



Homology bounds for hyperbolic orbifolds

Hartwig Senska¹

Received: 12 June 2021 / Accepted: 30 September 2022
© The Author(s) 2022

Abstract

We will provide bounds on both the Betti numbers and the torsion part of the homology of hyperbolic orbifolds. These bounds are linear in the volume and are a direct consequence of an efficient simplicial model of the thick part, which we will construct as well. The homology statements complement previous work of Bader, Gelander and Sauer (torsion homology of manifolds), Samet (Betti numbers of orbifolds) and a classical theorem of Gromov (Betti numbers of manifolds). For arithmetic, non-compact hyperbolic orbifolds—i.e. in the case of arithmetic, non-uniform lattices in $\text{Isom}(\mathbb{H}^n)$ —the strongest results will be obtained.

Keywords Orbifolds · Hyperbolic geometry · Homology · Betti numbers · Torsion in homology · Thick-thin decomposition · Arithmetic lattices

Mathematics Subject Classification 55N99 · 57R18

1 Introduction

An interesting feature of negative curvature is that the topology (e.g. in terms of the homology) of sufficiently well-behaved spaces like manifolds or orbifolds can in some sense be controlled by their volume. The following result of Gromov [2] can be seen as a starting point for studies into this theme: for a Hadamard n -manifold X with pinched negative sectional curvature¹ $-1 \leq K \leq a < 0$ (for some $a < 0$) and a torsion-free lattice $\Gamma < \text{Isom}(X)$, the Betti numbers—that is, the free part of the homology—of the quotient manifold X/Γ is linearly bounded by the volume, i.e.

$$b_k(X/\Gamma; \mathbb{K}) \leq C \cdot \text{Vol}(X/\Gamma).$$

This holds for all degrees $k = 0, \dots, n$ and arbitrary coefficient fields \mathbb{K} . Here, $C = C(n) > 0$ is a constant depending only on the dimension n . Note that in this and all following statements,

¹ In the analytic case, this extends to non-positive curvature $-1 \leq K \leq 0$, given that there are no Euclidean de Rham factors in X .

The author acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 281869850 (RTG 2229).

✉ Hartwig Senska
hartwig.senska@alumni.kit.edu

¹ Institute for Algebra and Geometry, Karlsruhe Institute of Technology, 76131 Karlsruhe, Germany

when we speak about homology, we simply mean the singular homology of the topological quotient space X/Γ ; we make no statements about the group homology² of Γ .

Using a similar Morse theoretic argument to the one of Gromov, Samet [7] showed that this can be extended to general lattices $\Gamma < \text{Isom}(X)$, i.e. where X/Γ is an orbifold and not necessarily a manifold. Again, the Betti numbers of the quotient orbifold satisfy

$$b_k(X/\Gamma; \mathbb{K}) \leq C \cdot \text{Vol}(X/\Gamma),$$

where $k = 0, \dots, n$ and $C = C(n) > 0$ depends only on the dimension n ; here, the coefficient field \mathbb{K} has to have characteristic 0.

As the Betti numbers constitute only one part of the homology, it is a natural question whether the other part—the torsion in the homology—might admit similar bounds. Bader, Gelander and Sauer [1] settled this question positively for the case of negatively pinched manifolds (i.e. torsion-free lattices again), showing that

$$\log |\text{tors } H_k(X/\Gamma; \mathbb{Z})| \leq C \cdot \text{Vol}(X/\Gamma)$$

holds for all $k = 0, \dots, n$, where $C = C(n) > 0$ is a constant depending only on the dimension n . As a special case, the statement for degree $k = 1$ in dimension $n = 3$ has to be excluded; using Dehn surgery, [1] gives an explicit counterexample for that situation. Unlike Gromov [2] and Samet [7], Bader, Gelander and Sauer [1] don't employ Morse theory to prove the statement above; instead, they construct an efficient simplicial decomposition of the thick part of X/Γ , which yields the torsion homology result as a direct consequence. This decomposition would also imply another proof for Gromov's theorem above. The curvature conditions in [1] can be relaxed to negatively curved visibility manifolds, as was shown by the author of the present paper in [8].

We will provide answers for the remaining case of torsion homology for orbifolds (i.e. general, not necessarily torsion-free lattices) in the hyperbolic setting. Similar to [1], we will achieve this by first constructing an efficient simplicial decomposition of the thick part of the orbifold. To fix notation, $\Gamma < \text{Isom}(\mathbb{H}^n)$ will denote a lattice, $\eta \in \mathbb{N}$ an upper bound on the order of finite subgroups of Γ (i.e. $|G| \leq \eta$ for all $G < \Gamma$ finite) and $\nu > 0$ a lower bound on the displacement of hyperbolic elements³ of Γ (both η and ν always exist). We let $M := \mathbb{H}^n/\Gamma$ denote the quotient orbifold and write M_+ for its thick part. Our main result now states:

Theorem (see Theorem 3) *There are constants $C = C(n, \eta, \nu) > 0$ and $D = D(n, \nu) > 0$, such that for any such orbifold M , the pair $(M_+, \partial M_+)$, i.e. the thick part and its boundary, is as a pair homotopy equivalent to a simplicial pair (S, S') , where the number of vertices of S is bounded by $C \cdot \text{Vol}(M)$ and the degree at the vertices of S is universally bounded by D . For arithmetic, non-uniform Γ , the constants C and D will only depend on the dimension n .*

By a straightforward Mayer-Vietoris argument, our simplicial decomposition yields another proof for the linear bounds on the Betti numbers as a byproduct.

² The reason is that in general, the group homology of Γ will not coincide with the (singular) homology of X/Γ . This is especially true for the case of torsion in Γ . Despite this, recall that in the situation of Gromov's statement the space X/Γ would be a $B\Gamma$, thus leading to the group homology and the singular to homology to coincide in that particular example.

³ In the manifold setting, this would be equivalent to a lower bound on the length of geodesic loops in M ; the interpretation is only slightly more complicated in our orbifold situation.

Theorem (see Theorem 4) *There is a constant $E = E(n, \eta, \nu) > 0$ such that for any such orbifold M , we have*

$$b_k(M; \mathbb{K}) \leq E \cdot \text{Vol}(M)$$

for all $k \in \mathbb{N}_0$ and arbitrary coefficient field \mathbb{K} . For arithmetic, non-uniform Γ , the constant E will only depend on the dimension n .

While Samet’s [7] result above already provided a linear bound for the Betti numbers even under more general curvature assumptions, it is restricted to coefficient fields of characteristic 0; our statement has the advantage that it is valid for arbitrary coefficients.

The application of central interest is the torsion of the homology, where we will show:

Theorem (see Theorem 5) *There is a constant $F = F(n, \eta, \nu) > 0$ such that for any such orbifold M , we have*

$$\log |\text{tors } H_k(M; \mathbb{Z})| \leq F \cdot \text{Vol}(M)$$

for all $k \in \mathbb{N}_0$. For arithmetic, non-uniform Γ , the constant F will only depend on the dimension n .

Note that we did not exclude the case of degree $k = 1$ in dimension $n = 3$, in contrast to the similar statement of Bader, Gelander and Sauer [1] above; our extra restrictions regarding ν enable us to do so.

Although the general situation of our statements uses assumptions on η and ν , the results are as good as possible for the arithmetic, non-uniform case. Recall that arithmetic, non-uniform lattices in $\text{Isom}(\mathbb{H}^n)$ correspond to arithmetic, non-compact hyperbolic orbifolds. They form an interesting and widely studied class of orbifolds, with the maybe most prominent examples given by the Bianchi orbifolds (in dimension $n = 3$).

Our restriction to the hyperbolic setting is mainly for technical reasons, as we will use convexity arguments that only hold in constant curvature to prove that the cover we construct for the thick part is indeed a good cover; the goodness of the cover is needed to be able to use (some appropriate modification of) the Nerve lemma to obtain the desired simplicial model.

Given the bound η on the order of finite subgroups of Γ that we require, it is possible to linearly bound the number of components of the thin part (see our thick-thin decomposition Theorem 2) using a fairly straightforward volume argument. This in turn is crucial for controlling the homology. Note that using a more subtle approach, Samet [7] was able to bound the number of thin components without imposing such an η . This might suggest that it would be possible to dispose of η in our case as well.

Similarly, we suspect that ν , which was needed for technical reasons only, could be eliminated from all the above statements. In fact, more general versions of the results presented here (see [9]) already no longer rely on ν ; but as of now, this comes at the price of having bounds polynomial in the volume⁴ instead of the linear ones presented in this paper. As these slight generalizations need significantly more technical work, we refrain from presenting them here.

1.1 Structure of the paper

The following sect. 2 summarizes some well-known facts utilized throughout this paper. As a next step, in Sect. 3 we state the thick-thin decomposition for orbifolds in a fairly general

⁴ Essentially in the form of $C \cdot \text{Vol}(M)^{k+1}$ for a constant $C = C(n, \eta) > 0$ and homology degree k .

setting. In Sect. 4, we restrict ourselves to the hyperbolic case and prove the main result, namely the efficient simplicial model for the thick part of hyperbolic orbifolds. The main applications for the homology of such orbifolds are contained in the final Sect. 5.

1.2 Acknowledgement

The results presented here are part of my doctoral thesis [9], which was written under the supervision of Prof. Roman Sauer, to whom I am grateful for supporting my work. The author acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 281869850 (RTG 2229).

2 Preliminaries

First, let us fix notation and state some useful facts. General references for the concepts covered here are [2] and [3].

We will always let X denote a Hadamard manifold with curvature $-1 \leq K \leq a < 0$ for some $a < 0$. For a lattice $\Gamma < \text{Isom}(X)$, $M := X/\Gamma$ will be the finite-volume orbifold to be studied (M is a manifold if and only if Γ is torsion-free). The boundary at infinity of X will be denoted by $X(\infty)$ or sometimes ∂X .

Every isometry $\gamma \in \text{Isom}(X)$ gives rise to a displacement function

$$d_\gamma : X \rightarrow [0, \infty), \quad x \mapsto d_\gamma(x) := d(x, \gamma x),$$

which can be used to classify the nontrivial isometries of $\text{Isom}(X)$: γ is elliptic if d_γ has minimum 0; it is hyperbolic if d_γ has minimum > 0 ; and it is parabolic if d_γ has no minimum. Elliptic isometries are precisely the torsion elements of Γ ; if M is non-compact, there has to be a parabolic $\gamma \in \Gamma$. The different isometry types are stable under taking powers (with powers $\neq 0$). Hyperbolic isometries have precisely two fixed points in $X(\infty)$, whereas parabolic isometries have precisely one.

Since the distance function of X is convex, the same holds for the displacement functions. Hence the sublevel sets $\{d_\gamma < a\}$ and $\{d_\gamma \leq a\}$ (for $a \geq 0$) are convex as well. Fixed point sets $\text{Fix}(\gamma)$ of elliptic isometries γ arise in the special case of $a = 0$; they are complete, totally geodesic submanifolds of codimension ≥ 1 (sometimes we treat id as an elliptic isometry, in which case $\text{Fix}(\text{id}) = X$ has codimension 0).

For a closed convex set $W \subseteq X$, there is a well-defined projection $\pi_W : X \rightarrow W$ sending a point $x \in X$ to the (unique) point $\pi_W(x) \in W$ of smallest distance to x ; we will call $\pi_W(x)$ the projection point or foot point of x in W . This projection is equivariant under isometries preserving W , i.e. if $\gamma \in \text{Isom}(X)$ with $\gamma W = W$, then $\pi_W(\gamma x) = \gamma \pi_W(x)$ for all $x \in X$.

To denote the open r -neighborhood of a subset $A \subseteq X$ (and similarly for subsets of M), we will write $(A)_r := \{x \in X : d(x, A) < r \text{ for some } a \in A\}$.

We will adopt some notation from [7] regarding singular submanifolds. For $G < \text{Isom}(X)$, let $\text{Fix}(G) := \bigcap_{g \in G} \text{Fix}(g)$. Now for $\Delta < \text{Isom}(X)$, define

$$\begin{aligned} \Sigma(\Delta) &:= \{\text{Fix}(G) : G < \Delta \text{ finite}\}, \\ \Sigma_i(\Delta) &:= \{Y \in \Sigma(\Delta) : \dim(Y) = i\}, \text{ and} \\ \Sigma_{<i}(\Delta) &:= \{Y \in \Sigma(\Delta) : \dim(Y) < i\}. \end{aligned}$$

Moreover,

$$S_i(\Delta) := \bigcup_{Y \in \Sigma_i(\Delta)} Y \quad \text{and} \quad S_{<i}(\Delta) := \bigcup_{Y \in \Sigma_{<i}(\Delta)} Y.$$

We will often omit Δ in the notation if it is obvious from the context. Note that since $\text{id} \in \Delta$ (for arbitrary $\Delta < \text{Isom}(X)$), we always have $X = \text{Fix}(\{\text{id}\}) \in \Sigma(\Delta)$.

To prove our main results, we will restrict ourselves to the hyperbolic space $X = \mathbb{H}^n$ and use the upper half space model:

$$\mathbb{H}^n = \{(x, t) = (x_1, \dots, x_{n-1}, t) \in \mathbb{R}^{n-1} \times \mathbb{R} \mid t > 0\} = \mathbb{R}^{n-1} \times \mathbb{R}_{>0} \subseteq \mathbb{R}^n$$

with the usual hyperbolic metric. In \mathbb{H}^n , a nonempty, closed subset is convex if and only if it is the intersection of all its (closed) supporting half spaces ([5] Proposition II.1.4.1). For the upper half space model, the distances are given by the following formulas:

- $d((x, t), (y, s)) = 2 \cdot \text{artanh} \left(\sqrt{\frac{\|x-y\|^2 + (t-s)^2}{\|x-y\|^2 + (t+s)^2}} \right),$
- $d((x, t), (x, s)) = \left| \ln \left(\frac{t}{s} \right) \right|,$
- $d((x, t), (y, t)) = 2 \cdot \text{arsinh} \left(\frac{\|x-y\|}{2t} \right).$

Here, $\| \cdot \|$ denotes the usual Euclidean norm on the \mathbb{R}^{n-1} -factor.

3 Thick-thin decomposition

The general idea behind the thick-thin decomposition is most obvious in the manifold case: the manifold can be decomposed into two parts-thick part and thin part-, with the thick part being characterized by a uniform lower bound on the injectivity radius; this makes its geometry easy to control. On the other hand, the thin part turns out to consist of only two types of components, tubes (which are ball bundles over the circle) and cusps (which are products of a compact manifold with a ray). Tubes correspond to sublevel sets of (the displacement function of) hyperbolic isometries in the universal cover, while cusps similarly correspond to sublevel sets of parabolic isometries. This relies on the absence of elliptic isometries in the lattice Γ . While the general situation is similar in the orbifold case, the elliptic isometries occurring now might complicate the picture. As we will see, using our slightly different construction, we can essentially remove the influence elliptic isometries might have: the thin part will be given by the sublevel sets of parabolic and hyperbolic isometries only.

Arguably the most essential tool in the thick-thin decomposition is the Margulis lemma:

Theorem 1 (Margulis lemma; [7] Theorem 2.1) *There are constants $\varepsilon(n) > 0$ and $m(n) \in \mathbb{N}$ depending only on n , such that if X is an n -dimensional Hadamard manifold with sectional curvature $-1 \leq K \leq 0$, then for every discrete group $\Gamma < \text{Isom}(X)$, every $x \in X$ and every $\varepsilon \leq \varepsilon(n)$, the group*

$$\Gamma_\varepsilon(x) := \{\gamma \in \Gamma : d_\gamma(x) < \varepsilon\}$$

contains a nilpotent normal subgroup N of index $\leq m(n)$. If $\Gamma_\varepsilon(x)$ is finite, then N is abelian.

The constants $\varepsilon(n)$ and $m(n)$ in Theorem 1 will be called Margulis ε and Margulis index constant, respectively.

In the standard thick-thin decompositions (see e.g. [2] chapter 10, [3] chapter D; moreover, [4] chapter 3.5 is interesting as it explicitly covers the orbifold case), all isometries are treated

equally: the sublevel sets in question are given by $\{d_\gamma < \varepsilon\}$ for some fixed $\varepsilon \in (0, \varepsilon(n))$. [1] already introduced the concept of varying levels ε_γ for different γ for manifolds (see also [8] chapter 3 for a proof). Our definition extends this to the orbifold case.

Let $\varepsilon \in (0, \varepsilon(n)/2]$ be arbitrary, but fixed⁵, and let $\Gamma \ni \gamma \mapsto \varepsilon_\gamma \in [\varepsilon, \varepsilon(n)/2]$ be a conjugation-invariant choice of levels, i.e. $\varepsilon_{\gamma\gamma^{-1}} = \varepsilon_\gamma$ for all $\gamma, \gamma' \in \Gamma$. Define

$$\Gamma_{\varepsilon_\Gamma}(x) := \{\gamma \in \Gamma : d_\gamma(x) < \varepsilon_\gamma\}$$

and consequently the **thick part** X_+ of X as

$$X_+ := \{x \in X : \Gamma_{\varepsilon_\Gamma}(x) \text{ is finite}\}.$$

The **thin part** X_- is its complement, i.e.

$$X_- := \{x \in X : \Gamma_{\varepsilon_\Gamma}(x) \text{ is infinite}\}.$$

We define the **thick part** M_+ and the **thin part** M_- of $M = X/\Gamma$ as the quotient of the thick part and the thin part of X , respectively, by the group action of Γ , i.e.

$$M_+ := X_+/\Gamma \quad \text{and} \quad M_- := X_-/\Gamma.$$

By conjugation-invariance of $\gamma \mapsto \varepsilon_\gamma$, this is well-defined.

An important fact is that in the present situation, for a fixed lattice Γ , there always exists an upper bound $\eta := \eta(\Gamma) \in \mathbb{N}$ on the order of finite subgroups of Γ (see [4] Proposition 5.4.2), i.e. $|G| \leq \eta$ for all finite $G < \Gamma$.

As the lengthy proof of the thick-thin decomposition is of limited benefit, we will omit it. The interested reader might look up the details in [9] chapter 3.2⁶. Eventually, we arrive at the following result:

Theorem 2 *We have:*

1. M_+ is a compact orbifold with boundary.
2. M_+ is connected for dimension $n > 2$.
3. The number of connected components of M_- is bounded by $C \cdot \text{Vol}(M)$, where $C = C(n, \eta) > 0$ is a constant only depending on n and $\eta = \eta(\Gamma) \in \mathbb{N}$. In dimension $n = 2$, C is independent of η .
4. The connected components U of M_- are of one of the following two shapes:
 - Tubes (bounded components), i.e. (type 1) U is homeomorphic to a (D^{n-1}/E) -bundle over \mathbb{S}^1 with $E < O(n-1)$ finite; or (type 2) U is homeomorphic to

$$D' \times [0, 1] \quad \text{or} \quad D' \times (0, 1),$$

where $D' := (D^{n-1}/E)/\mathbb{Z}_2$ with finite $E < O(n-1)$, and the \mathbb{Z}_2 -action on D^{n-1}/E might be trivial.

Type 1 tubes or homotopy equivalent to \mathbb{S}^1 , whereas type 2 tubes are contractible

- Cusps (unbounded components), i.e. U is homeomorphic to $V \times (0, \infty)$ for some compact $(n-1)$ -dimensional orbifold V . If ∂U is the boundary of U in M , then there is a strong deformation retraction of U onto ∂U .

⁵ We will generally choose ε to be some fixed fraction of $\varepsilon(n)$ to make sure all constants only depend on the dimension n .

⁶ The proof is—as a gross simplification—the combination of the ideas of the thick-thin decomposition for orbifolds with constant levels ε as in [4] and the thick-thin decomposition with variable levels ε_γ (for manifolds) as in [1] and [8].

In particular, M is homotopy equivalent to the compact orbifold M_C with boundary, which is constructed out of M by contracting the cusps onto their common boundary with M_+ . Equivalently, M_C is the union of M_+ with the finitely many tubes.

It can be shown that using $M(n) := 2m(n) + 1$ (with $m(n)$ the Margulis index constant), if the group

$$\Gamma_{\varepsilon/M(n)}(x) = \langle \gamma \in \Gamma : d_\gamma(x) < \varepsilon/M(n) \rangle$$

is infinite (where ε s.t. $0 < \varepsilon/M(n) \leq \varepsilon(n)$), there already has to be some $\gamma \in \Gamma_{\varepsilon/M(n)}(x)$ of infinite order with $d_\gamma(x) < \varepsilon$. Using this fact, we can deduce the following corollary.

Corollary 1 *If for some fixed $\varepsilon' \in (0, \varepsilon(n)/2]$ we choose*

$$\varepsilon_\gamma := \begin{cases} \varepsilon' & \text{if } \gamma \text{ hyperbolic or parabolic,} \\ \varepsilon := \frac{\varepsilon'}{M(n)} & \text{if } \gamma \text{ elliptic,} \end{cases}$$

with $M(n) \in \mathbb{N}$ as above, then the constant $C > 0$ in Theorem 2. depends only on ε', n and η , i.e. $C = C(\varepsilon', n, \eta)$, and we have

$$\begin{aligned} X_- &= \{x \in X : \text{There is a hyperbolic or parabolic } \gamma \in \Gamma \text{ with } d_\gamma(x) < \varepsilon_\gamma\} \\ &= \bigcup_{\gamma \in \Gamma'} \{d_\gamma < \varepsilon_\gamma\}, \quad \text{where } \Gamma' = \{\gamma \in \Gamma : \gamma \text{ hyperbolic or parabolic}\}. \end{aligned}$$

Remark 1 Note that for a fixed lattice Γ , there always exists a lower bound $\nu := \nu(\Gamma) > 0$ on the displacement of hyperbolic isometries, i.e. $d_\gamma(x) > \nu$ for all $x \in X$ and $\gamma \in \Gamma$ hyperbolic. This can be used to move the tubes to the thick part; one way of doing this is to simply assume that the Margulis ε is smaller than ν , which effectively means replacing every occurrence of $\varepsilon(n)$ by $\tilde{\varepsilon}(n)$, where $\tilde{\varepsilon}(n) := \min(\varepsilon(n), \nu)$. Hence every constant depending on $\varepsilon(n)$, in particular the C from the previous statements, will depend on ν as well. From now on we will always assume this, i.e. every time we use the Margulis ε , we will implicitly take its minimum with ν . The benefit is that M_- will then consist of cusps only (so $X_- = \bigcup_{\gamma \in \Gamma_p} \{d_\gamma < \varepsilon_\gamma\}$, where $\Gamma_p = \{\gamma \in \Gamma : \gamma \text{ parabolic}\}$). In particular, $M_+ = M_C$ and thus M itself will be homotopy equivalent to its thick part.

4 Efficient simplicial model

In order to achieve the desired bounds on the homology of the orbifold M , we will show that M admits an efficient simplicial model, i.e. there is a suitable homotopy equivalence to a simplicial complex with bounded complexity. This is done in several steps. First, contracting the cusps is a straightforward way to get a homotopy equivalence between M and its shrunken thick part M'_+ defined below⁷, so it suffices to construct a nice simplicial model for M'_+ . For that, we will essentially use the nerve construction: a space with a good cover (i.e. contractible covering sets with contractible intersections) is homotopy equivalent to its nerve (simplicial) complex. The tricky part is constructing this good cover of M'_+ . Due to the singularities in the orbifold M , we are no longer free to choose the position of possible covering sets. In fact,

⁷ Again, using our construction in which the thin part consists solely of cusps; of course the presence of tubes in the thin part would further complicate the matter. This more general situation is dealt with in [9], yielding only slightly better results than the ones stated here.

just to achieve a good cover of M'_+ as a subspace of M , the covering sets have to conform to the positions of the singular submanifolds of M (to that end, we will use *foldable sets* similar to those in [7]). This rigidity in the position of the covering sets leads to a complicated boundary of the covering space (as a subspace of M), which will in general no longer be homotopy equivalent to M'_+ . To remedy this, we will fill in possible gaps at the boundary of that covering space by using new sets called *stretched balls*, eventually obtaining a good cover of M'_+ homotopy equivalent to M'_+ .

4.1 Defining the flow

We often utilize that the orbifolds we study are homotopy equivalent to their thick parts. While there are many possible ways to contract onto the thick part, there is a very natural one which turns out to have many useful properties; in particular, it will define a homotopy equivalence between the thick part and the shrunken thick part as well. From now on, we will use the thick-thin decomposition (Theorem 2 and Corollary 1) with $\varepsilon' := \varepsilon(n)/2$. Here and henceforth, the **shrunken thick part** will be defined by

$$X'_+ := X \setminus (X_-)_{\varepsilon(n)/32} \quad \text{and} \quad M'_+ := M \setminus (M_-)_{\varepsilon(n)/32},$$

respectively. Recall that in our case, the thin part consists of cusps only. Now for every cusp of M , there is a parabolic fixed point $z \in X(\infty)$. Hence for the region around a preimage of such a cusp, we can define the flow (in X) by flowing along the geodesics to/from z . As this turns out to be Γ -equivariant, we get a similar flow in M . This procedure can be repeated for every single cusp separately, and since the cusps have a uniform minimal distance from each other, we can stick these flows together to construct a global flow for all cusps simultaneously. Using the notation from the thick-thin decomposition (Theorem 2), the precise statement is as follows:

Lemma 1 *The map $F : X_+ \times [0, 1] \rightarrow X_+$, given by flowing away from the parabolic fixed points $z \in X(\infty)$ along the geodesics to them, defines a strong deformation retraction of X_+ onto X'_+ , which is Γ -equivariant and at time 1 induces a homeomorphism $F(\cdot, 1)|_{\partial X_+}$ between ∂X_+ and $\partial X'_+$.*

Consequently, flowing along the images (under the projection $\pi : X \rightarrow X/\Gamma = M$) of these geodesics yields a map $f : M_+ \times [0, 1] \rightarrow M_+$, which is a strong deformation retraction of M_+ onto M'_+ and at time 1 induces a homeomorphism $f(\cdot, 1)|_{\partial M_+}$ between ∂M_+ and $\partial M'_+$.

Proof The desired properties of the flow can be shown in a similar way as in [8] chapter 4.1. A more detailed discussion can also be found in [9] chapter 3. □

4.2 Stretched and foldable sets

Our special tools to arrive at a good covering of the thick part in the orbifold case will be the foldable sets and stretched balls mentioned earlier. While foldable sets can be used in the more general negative curvature setting, the stretched balls are only useful in the hyperbolic case⁸, hence from now on we will always assume $X = \mathbb{H}^n$.

⁸ In fact, stretched balls might also be helpful in the more general setting, but it will be much harder to show that. A crucial feature of the covering sets is that they and their intersections have to be contractible. In the hyperbolic case, it is easy to show that stretched balls are convex, immediately giving the desired properties.

Foldable sets

Foldable sets were already introduced in [7]. The motivation behind them is that in order to have a contractible image in the quotient orbifold, the shape and position of a set in the universal cover has to be compatible with the singular submanifolds⁹.

Definition 1 Let $U \subseteq X$ be open and $Y \subseteq X$ be a convex, complete, totally geodesic submanifold. U is **Y -foldable**, if it has the following properties:

1. U is convex and precisely invariant under Γ , i.e. for $\gamma \in \Gamma$ we always have $\gamma U = U$ or $\gamma U \cap U = \emptyset$.
2. Y is fixed pointwise by $\Gamma_U = \{\gamma \in \Gamma : \gamma U = U\}$.
3. $\pi_Y(U) \subseteq U$, where $\pi_Y : X \rightarrow Y$ projects to the closest point in Y .
4. The image $\pi(U \cap Y)$ of $U \cap Y$ in X/Γ is contractible.

If U is a Y -foldable set for suitable Y , we will call U **foldable**. In that case, the image $\pi(U) \subseteq X/\Gamma$ is a **folded set**.

Note that we altered the definition slightly: we only assume that $\pi(U \cap Y)$ is contractible, but not necessarily convex. Later on, Y will always be given by the fixed set of a finite subgroup of Γ , hence the assumptions on Y will be fulfilled automatically.

Just as in [7], we get the following statements.

Lemma 2 ([7] Proposition 4.9) *Folded sets are contractible.*

Lemma 3 *If U and Y are subsets of X , such that U is precisely invariant under Γ and Y is fixed pointwise by Γ_U , then $U \cap Y$ is mapped injectively into X/Γ . In particular, $U \cap Y$ is homeomorphic to $\pi(U \cap Y) \subseteq X/\Gamma$.*

We see that the above Lemma 3 holds in particular for Y -foldable sets U . Conversely, we can also use it to prove that a set is foldable: if properties 1. and 2. hold, then for property 4. the situation in X/Γ can be reduced to the one in X , which—in general—is less complicated.

The following lemma tells us, in which situations *ordinary* balls (and their intersections) are foldable and when an intersection of folded balls remains folded (and thus contractible); as it is a technical lemma concerned with the relationship between foldable sets and singular submanifolds, it might be helpful to revisit the Σ - S -notation regarding singular submanifolds (of or up to some dimension i) outlined in section 2. In a next step, we will extend the results to our special covering sets, the *stretched* balls.

Lemma 4 [[7] Proposition 4.10] *Let $Y \in \Sigma_i$ (where $i = n$, i.e. $Y = X$, is also possible).*

This uses the fact that in \mathbb{H}^n , every half space is convex; but on the other hand—due to a classical result of Cartan—a space where for every point x and every tangent plane Σ_x at that point there is a totally geodesic submanifold tangent to Σ_x (which holds true under the convex half space condition) already has to have constant curvature, restricting us to the case of \mathbb{H}^n . Hence in general negative curvature, other arguments than convexity would have to be used to show that stretched balls and their intersections are contractible.

⁹ As an example, let the fixed point $p \in \mathbb{H}^2$ of a suitable rotation γ be the singular submanifold. The quotient $M := \mathbb{H}^2/\langle\gamma\rangle$ will look like a cone, where $\pi(p)$ is the cone point. A ball far away from p will still have an image in M that looks like a ball, which thus is contractible. But as we let the ball move towards p , at some point the rotation will glue opposite parts of the ball together, leading to a non-contractible image homotopic to \mathbb{S}^1 , no matter how small the radius of the ball is. Note that if the ball was centered in p , the image would always be contractible, as it would just be a ball around the cone point. So the moral is that (in order to have a contractible image) balls would either have to lie far away from the singular submanifolds, or—if they lie close to them – already be centered in them.

- (a) If $y \in Y \setminus S_{<i}$ and $\mu > 0$ such that Y is fixed pointwise by $\Gamma_{4\mu}(y)$, then $B_\mu(y)$ is Y -foldable.
- (b) Let U_1 be a Y -foldable ball as in (a) and U_2, \dots, U_k foldable balls with centers $y_2, \dots, y_k \in Y$; note that we don't assume $y_2, \dots, y_k \in Y \setminus S_{<i}$ and hence the U_2, \dots, U_k will in general not be Y -foldable. If $U := \bigcap_{j=1}^k U_j \neq \emptyset$, then:

1. U is Y -foldable.

2. If the radii μ_j of the U_j are chosen in a way such that y_j is fixed by $\Gamma_{4\mu_j}(y_j)$ ($j = 1, \dots, k$), then $\pi(U) = \bigcap_{j=1}^k \pi(U_j)$.

Hence the intersection $\bigcap_{j=1}^k \pi(U_j)$ of the $\pi(U_j)$ is folded and thus contractible.

Stretched balls

We will now introduce stretched balls, which are needed later on to fill in gaps in the good cover of the thick part, which could appear if we used ordinary balls exclusively. Stretched balls only need to be defined near the common boundary of a cusp and the thick part. Let $z \in X(\infty) = \partial\mathbb{H}^n$ be the corresponding parabolic fixed point of the cusp and take the upper half space model of $X = \mathbb{H}^n$ with z as ∞ .

Denote the maximal parabolic subgroup of Γ corresponding to z by G' (i.e. $G' = \Gamma_z$). If $G < \Gamma$ is a finite group fixing z and $\text{Fix}(G) = \bigcap_{g \in G} \text{Fix}(g)$ its singular submanifold, we have $\text{Fix}(G) \in \Sigma(G')$. Note that $\text{Fix}(G)$ is either equal to \mathbb{H}^n (namely if $G = \{\text{id}\}$) or given by the (non-empty) intersection of Euclidean affine hyperplanes perpendicular to the boundary $\partial\mathbb{H}^n = \mathbb{R}^{n-1} \times \{0\}$. In the present situation we say that the singular submanifold $\text{Fix}(G)$ contains the parabolic fixed point z .

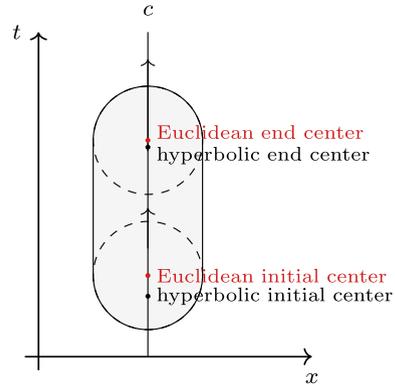
Let $c : \mathbb{R} \rightarrow \mathbb{R}^n$ be the hyperbolic geodesic with endpoints $c(\infty) = z$ and $c(-\infty) = (x^{(0)}, 0)$ for some $x^{(0)} \in \mathbb{R}^{n-1}$, so $c([t'_0, t'_1])$ (for $t'_1 \geq t'_0 > 0$) is a geodesic section. A **stretched ball** U of Euclidean radius $r > 0$ along this geodesic section is then defined by

$$U := \{(x, t) \in \mathbb{H}^n : d_{\text{eucl}}((x, t), c(t')) < r \text{ for some } t' \in [t'_0, t'_1]\},$$

where d_{eucl} denotes the Euclidean distance in \mathbb{R}^n . Since Euclidean balls are just hyperbolic balls (with different radius and center), we can also see such a U as a union of hyperbolic balls around points of c . Of those points, the one with the smallest t -coordinate will be the (hyperbolic) initial center, whereas the one with the largest t -coordinate will be the (hyperbolic) end center. Obviously, the hyperbolic radius decreases monotonically while going from the initial center to the end center. The (hyperbolic) radius at the (hyperbolic) initial center will be called the (hyperbolic) initial radius, and similarly we get the (hyperbolic) end radius; the corresponding balls will be the (hyperbolic) initial ball and the (hyperbolic) end ball, respectively. Note that the hyperbolic initial and end center do not coincide with the Euclidean initial and end center $c(t'_0)$ and $c(t'_1)$, respectively; despite this, the hyperbolic initial and end balls are the same as the Euclidean initial and end balls. For a better understanding, the construction is pictured in Fig. 1. As a union of open balls, U is itself open. U is also convex, as it is the intersection of its supporting half spaces (see [5] Proposition II.1.4.1). Finally, note that ordinary balls are special cases of stretched balls for $t'_0 = t'_1$.

Recall that G' denotes the maximal parabolic subgroup of Γ corresponding to z . We will always assume that U is entirely contained in the $\varepsilon(n)$ -thin part of \mathbb{H}^n with respect to G' , i.e. $G'_{\varepsilon(n)}(x) = \{g \in G' : d_g(x) < \varepsilon(n)\}$ is infinite for all $x \in U$. This way, we can reduce the group action of Γ on and around U to the action of G' , which behaves nicely with respect to the construction of U : since the geodesics going to z , which fiber U , are permuted by

Fig. 1 Construction of a stretched ball in the upper half space model of \mathbb{H}^n . The dashed balls are the initial ball (bottom) and the end ball (top)



the elements of $g \in G'$ and G preserves the horospheres around z , we see that gU is also a stretched ball. Now gU has the same radius as U ; additionally, if y denotes the initial center of U , then gy is the initial center of gU , and y and gy lie in the same horosphere around z .

From now on and if not explicitly stated otherwise, when talking about initial/end balls and initial/end centers, we will always mean the hyperbolic ones. The following lemma tells us that if the hyperbolic initial centers of two intersecting stretched balls lie in the same horosphere, we get the following estimate on their distance.

Lemma 5 *Let U and U' be stretched balls with initial balls $B_\mu(y)$ and $B_{\mu'}(y')$. If y and y' lie in the same horosphere around z and $U \cap U' \neq \emptyset$, then $d(y, y') < 2\mu + 2\mu'$.*

Proof Let r denote the Euclidean radius in the construction of U and $y = (x_y, t_y)$; the horosphere HS around z containing U is thus given by $\mathbb{R}^{n-1} \times \{t_y\}$. The hyperbolic initial ball $B_\mu(y)$ of U can also be seen as a Euclidean ball of radius r around some point $\tilde{y} = (x_{\tilde{y}}, t_{\tilde{y}})$, where $x_{\tilde{y}} = x_y$ and $t_{\tilde{y}} > t_y$. With the usual distance formulas, we see that the hyperbolic initial radius μ is given by

$$\mu \stackrel{!}{=} d((x_y, t_{\tilde{y}} + r), (x_y, t_y)) = \ln \left(\frac{t_{\tilde{y}} + r}{t_y} \right),$$

so $t_{\tilde{y}} + r = t_y \cdot e^\mu$. Using $t_{\tilde{y}} > t_y$ we conclude

$$r = t_y \cdot e^\mu - t_{\tilde{y}} < t_y \cdot e^\mu - t_y = t_y \cdot (e^\mu - 1).$$

Recall that for a point $(x, t_y) \in HS$, we have

$$d((x, t_y), (x_y, t_y)) = 2 \cdot \operatorname{arsinh} \left(\frac{\|x - x_y\|}{2 \cdot t_y} \right).$$

If (x, t_y) also lies in the Euclidean r -ball E around y , then $\|x - x_y\| < r < t_y \cdot (e^\mu - 1)$ and thus

$$d((x, t_y), (x_y, t_y)) < 2 \cdot \operatorname{arsinh} \left(\frac{t_y \cdot (e^\mu - 1)}{2 \cdot t_y} \right) = 2 \cdot \operatorname{arsinh} \left(\frac{e^\mu - 1}{2} \right).$$

As $\mu > 0$, also $(e^\mu - 1)/2 < (e^\mu - e^{-\mu})/2 = \sinh(\mu)$ holds. By monotonicity of arsinh , we get $\operatorname{arsinh}((e^\mu - 1)/2) < \operatorname{arsinh}(\sinh(\mu)) = \mu$, so

$$d((x, t_y), (x_y, t_y)) < 2 \cdot \operatorname{arsinh} \left(\frac{e^\mu - 1}{2} \right) < 2 \cdot \mu.$$

Fig. 2 Situation in the proof of Lemma 5. The dashed balls are the Euclidean r - and r' -balls E and E' around the hyperbolic initial centers y and y' , respectively

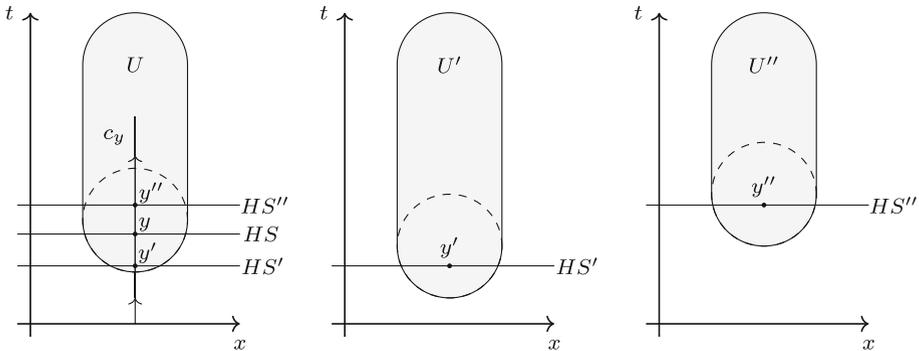
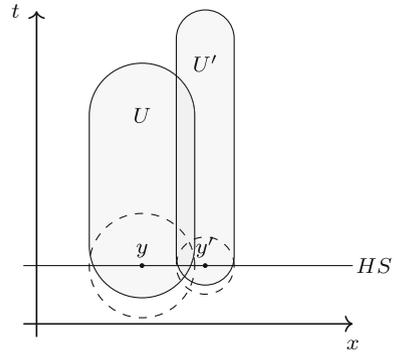


Fig. 3 In the left picture, we see the stretched ball U with initial center $y \in HS$ as well as the intersections $y' \in HS'$ and $y'' \in HS''$ of c_y ; these are used as the initial centers of the comparison balls U' and U'' , respectively. The dashed balls are the initial balls of U , U' and U''

Hence every point of $HS \cap E$ has a hyperbolic distance $< 2\mu$ to y .

The same arguments hold for U' , thus every point in the Euclidean r' -ball E' around y' , which also lies in the same horosphere HS' around z as y' , has hyperbolic distance $< 2\mu'$ to y' ; here, r' denotes the Euclidean radius in the construction of U' .

Now observe that since $U \cap U' \neq \emptyset$, we have $E \cap E' \neq \emptyset$. By assumption, $HS \ni y$ and $HS' \ni y'$ coincide, hence there is a point $y'' \in HS \cap E \cap E'$ (also see Fig. 2). By the above arguments we have $d(y, y'') < 2\mu$ and $d(y', y'') < 2\mu'$, thus

$$d(y, y') < 2\mu + 2\mu'.$$

□

The above Lemma 5 assumes that the initial centers lie in the *same* horosphere, so we will need another construction to be able to compare stretched balls with initial centers in *different* horospheres. If $B_\mu(y)$ is the initial ball of a stretched ball U and r the Euclidean radius of U , let c_y be the geodesic from y to the parabolic fixed point z . For a horosphere HS' other than the horosphere $HS \ni y$, let y' be the unique intersection of c_y with HS' . The **comparison ball** U' of U at height HS' is defined as the stretched ball with (hyperbolic) initial center y' , the same end ball as U and using the same Euclidean radius r in the construction, see Fig. 3.

In other words, we just choose the parameter t'_0 differently. In our later applications, the comparison ball U' will always be well-defined, i.e. the new initial center y' will always be farther away from z than the end center (which coincides with the end center of U).

We immediately see that the comparison ball U' intersects precisely those geodesics to z that also intersect with U . Since the distance between these geodesics decreases on its way to z , we get the following estimates for the hyperbolic initial radius μ' of U' : if HS' is closer to z than HS , we have $\mu' \leq \mu$; whereas if HS' is farther away from z than HS , we have $\mu' \geq \mu$. Also note that U intersects another stretched ball U'' with initial center $y'' \in HS''$ if and only if the comparison ball U' of U at height HS'' intersects U'' .

Similar to the case of ordinary balls (compare Lemma 4), suitably chosen stretched balls are foldable:

Lemma 6 *Let $Y \in \Sigma_i$ be a singular submanifold containing the parabolic fixed point z . Moreover, let $y \in Y \setminus S_{< i}$ and $\mu > 0$ be sufficiently small, such that Y is fixed pointwise by $\Gamma_{4\mu}(y)$. If U denotes a stretched ball with (hyperbolic) initial ball $B_\mu(y)$, then U is Y -foldable.*

Proof We have already seen that stretched balls are open and convex, so let's turn to the precise invariance of U . Since we always assume that U lies in the $\varepsilon(n)$ -thin part of \mathbb{H}^n w.r.t. G (where G is the maximal parabolic subgroup corresponding to the parabolic fixed point z), and that part is (seen as a component of the $\varepsilon(n)$ -thin part w.r.t. Γ) precisely invariant under Γ , it only remains to show the precise invariance of U w.r.t. elements of G . Hence, we have to check if for all $g \in G$, the condition $gU \cap U \neq \emptyset$ already implies $gU = U$. As $g \in G$, we know that gU is a stretched ball with same initial radius as U and initial center gy in the same horosphere around z as y . Since $gU \cap U \neq \emptyset$, we can apply Lemma 5 and get

$$d_g(y) = d(y, gy) < 2\mu + 2\mu = 4\mu,$$

i.e. $g \in \Gamma_{4\mu}(y)$. By assumption, this means that g fixes Y pointwise, thus $gy = y$ and hence $gU = U$.

As a next step, we will prove that Y is fixed pointwise by Γ_U . Note that in our situation, we have $\Gamma_U = G_U$, where $G_U = \{g \in G : gU = U\}$. By the above arguments we know that $g \in G_U$ already implies $d_g(y) < 4\mu$, so $g \in \Gamma_{4\mu}(y)$; hence by the assumption, g fixes Y pointwise.

We will now prove $\pi_Y(U) \subseteq U$. As every $u \in U$ lies in a suitable hyperbolic ball $B_{\mu_0}(y_0) \subseteq U$ of radius μ_0 around some $y_0 \in Y$ (recall that the geodesic c in the construction of U is entirely contained in Y), we get (using that the projection to Y is distance-decreasing, compare [2] chapter 1.6)

$$d(\pi_Y(u), \pi_Y(y_0)) \leq d(u, y_0) < \mu_0.$$

Since $\pi_Y(y_0) = y_0$, this implies $\pi_Y(u) \in B_{\mu_0}(y_0) \subseteq U$.

In the last step, we have to show that $\pi(U \cap Y) \subseteq \mathbb{H}^n/\Gamma$ is contractible; by Lemma 3, this is equivalent to $U \cap Y \subseteq \mathbb{H}^n$ being contractible. Now recall that U and Y are convex, so $U \cap Y$ is convex and thus contractible. □

We also need a similar statement for the intersection of (the images of) several stretched balls:

Lemma 7 *Let U_1 be a Y -foldable stretched ball as in Lemma 6, i.e. with initial center $y_1 \in Y \setminus S_{< i}$, and let U_2, \dots, U_k be foldable stretched balls with initial centers $y_2, \dots, y_k \in Y$; note that we don't assume $y_2, \dots, y_k \in Y \setminus S_{< i}$ and hence the U_2, \dots, U_k will in general not be Y -foldable. If $U := \bigcap_{j=1}^k U_j \neq \emptyset$, then:*

1. U is Y -foldable.

2. If the hyperbolic initial radii μ_j of the U_j at the initial centers y_j are chosen in a way such that y_j is fixed by $\Gamma_{8\mu_j}(y_j)$ ($j = 1, \dots, k$), then $\pi(U) = \bigcap_{j=1}^k \pi(U_j)$. Hence the intersection $\bigcap_{j=1}^k \pi(U_j)$ of the $\pi(U_j)$ is folded and thus contractible.

Proof Many ideas are similar to those in the proof of [7] Proposition 4.10 and its preceding text.

1. As an intersection of open, convex and precisely invariant sets, U itself is open, convex and precisely invariant. Since Γ_{U_1} fixes Y pointwise (because U_1 is Y -foldable), we only need to show that Γ_U is a subgroup of Γ_{U_1} to conclude that Γ_U fixes Y pointwise. Let $\gamma \in \Gamma_U$. As $\gamma U = U$ and $U \subseteq U_1$, we see that

$$\gamma U_1 \supseteq \gamma U = U \subseteq U_1,$$

so $\gamma U_1 \cap U_1 \neq \emptyset$ (because $U \neq \emptyset$). Knowing that U_1 is precisely invariant, this leads to $\gamma U_1 = U_1$, hence $\gamma \in \Gamma_{U_1}$.

Next, let us check if $\pi_Y(U) \subseteq U$. Let c_j denote the geodesic from the initial center $y_j \in Y$ of U_j to the parabolic fixed point z ; note that all the c_j are entirely contained in Y . Hence the U_j are constructed as the union of hyperbolic balls around points of $c_j \subseteq Y$. So if $u \in U$, there are points $y_0^{(j)} \in Y$ and radii $\mu_0^{(j)}$ ($j = 1, \dots, k$) such that u lies in every ball $B_{\mu_0^{(j)}}(y_0^{(j)}) =: B_j \subseteq U_j$. Just as in the proof of Lemma 6 we conclude that $\pi_Y(u) \in B_j$ for all $j = 1, \dots, k$. Thus

$$\pi_Y(u) \in \bigcap_{j=1}^k B_j \subseteq \bigcap_{j=1}^k U_j = U,$$

proving $\pi_Y(U) \subseteq U$.

It remains to show that $\pi(U \cap Y)$ is contractible. Again, this is equivalent to $U \cap Y$ being contractible (by Lemma 3), which itself is a consequence of the convexity of U and Y .

2. Similar to [7] Proposition 4.10, we see that a preimage of $\bigcap_{j=1}^k \pi(U_j)$ in \mathbb{H}^n is just a union of intersections $\bigcap_{j=1}^k \gamma_j U_j$ ($\gamma_j \in \Gamma$), so it only remains to show that every such non-empty intersection $\bigcap_{j=1}^k \gamma_j U_j$ arises as a translate of $\bigcap_{j=1}^k U_j$ under a suitable element of Γ . Let r_j be the Euclidean radius in the construction of U_j and choose $j_0 \in \{1, \dots, k\}$ such that $r_{j_0} \leq r_j$ for all $j = 1, \dots, k$. After a possible translation of $\bigcap_{j=1}^k \gamma_j U_j$ by $\gamma_{j_0}^{-1}$, we can assume that $\gamma_{j_0} = \text{id}$. As $\bigcap_{j=1}^k \gamma_j U_j \neq \emptyset$ by assumption, U_{j_0} intersects all the other $\gamma_j U_j$. We let $U_{j_0}^{(j)}$ denote the comparison ball of U_{j_0} at the height of the horosphere $HS_j \ni y_j$ (for $j = 1, \dots, k$), and $\mu_{j_0}^{(j)}$ the (hyperbolic) initial radius around the (hyperbolic) initial center $y_{j_0}^{(j)}$. Since U_{j_0} intersects U_j and $\gamma_j U_j$, also $U_{j_0}^{(j)}$ intersects U_j and $\gamma_j U_j$ (see definition of the comparison balls). As the Euclidean radius r_{j_0} was chosen to be minimal—and the Euclidean radius of the comparison ball coincides with r_{j_0} —, the hyperbolic initial radius $\mu_{j_0}^{(j)}$ also has to satisfy $\mu_{j_0}^{(j)} \leq \mu_j$ for all $j = 1, \dots, k$. Using Lemma 5, we get $d(y_j, y_{j_0}^{(j)}) < 2\mu_j + 2\mu_{j_0}^{(j)} \leq 4\mu_j$ and $d(\gamma_j y_j, y_{j_0}^{(j)}) < 2\mu_j + 2\mu_{j_0}^{(j)} \leq 4\mu_j$, thus

$$d_{\gamma_j}(y_j) = d(y_j, \gamma_j y_j) < 8\mu_j$$

for every j . Hence $\gamma_j \in \Gamma_{8\mu_j}(y_j)$ and so by assumption, γ_j fixes y_j . We conclude $\gamma_j U_j = U_j$ and thus

$$\bigcap_{j=1}^k \gamma_j U_j = \bigcap_{j=1}^k U_j.$$

So $\bigcap_{j=1}^k \pi(U_j) = \pi(U)$, the latter set being the image of the (by 1.) Y -foldable set U . Hence $\bigcap_{j=1}^k \pi(U_j)$ is folded and thus contractible.

□

4.3 Constructing the cover

Our next goal is to define a suitable cover of the thick part M_+ . To this end, we will extend the construction of Samet [7] Theorem 4.2: while the cover given there is indeed a good cover and contains M_+ , it is far from being homotopy equivalent to M_+ ; it goes well beyond M_+ and, in general, will have gaps outside M_+ . By gaps we mean that coming from the thin part (in a suitable way), we might enter and leave the cover several times before entering it for a last time and staying in M_+ . To fill these gaps, we will use the previously defined stretched balls, eventually giving us a good cover that is also homotopy equivalent to M_+ .

To achieve all the said properties, a very delicate choice of positions and sizes of the covering sets will be required; the following two lemmas are a major tool for this.

Lemma 8 ([7] Proposition 4.6) *For every $\varepsilon_1 > 0$ there is $\varepsilon_2 = \varepsilon_2(\varepsilon_1) > 0$ with the following property. Let $Y_1, Y_2 \in \Sigma(\Gamma)$ with $i = \dim(Y_2) \leq \dim(Y_1)$ and $y_j \in Y_j \cap X_+$ ($j = 1, 2$), such that $d(y_1, S_{<i}(\Gamma)) \geq \varepsilon_1$ or $d(y_2, S_{<i}(\Gamma)) \geq \varepsilon_1$. If $d(y_1, y_2) < \varepsilon_2$, then already $Y_2 \subseteq Y_1$.*

Lemma 9 ([7] Proposition 4.7) *For every $\varepsilon_1 > 0$ there is $\varepsilon_3 = \varepsilon_3(\varepsilon_1) > 0$ with the following property. Let $Y \in \Sigma(\Gamma)$, $y \in Y \cap X_+$ and $i = \dim(Y)$. If $d(y, S_{<i}(\Gamma)) > \varepsilon_1$, then Y is fixed pointwise by every element of $\Gamma_{\varepsilon_1}(y)$.*

We have let our X_+ take the role of $X_{\geq \varepsilon, m}$ in [7], so we will always have to assume (without restriction) that $\varepsilon_2(\cdot), \varepsilon_3(\cdot) \leq \varepsilon(n)/(2M(n))$; recall that $X_+ \subseteq X_{\geq \varepsilon, \eta}$ with $\varepsilon = \varepsilon(n)/(2M(n))$. Define

$$\mu_{-1} := \min \left(\frac{\varepsilon(n)}{64}, \nu \right),$$

where ν is the minimal displacement of hyperbolic isometries of Γ . We then iteratively define $\mu_{-1} > \mu_0 > \dots > \mu_n$ by

$$\mu_{i+1} := \min \left(\frac{\varepsilon_2(\mu_i)}{12}, \frac{\varepsilon_3(\mu_i)}{24}, \frac{\mu_i}{12} \right),$$

with $\varepsilon_2(\cdot)$ and $\varepsilon_3(\cdot)$ given by Lemma 8 and 9. Furthermore, let

$$\begin{aligned} \mathcal{D}_0 &:= \text{maximal } \mu_0\text{-discrete subset in } \overline{(M'_+)_8\mu_0} \cap \pi(S_0), \quad \text{and} \\ \mathcal{D}_i &:= \text{maximal } \mu_i\text{-discrete subset in } \left(\overline{(M'_+)_8\mu_i} \cap \pi(S_i) \right) \setminus \bigcup_{j < i} (\pi(S_j))_{\mu_j} \end{aligned}$$

for $i > 0$. Without restriction, we can assume that \mathcal{D}_0 contains a maximal μ_0 -discrete subset of

$$\partial(M'_+)^{8\mu_0} \cap \pi(S_0) \subseteq \overline{(M'_+)^{8\mu_0}} \cap \pi(S_0),$$

and similarly \mathcal{D}_i (for $i > 0$) contains a maximal μ_i -discrete subset of

$$(\partial(M'_+)^{8\mu_i} \cap \pi(S_i)) \setminus \bigcup_{j < i} (\pi(S_j))_{\mu_j} \subseteq \overline{((M'_+)^{8\mu_i} \cap \pi(S_i))} \setminus \bigcup_{j < i} (\pi(S_j))_{\mu_j},$$

because maximal μ_i -discrete subsets in the sets on the left hand side can be extended to maximal μ_i -discrete subsets in the sets on the right hand side. Denote by

$$\mathcal{D} := \bigcup_{i=0}^n \mathcal{D}_i$$

the set of all centers. In a first step, let

$$\mathcal{B}' := \{B_{3\mu_i}^M(x) : x \in \mathcal{D}_i \text{ for some } i \in \{0, \dots, n\}\}$$

be the set of all ordinary balls (in M) around the points of \mathcal{D} . The following lemma says that this already covers the shrunken thick part.

Lemma 10 *The sets of \mathcal{B}' form a cover of M'_+ , i.e.*

$$M'_+ \subseteq \bigcup_{B \in \mathcal{B}'} B.$$

Proof The proof contains ideas of [7] Theorem 4.2 step 1). Let $x \in M'_+$ and choose $i \in \{0, \dots, n\}$ minimal, such that $d(x, \pi(S_i)) < 2\mu_i$; since $S_n = X$, i.e. $\pi(S_n) = M$, this is always possible. Hence there is $y \in \pi(S_i)$ with $d(x, y) < 2\mu_i$.

We claim that $y \notin \bigcup_{j < i} (\pi(S_j))_{\mu_j}$. Assume the contrary, then there would be $j < i$ and $z \in \pi(S_j)$ with $d(y, z) < \mu_j$. As $\mu_i \leq \mu_j/12$ (because $j < i$), this leads to

$$d(x, z) \leq d(x, y) + d(y, z) < 2\mu_i + \mu_j < 2\mu_j,$$

i.e. $d(x, \pi(S_j)) < 2\mu_j$; since $j < i$, this contradicts the minimality of i . Thus $y \notin \bigcup_{j < i} (\pi(S_j))_{\mu_j}$.

Using $d(x, y) < 2\mu_i < 8\mu_i$, we also get $d(y, M'_+) \leq d(y, x) < 8\mu_i$, so $y \in \overline{(M'_+)^{8\mu_i}}$ and hence

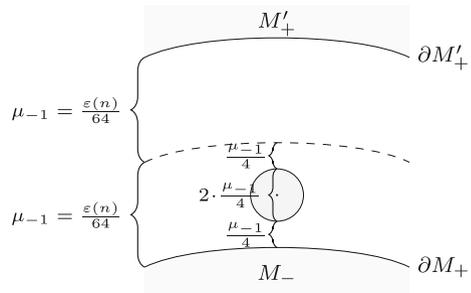
$$y \in \overline{((M'_+)^{8\mu_i} \cap \pi(S_i))} \setminus \bigcup_{j < i} (\pi(S_j))_{\mu_j}.$$

By definition, \mathcal{D}_i lies in that set as a maximal μ_i -discrete subset, so there must be a $y' \in \mathcal{D}_i$ with $d(y, y') < \mu_i$. We conclude

$$d(x, y') \leq d(x, y) + d(y, y') < 2\mu_i + \mu_i = 3\mu_i,$$

i.e. $x \in B_{3\mu_i}^M(y') \in \mathcal{B}'$, which finishes the proof. □

Fig. 4 Choosing a suitable stretching length. It is enough to specify the position of the end center (here: the center of the gray ball). Ordinary balls of \mathcal{B}' (equivalently: their preimages in X) can not exceed the dashed line, while end balls of stretched balls will always exceed the dashed line (see Lemma 12)



Defining the stretching

To fill in possible gaps of the cover outside of M_+ (or using another interpretation: to make the cover stable¹⁰ under the flow to M_+), we will have to stretch some of the balls of the cover; this simply means that we replace that ordinary ball by (the image of) a suitable stretched ball with that ball as its initial ball. Recall that in the present situation, $\pi(B_r^X(y)) = B_r^M(\pi(y))$ for all $y \in X$ and $r > 0$, so for the sake of simplicity, the image of a stretched ball in X will also be called a stretched ball (in M).

Defining the stretching will amount to two tasks: first, we have to state which balls should be stretched; and second, we have to say for how long to stretch these balls, i.e. define the stretching length¹¹. The general situation and strategy is summarized in Fig. 4.

- *Balls to be stretched:* If $x \in \mathcal{D}_i \cap \partial(M'_+)_8\mu_i$, then (by choice of M'_+) $B := B_{3\mu_i}^M(x)$ is contained in the $\varepsilon(n)$ -thin part of M . We will replace such a ball by a corresponding stretched ball with initial center x .
- *Stretching length:* Let x be the center of a ball $B_{3\mu_i}^M(x) \in \mathcal{B}'$ which should be stretched; to specify the stretching length, it is enough to define the position of the end center¹². Let c_x be the flow geodesic of x ; we will flow along c_x in direction of ∂M_+ until $c_x(t)$ is precisely $\mu_{-1}/2$ away from ∂M_+ , but still lies in M_+ , i.e. for this t we have $c_x(t) \in M_+$ and $d(c_x(t), \partial M_+) = \mu_{-1}/2$.

With the stretching given as above, we will turn the set \mathcal{B}' into our final set of covering sets \mathcal{B} . Observe that the set $\mathcal{D}_i \cap \partial(M'_+)_8\mu_i$ of initial centers of the stretched balls contains precisely the points of \mathcal{D}_i with maximal distance to M'_+ (namely $8\mu_i$).

Effects of the stretching

Before proving the desired properties of the cover, we will further investigate the effects that the stretching had on the balls that were stretched. The following lemma states that with this choice of stretching, we have achieved a monotonicity of the radii of the (stretched) balls; in simple terms, it means that the initial centers of stretched balls with large radius are closer

¹⁰ We say that a set is stable under the flow if the flow does not leave that set after entering it for the first time.

¹¹ As measured by the (hyperbolic) length of the (Euclidean) line segment in the construction of the stretched balls or equivalently by the (hyperbolic) distance between the initial center and the end center.

¹² We stated the construction of stretched balls only in X , but the construction in M can be translated to X by lifting the end center.

to the thin part than the (initial) centers of arbitrary balls of smaller radius, if these balls intersect¹³:

Lemma 11 *Let $y_i \in \widetilde{\mathcal{D}}_i$ be an initial center of a stretched ball $B_i \in \widetilde{\mathcal{B}}$ with initial ball $B_{3\mu_i}^X(y_i)$ and $B_j \in \widetilde{\mathcal{B}}$ be an arbitrary (i.e. ordinary or stretched) ball with (initial) center $B_{3\mu_j}^X(y_j)$ for $y_j \in \widetilde{\mathcal{D}}_j$. By $HS_i \ni y_i$ and $HS_j \ni y_j$ we denote the corresponding horospheres around z . If $B_i \cap B_j \neq \emptyset$ and $3\mu_i > 3\mu_j$ (so equivalently: $i < j$), then HS_i is closer to z than HS_j .*

Proof We already see from the formulation of the statement that we will treat the situation in X , where $\widetilde{\mathcal{D}}_i$ and $\widetilde{\mathcal{B}}$ denote the lifts (to X) of the corresponding sets in M .

By the choice of centers, we have $d(y_i, X'_+) = 8\mu_i$; moreover, $d(y_j, X'_+) \leq 8\mu_j$. Let $x \in B_i \cap B_j$. If the (initial) balls $B_{3\mu_i}^X(y_i)$ and $B_{3\mu_j}^X(y_j)$ would intersect (in x , without restriction), then we would reach the contradiction

$$8\mu_i = d(y_i, X'_+) \leq d(y_i, x) + d(x, y_j) + d(y_j, X'_+) < 3\mu_i + 3\mu_j + 8\mu_j < 4\mu_i$$

(note that $12\mu_j \leq \mu_i$, since $i < j$). Let c_x be the flow geodesic of x , with parametrization $c_x(0) = x$ and $c_x(-\infty) = z$ (i.e. flowing to X'_+). We've already seen that the (initial) balls can not intersect, but c_x has to intersect both (initial) balls. We will distinguish between the two possible cases:

In the first case, c_x leaves the smaller (initial) ball $B_{3\mu_j}^X(y_j)$ before entering the larger initial ball $B_{3\mu_i}^X(y_i)$. Observe that all points of $B_{3\mu_j}^X(y_j)$ have distance $< 3\mu_j + 8\mu_j < 12\mu_j \leq \mu_i$ to X'_+ , because $d(y_j, X'_+) \leq 8\mu_j$. As the distance to X'_+ decreases along the flow, this means that the entry point of c_x in $B_{3\mu_i}^X(y_i)$ also has distance $< \mu_i$ to X'_+ . This is a contradiction, since $d(y_i, X'_+) = 8\mu_i$ already implies that all points $y \in B_{3\mu_i}^X(y_i)$ satisfy $d(y, X'_+) \geq 8\mu_i - 3\mu_i = 5\mu_i$. Consequently, this case can not happen.

In the second case, c_x leaves the larger initial ball $B_{3\mu_i}^X(y_i)$ before entering the smaller (initial) ball $B_{3\mu_j}^X(y_j)$. Using the above estimates for the distance to X'_+ , we know that along c_x , we have to flow at least $5\mu_i - \mu_i = 4\mu_i$ after leaving $B_{3\mu_i}^X(y_i)$ before we enter $B_{3\mu_j}^X(y_j)$. If $x_i = c_x(t_i)$ denotes the exit point of $B_{3\mu_i}^X(y_i)$ and $x_j = c_x(t_j)$ the entry point in $B_{3\mu_j}^X(y_j)$, we see that $t_j \geq t_i + 4\mu_i$

We will now use the upper half space model (again with z as point ∞), so horospheres around z are (Euclidean) hyperplanes $\mathbb{R}^{n-1} \times \{t\} \subseteq \mathbb{H}^n$ and the flow geodesics are (Euclidean) lines going away from ∞ and perpendicular to $\mathbb{R}^{n-1} \times \{0\}$. So the exit point x_i is on the lower half of the initial ball $B_{3\mu_i}^X(y_i)$ (seen as a Euclidean ball), i.e. it has a t -coordinate smaller than the t -coordinate of the Euclidean center of $B_{3\mu_i}^X(y_i)$. Assuming $y_i = (0, \dots, 0, 1)$ (without restriction), by similar arguments as in the proof of Lemma 5 and using the usual distance formulas we can deduce that the t -coordinate of the Euclidean center of $B_{3\mu_i}^X(y_i)$ is $\cosh(3\mu_i) > 1$. Similarly, the t -coordinate of a point that is reached after flowing $4\mu_i$ (in hyperbolic length)—and starting from a t -coordinate of $\cosh(3\mu_i)$ —can be computed to be $e^{-\mu_i} < 1$ (again using the usual distance formulas). By the above arguments, this value is an upper bound on the t -coordinate of any entry point of the ball $B_{3\mu_j}^X(y_j)$; so an arbitrary entry point of c_x into $B_{3\mu_j}^X(y_j)$ has a t -coordinate < 1 . Since the t -coordinate of the center y_j has

¹³ Note that we will only have to prove this monotonicity for intersecting balls. This is also the reason why we defined the stretching before stating this lemma: in order to see which balls (stretched and/or ordinary) intersect, we have to know which balls to stretch and how to stretch them.

an even smaller t -coordinate than the entry point, y_j has a t -coordinate < 1 . Thus y_j lies in a horosphere $\mathbb{R}^{n-1} \times \{t\}$ for some $t < 1$, whereas y_i lies in the horosphere $\mathbb{R}^{n-1} \times \{1\}$. \square

Although stretching has the effect of extending the cover in the direction of the thin part, the resulting cover will still be contained in the (ordinary) thick part; furthermore, there is some level which ordinary balls will never exceed, but stretched balls will always exceed:

Lemma 12 *With the above choice of stretching, we have:*

1. *All the sets of \mathcal{B} are contained in the (ordinary) thick part, i.e.*

$$\bigcup_{B \in \mathcal{B}} B \subseteq M_+.$$

2. *If B is the end ball of a stretched ball in \mathcal{B} , then every point of B has distance less than μ_{-1} from ∂M_+ .*
 3. *If B' is an ordinary ball (i.e. non-stretched) in \mathcal{B} , then every point of B' has distance more than μ_{-1} from ∂M_+ .*

Proof We first remark that by the monotonicity $\mu_i > \mu_{i+1}$, the balls of \mathcal{B}' with centers $x \in \mathcal{D}_0$ are those which can lie the furthest away from M'_+ ; for these x we have $d(x, M'_+) \leq 8\mu_0$, so for arbitrary $y \in B_{3\mu_0}^M(x)$ we get

$$d(y, M'_+) < 8\mu_0 + 3\mu_0 = 11\mu_0 < 12\mu_0 \leq \mu_{-1},$$

where $\mu_{-1} = \varepsilon(n)/64$ is half the size $\varepsilon(n)/32$ of the shrinking of M_+ onto M'_+ . Hence every point of $\bigcup_{B \in \mathcal{B}'} B$ has a distance $> \mu_{-1}$ to the boundary ∂M_+ , which proves the third statement.

This also yields $\bigcup_{B \in \mathcal{B}'} B \subseteq M_+$, so it remains to show that this property is preserved when stretching the balls, i.e. when going from \mathcal{B}' to \mathcal{B} . Let y be the center of a ball $B_{3\mu_i}^M(y) \in \mathcal{B}'$ that will be stretched; in particular, y will be the initial center of the corresponding stretched ball. Observe that the hyperbolic radius at the end center is smaller than the hyperbolic radius at the initial center, and since

$$3\mu_i \leq 3\mu_0 \leq \frac{3\mu_{-1}}{12} = \frac{\mu_{-1}}{4},$$

the latter is bounded by $\mu_{-1}/4$. Hence the hyperbolic end radius is also smaller than $\mu_{-1}/4$. Recall that by the definition of the stretching length, the end center will have distance $\mu_{-1}/2$ to ∂M_+ . Thus every point of the end ball has distance $\geq \mu_{-1}/2 - \mu_{-1}/4 = \mu_{-1}/4$ from ∂M_+ , so the end ball itself is completely contained in M_+ . Consequently, the entire stretched ball lies inside M_+ , which yields the first statement. Also note that the end ball has distance $\geq \mu_{-1}/4$ to any ordinary ball of \mathcal{B} , because the distance of points of the ordinary balls to ∂M_+ is always $\geq \mu_{-1}$.

For the second statement, let y' be the end center of the stretched ball and y'' an arbitrary point in the end ball. Using the above bound on the hyperbolic end radius, we get $d(y', y'') < \mu_{-1}/4$; by the definition of the stretching length, we also have $d(y', \partial M_+) = \mu_{-1}/2$. This yields $d(y'', \partial M_+) < \mu_{-1}/4 + \mu_{-1}/2 < \mu_{-1}$, proving the second statement. \square

The stretching length was originally defined via the distance to ∂M_+ ; this means that so far, we only know when to stop stretching, but not how long we have stretched in absolute terms. To get such a value for the stretching length in absolute terms, we have to check how long we have to flow to realize said distance to ∂M_+ .

Lemma 13 *Every stretched ball $B \in \mathcal{B}$ with initial ball $B_{3\mu_i}^M(y)$ is contained in a ball $B_{R(n,\nu)}^M(y)$, where $R(n,\nu) > 0$ is a constant only depending on the dimension n and the constant ν .*

Proof We will omit any further mentioning of ν in the proof, as the dependency of R on ν only comes from the definition of μ_{-1} as the minimum of $\varepsilon(n)/64$ and ν ; here, the dependency on $\varepsilon(n)/64$ contains the interesting information on the correct size of R , whereas ν is only needed to account for the fact that we defined our thick part using ν in order to have no tubes in the thin part.

Let c_y be the flow geodesic of $y = c_y(0)$. By the definition of the stretching length, we have seen that the end center $c_y(t)$ of B is characterized by $c_y(t) \in M_+$ and $d(c_y(t), \partial M_+) = \mu_{-1}/2$, for suitable $t > 0$. Given the construction of the stretched balls, every point $y' \in B$ is contained in a ball around some $c_y(t')$, where $0 \leq t' \leq t$; the radius of that ball is at most $3\mu_i$ (which is the initial radius). The latter value is bounded by μ_{-1} , so using $d(y, c_y(t')) = t'$ we get

$$B \subseteq B_{t+\mu_{-1}}^M(y).$$

As μ_{-1} only depends on n , it remains to show that t can also be controlled by only n . Recall that y is not contained in M'_+ (because it is the initial center of a stretched ball), but in the piece between $\partial M'_+$ and ∂M_+ , and similarly for the end center. If we could bound the length L of c_y between $\partial M'_+$ and ∂M_+ , we would thus get the desired bound on the above t . Note that equivalently, we can try to control the length of the lift of c_y (which we also denote by c_y) in X between $\partial X'_+$ and ∂X_+ .

Let $y_1 \in \partial X'_+$ and $y_2 \in \partial X_+$ be the intersections of c_y with $\partial X'_+$ and ∂X_+ . We reparametrize c_y such that $c_y(0) = y_1$, so we only have to find an upper bound for $t_2 > 0$, where $c_y(t_2) = y_2$. By definition of X'_+ , we have that $y_1 \in \overline{\{d_\gamma < \varepsilon(n)/2\}}_{\varepsilon(n)/32}$ for some parabolic isometry $\gamma \in \Gamma$ with fixed point $z = c_y(\infty)$. Let $y_3 = c_y(t_3)$ be the intersection of c_y with $\partial\{d_\gamma < \varepsilon(n)/2\}$, so $t_3 \geq t_2$ holds¹⁴.

Thus it remains to show: if we let $y_1 = c_y(0)$ denote the entry point of c_y in $\overline{\{d_\gamma < \varepsilon(n)/2\}}_{\varepsilon(n)/32}$ and $y_3 = c_y(t_3)$ is the entry point of c_y in $\overline{\{d_\gamma < \varepsilon(n)/2\}}$, then $t_3 \geq 0$ is bounded from above by a constant only depending on n .

Again, we choose the upper half space model with $z = c_y(\infty)$ as the point ∞ , and $y_1 = (0, 1) \in \mathbb{R}^{n-1} \times \mathbb{R}_{>0}$ without restriction. Note that γ acts on $\mathbb{R}^{n-1} \times \mathbb{R}_{>0} = \mathbb{H}^n$ as a Euclidean motion and on $\mathbb{R}_{>0}$ as the identity; we will denote the restricted action of γ on \mathbb{R}^{n-1} also by γ . By the triangle inequality we already have

$$\varepsilon' := d_\gamma((0, 1)) \leq \frac{\varepsilon(n)}{2} + 2 \cdot \frac{\varepsilon(n)}{32} < \varepsilon(n).$$

With the usual distance formulas we deduce

$$\varepsilon' = d((0, 1), (\gamma(0), 1)) = 2 \cdot \operatorname{arsinh}\left(\frac{\|0 - \gamma(0)\|}{2 \cdot 1}\right) = 2 \cdot \operatorname{arsinh}\left(\frac{\|\gamma(0)\|}{2}\right),$$

hence $\|\gamma(0)\| = 2 \sinh(\varepsilon'/2)$. Since c_y is parametrized by arc length, we get $c_y(t) = (0, e^t)$; so similar to the above, the time t displacement is $d_\gamma(c_y(t)) = d_\gamma((0, e^t)) =$

¹⁴ Either y_3 is the intersection of c_y with ∂X_+ (i.e. $t_3 = t_2$), or c_y entered ∂X_+ earlier (at the boundary of another sublevel set, i.e. $t_2 \leq t_3$).

$2 \operatorname{arsinh}(\|\gamma(0)\|/(2e^t))$. Using the above value for $\|\gamma(0)\|$ and the bound on ε' , we obtain

$$d_\gamma(c_\gamma(t)) = 2 \cdot \operatorname{arsinh} \left(\frac{\sinh \left(\frac{\varepsilon'}{2} \right)}{e^t} \right) < 2 \cdot \operatorname{arsinh} \left(\frac{\sinh \left(\frac{\varepsilon(n)}{2} \right)}{e^t} \right).$$

Hence if t is so large that the term on the right hand side is at most $\varepsilon(n)/2$, then $c_\gamma(t)$ already lies in $\{d_\gamma < \varepsilon(n)/2\}$; thus we have to solve

$$2 \cdot \operatorname{arsinh} \left(\frac{\sinh \left(\frac{\varepsilon(n)}{2} \right)}{e^t} \right) \stackrel{!}{\leq} \frac{\varepsilon(n)}{2}$$

for t , i.e.

$$\ln \left(\frac{\sinh \left(\frac{\varepsilon(n)}{2} \right)}{\sinh \left(\frac{\varepsilon(n)}{4} \right)} \right) \leq t.$$

So the left hand term tells us the maximal amount of time needed for c_γ to enter $\{d_\gamma < \varepsilon(n)/2\}$; as it only depends on n , this finishes the proof. \square

Properties of the covering

We will now prove the desired properties of the covering.

Lemma 14 *The sets of \mathcal{B} form a covering of M'_+ , i.e.*

$$M'_+ \subseteq \bigcup_{B \in \mathcal{B}} B,$$

which is stable under the flow f in the sense that if a flow geodesic enters $\bigcup_{B \in \mathcal{B}} B$, then it remains inside $\bigcup_{B \in \mathcal{B}} B$ until it meets $\partial M'_+$.

Proof Using Lemma 10, $\bigcup_{B \in \mathcal{B}'} B \subseteq \bigcup_{B \in \mathcal{B}} B$ yields the first statement.

As every flow geodesic meets $\partial M'_+$ at some point, the stability is equivalent to the following property: a flow geodesic of a point $x \in \partial M'_+$ (flowing to the thin part, i.e. in the opposite direction), is contained in $\bigcup_{B \in \mathcal{B}} B$ until it leaves this set at some point and never enters it again.

Let $x \in \partial M'_+$ and $i \in \{0, \dots, n\}$ be minimal such that

$$x \in (\partial M'_+ \cap (\pi(S_i))_{2\mu_i}) \setminus \bigcup_{j < i} (\pi(S_j))_{3\mu_j/2}.$$

This is always possible: if $i = n$, then by $S_n = M$ we get $\partial M'_+ \cap (\pi(S_n))_{2\mu_n} = \partial M'_+$, i.e. x lies in $\partial M'_+ \setminus \bigcup_{j < n} (\pi(S_j))_{3\mu_j/2}$ —proving this statement —, or $x \in \partial M'_+ \cap \bigcup_{j < n} (\pi(S_j))_{3\mu_j/2}$. In the latter case, we can chose $j < n$ minimally such that $x \in (\pi(S_j))_{3\mu_j/2}$. Since $(\pi(S_j))_{2\mu_j} \supseteq (\pi(S_j))_{3\mu_j/2}$, this means $x \in \partial M'_+ \cap (\pi(S_j))_{2\mu_j}$, hence $x \notin \bigcup_{j' < j} (\pi(S_{j'}))_{3\mu_{j'}/2}$ by minimality of j . So

$$x \in (\partial M'_+ \cap (\pi(S_j))_{2\mu_j}) \setminus \bigcup_{j' < j} (\pi(S_{j'}))_{3\mu_{j'}/2}.$$

Thus we can always find a minimal i as described above.

Let c_x be the flow geodesic of x in direction of M_- , where $c_x(0) = x$. By the above arguments, there is $x' \in \pi(S_i)$ with $d(x, x') < 2\mu_i$. Denote the corresponding flow geodesic (again in direction of M_-) of x' by $c_{x'}$, where $c_{x'}(0) = x'$. Then $d(c_x(t), c_{x'}(t)) < 2\mu_i$ for $t > 0$, because the distance between the geodesics decreases when flowing towards M_- . Note that $c_{x'}$ is entirely contained in $\pi(S_i)$, so $c_x(t) \in (\pi(S_i))_{2\mu_i}$ for $t > 0$. As $d(c_{x'}(0), M'_+) = d(x', M'_+) \leq d(x', x) < 2\mu_i$, we see that $c_{x'}(t)$ has not entered $\partial(\overline{M'_+})_{8\mu_i}$ for sufficiently small $t > 0$; moreover, $c_x(t)$ does not lie in $(\pi(S_j))_{3\mu_j/2}$ for sufficiently small $t > 0$ and $j < i$, as (similar to the above arguments) this would contradict the minimality of i . With increasing t we will now flow towards M_- (equivalently: towards ∂M_+). If $c_{x'}(t)$ meets the boundary $\partial(\overline{M'_+})_{8\mu_i}$, we say that event I happened; if on the other hand $c_x(t)$ enters some $(\pi(S_j))_{3\mu_j/2}$ for $j < i$, we say that event II happened.

As $\partial M'_+ \subseteq M'_+$, x is already contained in some $B \in \mathcal{B}$. Thus for sufficiently small $t > 0$, we know that $c_x(t)$ lies inside $\bigcup_{B \in \mathcal{B}} B$ and neither event I nor II has happened.

We will now show that for increasing t , $c_x(t)$ will remain inside $\bigcup_{B \in \mathcal{B}} B$ if neither of these events happen. To this end note that $c_{x'}(t) \notin \bigcup_{j < i} (\pi(S_j))_{\mu_j}$ has to hold, because otherwise, for some $j < i$, we would find $x'' \in \pi(S_j)$ with $d(x'', c_{x'}(t)) < \mu_j$. By

$$d(c_x(t), \pi(S_j)) \leq d(c_x(t), x'') \leq d(c_x(t), c_{x'}(t)) + d(c_{x'}(t), x'') < 2\mu_i + \mu_j < \frac{3}{2}\mu_j$$

(recall $2\mu_i \leq \mu_j/6$, as $j < i$) this would mean $c_x(t) \in (\pi(S_j))_{3\mu_j/2}$, contradicting the assumption that event II hasn't happened. Since we further assumed that event I hasn't happened—i.e. $c_{x'}(t)$ still lies inside $(M'_+)_{8\mu_i}$ —, we conclude

$$c_{x'}(t) \in \left(\overline{(M'_+)_{8\mu_i}} \cap \pi(S_i) \right) \setminus \bigcup_{j < i} (\pi(S_j))_{\mu_j}.$$

By definition, this set contains \mathcal{D}_i as a maximal μ_i -discrete subset, so there is $y \in \mathcal{D}_i$ with $d(y, c_{x'}(t)) < \mu_i$. This yields

$$d(c_x(t), y) \leq d(c_x(t), c_{x'}(t)) + d(c_{x'}(t), y) < 2\mu_i + \mu_i = 3\mu_i,$$

i.e. $c_x(t) \in B_{3\mu_i}^M(y) \subseteq \bigcup_{B \in \mathcal{B}} B$, what we wanted to show.

So it only remains to check what occurs if event I or event II happen. Let us begin with the case that event II happens first. Hence we can chose $j < i$ minimally such that $c_x(t) \in (\pi(S_j))_{3\mu_j/2} \subseteq (\pi(S_j))_{2\mu_j}$, where—by minimality of j —also $c_x(t) \notin \bigcup_{j' < j} (\pi(S_{j'}))_{3\mu_{j'}/2}$. As event II happened before event I, we get

$$d(c_x(t), M'_+) \leq d(c_x(t), c_{x'}(t)) + d(c_{x'}(t), M'_+) < 2\mu_i + 8\mu_i < \mu_j < 8\mu_j$$

(recall $j < i$), so

$$c_x(t) \in \left(\overline{(M'_+)_{8\mu_i}} \cap (\pi(S_j))_{2\mu_j} \right) \setminus \bigcup_{j' < j} (\pi(S_{j'}))_{3\mu_{j'}/2}.$$

Similar to the definition of $c_{x'}$, we can find a flow geodesic $c_{x''}$ which lies inside $\pi(S_j)$ and which fulfills $d(c_{x''}(t), c_x(t)) < 2\mu_j$ after event II happened. As

$$d(c_{x''}(t), M'_+) \leq d(c_{x''}(t), c_x(t)) + d(c_x(t), M'_+) < 2\mu_j + \mu_j < 8\mu_j,$$

the analogously defined event I for this index j and the geodesics $c_x, c_{x''}$ hasn't happened yet; by minimality of j , the analogously defined event II for this data has also not happened. So replacing i by j and $c_{x'}$ by $c_{x''}$ in the above paragraphs, we can repeat the corresponding arguments and deduce that $c_x(t)$ still lies inside $\bigcup_{B \in \mathcal{B}} B$.

Let us now assume that event I happens first; hence $c_{x'}(t) \in \partial(M'_+)_{8\mu_i}$. Again, $c_{x'}(t) \notin \bigcup_{j < i} (\pi(S_j))_{\mu_j}$, because

$$d(c_x(t), \pi(S_j)) \leq d(c_x(t), c_{x'}(t)) + d(c_{x'}(t), \pi(S_j)) < 2\mu_i + \mu_j < \frac{3}{2}\mu_j$$

(for suitable $j < i$) would once more mean that event II has already happened, contradicting our assumption. Hence

$$c_{x'}(t) \in (\partial(M'_+)_{8\mu_i} \cap \pi(S_i)) \setminus \bigcup_{j < i} (\pi(S_j))_{\mu_j}.$$

By definition, \mathcal{D}_i contains a maximal μ_i -discrete subset inside this set, so there is some $y \in \mathcal{D}_i \cap \partial(M'_+)_{8\mu_i}$ with $d(y, c_{x'}(t)) < \mu_i$. Again, $d(c_x(t), y) < 3\mu_i$, thus $c_x(t) \in B_{3\mu_i}^M(y) \subseteq \bigcup_{B \in \mathcal{B}} B$, what we wanted to show. Note that by our choice of stretched balls, $B_{3\mu_i}^M(y)$ is the initial ball of some stretched ball U ; consequently, $c_x(t)$ lies inside U from this point on. Using the second and third statement of Lemma 12, we see that we can flow along c_x (and thus remain inside $\bigcup_{B \in \mathcal{B}} B$) until we are in a distance of $< \mu_{-1}$ from ∂M_+ , where the only remaining sets of the cover \mathcal{B} are stretched balls, and no ordinary balls will be able to intersect c_x anymore on its way to the thin part. Since the stretched balls were constructed using precisely such flow geodesics as c_x , we conclude that c_x will either remain outside of $\bigcup_{B \in \mathcal{B}} B$ after leaving U , or already lie inside some other stretched ball U' ; in the latter case, we can repeat the argument, proving the statement. \square

The following lemmata state that the covering sets and their intersections are contractible.

Lemma 15 *Elements of \mathcal{B} are folded sets and thus contractible.*

Proof The proof is similar to [7] Theorem 4.2 step 2). Let $x \in \mathcal{D}_i$ and $\tilde{x} \in X$ be a preimage of x in X ; moreover, let $Y \in \Sigma_i$ be the singular submanifold containing \tilde{x} . By choice of \mathcal{D}_i we have $d(x, \pi(S_{<i})) > \mu_{i-1}$ and thus $d(\tilde{x}, S_{<i}) > \mu_{i-1}$. Note that since $x \in \mathcal{D}_i \subseteq M_+$, also $\tilde{x} \in X_+$, hence $\tilde{x} \in Y \cap X_+$.

By the definition of the μ_i we have $24\mu_i \leq \varepsilon_3(\mu_{i-1})$, so using Lemma 9 we conclude that $\Gamma_{24\mu_i}(\tilde{x})$ fixes Y pointwise. Hence the assumptions of Lemma 4 and 6 are fulfilled; using these, we see that the respective ball with (initial) center \tilde{x} is Y -foldable, which proves the statement. \square

Lemma 16 *Nonempty intersections of the elements of \mathcal{B} are folded sets and thus contractible.*

Proof The main idea is similar to [7] Theorem 4.2 step 3), although we have to work considerably harder because of the presence of stretched balls.

Let $x_1, \dots, x_k \in \mathcal{D}$, where $x_j \in \mathcal{D}_{n_j}$ for $j = 1, \dots, k$, such that the corresponding sets U_j of \mathcal{B} have a non-empty intersection. Note that these balls are either all ordinary balls, or we have an intersection of stretched balls and possibly some ordinary balls. The proof for the case of ordinary balls only is essentially the same as in [7] Theorem 4.2 step 3), so we will skip it. Without restriction, let n_1 be the maximum of the n_j ($j = 1, \dots, k$).

In a first step, we will treat the case where all the balls are stretched; so let $B_{3\mu_{n_j}}^X(\tilde{x}_j)$ be the corresponding initial balls of the U_j and note that $\tilde{x}_j \in X_+$ for all $j = 1, \dots, k$. Denote the singular submanifold containing \tilde{x}_j by $Y_j \in \Sigma_{n_j}$. By definition of \mathcal{D}_j , we have $d(\tilde{x}_j, S_{<n_j}) \geq \mu_{n_j-1}$. Observe that by maximality of n_1 , the point \tilde{x}_1 is the farthest away from the thin part among all \tilde{x}_j : for $j \in \{1, \dots, k\}$ with $n_1 > n_j$ (i.e. $\mu_{n_j} > \mu_{n_1}$) this follows from Lemma 11; and for those j with $n_1 = n_j$ (i.e. $\mu_{n_1} = \mu_{n_j}$) the point \tilde{x}_1 can be chosen

to be the one maximizing the distance to the thin part (among all \tilde{x}_j with $n_j = n_1$), without restriction. Let $U_1^{(j)}$ be the comparison ball of U_1 at the height of \tilde{x}_j , with initial center $\tilde{x}_1^{(j)}$ and initial radius $3\mu_{n_1}^{(j)}$. Note that $\tilde{x}_1^{(j)}$ lies in the same singular submanifolds as \tilde{x}_1 , so in particular $\tilde{x}_1^{(j)} \in Y_1$. Since \tilde{x}_1 had the largest distance to the thin part, we get $3\mu_{n_1}^{(j)} \leq 3\mu_{n_1}$ for all $j = 2, \dots, k$; by maximality of n_1 , we moreover have $\mu_{n_1} \leq \mu_{n_j}$ for all j , hence $3\mu_{n_1}^{(j)} \leq 3\mu_{n_j}$. The construction of the comparison balls now allow the application of Lemma 5, which yields

$$d(\tilde{x}_j, \tilde{x}_1^{(j)}) < 2 \cdot 3\mu_{n_j} + 2 \cdot 3\mu_{n_1}^{(j)} \leq 12\mu_{n_j}.$$

Since $12\mu_{n_j} \leq \varepsilon_2(\mu_{n_j-1})$ and $d(\tilde{x}_j, S_{<n_j}) \geq \mu_{n_j-1}$, we can use Lemma 8 and deduce $Y_j \subseteq Y_1$. Hence all initial centers lie in the same singular submanifold Y_1 , which already contained the center of the Y_1 -foldable stretched ball U_1 . As $\Gamma_{24\mu_{n_j}}(\tilde{x}_j)$ (for all $j = 1, \dots, k$) fixes \tilde{x}_j (see the proof of the previous Lemma 15), the statement follows after applying Lemma 7.

It remains to examine what happens if we have a mixed intersection of ordinary and stretched balls. This situation can be reduced to the case of the intersection of a single ordinary ball and a single stretched ball: if U_1, \dots, U_k are stretched balls as above such that $n_1 \geq n_j$ for all $j = 1, \dots, k$, we saw that the initial centers of the U_j also lie in the singular submanifold Y_1 , and a similar statement holds for ordinary balls U'_1, \dots, U'_k (so the centers of the U'_j lie in Y'_1 , which contained the center of U'_1); so if we can show that either $Y_1 \subseteq Y'_1$ or $Y'_1 \subseteq Y_1$, then all (initial) centers of the mixed intersection would be contained in the same singular submanifold. Applying Lemma 7, this would yield the statement (the additional assumption that $\Gamma_{24\mu_{n_j}}(\tilde{x}_j)$ fixes the centers \tilde{x}_j can again be taken from the proof of Lemma 15).

So let U_1 be a stretched ball with initial ball $B_{3\mu_{n_1}}^X(\tilde{x}_1)$ and $U_2 = B_{3\mu_{n_2}}^X(\tilde{x}_2)$ be an ordinary ball; denote the singular submanifolds containing \tilde{x}_1 and \tilde{x}_2 by $Y_1 \in \Sigma_{n_1}$ and $Y_2 \in \Sigma_{n_2}$, respectively. We want to show that one of $Y_1 \subseteq Y_2$ or $Y_2 \subseteq Y_1$ always holds.

By the choice of centers, U_1 can not intersect an ordinary ball of strictly smaller radius (compare the proof of Lemma 11), so we already have $3\mu_{n_1} \leq 3\mu_{n_2}$, i.e. $n_1 \geq n_2$. If the initial ball of U_1 intersects (the ordinary ball) U_2 , then the above argument for the intersection of ordinary balls can be used, giving $Y_2 \subseteq Y_1$, which proves the statement. Thus without restriction, we can assume that the intersection of U_1 and U_2 is outside the initial ball of U_1 . We will distinguish between two cases, depending on the proximity of \tilde{x}_1 and \tilde{x}_2 to the thin part.

- In the first case, \tilde{x}_1 is closer to the thin part than \tilde{x}_2 . In a first step, assume $n_1 = n_2$, i.e. $3\mu_{n_1} = 3\mu_{n_2}$. Note that since U_1 and U_2 intersect, U_1 also intersects the stretched ball U'_2 with initial ball U_2 ; let U''_2 be the comparison ball of U'_2 at height \tilde{x}_1 . As \tilde{x}_1 is closer to the thin part than \tilde{x}_2 , the initial radius $3\mu''_{n_2}$ of U''_2 satisfies $3\mu''_{n_2} \leq 3\mu_{n_2} (= 3\mu_{n_1})$. If \tilde{x}''_2 denotes the initial center of U''_2 (observe that $\tilde{x}''_2 \in Y_2$), then using Lemma 5 we deduce

$$d(\tilde{x}_1, \tilde{x}''_2) < 2 \cdot 3\mu_{n_1} + 2 \cdot 3\mu''_{n_2} \leq 12\mu_{n_1}.$$

By definition, $12\mu_{n_1} \leq \varepsilon_2(\mu_{n_1-1})$ and $d(\tilde{x}_1, S_{<n_1}) \geq \mu_{n_1-1}$, so Lemma 8 yields $Y_2 \subseteq Y_1$, the statement.

Assume now that $n_1 > n_2$, i.e. $3\mu_{n_1} < 3\mu_{n_2}$. Let $x \in U_1 \cap U_2$ and c_x be the flow geodesic towards the parabolic fixed point z , with parametrization $c_x(0) = x$ and $c_x(-\infty) = z$;

moreover, let HS_1 and HS_2 be the horospheres around z containing \tilde{x}_1 and \tilde{x}_2 , respectively. As \tilde{x}_1 is closer to the thin part and the initial ball of U_1 does not intersect U_2 – i.e. c_x (coming from z) leaves U_2 before entering the initial ball of U_1 –, we have $c_x(t_1) \in HS_1$ and $c_x(t_2) \in HS_2$ for suitable $0 < t_1 < t_2$, where $t_2 < 3\mu_{n_2}$ (recall that $x \in U_2 = B_{3\mu_{n_2}}^X(\tilde{x}_2)$). Note that on its way from x to $c_x(t_1)$, c_x has to meet the initial ball $B_{3\mu_{n_1}}^X(\tilde{x}_1)$ of U_1 (compare the proof of Lemma 11). Thus there is some t' with $0 < t' < t_1$ (hence $t' < 3\mu_{n_2}$) such that $c_x(t') \in B_{3\mu_{n_1}}^X(\tilde{x}_1)$. We deduce

$$\begin{aligned} d(\tilde{x}_1, \tilde{x}_2) &\leq d(\tilde{x}_1, c_x(t')) + d(c_x(t'), c_x(0)) + d(c_x(0), \tilde{x}_2) \\ &= d(\tilde{x}_1, c_x(t')) + t' + d(x, \tilde{x}_2) \\ &< 3\mu_{n_1} + t' + 3\mu_{n_2} \\ &< 3\mu_{n_1} + 3\mu_{n_2} + 3\mu_{n_2} \\ &< 7\mu_{n_2}. \end{aligned}$$

Since $7\mu_{n_2} < 12\mu_{n_2} \leq \varepsilon_2(\mu_{n_2-1})$, we conclude $d(\tilde{x}_1, \tilde{x}_2) < \varepsilon_2(\mu_{n_2-1})$. By definition of \mathcal{D}_i , we also have $d(\tilde{x}_2, S_{<n_2}) \geq \mu_{n_2-1}$; applying Lemma 8 gives $Y_2 \subseteq Y_1$, the desired statement.

- In the other case we assume that \tilde{x}_1 is farther away from the thin part than \tilde{x}_2 . So the comparison ball U'_1 of U_1 at the height of \tilde{x}_2 satisfies $3\mu'_{n_1} \leq 3\mu_{n_1}$, where $B_{3\mu'_{n_1}}^X(\tilde{x}'_1)$ denotes the initial ball of U'_1 (note that $\tilde{x}'_1 \in Y_1$, as above). Using Lemma 5 and $3\mu_{n_1} \leq 3\mu_{n_2}$, we deduce

$$d(\tilde{x}'_1, \tilde{x}_2) < 2 \cdot 3\mu'_{n_1} + 2 \cdot 3\mu_{n_2} \leq 12\mu_{n_2}.$$

Since $12\mu_{n_2} \leq \varepsilon_2(\mu_{n_2-1})$ and $d(\tilde{x}_2, S_{<n_2}) \geq \mu_{n_2-1}$, applying Lemma 8 again yields $Y_2 \subseteq Y_1$, which finishes the proof. □

Our construction guarantees that we can control the number of covering sets and the number of nonempty intersections between such sets linearly by the volume:

Lemma 17 *There are constants $C = C(n, \eta, \nu)$, $D = D(n, \nu) > 0$ satisfying the following statements.*

1. We have $|\mathcal{B}| \leq C \cdot \text{Vol}(M)$.
2. A set of \mathcal{B} intersects at most D other sets of \mathcal{B} .

Proof The proof of the first statement is a standard argument which is basically identical to [7] Theorem 4.2 step 4). Note that our η corresponds to the m in [7], which explains the dependency of C on η ; the dependency of C on ν stems from the fact that we defined our thick part using ν as a lower bound on the hyperbolic displacement (so the thin part consisted only of cusps).

Note that in order to estimate the maximal number of intersecting sets, we can replace the stretched balls by larger ordinary balls as in Lemma 13 and use the resulting value as an upper bound. But for the (larger) ordinary balls, the argument is essentially the same as in [7] Theorem 4.2, this time step 5)¹⁵.

Detailed proofs of these statements not omitting these details can also be taken from [9] Lemma 3.50. □

¹⁵ The usual proof shows that there is no dependency of D on η ; the dependency on ν again comes from the fact that we needed ν to put the tubes in the thick part.

4.4 Nerve construction

Our goal is now to build the desired simplicial complex out of the covering via the nerve construction. Let $\mathcal{B}_g \subseteq \mathcal{B}$ be the set of all stretched balls in \mathcal{B} ; define

$$N_+ := \bigcup_{B \in \mathcal{B}} B \quad \text{and} \quad N_0 := \bigcup_{B \in \mathcal{B}_g} B.$$

We denote the nerve complexes (see also [1] section 2.2) associated to \mathcal{B} and \mathcal{B}_g by $N(\mathcal{B})$ and $N(\mathcal{B}_g)$, respectively; obviously, we can think of $N(\mathcal{B}_g)$ as a subcomplex of $N(\mathcal{B})$. The following lemma will be needed to prove the homotopy equivalence between the nerve complexes and the thick part.

Lemma 18 *There is a homotopy equivalence $F : M_+ \xrightarrow{\cong} N_+$ which induces a commutative diagram, with vertical maps given by the inclusions $\partial M_+ \hookrightarrow M_+$ and $N_0 \hookrightarrow N_+$:*

$$\begin{array}{ccc} \partial M_+ & \xrightarrow[\cong]{F|_{\partial M_+}} & N_0 \\ \downarrow & & \downarrow \\ M_+ & \xrightarrow[\cong]{F} & N_+. \end{array}$$

Proof Let M''_+ be the shrinking of the thick part M_+ by $3\mu_{-1}/4$, i.e. $M''_+ := M \setminus (M_-)_{3\mu_{-1}/4}$. By the choice of the stretching length we know that every stretched ball $B \in \mathcal{B}_g$ intersects $\partial M''_+$; on the other hand, no ordinary ball of \mathcal{B} meets the boundary $\partial M''_+$, see Fig. 5.

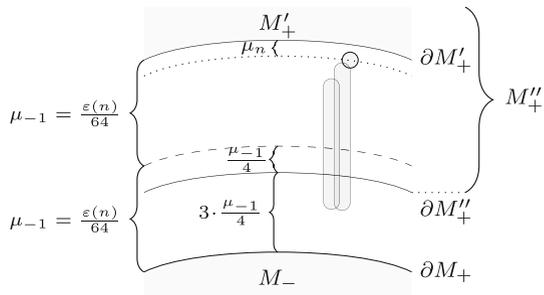
Similar to Lemma 1, the flow away from the thin part up to $\partial M''_+$ (where we stop flowing) induces a commutative diagram

$$\begin{array}{ccc} \partial M_+ & \xrightarrow{\cong} & \partial M''_+ \\ \downarrow & & \downarrow \\ M_+ & \xrightarrow{\cong} & M''_+. \end{array}$$

Moreover, N_+ is stable under this flow (compare Lemma 14); by the choice of the shrinking $3\mu_{-1}/4$, the same is true for N_0 . Let N'_+ and N'_0 denote the images of N_+ and N_0 under this flow up to $\partial M''_+$, then we get a commutative diagram

$$\begin{array}{ccc} N_0 & \xrightarrow{\cong} & N'_0 \\ \downarrow & & \downarrow \\ N_+ & \xrightarrow{\cong} & N'_+. \end{array}$$

Fig. 5 Simplified depiction of the position of the covering sets. Ordinary balls do not exceed the dashed line in the middle and thus can not intersect $\partial M''_+$. Stretched balls always intersect $\partial M''_+$, but do not meet the upper dotted line. N_+ consists of all balls (ordinary and stretched), whereas $N_0 \subseteq N_+$ is made up of all stretched balls. Observe that $M''_+ \subseteq N_+$



Observe that $N'_+ = M''_+$. On the other hand—by the construction of the stretched balls $\bar{\cdot}$, N'_+ and N'_0 are also stable under the flow in opposite direction (i.e. flowing towards M_-) up to $\partial M''_+$; this yields a homotopy equivalence $N'_+ \simeq M''_+$, which itself induces a homotopy equivalence $N'_0 \simeq \partial M''_+$. Consequently, we get the commutative diagram

$$\begin{CD} N'_0 @>\simeq>> \partial M''_+ \\ @VVV @VVV \\ N'_+ @>\simeq>> M''_+ \end{CD}$$

The desired diagram of the statement is now obtained by composing the above diagrams, where—if needed—the horizontal arrows can be reversed by taking the respective homotopy inverses. □

Our main result will depend on the following homotopy equivalence.

Lemma 19 $(M_+, \partial M_+)$ is as a pair homotopy equivalent to $(N(\mathcal{B}), N(\mathcal{B}_g))$.

Proof By Lemma 15 and 16 we know that the (open) covers of N_+ and N_0 by \mathcal{B} and \mathcal{B}_g , respectively, are good covers¹⁶. Using [1] Theorem 2.7 we get a commutative diagram

$$\begin{CD} N_0 @<\simeq<< (N_0)_{\mathcal{B}_g} @>\simeq>> N(\mathcal{B}_g) \\ @VVjV @VVkV @VVV \\ N_+ @<\simeq<< (N_+)_{\mathcal{B}} @>\simeq>> N(\mathcal{B}), \end{CD}$$

where the vertical maps are given by the respective inclusions; here, the spaces $(N_0)_{\mathcal{B}_g}$ and $(N_+)_{\mathcal{B}}$ are the generalized nerve spaces as in [1] section 2.2. Attaching the diagram of Lemma 18 on the left hand side yields the commutative diagram

$$\begin{CD} \partial M_+ @>\simeq>> N_0 @<\simeq<< (N_0)_{\mathcal{B}_g} @>\simeq>> N(\mathcal{B}_g) \\ @VVV @VVjV @VVkV @VVV \\ M_+ @>\simeq>> N_+ @<\simeq<< (N_+)_{\mathcal{B}} @>\simeq>> N(\mathcal{B}), \end{CD} \tag{*}$$

where the outer vertical maps are cofibrations, because $\partial M_+ \hookrightarrow M_+$ is the inclusion of the boundary and $N(\mathcal{B}_g) \hookrightarrow N(\mathcal{B})$ the inclusion of a subcomplex.

The diagram (*) can take the role of diagram (11) in step 4 of the proof of [1] Theorem 3.1. Note that the second column of that diagram (11) has no counterpart in our diagram (*); this is no issue, as that column was only needed to construct a diagram similar to the one in our Lemma 18. Repeating the arguments of [1] Theorem 3.1 step 4 with our diagram (*) replacing diagram (11) there now yields the statement. □

Of course the impact of Lemma 19 depends on if we can control the complexity of the simplicial pair $(N(\mathcal{B}), N(\mathcal{B}_g))$; here, Lemma 17 will come into play. Before summarizing all these statements in our main result, we will see that the case of arithmetic, non-uniform lattices Γ is particularly nice, as all constants will only depend on the dimension n (and not on the other constants η and ν):

Lemma 20 Let \mathfrak{A}_n be the class of arithmetic, non-uniform lattices in $\text{Isom}(\mathbb{H}^n)$, then:

¹⁶ In the sense that the covering sets and their non-empty intersections are contractible.

1. There is a constant $\eta = \eta(n) \in \mathbb{N}$ only depending on n , such that for all $\Gamma \in \mathfrak{A}_n$, the order of any finite subgroup G of Γ is bounded by η , i.e. $|G| \leq \eta$.
2. There is a constant $\nu = \nu(n) > 0$ only depending on n , such that for all $\Gamma \in \mathfrak{A}_n$ and every hyperbolic $\gamma \in \Gamma$, the minimal displacement of γ is at least ν , i.e. $d_\gamma(x) \geq \nu$ for all $x \in \mathbb{H}^n$.

Proof These statements are straightforward consequences of [6] Lemma 13.1 and [6] Remark 5.7. □

We will now state the main result in its general form; recall that a pair (S, S') with a simplicial complex S and (possibly empty) subcomplex $S' \subseteq S$ is an (A, B) -**simplicial pair**, if S has at most B vertices and the degree at every vertex is bounded by A .

Theorem 3 *As usual, let $\eta \in \mathbb{N}$ be an upper bound on the order of finite subgroups of Γ and $\nu > 0$ be a lower bound on the displacement of hyperbolic elements of Γ . Then there are constants $C = C(n, \eta, \nu)$ and $D = D(n, \nu)$, such that $(M_+, \partial M_+)$ is as a pair homotopy equivalent to a $(D, C \cdot \text{Vol}(M))$ -simplicial pair. For $\Gamma \in \mathfrak{A}_n$ (i.e. Γ is arithmetic, non-uniform) the constants C and D will only depend on the dimension n .*

Proof This is a combination of Lemma 19 and 17, as well as 20 in the arithmetic, non-uniform case. □

Remark 2 Using the formulas given for the construction of the constant C of Theorem 3, it can be shown that $C = C(n, \eta, \nu)$ grows exponentially in η . The dependency on ν can not be deduced that easily, as it was also used for the definition of μ_{-1} , and thus its influence would have to be traced along the iterative construction of all the μ_i up to μ_n .

5 Applications

Using the main result Theorem 3 (and its notation), bounds on the homology of hyperbolic orbifolds are an immediate consequence.

Theorem 4 *Let \mathbb{K} be an arbitrary field and let $b_k(M; \mathbb{K}) = \dim_{\mathbb{K}} H_k(M; \mathbb{K})$ denote the k -th Betti number of M with coefficients in \mathbb{K} . Then there is a constant $E = E(n, \eta, \nu) > 0$ such that*

$$b_k(M; \mathbb{K}) \leq E \cdot \text{Vol}(M)$$

for all $k \in \mathbb{N}_0$. In the arithmetic, non-uniform case, the constant E will only depend on the dimension n .

Proof With our main result Theorem 3, the proof is a standard argument utilizing the Mayer-Vietoris sequence (see e.g. [8] Theorem 4.11 or [9] Satz 3.59); E will be given by $E := (D^{n-1} + D^n + 1) \cdot C$. □

The similar result for the torsion part of the homology is given as follows.

Theorem 5 *There is a constant $F = F(n, \eta, \nu) > 0$ such that*

$$\log |\text{tors } H_k(M; \mathbb{Z})| \leq F \cdot \text{Vol}(M)$$

for all $k \in \mathbb{N}_0$. In the arithmetic, non-uniform case, the constant F will only depend on n .

Proof By the main result Theorem 3, $M \simeq M_+$ is homotopy equivalent to a $(D, C \cdot \text{Vol}(M))$ -simplicial complex. Now [1] Lemma 5.2 yields the result. Note that—following the proof of [1] Lemma 5.2 and assuming $D \geq 1$ (without restriction)—the constant F will be given by $F := D^n \cdot \log(n + 2) \cdot C$. \square

As we see from the previous two theorems, the constants for the homology bounds depend on the dimension n , the maximal order η of finite subgroups of the lattice and the minimal hyperbolic displacement ν . Using a more general approach than the one presented in this paper, the dependence on ν can be relaxed – yielding bounds independent of ν , but polynomial in $\text{Vol}(M)$ – and we suspect that it might be removed altogether. This more general approach basically consists of replicating the construction of the stretched balls also near the tubes (which we basically ignored by putting them into the thick part, using ν); a detailed description of this procedure and the corresponding proofs are contained in [9] and might be a starting point for future improvements of the results given here.

Acknowledgements The results presented here are part of my doctoral thesis [9], which was written under the supervision of Prof. Roman Sauer, to whom I am grateful for supporting my work.

Data availability Data sharing not applicable to this article as no datasets were generated or analysed during the current study.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Bader, U., Gelander, T., Sauer, R.: Homology and homotopy complexity in negative curvature. *J. Eur. Math. Soc.* **22**(8), 2537–2571 (2020)
2. Ballmann, W., Gromov, M., Schroeder, V.: *Manifolds of Nonpositive Curvature*. Springer Science + Business Media, New York (1985)
3. Benedetti, R., Petronio, C.: *Lectures on Hyperbolic Geometry*. Springer-Verlag, Berlin, Heidelberg (1992)
4. Bowditch, B.H.: Geometrical finiteness with variable negative curvature. *Duke Math. J.* **77**(1), 229–274 (1995)
5. Canary, R.D., Epstein, D., Marden, A.: *Fundamentals of Hyperbolic Geometry*. Cambridge University Press, Cambridge (2006)
6. Gelander, T.: Homotopy type and volume of locally symmetric manifolds. *Duke Math. J.* **124**(3), 459–515 (2004)
7. Samet, I.: Betti numbers of finite volume orbifolds. *Geom. Topol.* **17**(2), 1113–1147 (2013)
8. Senska, H.: Topological complexity of visibility manifolds. *arXiv* (2019), <https://arxiv.org/abs/1910.06087>
9. Senska, H.: *Topologie und Volumen in negativer Krümmung*. Dissertation, Karlsruhe Institute of Technology (2019). <https://doi.org/10.5445/IR/1000096617> (German)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.