigma) \cup \text{supp}(\sigma)$ is contained in some $O \in \overline{\mathcal{O}}_{k-i}$. In addition, if $\tau_0, \dots, \tau_{i+1}$ are the i -simplices in $\partial\sigma$, there exist $O_j \in \overline{\mathcal{O}}_{k-i}$ containing $\text{supp}(H_i\tau_j) \cup \text{supp}(\tau_j)$. Since O, O_0, \dots, O_{i+1} are pairwise-intersecting elements of $\overline{\mathcal{O}}_{k-i}$, their union is contained in some $\overline{O} \in \overline{\mathcal{O}}_{k-i-1}$ (for this, note that $\overline{\mathcal{O}}_{k-i}$ is a double super-refinement of $\overline{\mathcal{O}}_{k-i-1}$ and apply Remark 1.11). Now, \overline{O} contains the support of $(\phi_{i+1}\sigma - \sigma - H_i\partial\sigma)$, as required.

Finally, by Lemma 1.3, the cycle $(\phi_{i+1}\sigma - \sigma - H_i\partial\sigma)$ is the boundary of the cone over any of its vertices, which is an $(i + 2)$ -chain contained in $X_{n_{i+1}} \cap \overline{O}$. We define $H_{i+1}\sigma$ as this chain. It is clear from the construction that $\partial H_{i+1}\sigma + H_i\partial\sigma = \phi_{i+1}\sigma - \text{id}$ and that $\text{supp}(H_{i+1}\sigma) \cup \text{supp}(\sigma)$ is contained in \overline{O} , hence it is an $\overline{\mathcal{O}}_{k-i-1}$ -tiny set.

This concludes the whole proof, setting $\mathcal{O}' := \mathcal{O}_{k+1}$ and $f_2(n) := n_k$. \square

4.2. Main result

We are finally ready to prove Theorem 4.1, which is the main result of this section. It claims that, $\check{H}^i(X, \mathcal{R}) = \{0\}$ for a principal ideal domain \mathcal{R} and for $1 \leq i < k$, where $X = \partial_*G$.

First, combining Propositions 3.6 and 4.2, we quickly obtain that X satisfies Condition (DA_i) in low degrees, which was introduced in Subsection 1.2.

Proposition 4.9. *For every open cover \mathcal{O} of X , there exists a refinement \mathcal{O}' such that, for every $1 \leq i < k$, every \mathcal{O}' -fine discrete i -cycle is the boundary of an \mathcal{O} -fine discrete chain.*

Proof. Let λ_n be a Lebesgue number for the cover $\{O \cap X_n \mid O \in \mathcal{O}\}$ of X_n . Let f_1, g_1, f_2 be the functions provided by Propositions 3.6 and 4.2. Let \mathcal{O}' be the refinement of \mathcal{O} provided by Proposition 4.2. Set $\delta := g_1(\lambda_{f_1(f_2(n))}, f_2(n))$.

Let c be an \mathcal{O}' -fine cycle. By Proposition 4.2, c is \mathcal{O} -homologous to a δ -fine cycle c' with $\text{supp}(c') \subseteq X_{f_2(n)}$. By Proposition 3.6, we have $c' = \partial d'$ for a $\lambda_{f_1(f_2(n))}$ -fine chain d' with $\text{supp}(d') \subseteq X_{f_1(f_2(n))}$. In particular, d' is \mathcal{O} -fine, showing that c is the boundary of an \mathcal{O} -fine chain. \square

Proof of Theorem 4.1. By Proposition 1.12, the topological space X is super-refinable. Proposition 4.9 shows that X satisfies Condition (DA_i) for $1 \leq i < k$, so Corollary 1.10 yields $\check{H}^i(X, \mathcal{R}) = \{0\}$ in the same range. \square

5. Non-vanishing

Throughout Section 5, we again consider a relatively hyperbolic group G , its Morse boundary $X := \partial_* G$ and the filtration $X = \varinjlim X_n$. Notation and assumptions are exactly the same as in the previous sections (see e.g. Subsection 3.2). In particular, the Bowditch boundary \mathcal{S} of (G, \mathcal{P}) is homeomorphic to a $(k + 1)$ -dimensional sphere S^{k+1} .

The goal of this section is to prove the second part of Theorem B (which implies the second part of Theorem A). To do so, we show that for any finite set of parabolic points F , the rank of the singular homology $H_k(\mathcal{S} - F, \mathbb{Z})$ (as an abelian group) is a lower bound to the dimension of $\check{H}^k(\partial_* G, \mathbb{F})$, where \mathbb{F} is a field of characteristic 0. Recall that the rank of an abelian group A , which we denote by $\text{rk } A$, is the maximal cardinality of a linearly independent subset.

The following is the crucial lemma in this section. Roughly speaking, it shows that, if there are many discrete cycles that are not boundaries, then the Čech cohomology is large. Note however that we have to be rather careful with how fine the various chains are.

Lemma 5.1. *Let \mathbb{F} be a field of characteristic 0. Every open cover \mathcal{O} of X admits a refinement \mathcal{O}' such that*

$$\begin{aligned} \dim \check{H}^k(X, \mathbb{F}) \\ \geq \text{rk} \left(\{ \mathcal{O}'\text{-fine } k\text{-cycles} \} / \{ \mathcal{O}'\text{-fine boundaries of } \mathcal{O}\text{-fine } (k + 1)\text{-chains} \} \right). \end{aligned}$$

Proof. Let $\mathcal{O}_2 \ll \mathcal{O}_1 \ll \mathcal{O}$ be a chain of super-refinements and let $\mathcal{O}' < \mathcal{O}_2$ be the refinement provided by Proposition 4.2. The following diagram is an overview of the various covers and chain maps that we will introduce during the proof.

$$\begin{array}{ccccc} \mathcal{C}_*(X, \mathcal{O}') & \subseteq & \mathcal{C}_*(X, \mathcal{O}_2) & \xrightarrow{g_*} & C_*(N(\mathcal{O}_1)) & \xrightarrow{g'_*} & \mathcal{C}_*(X, \mathcal{O}) \\ \cup & & & & \uparrow h_* & & \\ \mathcal{C}_*(X, \mathcal{U}') & \xrightarrow{h'_*} & & & C_*(N(\mathcal{U})) & & \end{array}$$

Let $g: X \rightarrow \mathcal{O}_1$ be an $(\mathcal{O}_2, \mathcal{O}_1)$ -parent map and $g': \mathcal{O}_1 \rightarrow X$ be a child map. Recalling the chain map g_* from discrete chains to chains in nerves provided by Lemma 1.8, we can consider the map $\Phi: \{ \mathcal{O}'\text{-fine } k\text{-cycles} \} \rightarrow H_k(N(\mathcal{O}_1), \mathbb{Z})$ given by $c \mapsto [g_*(c)]$.

If $[g_*(c)] = [g_*(c')]$ there exists some $d \in C_{k+1}(N(\mathcal{O}_1))$ with $\partial d = g_*(c) - g_*(c')$. Lemma 1.8 implies that $g'_* \circ g_*: \mathcal{C}_*(X, \mathcal{O}_2) \rightarrow \mathcal{C}_*(X, \mathcal{O})$ is a chain map homotopic to the standard inclusion and that $g'_*(d)$ is \mathcal{O} -fine. Thus $c - c' = \partial g'_*(d) + \partial d'$ for some \mathcal{O} -fine

chain d' . This shows that $\text{Ker}(\Phi) \subseteq \{\text{boundaries of } \mathcal{O}\text{-fine } (k + 1)\text{-chains}\}$. Define $A = \text{Im}(\Phi)$.

From now on we identify $H^k(\cdot, \mathbb{F})$ with $\text{Hom}(H_k(\cdot, \mathbb{Z}), \mathbb{F})$, and denote it simply by $H^k(\cdot)$. We are only left to show that $\dim \check{H}^k(X, \mathbb{F}) \geq \text{rk}(A)$.

Claim. *There exists a linear subspace V of $H^k(N(\mathcal{O}_1))$ of dimension $\text{rk}(A)$ and such that, for all $\psi \in V - \{0\}$, we have $\psi|_A \neq 0$.*

Proof of Claim. Let $\{x_i\}$ be a maximal linearly independent subset of A , and extend it to a maximal linearly independent subset $\{x_i\} \sqcup \{y_j\}$ of $H_k(N(\mathcal{O}_1), \mathbb{Z})$. We claim that we can take V to be the subspace of $H^k(N(\mathcal{O}_1))$ of all elements that vanish on all y_j . Indeed, we now argue that any map $f: \{x_i\} \sqcup \{y_j\} \rightarrow \mathbb{F}$ extends uniquely to a homomorphism h_f in $H^k(N(\mathcal{O}_1))$; this gives both the dimension bound and the non-vanishing property from the claim.

To define h_f , we note that all $x \in H_k(N(\mathcal{O}_1), \mathbb{Z})$ have a multiple $n_x x$ that lies in $\bar{L} = \langle \{x_i\} \sqcup \{y_j\} \rangle$ for some integer $n_x \neq 0$, by maximality. By linear independence, f extends to a homomorphism \bar{f} defined on \bar{L} , which we can then further extend by setting $h_f(x) = \bar{f}(n_x x)/n_x$. It is readily checked that h_f is a well-defined homomorphism, and the only one that restricts to f on $\{x_i\} \sqcup \{y_j\}$. \square

We now show that V as in the claim embeds in $\check{H}^k(X, \mathbb{F})$. Indeed, consider any $\alpha \in V - \{0\}$ and consider any refinement $\mathcal{U} < \mathcal{O}_1$ and any spouse map $h: \mathcal{U} \rightarrow \mathcal{O}_1$. We have to show $h^*(\alpha) \neq 0$ for the induced map $h^*: H^k(N(\mathcal{O}_1)) \rightarrow H^k(N(\mathcal{U}))$.

Since $\alpha \in V - \{0\}$ there exists a homology class $\tilde{c} \in A$ such that $\alpha(\tilde{c}) \neq 0$. By the definition of A , there exists a cycle $c \in \mathcal{C}_k(X, \mathcal{O}')$ such that $[g_*(c)] = \tilde{c}$. There exists some $n \in \mathbb{N}$ such that $c \subseteq X_n$. Let \mathcal{U}' be a super-refinement of \mathcal{U} that is also a refinement of \mathcal{O}' and let $h': X \rightarrow \mathcal{U}$ be a $(\mathcal{U}', \mathcal{U})$ -parent map. Let f be the function f_2 from Proposition 4.2 applied to the open cover \mathcal{O}_2 .

Let δ be a Lebesgue number for the open cover $\{U' \cap X_{f(n)} \mid U' \in \mathcal{U}'\}$ of the compact space $X_{f(n)}$. By Proposition 4.2 applied to \mathcal{O}_2 , δ and c , there exists a \mathcal{O}_2 -fine, δ -fine cycle $c' \subseteq X_{f(n)}$ which is \mathcal{O}_2 -homologous to c . Since $g_*: \mathcal{C}_*(X, \mathcal{O}_2) \rightarrow C_*(N(\mathcal{O}_1))$ is a chain map, we have that $[g_*(c)] = [g_*(c')]$. Also, by the choice of δ , the cycle c' is \mathcal{U}' -fine.

Note that both g and $h \circ h'$ are $(\mathcal{U}', \mathcal{O}_1)$ -parent maps, so Lemma 1.8 shows that their induced maps $\mathcal{C}_k(X, \mathcal{U}') \rightarrow C_k(N(\mathcal{O}_1))$ are chain-homotopic. Defining $c_{\mathcal{U}} = h'_*(c') \in C_*(N(\mathcal{U}))$, we get that $[h_*(c_{\mathcal{U}})] = [g_*(c)] = \tilde{c}$ in $H_k(N(\mathcal{O}_1))$. Finally, observe that $h^*(\alpha)([c_{\mathcal{U}}]) = \alpha(\tilde{c}) \neq 0$, as required. \square

We are now ready to prove that $\check{H}^k(X, \mathbb{F})$ is infinite dimensional (Theorem B(2)).

Theorem 5.2. *Let G be a finitely generated, relatively hyperbolic group with virtually nilpotent peripherals and Bowditch boundary homeomorphic to a sphere S^{k+1} , for some $k \geq 1$, and with at least one peripheral subgroup. Let \mathbb{F} be a field of characteristic 0. Then $\check{H}^k(\partial_* G, \mathbb{F})$ is infinite-dimensional.*

Proof. It suffices to show that, for every finite set of parabolic points $F \subseteq \mathcal{S}$, we have the inequality $\text{rk } H_k(\mathcal{S} - F, \mathbb{Z}) \leq \dim \check{H}^k(X, \mathbb{F})$, where $X = \partial_* G$.

Let \mathcal{U} be a convex cover of $\mathcal{S} - F$ as in Lemma 3.9. Set $\mathcal{O} := \mathcal{U}|_X$. To define this cover restriction, recall that there is a natural continuous and injective map $X \rightarrow \mathcal{S}$ (though it is not an open map).

Let $\mathcal{O}' < \mathcal{O}$ be the refinement provided by Lemma 5.1. Finally, let N be the integer associated to F by Proposition 3.8, and let δ be a Lebesgue number for the restriction $\mathcal{O}'|_{X_N}$.

It suffices to show that $H_k(\mathcal{S} - F, \mathbb{Z})$ embeds into

$$\{\mathcal{O}'\text{-fine } k\text{-cycles}\} / \{\mathcal{O}'\text{-fine boundaries of } \mathcal{O}\text{-fine } (k+1)\text{-chains}\}$$

(and hence has smaller rank) and apply Lemma 5.1. By Proposition 3.8, every element of $H_k(\mathcal{S} - F, \mathbb{Z})$ can be represented by a straight singular cycle c such that $\text{disc}(c)$ is δ -fine and contained in X_N , and hence \mathcal{O}' -fine. Since c is a straight cycle and \mathcal{U} is a convex cover, not only $\text{disc}(c)$ but also c itself is \mathcal{U} -fine. Lemma 3.9 concludes the proof. \square

References

- [1] G. Baumslag, *Lecture Notes on Nilpotent Groups*, Regional Conference Series in Mathematics, vol. 2, American Mathematical Society, Providence, R.I., 1971.
- [2] H. Barcelo, V. Capraro, J.A. White, Discrete homology theory for metric spaces, *Bull. Lond. Math. Soc.* 46 (5) (2014) 889–905.
- [3] M.R. Bridson, A. Haefliger, *Metric Spaces of Non-positive Curvature*, Grundlehren der Mathematischen Wissenschaften (Fundamental Principles of Mathematical Sciences), vol. 319, Springer-Verlag, Berlin, 1999.
- [4] M. Bestvina, G. Mess, The boundary of negatively curved groups, *J. Am. Math. Soc.* 4 (3) (1991) 469–481.
- [5] B.H. Bowditch, Relatively hyperbolic groups, *Int. J. Algebra Comput.* 22 (3) (2012) 1250016.
- [6] R. Bott, L.W. Tu, *Differential Forms in Algebraic Topology*, Graduate Texts in Mathematics, vol. 82, Springer-Verlag, New York-Berlin, 1982.
- [7] R. Charney, M. Cordes, A. Sisto, Complete topological descriptions of certain Morse boundaries, *Groups Geom. Dyn.* 17 (1) (2023) 157–184.
- [8] M. Cordes, M.G. Durham, Boundary convex cocompactness and stability of subgroups of finitely generated groups, *Int. Math. Res. Not.* 6 (2019) 1699–1724.
- [9] M. Cordes, D. Hume, Stability and the Morse boundary, *J. Lond. Math. Soc.* (2) 95 (3) (2017) 963–988.
- [10] C.H. Cashen, J.M. Mackay, A metrizable topology on the contracting boundary of a group, *Trans. Am. Math. Soc.* 372 (3) (2019) 1555–1600.
- [11] M. Cordes, Morse boundaries of proper geodesic metric spaces, *Groups Geom. Dyn.* 11 (4) (2017) 1281–1306.
- [12] R. Charney, H. Sultan, Contracting boundaries of CAT(0) spaces, *J. Topol.* 8 (1) (2015) 93–117.
- [13] C. Druțu, M. Kapovich, *Geometric Group Theory*, American Mathematical Society Colloquium Publications., vol. 63, American Mathematical Society, Providence, RI, 2018, With an appendix by Bogdan Nica.
- [14] M. Gromov, Hyperbolic groups, in: *Essays in Group Theory*, in: *Math. Sci. Res. Inst. Publ.*, vol. 8, Springer, New York, 1987, pp. 75–263.
- [15] D. Gruber, A. Sisto, Infinitely presented graphical small cancellation groups are acylindrically hyperbolic, *Ann. Inst. Fourier (Grenoble)* 68 (6) (2018) 2501–2552.
- [16] J.R. Isbell, On finite-dimensional uniform spaces, *Pac. J. Math.* 9 (1959) 107–121.
- [17] R. Karidi, Geometry of balls in nilpotent Lie groups, *Duke Math. J.* 74 (2) (1994) 301–317.

- [18] C.J. Leininger, S. Schleimer, Hyperbolic spaces in Teichmüller spaces, *J. Eur. Math. Soc.* 16 (12) (2014) 2669–2692.
- [19] J.M. Mackay, A. Sisto, Quasi-hyperbolic planes in relatively hyperbolic groups, *Ann. Acad. Sci. Fenn., Math.* 45 (1) (2020) 139–174.
- [20] J. Mackay, A. Sisto, Maps between relatively hyperbolic spaces and between their boundaries, *Trans. Am. Math. Soc.* 377 (2) (2024) 1409–1454.
- [21] R. Schultz, Notes on topological dimension theory, Available at <https://math.ucr.edu/~res/miscpapers/top-dimension-theory.pdf>, 2012.
- [22] A. Sisto, Projections and relative hyperbolicity, *Enseign. Math.* (2) 59 (1–2) (2013) 165–181.
- [23] A. Sisto, Quasi-convexity of hyperbolically embedded subgroups, *Math. Z.* 283 (3–4) (2016) 649–658.
- [24] E.H. Spanier, *Algebraic Topology*, McGraw-Hill Book Co., New York-Toronto, Ont.-London, 1966.
- [25] M. Zarichnyi, Regular linear operators extending metrics: a short proof, *Bull. Pol. Acad. Sci., Math.* 44 (3) (1996) 267–269.
- [26] S. Zbinden, Morse boundaries of 3-manifold groups, arXiv preprint arXiv:2212.08368, 2022.
- [27] S. Zbinden, Morse boundaries of graphs of groups with finite edge groups, *J. Group Theory* 26 (5) (2023) 969–1002.