

HALF-DIMENSIONAL COLLAPSE OF ENDS OF MANIFOLDS OF NONPOSITIVE CURVATURE

GRIGORI AVRAMIDI AND T. TÂM NGUYỄN-PHAN



Abstract. This paper accomplishes two things. First, we construct a geometric analog of the rational Tits building for general noncompact, complete, finite volume n -manifolds M of bounded nonpositive curvature. Second, we prove that this analog has dimension less than $\lfloor n/2 \rfloor$.

1 Introduction

Let M be a noncompact, complete Riemannian n -manifold with bounded nonpositive sectional curvature $-1 < K \leq 0$ and finite volume. We also assume that M does not have arbitrarily small geodesic loops, that is, there is a positive lower bound to the length of any geodesically embedded circle. Good examples to think about are locally symmetric spaces of noncompact type, such as hyperbolic manifolds, products of surfaces, and the usual beloved $K \backslash G / \Gamma$. Sometimes taking $G = \mathrm{SL}_m \mathbb{R}$ and $K = \mathrm{SO}_m$ can be as satisfactory as any other semisimple Lie groups. This sentiment holds true in terms of examples to keep in mind as one reads since our approach throughout this paper is purely geometric/topological but can be demonstrated by thinking about these concrete examples the right way.

The condition that M has no arbitrarily small geodesic loops holds when M is negatively curved, i.e. $-1 < K < 0$, or when M is locally symmetric. One reason we need this condition in the general setting of bounded nonpositive curvature is in order to ensure, by a theorem of Gromov–Schroeder, that M is *tame* in the sense that the thin part $M_{<\varepsilon}$ has finitely many components and each component is topologically a product of a closed $(n - 1)$ -manifold with a ray. The mechanism for tameness is that the injectivity radius function on M does not have any critical point outside a compact set which can be taken to be the thick part $M_{\geq\varepsilon}$ for some small $\varepsilon > 0$ (see also Appendix 2 of [BGS85] for a generalization). Let $\widetilde{M}_{<\varepsilon}$ be the preimage of the thin part $M_{<\varepsilon}$ under the universal covering map. We will call $\widetilde{M}_{<\varepsilon}$ the thin part of \widetilde{M} . It is the topology of $\widetilde{M}_{<\varepsilon}$ that we would like to describe.

When M is locally symmetric (and arithmetic), $\widetilde{M}_{<\varepsilon}$ is homotopy equivalent to the rational Tits building, which is a $(k - 1)$ -dimensional complex, where k is the rational rank of M . The rational Tits building of M can be realized as a subset of the visual boundary ∂_∞ of \widetilde{M} and can be thought of as the set of points at infinity

that can be reached if one moves only within $\widetilde{M}_{<\varepsilon}$. The rank k is at most $n/2$ with $k = n/2$ when M is a product of non-compact surfaces. The main purpose of this paper is to show that the $n/2$ bound on the dimension of the rational Tits building is no arithmetic coincidence but in a slightly weaker sense.

In the general nonpositively curved setting, for an n -dimensional manifold M satisfying the conditions described above, we define a map

$$\rho: \partial\widetilde{M}_{<\varepsilon} \rightarrow \partial_\infty$$

that is an analog of a rational Tits building in the sense that ρ encodes all the directions to infinity necessary to push any topological feature (e.g. homology cycles, maps) in $\widetilde{M}_{<\varepsilon}$ without it leaving $\widetilde{M}_{<\varepsilon}$. We then prove that the image of ρ has dimension at most $(\lfloor n/2 \rfloor - 1)$, where $\lfloor n/2 \rfloor$ is the greatest integer less than or equal to $n/2$.

Theorem 1. *Let M be a noncompact, complete, Riemannian manifold with bounded nonpositive sectional curvature $-1 < K \leq 0$ and finite volume. Assume that M has no arbitrarily small geodesic loops. Let $\varepsilon > 0$ be smaller than the Margulis constant and small enough so that $M_{<2\varepsilon}$ is topologically a product with a ray. Then there is a $\pi_1(M)$ -equivariant, Lipschitz map $\rho: \partial\widetilde{M}_{<\varepsilon} \rightarrow \partial_\infty$, defined on a triangulation of $\partial\widetilde{M}_{<\varepsilon}$, with the following properties.*

- (a) *For each $x \in \partial\widetilde{M}_{<\varepsilon}$, the unique unit speed geodesic ray $[x, \rho(x))$ connecting x to $\rho(x)$ stays in $\widetilde{M}_{<2\varepsilon}$. Moreover, the projection of $[x, \rho(x))$ to M escapes every compact set in M .*
- (b) *If σ is a simplex in $\partial\widetilde{M}_{<\varepsilon}$, then $\rho(\sigma)$ has dimension less than $\lfloor n/2 \rfloor$.*

We can use ρ to show that any polyhedron in $\widetilde{M}_{<\varepsilon}$ can be homotoped within $\widetilde{M}_{<\varepsilon}$ to one with dimension at most $(\lfloor n/2 \rfloor - 1)$.

Theorem 2. *Assume the hypotheses of Theorem 1. Let P be a finite polyhedron and let $\varphi: P \rightarrow \widetilde{M}_{\leq\varepsilon}$ be a continuous map. Then φ can be homotoped within $\widetilde{M}_{<\varepsilon}$ to a map $\widehat{\varphi}: P \rightarrow \widetilde{M}_{<\varepsilon}$ whose image has dimension $\leq \lfloor n/2 \rfloor - 1$.*

This is done by pushing P toward $\rho(P)$ until it is deep enough in $\widetilde{M}_{<\varepsilon}$ that we can “collapse” P onto a close-by copy of $\rho(P)$. The following corollary is an almost immediate consequence.

COROLLARY 3. *The map $\widehat{\varphi}$ in Theorem 2 can be homotoped in $\widetilde{M}_{<\varepsilon}$ to factor through a polyhedron Q of dimension $\lfloor n/2 \rfloor - 1$. Consequently, the homology of $\widetilde{M}_{<\varepsilon}$ vanishes in dimension $\geq \lfloor n/2 \rfloor$, i.e.*

$$H_{\geq \lfloor n/2 \rfloor}(\widetilde{M}_{<\varepsilon}) = 0. \tag{1}$$

The idea behind all of this is that we push any topological feature, such as a polyhedron P , to infinity within $\widetilde{M}_{<\varepsilon}$. Of course, we can always push anything in $\widetilde{M}_{<\varepsilon}$ to infinity since M has tame ends, but we want to push P in such a way not to stretch it more than we absolutely have to. The number of degrees of freedom in stretching P is the dimension of $\rho(P)$. This is why we build ρ and why we need to make it as low dimensional as possible. In a way, the topology of $\widetilde{M}_{<\varepsilon}$ is very much determined by ρ .

REMARK. The upper bound on the dimension of Q is sharp by the following example. If M is the product of k punctured tori $S_1 \times \cdots \times S_k$, each with a complete hyperbolic metric with finite volume, then M has dimension $n = 2k$, so $\lfloor n/2 \rfloor - 1 = k - 1$. We also know that in this case $\widetilde{M}_{<\varepsilon}$ is homotopy equivalent to a wedge of $(k-1)$ -spheres. One can see this explicitly by covering $\widetilde{M}_{<\varepsilon}$ with convex sets and taking the nerve of this cover. The convex sets can be chosen as follows. For each S_i , let γ_i be the loop that shrinks to zero, and take the pre-image in \widetilde{M} of the sets $U_i = \{|\gamma_i| < \varepsilon\}$ of points in M where the loop γ_i is smaller than ε . For ε small enough, the collection $\{p^{-1}(U_i)\}_{i=1,2,\dots,k}$ covers $\widetilde{M}_{<\varepsilon}$. Each connected component of $p^{-1}(U_i)$ is a convex subset of \widetilde{M} . The collection of connected components of $\{p^{-1}(U_i)\}_{i=1,2,\dots,k}$ is the cover that we take. It is not hard to see that the nerve of this cover is the k -fold join of infinite discrete sets, which is homotopy equivalent to an infinite wedge of $(k-1)$ -spheres.

Optimal examples. In [AN17] we build, for each n , an n -manifold M satisfying the hypotheses of Theorem 1 which has

$$\overline{H}_k(\widetilde{M}_{<\varepsilon}) \neq 0 \quad \text{for all } k < \lfloor n/2 \rfloor.$$

So, unlike in the case of locally symmetric spaces, in general there are no low dimensional homology vanishing results complementing the high dimensional homology vanishing (1).

Some other forms of collapse. Of course in some situations we expect to be able to do better than $\lfloor n/2 \rfloor$. For example, in the simple case of a finite volume hyperbolic manifold, every component of the thin part $\widetilde{M}_{<\varepsilon}$ collapses to a point. The feature responsible for this collapse is that the Tits boundary is discrete. In fact, the dimension of the Tits boundary $(\partial_\infty, \text{Td})$ of \widetilde{M} is one of the factors that control the topology of $\widetilde{M}_{<\varepsilon}$. This is reflected in the fact that the map ρ we construct is continuous (in fact, Lipschitz) in the Tits metric.

REMARK. More precisely, by *dimension* we mean the *compact topological dimension*, which is defined as the maximum of the covering dimensions of compact subspaces. See [Kle99] for a comparison of different notions of dimension for the Tits boundary.

In addition, ρ is constructed in such a way that it factors through a complex built out of virtual equivalence classes of certain Abelian subgroups of $\pi_1 M$. So, the topology

of this complex is another factor that controls the topology of $\widetilde{M}_{<\varepsilon}$. This leads to two additional forms of collapse via ranks of Abelian groups and dimension of the Tits boundary. To express it, let

$$d = \min\{\text{rank}_{Ab}(\pi_1 M) - 1, \quad \dim(\partial_\infty, \text{Td}), \quad \lfloor n/2 \rfloor - 1\}$$

where

- $\text{rank}_{Ab}(\pi_1 M)$ is the maximum rank of an Abelian subgroup of $\pi_1 M$,
- $\dim(\partial_\infty, \text{Td})$ is the dimension of the Tits boundary, and
- $\lfloor n/2 \rfloor$ still denotes the greatest integer less than or equal to $n/2$.

Theorem 4. *Theorem 2 is true if we replace “ $\lfloor n/2 \rfloor - 1$ ” by “ d ”.*

REMARK. If M is a noncompact, finite volume hyperbolic n -manifold, with several cusps then the fundamental group Γ contains a parabolic Abelian subgroup of rank $n - 1$. Doubling such a manifold along one of its cusps gives a non-positively curved manifold containing an Abelian subgroup of rank $n - 1$ that consists of hyperbolic elements, in which case $\dim(\partial_\infty, \text{Td}) \geq n - 2$. So, the bounds via ranks of Abelian groups and via the dimension of the Tits boundary are situational. But, the half dimension bound $d \leq \lfloor n/2 \rfloor - 1$ is something we always have.

Low dimensional collapse. Now let us turn to the special situation when d is small. Notice that $d \leq 1$ if

- $\pi_1 M$ does not contain \mathbb{Z}^3 , or
- $\dim(\partial_\infty, \text{Td}) \leq 1$, or
- $\dim M \leq 5$.

In any of these situations we get the following immediate corollary.

COROLLARY 5. *If $d \leq 1$ then each component of the end $M_{<\varepsilon}$ is aspherical.*

Once we know that each component of the end is aspherical, it is natural to ask “how bad” this aspherical manifold can be, i.e. how exotic its fundamental group is. We prove the following amplification of Corollary 5. Recall that a countable group F is *locally free* if every finitely generated subgroup is free.

Theorem 6. *If $d \leq 1$ then for each component C of $M_{<\varepsilon}$ the fundamental group is an extension*

$$1 \rightarrow F \rightarrow \pi_1 C \rightarrow \pi_1 M$$

of a subgroup of $\pi_1 M$ by a locally free group F .

In proving the main theorems we obtain the following.

Technical byproduct of independent interest. Let G be a group that acts on a Hadamard n -manifold \widetilde{M} via covering space transformations. Note that the sectional curvatures of \widetilde{M} need not be bounded from below. Suppose that G preserves horospheres centered at $\{z_i\}_{i=0,\dots,k} \subset \partial_\infty$. We ask the following questions.

1. Can one connect in ∂_∞ the z_i 's through points the horospheres centered at which are preserved by G ?
2. If so, then the z_i 's can be filled by a k -simplex Δ^k . That is, one gets a map

$$\sigma: \Delta^k \rightarrow \partial_\infty,$$

such that horospheres centered at each point in $\sigma(\Delta^k)$ are preserved. How nice (e.g. continuous, Hoelder, Lipschitz, etc) can this map be?

3. Let $\text{Fix}^0(G)$ be the set points in ∂_∞ whose horospheres are preserved by G . What is the relation between the dimension of $\text{Fix}^0(G)$ and the dimension of G ?

We answer these questions in the case when the vertices z_i are mutually a Tits distance $\leq \pi/2$ apart in Sections 9 and 12. In short, the answer to the first question is yes, the answer to the second question is Lipschitz—we construct such a map σ , which we call a *Busemann simplex* (defined in 9.2)—and the answer to the third question is the following.

Theorem 7. *If the vertices z_i span a non-degenerate Busemann k -simplex, then the homological dimension of G is less than $(n - k)$.*

On Section 2 and Subsection 12.1. We would like to warn the readers about the unorthodox way these sections are written and that they are written *not* to sketch the proof of the main theorems but to communicate what is not captured in the proof. The purpose of Section 2 and Subsection 12.1 is to discuss the *problems*, that is, to point the readers to where the difficulties lie. Their role is complementary to the proof. The proof of the main theorem in this paper consists of a mixture of steps and lemmas, some of which are trivial, some of which are just technicalities and some of which deal with actual difficulties. Section 2 and Subsection 12.1 are our attempt to pin-point where the difficulties lie in a coherent, though imprecise, way.

We “motivate” the proof of Theorem 2 in Section 2. Theorem 1 will be attained along the way. We will try to explain why things are done the way they are through an iteration of “what is the simplest thing to do?” and “what are the problems to overcome?” until there are no more problems. Good examples to think about are locally symmetric spaces, in particular $\text{SO}_3 \setminus \text{SL}_3 \mathbb{R} / \text{SL}_3 \mathbb{Z}$ and products of surfaces.

Some readers might find Section 2 and Subsection 12.1 “madness”, in which case they are encouraged to skip them (or read them afterwards for perspective) and proceed to the precise formulation given in the rest of the paper.

2 Problems and Solutions

Let P be as in Theorem 2. One can naively take a length minimizing geodesic ray γ in M starting at a point in $\varphi(P)$, take a lift $\tilde{\gamma}$ of γ , and then push P toward $\tilde{\gamma}(\infty) \in \partial_\infty$ with unit speed. Then the diameter of P will stay bounded, so once it is far enough to infinity it will be contained in a ball that is contained in $\tilde{M}_{<\varepsilon}$. We

then can contract P to a point within this ball. However, there is a problem with this approach, which is that as we push P toward $\tilde{\gamma}(\infty)$ it might slide off $\widetilde{M}_{<\varepsilon}$ for some time during this process. This problem does not have a solution for otherwise one could contract any such P to a point within $\widetilde{M}_{<\varepsilon}$, which is not true if M is a product of noncompact surfaces.

So we need to find a way to push P to infinity without it leaving $\widetilde{M}_{<\varepsilon}$. The strategy is that we push different points of P to different points in the visual boundary ∂_∞ along geodesics and we keep track of the amount of directions to infinity we need. The set of points in ∂_∞ to which we push P tells us how much P “expands” as we push it to infinity. It also gives a complex $c_t(P)$ in $\widetilde{M}_{<\varepsilon}$ to which we can “collapse” P onto. We then bound the dimension of this complex to be less than $\lfloor n/2 \rfloor$ by trying to make this process as efficient (in terms of how many degrees of freedom are needed as P expands) as possible.

Keeping the homotopy within the thin part $\widetilde{M}_{<\varepsilon}$. For each point x of the polyhedron P , we find in ∂_∞ a point $\rho(x)$ to which we push x with unit speed along the geodesic connecting x with $\rho(x)$ as illustrated in Figure 1. We call this homotopy

$$\varphi_t: P \rightarrow \widetilde{M}_{<\varepsilon}, \quad \text{for } t \in [0, \infty), \quad \text{with } \varphi_0 = \varphi.$$

We will define $\rho: P \rightarrow \partial_\infty$ systematically, skeleton by skeleton. Start with the vertices of P and suppose that x is a vertex of P . Since $\varphi(x)$ is in $\widetilde{M}_{<\varepsilon}$, there is a parabolic isometry γ_x that moves $\varphi(x)$ by a *small* amount, where *small* means less than ε , so that the group Γ_x generated by such γ_x is virtually nilpotent by the Margulis lemma. Therefore, a natural choice for $\rho(x)$ is the center of a horosphere preserved by Γ_x because if we push x to $\rho(x)$ along a geodesic $\varphi_t(x)$, then the *small* elements in Γ_x will remain *small* along $\varphi_t(x)$, so $\varphi_t(x)$ will stay in $\widetilde{M}_{<\varepsilon}$.

Next, we extend ρ to the edges and higher dimensional simplices of P . Just like for the vertices, the way to go to infinity is to “follow the shrinking small loops”. Let e be the edge connecting vertices x_0 and x_1 of P . Clearly a problem is that there is no clear “transition” in terms of *small* loops at $\varphi(x_0)$ to *small* loops at $\varphi(x_1)$. But there is a solution, which is to take a fine enough subdivision of P at the beginning. That is, we take a subdivision of P in which the diameter of each simplex is tiny enough so that if σ is a k -simplex of P with vertices x_0, x_1, \dots, x_k , then for each point $y \in \sigma$, some parabolic isometry that is *small* at $\varphi(x_i)$ is still *small* at $\varphi(y)$. We can always take such a subdivision of P , so we can harmlessly assume that P is triangulated in such a way at the beginning.

This gives us a way to assign to each simplex a nontrivial nilpotent group as follows. Let $\Gamma = \pi_1(M)$. For a vertex x of P , let

$$S_x = \{\gamma \in \Gamma \mid d(x, \gamma(x)) < \varepsilon\}$$

be the set of ε -small parabolic isometries at $\varphi(x)$. Then by the Margulis lemma, the group

$$\Gamma_x = \langle \gamma \mid \gamma \in S_x \rangle,$$

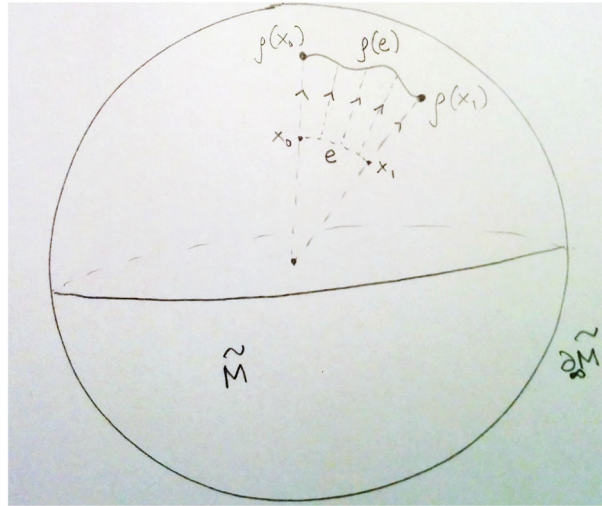


Figure 1: Pushing to infinity.

has a nilpotent subgroup of index less than a constant I_n that depends only on n . We assign to x the following nontrivial nilpotent subgroup

$$N_x = \langle \gamma^{I_n!} \mid \gamma \in \Gamma_x \rangle$$

of Γ_x . Note that N_x is characteristic and thus normal in Γ_x . One might worry with this choice of nilpotent group at x the group N_x may not contain any *small* parabolic isometry. This can easily be fixed by making ε small enough at the beginning. For a k -simplex $\sigma = x_0 * x_1 * \dots * x_k$, let

$$\Gamma_\sigma = \langle \Gamma_{x_0}, \dots, \Gamma_{x_k} \rangle.$$

We assign to σ the nilpotent group

$$N_\sigma = \langle \gamma^{I_n!} \mid \gamma \in \Gamma_\sigma \rangle.$$

Let $Z_\sigma = Z_{N_\sigma}$ be the center of N_σ . Note that Z_σ is normal in Γ_σ .

Since nilpotent groups are at times harder to deal with than Abelian groups—life is hard enough already—we try to make things as easy as we can by assigning to the simplex σ the Abelian group $A_\sigma = Z_\sigma$. Again, if one is worried, one can pick ε to be small enough at the beginning so that A_σ has *small* elements.

Now the problem with this is that this does not quite give us a way to define ρ on an edge e connecting x_0 and x_1 because what we get from A_e is the center of a horosphere preserved by A_e to which we can push e without it leaving $\widetilde{M}_{<\varepsilon}$. However, this center is only one point while what we are looking for is a path connecting $\rho(x_0)$ and $\rho(x_1)$. Nevertheless, what we have obtained is a way to define ρ on the vertices of the first barycentric subdivision P_1 of P . Note that in this assignment of Abelian groups adjacent vertices in P_1 have commuting Abelian groups.

There will be a solution to the above problem if the distance between adjacent vertices of P_1 in the Tits metric Td on ∂_∞ is less than π because then there will be a unique geodesic in $(\partial_\infty, \text{Td})$ connecting them, so we can use this to define ρ on the edges of P_1 . This turns out to be true if we are picky enough when we pick where ρ sends vertices of P_1 . In fact, we can, for each simplex, make ρ send all of its vertices to a set of Tits-diameter $\leq \pi/2$ in ∂_∞ .

For each Abelian group A_x , where x is a vertex of P_1 , let $\text{Fix}(A_x)$ be the set of fixed points of A_x in ∂_∞ . There is a canonical way (see Section 5) to define a unique “Center of Mass” $\xi_{A_x} \in \text{Fix}(A_x)$ such that

- any isometry γ that normalizes A_x fixes ξ_{A_x} , and
- all points in $\text{Fix}(A_x)$ are within a Tits distance of $\pi/2$ from ξ_{A_x} .

Define $\rho(x) = \xi_{A_x}$. Then by the first property above, adjacent Abelian groups in P_1 fix each other’s Centers of Mass. By the second property, they are within a Tits distance of $\pi/2$ from each other. It follows that any two adjacent vertices x_0 and x_1 in P_1 can be connected by a *unique* geodesic $\rho(e)$ in ∂_∞ connecting $\rho(x_0)$ and $\rho(x_1)$. Parametrize both e and $\rho(e)$ by constant speed and use this to define ρ on e the obvious way. We can extend ρ to higher dimensional skeleta via geodesic triangles in the obvious way.

Now that we have found a way to define ρ on P_1 , we need to check that the homotopy φ_t does not push P off the thin part $\widetilde{M}_{<\varepsilon}$. For each k -simplex σ in P_1 , we have a chain $\Gamma_{x_0} \leq \dots \leq \Gamma_{x_k}$. The “bottom” group Γ_{x_0} normalizes A_{x_0}, \dots, A_{x_k} and therefore fixes $\rho(x_i) = \xi_{A_{x_i}}$ for all $i = 0, 1, \dots, k$. It follows that Γ_{x_0} fixes $\rho(\sigma)$ *pointwise*. Remember that there is an element γ of Γ_{x_0} that is *small* at x_0 and that is still *small* at all points $y \in \sigma$. Since γ fixes $\rho(y)$, it will stay *small* on $\varphi_t(y)$ for all $t > 0$, and therefore, the homotopy φ_t does not move P off $\widetilde{M}_{<\varepsilon}$. Note that this shows we did not have to go back and manually make ε smaller at the beginning as one might have worried before. Now, we can move on to the next task.

“Collapsing” P within $\widetilde{M}_{<\varepsilon}$. Now that we have defined ρ of P and made sure that pushing P to $\rho(P)$ does not leave $\widetilde{M}_{<\varepsilon}$, we want to find a copy of $\rho(P)$ in \widetilde{M} to which we can “collapse” P within $\widetilde{M}_{<\varepsilon}$. Take a point $c_0 \in \widetilde{M}$ and take the geodesic cone on $\rho(P)$ with cone point c_0 . For $t \geq 0$ and $x \in P$, let $c_t(x)$ be the point obtained by flowing for time t along the geodesic ray from c_0 to $\rho(x)$. Then $c_t(P)$ homeomorphic to $\rho(P)$ because geodesic retractions are homeomorphisms. Also, it is not hard to see that the distance between $\varphi_t(P)$ and $c_t(P)$ is bounded by some number R that does not depend on $t \geq 0$ (but depends on P and c_0). Thus, we can “collapse” $\varphi_t(P)$ onto $c_t(P)$ in an R -neighborhood of $\varphi_t(P)$.

There is a problem, which is that the “collapse” might leave $\widetilde{M}_{<\varepsilon}$, which could happen if the R -neighborhood of $\varphi_t(P)$ is large enough that it contains points outside $\widetilde{M}_{<\varepsilon}$. But there is a solution if we can show that for t large enough, $\varphi_{t_{\text{large}}}(P)$ is deep enough in $\widetilde{M}_{<\varepsilon}$ that an R -neighborhood of $\varphi_{t_{\text{large}}}(P)$ is contained in $\widetilde{M}_{<\varepsilon}$, so when we collapse $\varphi_{t_{\text{large}}}(P)$ to $c_{t_{\text{large}}}(P)$ it will not leave $\widetilde{M}_{<\varepsilon}$ during this process. Therefore,

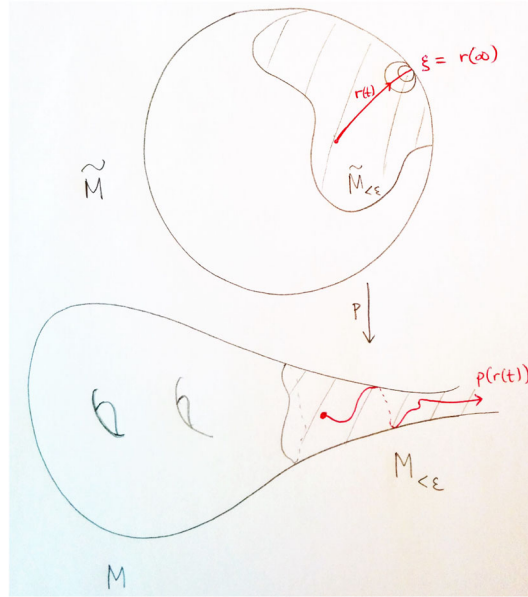


Figure 2: Divergent ray.

in addition to making sure that $\varphi_t(\sigma)$ stays in $\widetilde{M}_{<\varepsilon}$ for all $t > 0$, we also need its projection under the covering space projection $p: \widetilde{M} \rightarrow M$ to be *divergent* in M . That is, $p(\varphi_t(P))$ leaves all compact sets in M as $t \rightarrow \infty$ (Figure 2). This is true by the following key lemma.

LEMMA 8 (Divergent Geodesic Ray). *Let A be a free Abelian group of isometries. Suppose the centralizer C_A preserves each horosphere centered at a point ξ in $\partial_\infty \widetilde{M}$. Then for any geodesic ray $r: [0, \infty) \rightarrow \widetilde{M}$ with end point $r(\infty) = \xi$ the projection $p(r(t))$ is divergent.*

More generally, if $r(\infty) = \eta$ for some η fixed by C_A and $\text{Td}(\eta, \xi) < \pi/2$ then $p(r(t))$ is also divergent.

Lemma 8 takes care of the above problem if for each Abelian group A_x above the horospheres centered at the Center of Mass ξ_{A_x} are preserved by the centralizer C_{A_x} . We prove that this is true in Section 5. If one is concerned that Lemma 8 might apply to only one single ray $\varphi_t(x)$ at a time but not uniformly to a family $\varphi_t(P)$ of rays, then one is absolutely right, but we take care of this in Proposition 14, which says that one needs not worry if P is bounded (and P is indeed bounded).

Bounding the dimension of $\rho(P)$. That the dimension of $\rho(P)$ is at most $\lfloor n/2 \rfloor - 1$ is due to two factors.

- First, for each simplex $\sigma = x_0 * x_1 * \dots * x_k$ in P_1 , we get a “biggest” Abelian group $A^\sigma = \langle A_{x_0}, \dots, A_{x_k} \rangle$. The group A^σ is Abelian because the groups A_{x_i} ’s commute with each other. Also, A^σ preserves horospheres centered at $\rho(x_i) =$

$\xi_{A_{x_i}}$ for $i = 0, \dots, k$ and therefore preserves their intersection. If $\rho(x_i)$'s span an l -dimensional simplex at infinity (for $l \leq k$), then the dimension of the intersection of the horospheres *should* be $n - (l + 1)$. This *should* mean that the rank of A^σ is less than or equal to $n - (l + 1)$.

- Second, if σ is a simplex in P_1 , we expect the dimension of $\rho(\sigma)$ to be less than the rank of A^σ . One reason is because virtually equivalent Abelian groups are too similar to demand different treatments, in particular, they should be assigned the same point at infinity.

Putting these two factors together we get that if A^σ has rank r , then

$$r \leq n - (l + 1) \quad \text{and} \quad r \geq l + 1.$$

Therefore, $l + 1 \leq \lfloor n/2 \rfloor$. So the dimension of $\rho(P)$ is at most $\lfloor n/2 \rfloor - 1$.

There are, of course, problems to overcome in both claims. There are two problems in the second claim. One is that virtually equivalent Abelian groups A_x and A_y need not have the same Centers of Mass, in which case the complex $\rho(P)$ might be higher dimensional than it should be. Even with little optimism one expects that if A_x and A_y share a finite index subgroup, there should be a point at infinity whose horospheres are preserved by both A_x and A_y . The solution is to construct, for each such Abelian group, a Center of Mass that is invariant under virtual equivalence and that has all the metric and invariance properties we mentioned above. This can be done and is done in Section 5.

The other problem in the second claim is that the rank of A^σ could be *strictly* less than the number of virtual equivalence classes of Abelian groups at the vertices. For example, σ is a triangle and the group at each vertex of σ is isomorphic to \mathbb{Z} and no two of them share a finite index subgroup, yet A^σ could be \mathbb{Z}^2 . A solution is to work with the second barycentric subdivision P_2 of P , instead of P_1 , right from the beginning. So we need to assign Abelian groups to vertices of P_2 . Each vertex x in P_2 that is not in P_1 is a point in the interior of a simplex τ of P_1 . Assign to x the Abelian group A^τ generated by the Abelian groups at the vertices of τ (i.e. let $A_x = A^\tau$). The pay off for working with P_2 is that for each k -simplex σ in P_2 , the Abelian groups at the vertices of σ form a chain $A_0 \leq \dots \leq A_k$. Another nice consequence is that the group generated by vertex groups is the biggest group A_k in the chain (so one can forget the upper index). It follows that the rank of A_k is greater than or equal to number of the number of virtual equivalence classes of Abelian groups at the vertices, which takes care of the problem.

The problem with the first claim is that it is not clear if the intersection of horospheres described above is a manifold of dimension $n - (l + 1)$. For example, it is unclear (to us) how to rule out the following situation (Figure 3).

Suppose that h_0, h_1 and h_2 are Busemann functions on \widetilde{M} . Let $z_i \in \partial_\infty$, for $i = 0, 1, 2$, be the center of the horosphere S_i defined as $h_i = 0$. Now, the intersection $S = S_1 \cap S_2 \cap S_3$ of three horospheres is an $(n - 3)$ -dimensional manifold if S_1, S_2 , and S_3 intersect transversely, i.e. the gradient vectors $\nabla h_0, \nabla h_1$ and ∇h_2 at each point

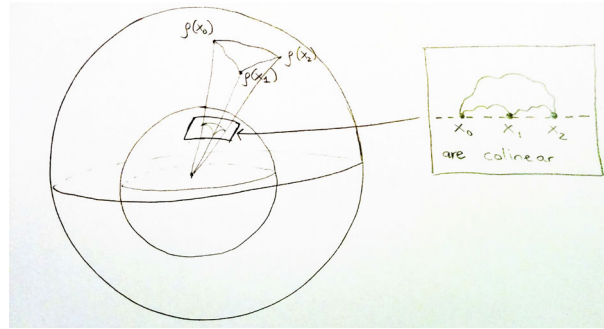


Figure 3: Colinear gradients.

in S are linearly independent. Suppose that z_0, z_1 and z_2 are not co-linear in ∂_∞ , i.e. none of the three points is on the geodesic connecting the other two, so they span a triangle in ∂_∞ . Since $\nabla h_0, \nabla h_1$ and ∇h_2 “point toward” z_0, z_1 and z_2 respectively, this strongly suggests that they *should* be linearly independent. However, there is no reason to relate the linear structure at a point to what happens at infinity, which is something obtained via a limiting process in terms of the metric. We are not sure if this is a real problem or the problem lies in our inability.

But we find a way around this problem and this is a solution, which requires a modification to how ρ is defined. This is the *last* modification we will make to ρ . We define ρ on the vertices of P_2 exactly as we did, but we will not use geodesics in $(\partial_\infty, \text{Td})$ to extend ρ to edges and higher dimensional simplices. Instead, we construct what we call *Busemann paths* and *Busemann simplices* and we use them in place of geodesics in the above process of defining ρ . A Busemann k -simplex

$$\sigma: \Delta^k \rightarrow \partial_\infty$$

is a singular k -simplex in ∂_∞ with vertices z_0, z_1, \dots, z_k and has the property (by construction) that if a parabolic isometry preserves the horospheres centered at z_i for $i = 0, 1, \dots, k$, then it will preserve horospheres centered at $\sigma(x)$ for *all* $x \in \Delta^k$. The construction of Busemann simplices uses convex combinations of Busemann functions, which explains the name, and is given in Section 9. The incentive for constructing Busemann simplices is to create more points at infinity whose horospheres are preserved so that we can use them in the case when the horospheres centered at the vertices do not intersect transversely.

It turns out, however, that even with a whole nondegenerate k -simplex of points at infinity whose horospheres are preserved by an Abelian group A_x we are unable to even prove existence of $(k + 1)$ points whose Busemann functions have linearly independent gradient vectors everywhere. Nevertheless, Busemann simplices are too good to waste and we manage to use them to show that the first claim is true if geodesic simplices are replaced by these. Busemann simplices are constructed as pointwise limits of singular simplices

$$\sigma_{R_i}: \Delta^k \rightarrow S(x_0, R_i) \subset \widetilde{M}$$

on larger and larger spheres centered at some fixed point $x_0 \in \widetilde{M}$. So for R large enough, σ_R “approximates” the Busemann simplex σ at infinity so well that non-degeneracy at infinity implies nondegeneracy of $\sigma_R(\Delta^k)$. The union of all $\sigma_R(\Delta^k)$ is called the “Busemann cone”, which is itself not a geodesic cone, serves as a parametrization space of intersections of horospheres centered at the vertices z_i , for $i = 0, 1, \dots, k$. We show that if an Abelian group A of rank r preserves horospheres centered at z_i 's, and if the Busemann simplex with vertices z_i 's is non-degenerate, then when we line up these intersections of horospheres over C the union has dimension at least $r + k + 1$, which gives the first claim. This is discussed in Section 12 and is hard enough to have its own “problems and solutions”.

Last but not least, all of the above effort will go to naught if Busemann simplices are space filling, in which case the Hausdorff dimension of a Busemann l -simplex could be greater than l . However, we show that Busemann simplices are Lipschitz and thus do not increase Hausdorff dimension. It follows that $\rho(P)$ has dimension at most $\lfloor n/2 \rfloor - 1$, and since $c_t(P)$ is homeomorphic to $\rho(P)$, the dimension of $c_t(P)$ is also at most $\lfloor n/2 \rfloor - 1$. This explains Theorems 1 and 2.

Finally, we can approximate $c_t(P)$ by a polyhedron Q of dimension at most $\lfloor n/2 \rfloor - 1$ in some small neighborhood of $c_t(P)$. So when we “collapse” $\varphi_t(P)$ to $c_t(P)$, we can “collapse” it to Q instead and not have to worry that the collapse will not leave $\widetilde{M}_{<\varepsilon}$ because Q is pointwise close to $c_t(P)$. This gives a map $\widehat{\varphi}$ that factors through Q as in the statement of Corollary 3.

There are no more problems.

3 Setup and Notation

3.1 Setup. In the rest of the paper, M is a complete, finite volume n -dimensional manifold of bounded non-positive curvature ($-1 \leq K \leq 0$) with fundamental group $\Gamma := \pi_1 M$ and universal cover $\widetilde{M} \rightarrow M$. Moreover, we assume that there are no arbitrarily small closed geodesics.

3.2 Margulis lemma. There are constants μ_n and I_n , depending only on the dimension n , for which the group $\langle \gamma \in \Gamma \mid d(x, \gamma x) < \mu_n \rangle$ generated by elements that move x less than μ_n is virtually nilpotent and contains a nilpotent subgroup of index $\leq I_n$. The constant μ_n is called the Margulis constant ([BGS85]).

3.3 Small ε . We fix a constant $\varepsilon > 0$ to be less than the Margulis constant and the length of the smallest closed geodesic in M . Then elements $\gamma \in \Gamma$ which have displacement $< \varepsilon$ at some point are parabolic. The “ ε -thin part”

$$\widetilde{M}_{<\varepsilon} := \{x \in \widetilde{M} \mid d(x, \gamma x) < \varepsilon \text{ for some } \gamma \in \Gamma \setminus \{1\}\}$$

is topologically (see [Gro78, BGS85]) a product $\partial \widetilde{M}_{<\varepsilon} \times [0, \infty)$. For each $x \in \widetilde{M}_{<\varepsilon}$, let

$$S_x := \{\gamma \in \Gamma \mid d_\gamma(x) < \varepsilon\},$$

$$\begin{aligned} \Gamma_x &:= \langle S_x \rangle, \\ N_x &:= \langle \gamma^{I_n!} \mid \gamma \in \Gamma_x \rangle. \end{aligned}$$

By the Margulis lemma, the group Γ_x is virtually nilpotent and N_x is a nilpotent subgroup of Γ_x . Moreover, N_x is *normal* in Γ_x and since Γ_x contains parabolic elements, so does N_x (Lemma 6.6 of [BGS85]).

3.4 Tiny δ . Fix another constant $\delta > 0$ so that $\varepsilon + 2\delta$ is still less than the Margulis constant and the length of the shortest closed geodesic in M . If $\sigma = x_0 * \dots * x_k$ is a k -simplex in \widetilde{M} of diameter $< \delta$, then at any point $x \in \sigma$ all the elements in the set

$$S_\sigma := S_{x_0} \cup \dots \cup S_{x_k}$$

have displacement $< \varepsilon + 2\delta$. This is less than the Margulis constant, so

$$\Gamma_\sigma := \langle S_\sigma \rangle = \langle \Gamma_{x_0}, \dots, \Gamma_{x_k} \rangle$$

is a virtually nilpotent group and

$$N_\sigma := \langle \gamma^{I_n!} \mid \gamma \in \Gamma_\sigma \rangle$$

is a normal nilpotent subgroup of Γ_σ containing all the groups N_τ for $\tau \subset \sigma$. Since N_σ contains parabolic elements, its center Z_σ does as well. (Section 7.6 of [BGS85].)

4 The Abelianization Map or “Mostly Abstract Nonsense with Little or No Content”

The goal of this section is to define a map μ from a triangulation of the boundary $\partial\widetilde{M}_{\leq\varepsilon}$ of $\widetilde{M}_{\leq\varepsilon}$ to an abstract complex $\Delta_{[pAb]}$ of virtual equivalence classes of Abelian subgroups of Γ . The existence of such a triangulation is addressed in “Appendix E”. Note that $\partial\widetilde{M}_{\leq\varepsilon}$ is the ordinary boundary, not the ideal boundary. Recall that two subgroups $A, B < \Gamma$ are *virtually equivalent* if they share a finite index subgroup.

The map μ is the zeroth step in defining a map $\rho: \partial\widetilde{M}_{\leq\varepsilon} \rightarrow \partial_\infty$. Then in the next section, we will construct for each virtual equivalence class of such Abelian subgroups a canonical center of mass in ∂_∞ and use it to define ρ on the vertices of $\partial\widetilde{M}_{\leq\varepsilon}$.

4.1 Complexes of Abelian and nilpotent groups. Let

$$pAb := \{A < \pi_1 M \mid A \text{ is an Abelian group containing a parabolic element}\}$$

be the set of Abelian groups in Γ containing parabolics. Also, let $[pAb]$ be the set of virtual equivalence classes of such things. Denote by Δ_{pAb} the complex whose vertices are elements of pAb and whose simplices are chains $A_0 < \dots < A_k$ of distinct such subgroups and define $\Delta_{[pAb]}$ similarly. In the same way, we define $pNil$ to be the set of nilpotent subgroups of Γ containing a parabolic element and we form a complex Δ_{pNil} . The group Γ acts on all these complexes by conjugation.

REMARK. Note that the simplices of Δ_{pAb} come with a natural ordering of vertices corresponding to inclusions of subgroups. When we take barycentric subdivisions, which we will do below, the barycentric subdivision also gets the standard ordering.

4.2 Labeling the thin part with Abelian groups. We assemble consequences of the Margulis lemma at points in the thin part in three steps. Let P be a Γ -equivariant triangulation of $\partial\widetilde{M}_{\leq\varepsilon}$ that is δ -fine, i.e. every simplex in the triangulation has diameter $< \delta$. Let P_1 be its barycentric subdivision.

- (1) Assign to each *vertex* τ of P_1 the nilpotent group

$$\mu'(\tau) := N_\tau.$$

This extends to a map

$$\mu' : P \rightarrow \Delta_{pNil}$$

because adjacent vertices in P_1 give inclusions of nilpotent groups. The nilpotent groups N_τ in the image contain parabolics because ε is small. Note that μ' is determined by the isometric Γ -action once we fix the Γ -equivariant triangulation P , so μ' is Γ -equivariant and

$$\Gamma_\tau \text{ fixes } \mu'(\tau).$$

- (2) Each vertex of the barycentric subdivision of Δ_{pNil} corresponds to a chain of nilpotent subgroups $N_0 < \cdots < N_k$. Assign to each vertex ($N_0 < \cdots < N_k$) of the barycentric subdivision of Δ_{pNil} the *group generated by the centers* Z_i of N_i ,

$$\zeta(N_0 < \cdots < N_k) := \langle Z_0, \dots, Z_k \rangle.$$

These groups are *Abelian* (because each Z_i centralizes the subgroup $\langle Z_0, \dots, Z_{i-1} \rangle$ of N_i) and contain parabolic elements. The assignment extends to a continuous map

$$\zeta : \Delta_{pNil} \rightarrow \Delta_{pAb},$$

because adding more groups to $N_0 < \cdots < N_k$ makes $\langle Z_0, \dots, Z_k \rangle$ bigger.

- (3) Pass to virtual equivalence classes via

$$\nu : \Delta_{pAb} \rightarrow \Delta_{[pAb]}.$$

Abelianization map. We call the composition $\mu := \nu \circ \zeta \circ \mu'$,

$$\mu : \partial\widetilde{M}_{\leq\varepsilon} \rightarrow \Delta_{[pAb]}.$$

the *Abelianization map*.

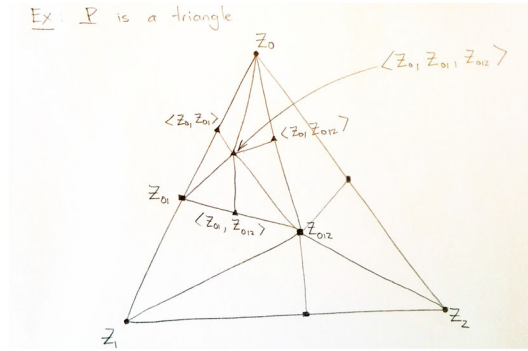


Figure 4: Labeling P_2 .

REMARK. In summary, the map μ was described on each simplex in terms of the *second barycentric subdivision* P_2 of P (Figure 4). The reason is that we can connect nilpotent groups N_0 and N_1 corresponding to vertices of a simplex in P by inclusions of nilpotent groups via

$$N_0 \leq N_{01} \geq N_1, \tag{2}$$

which is a path in the barycentric subdivision P_1 , and we can connect the centers Z_0 and Z_1 by inclusions of *Abelian* groups

$$Z_0 \leq \langle Z_0, Z_{01} \rangle \geq Z_{01} \leq \langle Z_{01}, Z_1 \rangle \geq Z_1, \tag{3}$$

which is a path in the second barycentric subdivision P_2 . It is necessary to pass from P_1 to P_2 when we go Abelian because unlike the situation with nilpotent subgroups, it is in general not true that $Z_0 \leq Z_{01} \geq Z_1$.

4.3 For later use (to stay thin). Let $\sigma = (\tau_0 \subset \dots \subset \tau_k)$ be an ordered k -simplex in the barycentric subdivision P_1 , where the order is by inclusion of simplices in the original complex. We noted above that Γ_{τ_i} fixes the vertex $\mu'(\tau_i)$. Since $\mu = \nu \circ \zeta \circ \mu'$ and the maps ν and ζ are obviously Γ -equivariant, we see that the “bottom” group Γ_{τ_0} in the chain $\Gamma_{\tau_0} < \dots < \Gamma_{\tau_k}$ fixes the entire simplex $\mu(\sigma) = \nu \circ \zeta \circ \mu'(\tau_0 \subset \dots \subset \tau_k)$. Moreover, there is an element in Γ_{τ_0} that is $(\varepsilon + 2\delta)$ -small everywhere on σ . This is because σ is in the subdivision of a simplex of P whose diameter is $\leq \delta$ and at least one of whose vertices x has a non-trivial element $\gamma \in S_x < \Gamma_{\tau_0}$ with $d_\gamma(x) \leq \varepsilon$. Thus $d_\gamma \leq \varepsilon + 2\delta$ everywhere on σ . In summary,

for every $x \in \partial \widetilde{M}_{\leq \varepsilon}$ there is $1 \neq \gamma \in \Gamma$ such that
 $d_\gamma(x) \leq \varepsilon + 2\delta$ and $\gamma(\mu(x)) = \mu(x)$.

5 Center of Mass for an Abelian Group with Parabolics

In this section we will build a map β on the set of *vertices* of the complex $\Delta_{[pAb]}$

$$\beta : \Delta_{[pAb]}^{(0)} \rightarrow (\partial_\infty, \text{Td}) \tag{4}$$

such that

- the map β is Γ -equivariant,
- horospheres centered at $\beta([A])$ are preserved by the centralizer C_A for any Abelian group A' virtually equivalent to A , and
- for any simplex σ in $\Delta_{[pAb]}$ we have in the Tits metric

$$\text{diam}(\beta(\sigma^{(0)})) < \pi/2, \quad (5)$$

where $\sigma^{(0)}$ is the set of vertices of σ . Then in later sections, we will discuss how to correctly fill in β with simplices at infinity in order to finally obtain a map $\rho: \partial\widetilde{M}_{\leq \varepsilon} \rightarrow \partial_\infty$.

Every Abelian group A containing a parabolic isometry has a nonempty fixed set $\text{Fix}(A)$ at infinity with a canonical *center of mass* ξ_A . This construction can be found in Appendix 3.B of [BGS85] and a variation is in 3.5 of [Ebe96]. We review it in ‘‘Appendix B’’ of this paper. It is, however, not invariant under virtual equivalences, so we cannot use it to define β . The plan for this section is to first recall properties of the classical center of mass, and then, inspired by this, construct a canonical center of mass that depends only on the virtual equivalence class of the parabolic Abelian group. We will prove similar properties for this new center of mass and use it to define β .

5.1 The classical center of mass for a parabolic Abelian group A The center of mass ξ_A has the crucial property that for every $y \in \text{Fix}(A)$,

$$\text{Td}(\xi_A, y) \leq \pi/2.$$

Also, since the construction of ξ_A is canonical, ξ_A is fixed by the normalizer of A . This implies that for a simplex $\sigma = (A_0 < \dots < A_k)$ in Δ_{pAb} each group A_i fixes all the points ξ_{A_j} , and from this the first property gives

$$\text{Td}(\xi_{A_i}, \xi_{A_j}) \leq \pi/2.$$

In addition, the following feature of ξ_A is fundamental and is crucial in the next section.

PROPOSITION 9 (Preserved horospheres). *The centralizer C_A preserves horospheres centered at ξ_A .*

Proof. Let h be a Busemann function centered at ξ_A . We need to show that h is γ -invariant under isometries $\gamma \in C_A$. Since $|\nabla h| = 1$ we have

$$\frac{|h(\gamma^n x) - h(x)|}{n} \leq \frac{d(\gamma^n x, x)}{n}.$$

Since $\gamma \in C_A$ we already know it fixes ξ_A , so the quantity on the left is independent of n and x and is equal to $|h(\gamma x) - h(x)|$. Letting $n \rightarrow \infty$ and using the well known formula for the infimum displacement of an isometry (Lemma 6.6.(2) in [BGS85])

$$|\gamma| := \inf_{x \in \widetilde{M}} d(\gamma x, x) = \lim_{n \rightarrow \infty} \frac{d(\gamma^n x, x)}{n}$$

we get

$$|h(\gamma x) - h(x)| \leq |\gamma|.$$

So we see that h is γ -invariant whenever $|\gamma| = 0$.

Now, suppose that $|\gamma| > 0$. Then, according to Karlsson-Margulis ([KM99], see also ‘‘Appendix C’’), there are geodesic rays $r_{\pm} = [x, \eta^{\pm})$ sublinearly tracking the positive and negative γ -orbits, i.e.

$$\lim_{n \rightarrow \infty} \frac{d(\gamma^n x, r_+(n|\gamma|))}{n} = 0,$$

and

$$\lim_{n \rightarrow \infty} \frac{d(\gamma^{-n} x, r_-(n|\gamma|))}{n} = 0.$$

It follows from this and the formula for $|\gamma|$ that

$$2|\gamma| = \lim_{t \rightarrow \infty} \frac{d(r_+(t|\gamma|), r_-(t|\gamma|))}{t}.$$

Reparametrizing and using the \angle -metric description on p. 36 of [BGS85] we get

$$1 = \lim_{t \rightarrow \infty} \frac{d(r_+(t), r_-(t))}{2t} = \sin\left(\frac{\angle(\eta^+, \eta^-)}{2}\right).$$

Therefore $\angle(\eta^+, \eta^-) = \pi$ which implies that

$$\text{Td}(\eta^+, \eta^-) \geq \pi.$$

On the other hand, since A commutes with γ it fixes the limit points $\lim_{n \rightarrow \infty} \gamma^n x = \eta^+$ and $\lim_{n \rightarrow \infty} \gamma^{-n} x = \eta^-$. Therefore

$$\text{Td}(\eta^{\pm}, \xi_A) \leq \pi/2.$$

Putting these two inequalities together we get

$$\pi \leq \text{Td}(\eta^+, \eta^-) \leq \text{Td}(\eta^+, \xi_A) + \text{Td}(\xi_A, \eta^-) \leq \pi$$

and consequently

$$\text{Td}(\eta^{\pm}, \xi_A) = \pi/2.$$

This implies that $\nabla h \cdot \dot{r}_+ \rightarrow 0$ along the geodesic ray r_+ (see 4.2 in [BGS85]) and consequently

$$\frac{h(r_+(n|\gamma|)) - h(x)}{n} = \frac{1}{n} \int_0^{n|\gamma|} |\nabla h \cdot \dot{r}_+| dt \rightarrow 0$$

as $n \rightarrow \infty$. Putting this together with

$$\frac{|h(\gamma^n x) - h(r_+(n|\gamma|))|}{n} \leq \frac{d(\gamma^n x, r_+(n|\gamma|))}{n} \rightarrow 0$$

we get that $\frac{h(\gamma^n x) - h(x)}{n} \rightarrow 0$. Since this quantity is actually constant and equal to $h(\gamma x) - h(x)$, we conclude that h is γ -invariant. \square

5.2 Dealing with finite index issues. Centers of mass of virtually equivalent Abelian groups might be different. Our goal in this subsection is to pick a single point at infinity that will play the role of the center of mass for the whole virtual equivalence class $[A]$ of A . We will do this by constructing a center of mass for the union of fixed point sets of all the groups virtually equivalent to A , which is also equal to

$$F_A := \bigcup_{n \in \mathbb{N}} \text{Fix}(n!A),$$

since every group virtually equivalent to A contains some $n!A$ as a subgroup. It is easy to see and occasionally useful to remember that

$$[A] = [A'] \implies F_A = F_{A'}, \text{ and} \quad (6)$$

$$B \leq A \implies F_B \supseteq F_A. \quad (7)$$

A two-step center of mass construction. Now we construct a center of mass for F_A that depends only on the virtual equivalence class of A . We do this in two steps.

(*Step 1*) First, we will show that there is a point ξ in F_A so that any other point of F_A is within $\pi/2$ of ξ . To show this, let

$$B_{n,A} := \{x \in \text{Fix}(A) \mid \text{Td}(x, y) \leq \pi/2 \text{ for all } y \in \text{Fix}(n!A)\}.$$

The sets $B_{n,A}$ are closed in the sphere topology (by the lower semicontinuity of Td , see “Appendix B”) and nested:

$$B_{0,A} \supseteq B_{1,A} \supseteq B_{2,A} \supseteq \cdots.$$

They are also non-empty because the center of mass $\xi_{n!A}$ of the fixed point set of $n!A$ is fixed by A and therefore $\xi_{n!A} \in B_{n,A}$. So there is a point ξ contained in the intersection $\bigcap_n B_{n,A}$. For this $\xi \in \text{Fix}(A)$ we have $\text{Td}(\xi, y) \leq \pi/2$ for all $y \in F_A$. This finishes the first step.

(*Step 2*) The point ξ constructed in Step 1 may not be unique, and we denote the set of all such points by

$$B_{[A]} := \{x \in F_A \mid \text{Td}(x, y) \leq \pi/2 \text{ for all } y \in F_A\}.$$

This set is our collection of *potential* centers of mass. The second step is to pick in a canonical way a single point from this set.

It is clear that $B_{[A]}$ has Tits diameter $\leq \pi/2$, so *if it was closed in the sphere topology* then one way to pick a unique point would be to take the center of $B_{[A]}$ (in the sense of “Appendix B”). Unfortunately, the set $B_{[A]}$ may not be closed in the sphere topology, so we first replace it by its closure in the sphere topology $\overline{B}_{[A]}$. It is easy to see (again, using lower semicontinuity of Td) that this closure is equal to

$$\overline{B}_{[A]} = \{x \in \overline{F}_A \mid \text{Td}(x, y) \leq \pi/2 \text{ for all } y \in \overline{F}_A\},$$

where \overline{F}_A is the closure of F_A in the sphere topology. In particular, $\overline{B}_{[A]}$ still has Tits diameter $\leq \pi/2$. Therefore, the function

$$\tau(\cdot) := \sup_{x \in \overline{B}_{[A]}} \text{Td}(\cdot, x)$$

has infimum

$$\inf \tau < \pi/2 - \alpha$$

for a positive constant $\alpha = \alpha_n > 0$ that only depends on the dimension n . This infimum is attained at a unique point in ∂_∞ which we will denote by $\xi_{[A]}$. In other words, there is a unique ball of smallest radius containing $\overline{B}_{[A]}$. This ball is centered at $\xi_{[A]}$ and has radius $\tau(\xi_{[A]})$. See “Appendix B” for everything in this paragraph. We call the point $\xi_{[A]}$ the *center of mass of $[A]$* .

What remains to be shown is that $\xi_{[A]}$ is actually contained in $B_{[A]}$. We prove this in the remainder of the subsection. We begin with

LEMMA 10. *The set $B_{[A]}$ is convex.*

Proof. Let $x_0, x_1 \in B_{[A]}$ and let x_t be a point on the geodesic segment in ∂_∞ connecting them. There are virtually equivalent groups $A_0, A_1 \in [A]$ fixing x_0 and x_1 , so the entire geodesic segment is fixed by the group $(A_0 \cap A_1) \in [A]$. Moreover, since ∂_∞ is CAT(1) we see for any $y \in F_A$ that $\text{Td}(x_t, y) \leq \pi/2$ by comparison with the round sphere. Thus $x_t \in B_{[A]}$. □

We will use this to show that $\xi_{[A]}$ is contained in $\overline{B}_{[A]}$. If the closure $\overline{B}_{[A]}$ was convex, then this would follow easily (see “Appendix B”). But, we only know that $\overline{B}_{[A]}$ is the closure in the sphere topology of a Td-convex set. So, we need the following lemma.

LEMMA 11. *Fix $\alpha > 0$. Suppose C is a convex set of diameter $\leq \pi/2$. Let \overline{C} be its closure in the sphere topology. Let $\xi \in \partial_\infty$ be a point for which*

$$\text{Td}(\xi, y) \leq \pi/2 - \alpha \quad \text{for every } y \in C.$$

Then there is a point $x \in \overline{C}$ so that

$$\text{Td}(x, y) \leq \text{Td}(\xi, y) \quad \text{for all } y \in C.$$

Proof. First, let $r = \inf_{y \in C} \text{Td}(\xi, y)$. Then there is a sequence of points $x_i \in C$ so that $\text{Td}(\xi, x_i) \rightarrow r$. After passing to a subsequence, we may assume that in the sphere topology $x_i \rightarrow x \in \overline{C}$. Now, pick a point $y \in C$. Since C is convex, the geodesic segment $[x_i, y]$ is contained in C . On this segment there is a unique closest point y_i to ξ . Since it is the closest point to ξ on the segment $[x_i, y]$, at this point the angles $\angle_{y_i}(x_i, \xi)$ and $\angle_{y_i}(\xi, y)$ are both obtuse (that is, greater or equal to $\pi/2$). Therefore, triangle comparison with obtuse triangles on the round sphere gives

$$\text{Td}(y_i, y) \leq \text{Td}(\xi, y),$$

and also

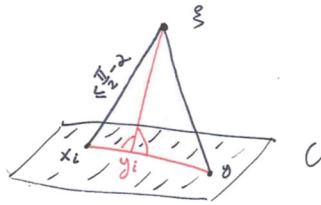
$$\mathrm{Td}(x_i, y_i) \leq \cos^{-1} \left(\frac{\cos(\mathrm{Td}(x_i, \xi))}{\cos(\mathrm{Td}(y_i, \xi))} \right).$$

As $i \rightarrow \infty$ the right hand side of this tends to zero because the denominator

$$\cos(\mathrm{Td}(y_i, \xi)) \geq \cos(\pi/2 - \alpha) > 0$$

doesn't approach zero and

$$r \leq \mathrm{Td}(y_i, \xi) \leq \mathrm{Td}(x_i, \xi) \rightarrow r.$$



Therefore, using lower semicontinuity of $\mathrm{Td}(\cdot, y)$ we get

$$\begin{aligned} \mathrm{Td}(x, y) &\leq \liminf_i \mathrm{Td}(x_i, y) \\ &\leq \liminf_i (\mathrm{Td}(x_i, y_i) + \mathrm{Td}(y_i, y)) \\ &\leq \mathrm{Td}(\xi, y). \end{aligned} \quad \square$$

So, since B_A is a convex set of diameter $\leq \pi/2$ and $\xi_{[A]}$ is the unique point at which ρ attains its infimum, we conclude that

$$\xi_{[A]} \in \overline{B}_A.$$

Therefore

$$\mathrm{Td}(\xi_{[A]}, y) \leq \pi/2 \quad \text{for all } y \in F_A, \text{ and} \quad (8)$$

$$\mathrm{Td}(\xi_{[A]}, y) \leq \pi/2 - \alpha \quad \text{for all } y \in B_{[A]}. \quad (9)$$

The set F_A is preserved by A so its center of mass $\xi_{[A]}$ is fixed by A , and therefore

$$\xi_{[A]} \in B_{[A]},$$

which is what we needed to show. This finishes the second step.

REMARK. At this point the reader can safely forget about the closures $\overline{B}_{[A]}$. While they appeared in the construction of $\xi_{[A]}$ they will never appear again.

5.3 The map $\beta : \Delta_{[pAb]}^{(0)} \rightarrow (\partial_\infty, \mathbf{Td})$. We set

$$\beta([A]) := \xi_{[A]}.$$

Let us verify that it has all the properties we promised at the beginning of Section 5. First, it follows from the construction that $\xi_{[\gamma A \gamma^{-1}]} = \gamma \xi_{[A]}$ so the map β is Γ -equivariant. Second, for any A' virtually equivalent to A , the centralizer $C_{A'}$ of A' fixes $\xi_{[A']} = \xi_{[A]}$, so the proof of Proposition 9 applies word-for-word and shows that $C_{A'}$ preserves horospheres centered at $\xi_{[A]}$. To prove the third bullet we proceed as follows. For any simplex $\sigma = ([A_0] < \dots < [A_k])$ in $\Delta_{[pAb]}$ we can pick representatives A_i so that $A_0 < \dots < A_k$. Then

$$F_{A_0} \supseteq \dots \supseteq F_{A_k},$$

so by (8) if $i \leq j$ we get $\text{Td}(\xi_{[A_i]}, y) \leq \pi/2$ for all $y \in F_{A_j}$. Since for all i, j each group A_j fixes all the points $\xi_{[A_i]}$ we also have $\xi_{[A_i]} \in F_{A_j}$. In summary

$$\xi_{[A_i]} \in B_{[A_j]} \quad \text{for all } i \leq j.$$

The upshot of these gymnastics is that (9) implies that for all i, j

$$\text{Td}(\xi_{[A_i]}, \xi_{[A_j]}) \leq \pi/2 - \alpha.$$

In other words, the diameter of $\beta(\sigma^{(0)})$ is strictly less than $\pi/2$, which is what we wanted to show.

6 A Criterion for and the Necessity of Being Divergent

Now that we have defined β (and thus $\rho := \beta \circ \mu$) on vertices, we need to extend it to each simplex. The extension must be canonical and satisfy a divergence property. The goal of this section is to carefully discuss this notion of a divergent simplex at infinity, give a criterion for when a simplex is divergent, and illustrate how it is useful in the context of the main theorem.

In this section we will need the observation that $M_{\geq \varepsilon}$ is compact. It is immediate that $M_{\geq \varepsilon}$ is closed. It is bounded because if not, one can pack infinitely many disjoint ε -balls into M , which means that M has infinite volume, but the volume of M is finite. So M is closed and bounded, and therefore compact.

6.1 Divergent rays, divergent simplices and divergent maps

6.1.1 *Divergent rays.* Recall that a geodesic ray in \widetilde{M} is *divergent* if its projection under the covering map p leaves all compact sets. It turns out, as we will see, that parabolic Abelian subgroups $A < \Gamma$ give geodesic rays $[x, \xi_A)$ which project to *divergent* rays in M , and thus determine distinguished directions to go to infinity in M .

PROPOSITION 12. Any ray $[x, \xi_{[A]})$ in \widetilde{M} projects to a divergent ray in M .

The key to obtaining results of this sort is the strong invariance of ξ_A established in Proposition 9. The centralizer C_A preserves horospheres centered at $\xi_{[A]}$, so Proposition 12 follows from Lemma 13 below. This is a way to produce divergent rays in M .

LEMMA 13. Suppose A is a subgroup of Γ . If C_A preserves horospheres centered at ξ and A fixes ξ , then any geodesic ray $r : [0, \infty) \rightarrow \widetilde{M}$ with endpoint $r(\infty) = \xi$ projects to a divergent ray in M .

Proof. We prove the contrapositive. If the projection of the geodesic ray r to M does not diverge then there is a sequence of times $t_i \rightarrow \infty$ and elements $g_i \in \Gamma$ so that $\{g_i r(t_i)\}$ converges to a point x_0 in \widetilde{M} . We will construct out of this an element in C_A that does not preserve a Busemann function h centered at ξ . This will prove the Proposition. Let $D := \sup_i d(x_0, g_i r(t_i))$.

CLAIM. After passing to a subsequence of $\{g_i\}$, we have $g_j^{-1} g_i \in C_A$.

Recall that for any isometry ρ , it follows from triangle inequality that

$$|d_\rho(x) - d_\rho(y)| \leq 2d(x, y).$$

This implies that for any element γ ,

$$|d_{g_i \gamma g_i^{-1}}(x_0) - d_{g_i \gamma g_i^{-1}}(g_i r(t_i))| \leq 2d(x_0, g_i r(t_i)) \tag{10}$$

$$\leq 2D. \tag{11}$$

If γ fixes $r(\infty) = \xi$, we also get

$$d_{g_i \gamma g_i^{-1}}(g_i r(t_i)) = d_\gamma(r(t_i)), \tag{12}$$

$$\leq d_\gamma(r(0)), \tag{13}$$

so that $\{d_{g_i \gamma g_i^{-1}}(x_0)\}_{i=1}^\infty$ is bounded. Thus, there are only finitely many different conjugates in the sequence $\{g_i \gamma g_i^{-1}\}_{i=1}^\infty$. After passing to a subsequence, we may assume that all the conjugates are the same, i.e. that

$$g_1 \gamma g_1^{-1} = g_2 \gamma g_2^{-1} = \dots, \tag{14}$$

and consequently that $g_j^{-1} g_i$ commutes with γ .

In the special case when $A = \langle \gamma_1, \dots, \gamma_r \rangle$ is a finitely generated group fixing ξ we can do the above argument for each one of the generators. So, after passing to subsequences finitely many times, we get a sequence $\{g_i\}$ for which $g_j^{-1}g_i$ commutes with the entire group $A = \langle \gamma_1, \dots, \gamma_r \rangle$.

In general, A is countable so we get the same result via diagonal argument.

CLAIM. For large enough j , the element $g_j^{-1}g_i$ does not preserve h .

Note that $d(g_j^{-1}g_i r(t_i), r(t_j)) = d(g_i r(t_i), g_j r(t_j))$ is bounded by $2D$, so

$$|h(g_j^{-1}g_i r(t_i)) - h(r(t_j))| \leq 2D.$$

On the other hand, as $j \rightarrow \infty$ we have

$$h(r(t_j)) = (h(r(0)) - t_j) \rightarrow -\infty. \tag{15}$$

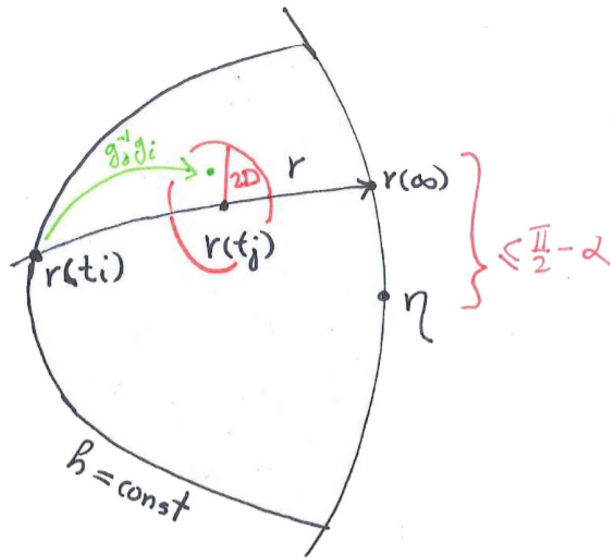
Therefore $\lim_{j \rightarrow \infty} h(g_j^{-1}g_i r(t_i)) = -\infty$. This implies that h is not $g_j^{-1}g_i$ -invariant for a fixed i and large enough j .

So we've found an element $g_j^{-1}g_i \in C_A$ that does not preserve horospheres centered at $r(\infty)$. This proves the proposition. \square

Wiggle room in the divergent ray argument. It is good to notice that the divergent ray argument is not delicate. There is quite a bit of "wiggle room" in the argument. If one looks through the proof, one sees that the assumptions can be weakened. We only need to know that

- (1) the group A fixes the point at infinity $r(\infty)$,
- (2) there is *some point* η such that C_A preserves horospheres at η , and
- (3) there is a positive constant $\alpha > 0$ such that

$$\text{Td}(\eta, r(\infty)) \leq \pi/2 - \alpha.$$



In other words, we can separate the point η from the endpoint of the geodesic ray $r(\infty)$ in the argument as illustrated in the figure above. The point $r(\infty)$ just needs to be fixed by A as long as the (much stronger) condition that horospheres are preserved by the entire centralizer is satisfied by some nearby point η . Having phrased things in this way, we note that *we can vary the endpoint $r(\infty)$ of the geodesic ray, as long as all the rays we use satisfy (1) and (3) for a single point η and a single constant $\alpha > 0$.* Finally, note that we can vary the startpoint of the geodesic ray r in a bounded set. So, we arrive at the following Proposition, which produces divergent sectors.

PROPOSITION 14. *Suppose A is a subgroup of Γ , B is a bounded subset of \widetilde{M} , and h is a C_A -invariant Busemann function centered at a point $\eta \in \partial_\infty$. Then for every $\varepsilon > 0$ there is a constant $T := T_{A,\varepsilon,\eta,\alpha,B}$ so that any geodesic ray $r : [0, \infty) \rightarrow \widetilde{M}$ with $r(0) \in B$ and $r(\infty)$ satisfying (1) and (3) has*

$$r(t) \in \widetilde{M}_{\leq \varepsilon} \quad \text{for all } t \geq T.$$

Proof. Suppose the conclusion is not true. Then there are times $t_i \rightarrow \infty$, elements $g_i \in \Gamma$, and rays r_i with $r_i(0) \in B$ and $r_i(\infty)$ satisfying (1) and (3) such that $\{g_i r_i(t_i)\}$ converges to a point $x_0 \in \widetilde{M}$. As before, using (1) we show that after passing to a subsequence we can assume $g_j^{-1} g_i \in C_A$ for all i, j . As before, h is a Busemann function whose level sets are horospheres centered at η and $|h(g_j^{-1} g_i r_i(t_i)) - h(r_j(t_j))| \leq 2D$. Condition (3) implies that $r'_j(t)$ and the tangent to the geodesic ray from $r_j(t)$ to η is at an angle at most $\pi/2 - \alpha$ for all t , which implies that

$$\begin{aligned} h(r_j(t_j)) &\leq h(r_j(0)) - t_j \cdot \sin \alpha, \\ &\rightarrow -\infty. \end{aligned}$$

So we again conclude that $\lim_{j \rightarrow \infty} h(g_j^{-1} g_i r_i(t_i)) \rightarrow -\infty$. This contradicts the assumption that h is C_A -invariant, so it proves the proposition. \square

6.1.2 *Divergent maps and divergent simplices.* In order to state our main application of Proposition 14, we introduce the following terminology. A family of maps $\{\varphi_t : X \rightarrow \widetilde{M}\}_{t \in \mathbb{R}^+}$ *diverges over M* if for any $\varepsilon > 0$ there is T so that

$$\text{Im}(\varphi_t) \subset \widetilde{M}_{\leq \varepsilon} \quad \text{for all } t \geq T.$$

Two families $\{\varphi_t, \psi_t : X \rightarrow \widetilde{M}\}_{t \geq 0}$ are *asymptotic* if for every compact set $K \subset X$, the distance $\sup_{x \in K, t \geq 0} d(\varphi_t(x), \psi_t(x))$ is finite. Next, fix a basepoint $z \in \widetilde{M}$, let $S_z(t)$ be the sphere of radius t centered at z , and let

$$c_t = c_t^z : \partial_\infty \rightarrow S_z(t)$$

be the geodesic retraction that sends $\xi \in \partial_\infty$ to the point $[z, \xi]_t$ obtained by flowing for a time t along the geodesic ray from z to ξ . We say that a singular simplex $\lambda : \Delta^k \rightarrow \partial_\infty$ *diverges over M* if the family $c_t \circ \lambda$ diverges over M . Whenever it

doesn't cause confusion, we will omit “over M ” and just say that the simplex λ diverges. A direct corollary of Proposition 14 is the following criterion for finding divergent simplices.

COROLLARY 15. *Suppose horospheres centered at $\eta \in \partial_\infty$ are C_A -invariant. For $\alpha > 0$, any simplex contained in $\text{Fix}(A) \cap B_{\pi/2-\alpha}(\eta)$ diverges over M .*

It is easy to see that the definition of divergence for simplices does not depend on the choice of basepoint z . The underlying reason is that for different basepoints z and z' the cone homotopies $c_t^z \circ \lambda$ and $c_t^{z'} \circ \lambda$ are asymptotic. Somewhat more generally we have the following lemma which will be useful in the next subsection.

LEMMA 16. *Suppose that two families $\{\varphi_t, \psi_t : X \rightarrow \widetilde{M}\}_{t \geq 0}$ are asymptotic. If φ_t diverges then for any compact subset $K \subset X$*

- for sufficiently large t , the geodesic homotopy between $\varphi_t|_K$ and $\psi_t|_K$ is inside $\widetilde{M}_{<\varepsilon}$, and in particular
- $\psi_t|_K$ diverges.

Proof. Let $D := \sup_{x \in K, t \geq 0} d(\varphi_t(x), \psi_t(x))$. This is finite because φ_t and ψ_t are asymptotic. We claim that since φ_t diverges in M and the injectivity radius function on M is proper, there is a time T so that for $t \geq T$ the closed D -neighborhood of $\varphi_t|_K$ is contained in $\widetilde{M}_{\leq \varepsilon}$. To see this, note that the injectivity radius being proper implies that there is $\varepsilon' < \varepsilon$ such that $d(M_{\geq \varepsilon}, M_{\leq \varepsilon'}) > D$. Then also $d(\widetilde{M}_{\geq \varepsilon}, \widetilde{M}_{\leq \varepsilon'}) > D$ and therefore once t is large enough so that φ_t is in $\widetilde{M}_{\leq \varepsilon'}$ its D -neighborhood will be in $\widetilde{M}_{\leq \varepsilon}$.

Since the geodesic homotopy between $\varphi_t|_K$ and $\psi_t|_K$ is in this D -neighborhood, we get the first bullet point. The second follows immediately from the first. \square

6.1.3 Filling β in with divergent simplices. We are now almost ready to establish our basic collapse result, Theorem 18 below. In order to do this, we will need to extend the “center of mass” map β to the entire complex $\Delta_{[pAb]}$ by filling it in with divergent simplices in a natural way. The resulting $\beta : \Delta_{[pAb]} \rightarrow (\partial_\infty, \text{Td})$ should be continuous, of course, and it should

- (1) be Γ -equivariant,
- (2) send a vertex $[A]$ to its center of mass ξ_A , and
- (3) for each simplex σ of $\Delta_{[pAb]}$, the simplex $\beta(\sigma)$ should diverge.

We will express (3) by saying “ β diverges on simplices”. Because of condition (2) and by Proposition 12, the *vertices* of a simplex σ are mapped to the endpoints of divergent rays by β . The extra condition (3) says that for all *points* in a simplex $p \in \sigma$ all the rays $[z, \beta(p))$ diverge, and they do so uniformly. In practice, if the map β doesn't distort things too much one can get (3) from (1) and (2):

LEMMA 17. *Suppose β satisfies (1) and (2), and let $\alpha > 0$. If $\beta(\sigma)$ is in the $(\pi/2-\alpha)$ -neighborhood of the vertex set $\beta(\sigma^{(0)})$, then $\beta(\sigma)$ diverges.*

Proof. Let σ be a simplex in $\Delta_{[pAb]}$ represented by a chain of Abelian groups $A_0 < \dots < A_k$. Then σ is fixed pointwise by each group A_i since all the A_i commute. Since β is equivariant, $\beta(\sigma)$ is fixed by A_i as well. Moreover, the group C_{A_i} preserves horospheres centered at $\beta([A_i])$. Therefore Corollary 15 implies

$$\beta(\sigma) \text{ diverges if for some vertex } v \text{ we have } \beta(\sigma) \subset B_{\pi/2-\alpha}(v). \quad (16)$$

Even if the entire simplex is not contained in the $(\pi/2 - \alpha)$ -neighborhood of a single vertex, applying Corollary 15 to a vertex shows that the portion of the simplex that lies in the $(\pi/2 - \alpha)$ -neighborhood of that vertex diverges. Doing this for each one of the vertices of σ proves the lemma. \square

We will describe different versions of β in Section 8 and use the lemma to check that for each of these versions all simplices $\beta(\sigma)$ diverge.

6.2 The perks of being a divergent simplex. Suppose we have, one way or another, got our hands on such a “divergent simplex” map β . Let us explain how it can be used, together with the Abelianization map μ constructed in section 4, to understand the topology of the thin part. The idea is that the composition $\rho := \beta \circ \mu$

$$\partial \widetilde{M}_{\leq \varepsilon} \xrightarrow{\mu} \Delta_{[pAb]} \xrightarrow{\beta} (\partial_\infty, \text{Td})$$

tells us how to push topological features to infinity while staying in the thin part. To make this precise, denote by

$$\begin{aligned} (\beta \circ \mu)_t : \partial \widetilde{M}_{\leq \varepsilon} &\rightarrow \widetilde{M} \\ x &\mapsto [x, \beta \circ \mu(x)]_t \end{aligned}$$

the map which sends a point $x \in \partial \widetilde{M}_{\leq \varepsilon}$ to the point obtained by going for a time t along the geodesic ray $[x, \beta \circ \mu(x)]$. Note that it is Γ -equivariant and that it “approaches $\beta \circ \mu$ ” in the sense that

$$(\beta \circ \mu)_t \text{ is asymptotic to the cone homotopy } c_t \circ \beta \circ \mu.$$

Two additional key features of this map (proved below) is that it stays in the thin part for all t and pushes further into the thin part for large t . This lets us collapse any compact subset in $\widetilde{M}_{\leq \varepsilon}$ to a subset of topological dimension less than or equal to $\dim(\text{Im}(\beta \circ \mu), \angle_x)$.

Theorem 18. *Let $\mu : \partial \widetilde{M}_{\leq \varepsilon} \rightarrow \Delta_{[pAb]}$ be the Abelianization map. Suppose there is a Γ -equivariant map $\beta : \Delta_{[pAb]} \rightarrow (\partial_\infty, \text{Td})$ which sends vertices $[A]$ to their centers of mass $\xi_{[A]}$ and which diverges on simplices. Then*

- $(\beta \circ \mu)_t : \partial \widetilde{M}_{\leq \varepsilon} \rightarrow \widetilde{M}$ is in $\widetilde{M}_{\leq \varepsilon+2\delta}$ for all $t \geq 0$, and
- $(\beta \circ \mu)_t$ diverges over M .

Denote the dimension of the image of $\beta \circ \mu$ in the sphere topology by

$$d := d_\beta = \dim(\text{Im}(\beta \circ \mu), \angle_x).$$

Then the inclusion of any compact subset $\varphi : K \hookrightarrow \partial\widetilde{M}_{\leq \varepsilon}$ can be homotoped in $\widetilde{M}_{\leq \varepsilon+2\delta}$ to a map $\hat{\varphi}$ with image of dimension $\leq d$.

We will prove this theorem at the end of this section. Next, we present several topological consequences of Theorem 18.

COROLLARY 19. *Assume the hypotheses of Theorem 18. Then*

$$H_{>d}(\widetilde{M}_{\leq \varepsilon}) = 0.$$

Proof. Let $\varphi : F \rightarrow \partial\widetilde{M}_{\leq \varepsilon}$ be a homology cycle of dimension $k > d$. By the previous theorem, it can be homotoped in $\widetilde{M}_{\leq \varepsilon+2\delta}$ to a map $\hat{\varphi}$ with d -dimensional image. Since M is tame, we can push the homotopy a little along the product direction of $\widetilde{M}_{\leq \varepsilon+2\delta} = \partial\widetilde{M}_{\leq \varepsilon+2\delta} \times [0, \infty)$ so that it stays in the ε -thin part $\widetilde{M}_{\leq \varepsilon}$. By a standard argument (recalled in ‘‘Appendix D’’) we can further homotope the map $\hat{\varphi}$ in $\widetilde{M}_{\leq \varepsilon}$ to a map $\bar{\varphi}$ whose image lands in the d -skeleton $\widetilde{M}_{\leq \varepsilon}^{(d)}$ of a triangulation of $\widetilde{M}_{\leq \varepsilon}$. Since $d < k$, we conclude that φ is zero on k -dimensional homology. \square

COROLLARY 20. *Assume the hypotheses of Theorem 18. If $d \leq 1$ then each component of $\widetilde{M}_{\leq \varepsilon}$ is aspherical.*

Proof. If $d \leq 1$ then, arguing as in the previous corollary, any map $\varphi : S^k \rightarrow \widetilde{M}_{\leq \varepsilon}$ can be homotoped in $\widetilde{M}_{\leq \varepsilon}$ to factor through a graph. Therefore each component of $\widetilde{M}_{\leq \varepsilon}$ is aspherical. \square

In order for these results to mean anything, we need some control over the dimension d . The first observation is that d is also equal to the dimension of the image of $\rho = \beta \circ \mu$ in the Tits metric.

PROPOSITION 21. *Assume the hypotheses of Theorem 18. Then*

$$d = \dim(\text{Im } \rho, \angle_x) = \dim(\text{Im } \rho, \text{Td}) \leq \dim(\partial_\infty, \text{Td}).$$

Proof. As a continuous image of a compact set, $\rho(\sigma)$ is compact in the Tits metric. Therefore, the identity map $(\rho(\sigma), \text{Td}) \rightarrow (\rho(\sigma), \angle_x)$ is a homeomorphism (since it is a continuous bijection from a compact space to a Hausdorff space) and thus it preserves topological dimensions, i.e. $\dim(\rho(\sigma), \angle_x) = \dim(\rho(\sigma), \text{Td})$. Now, since $\text{Im}(\rho)$ is a countable union of the images of simplices $\rho(\sigma)$, its dimension is equal, in either the \angle_x or the Td-metric, to the supremum of the dimensions of the simplices by the countable sum theorem in dimension theory (see III.2 in [HW48]). We conclude that

$$\dim(\text{Im } \rho, \angle_x) = \sup_\sigma \dim(\rho(\sigma), \angle_x) = \sup_\sigma \dim(\rho(\sigma), \text{Td}) = \dim(\text{Im } \rho, \text{Td})$$

which proves the proposition. \square

REMARK. We initially defined d via the sphere topology on ∂_∞ because this is the topology for which the cone map c_t is a homeomorphism onto its image. This is unsatisfying because the sphere topology does not reflect in any way the geometry of the universal cover. After all, the metric space $(\partial_\infty, \angle_x)$ is just a round $(n-1)$ -sphere. Proposition 21 is useful because the topological dimension of the Tits boundary is a geometrically meaningful quantity that can be (and often is) much smaller than $n-1$. For instance, for symmetric spaces it is one less than the dimension of a maximal flat.

We end this section by giving the proof of Theorem 18.

Proof of Theorem 18. The proof consists of several steps. First we prove the two properties of the homotopy $(\beta \circ \mu)_t$ mentioned in the bullets and then we explain how to use these properties to collapse K to a d -dimensional subset.

CLAIM. *The homotopy $(\beta \circ \mu)_t$ is in the $(\varepsilon + 2\delta)$ -thin part $\widetilde{M}_{\leq \varepsilon + 2\delta}$ for all t .*

Recall from 4.3 that for every point $x \in \partial \widetilde{M}_{\leq \varepsilon}$ there is a non-trivial element $\gamma \in \Gamma$ that is $(\varepsilon + 2\delta)$ -small at x and fixes $\mu(x)$. Because β is Γ -equivariant γ also fixes $\beta \circ \mu(x)$, so γ is $(\varepsilon + 2\delta)$ -small on the entire geodesic ray $[x, \beta \circ \mu(x))$. Since $(\beta \circ \mu)_t$ is defined by flowing along these geodesic rays for a time t , its image is in $\widetilde{M}_{\leq \varepsilon + 2\delta}$.

CLAIM. *The homotopy $(\beta \circ \mu)_t$ diverges.*

Let F be a compact fundamental domain for the Γ -action on $\widetilde{M}_{\leq \varepsilon}$. Since $\mu(F)$ is contained in a finite union of simplices of $\Delta_{[pAb]}$ and β diverges on simplices, we conclude that $c_t \circ \beta \circ \mu|_F$ diverges. It is asymptotic to $(\beta \circ \mu)_t|_F$ so Lemma 16 that $(\beta \circ \mu)_t|_F$ diverges. But since $(\beta \circ \mu)_t$ is Γ -equivariant and F is a fundamental domain, this actually implies that the entire $(\beta \circ \mu)_t$ diverges.

CLAIM. *Collapsing K to dimension d in the thin part.*

Since $(\beta \circ \mu)_t$ diverges and is asymptotic to $c_t \circ \beta \circ \mu$, for any compact subset K and large enough t the straight-line homotopy between $(\beta \circ \mu)_t|_K$ and $c_t \circ \beta \circ \mu|_K$ is *inside* $\widetilde{M}_{\leq \varepsilon}$ by Lemma 16. Thus for a compact K we can go along $(\beta \circ \mu)_t$ for a sufficiently large time and then take the straight line homotopy to $c_t \circ \beta \circ \mu$, and during this process the image of the set K will stay inside the $(\varepsilon + 2\delta)$ -thin part $\widetilde{M}_{\leq \varepsilon + 2\delta}$. Since c_t is a diffeomorphism, the topological dimension of the image of $c_t \circ \beta \circ \mu$ is equal to d . This finishes the proof of the theorem. \square

7 The Importance of being Lipschitz

As pointed out in the previous section, we need some control over the dimension $d = \dim(\text{Im}(\beta \circ \mu), \angle_x)$. The inconvenient truth that continuous maps can be space-filling means that if β is only continuous, then d can be as high as $(n-1)$ and

all information on the topology of $\partial\widetilde{M}_{\leq\varepsilon}$ will be lost. Therefore, we need β to be Lipschitz because Lipschitz maps do not raise dimensions, so that we will have

$$d \leq \dim(\Delta_{|pAb|}) \leq \text{rank}_{Ab}(\pi_1 M) - 1. \quad (17)$$

To get further constraints on the dimension d , it turns out to be important to understand non-degenerate simplices. A simplex $\lambda : \Delta^k \rightarrow X$ is *non-degenerate* if $\lambda(\Delta^k) \neq \lambda(\partial\Delta^k)$. Since $\text{Im}(\beta)$ is the union of all the non-degenerate simplices $\beta(\sigma)$, Lipschitzness of β will imply that

$$d \leq \dim(\text{Im}(\beta), \angle_x) \leq \max\{k \mid \text{there is a non-degenerate } k\text{-simplex } \beta(\sigma)\},$$

so understanding non-degenerate simplices may tell us something about d .

Some simplices are better adapted for this than others. We will discuss three possibilities in the next section. Of course, the third one is always the one to be chosen in the end. It will be named after Busemann.

8 Intermission and Flyers on Various Types of Simplices

Summary of previous sections. We have found a systematic way (given by the map β) of sending vertices of a sufficiently fine enough triangulation of $\partial\widetilde{M}_{\leq\varepsilon}$ to ∂_∞ . We need to fill in β with simplices in ∂_∞ that are Lipschitz and satisfy the criterion for divergence explained above.

So let us now turn to the problem of actually constructing divergent simplices for the map β . There are at least three different ways to do it. The easiest method is to use geodesic simplices so we will mention it first.

8.1 Geodesic simplices in $(\partial_\infty, \text{Td})$. To reassure the reader that the present discussion is not devoid of content, we note that one way to build β is using geodesic simplices. Recall that a geodesic simplex σ_k with (ordered set of) vertices v_0, \dots, v_k that are mutually $\leq \pi/2$ apart is defined inductively as the iterated geodesic join $\sigma_k = \sigma_{k-1} * v_k$. So, by definition, a geodesic simplex is contained in the convex hull of its vertices. If the set of vertices has diameter $< \pi/2 - \alpha$ then the geodesic simplex σ_k is inside the ball $B_{\pi/2-\alpha}(v_i)$ centered at any vertex. Therefore, by (16), if we form β using geodesic simplices, then the resulting simplices will diverge. It is also easy to see from the definition that the resulting map β will be Lipschitz and Γ -equivariant. So, this β will have all the properties listed in Subsection 6.1.3 and all the results of Subsection 6.2 and Section 7 apply to it. In particular, geodesic simplices are sufficient to establish the $\text{rank}_{Ab}(\pi_1 M)$ and $\dim(\partial_\infty, \text{Td})$ versions of Theorem 4. However it is difficult to say anything about non-degenerate geodesic simplices and we do not know how to get the half-dimensional bound of Theorems 1 and 2 using geodesic simplices.

8.2 Barycentric simplices in $(\partial_\infty, \mathbf{Td})$. These simplices were introduced in Section 4 of [Kle99]. Suppose the diameter of the set $\{v_0, \dots, v_k\}$ is $< \pi/2 - \alpha$. For each $t \in \Delta^k$, let $\lambda(t)$ be the unique minimum (see Lemma 4.22 of [Kle99]) of the function

$$f_t(\cdot) := \sum_i t_i \operatorname{Td}(\cdot, v_i)^2.$$

This defines a map $\lambda : \Delta^k \rightarrow \partial_\infty$ that is called the *barycentric simplex with vertices* v_0, \dots, v_k . Points x with $\operatorname{Td}(x, v_i) \geq \pi/2 - \alpha$ for all i , are not on the barycentric simplex λ . This is because any function of the form f_t has

$$f_t(x) \geq (\pi/2 - \alpha)^2 > f_t(v_i)$$

so it does not have a minimum at x . Therefore λ is contained in the $(\pi/2 - \alpha)$ -neighborhood of its vertex set $N_{\pi/2 - \alpha}(\lambda^{(0)})$. Barycentric simplices are Lipschitz and defined in an equivariant way, so we can use them to construct a “divergent simplex” map β . These simplicies are well adapted to understanding the Tits boundary with the Tits metric. Their key feature is that non-degenerate barycentric k -simplices must have $k \leq \dim(\partial_\infty, \mathbf{Td})$. In particular, we get from this that $d \leq \dim(\partial_\infty, \mathbf{Td})$, but we already knew that.

8.3 Busemann simplices in $(\partial_\infty, \mathbf{Td})$. The main focus of the rest of this paper is to introduce a new way of constructing simplices at infinity which we call *Busemann simplices* and which can also be used to build a “divergent simplex” map β . In this subsection, we give an informal account of Busemann simplices and their utility.

For a set of Busemann functions h_0, \dots, h_k centered at vertices $v_0, \dots, v_k \in \partial_\infty$ with $\operatorname{Td}(v_i, v_j) < \pi/2 - \alpha$, and a basepoint $x \in \widetilde{M}$ we let $\sigma_R(t)$ be the unique minimum of the function

$$f_t(\cdot) = \sum_i t_i h_i(\cdot)$$

on the sphere $S_R(x)$. This defines for each radius R a map $\sigma_R : \Delta^k \rightarrow S_R(x)$. Doing this for all the simplices in $\Delta_{[pAb]}$ gives a map

$$\beta_R : \Delta_{[pAb]} \rightarrow S_R(x).$$

As we will see in Section 9, this map is *Lipschitz* in the \angle_x -metric with a Lipschitz constant that *does not depend on R or x* . This allows us to find a convergent subsequence $\beta_{R_i} \rightarrow \beta : \Delta^k \rightarrow \partial_\infty$ converging to a Lipschitz map β . Such a limit map β is called “the” *Busemann map* and its restriction to each simplex σ is called a *Busemann simplex*.

REMARK. The map β depends on a sequence of scales $\{R_i \rightarrow \infty\}$ that we choose once and for all.

The limit map β does not depend on the choice of basepoint x and it follows from this that β is Γ -equivariant and Lipschitz in the Tits metric. The basepoint independence also leads, via (16), to divergence for Busemann simplices. Therefore, this β constructed out of Busemann simplices serves as a “divergent simplex” map and the results of Subsection 6.2 and Section 7 apply to it.

Busemann simplices are particularly well-adapted to studying the topology of the end. Let us say a bit about why this is the case. Let G be a discrete group which preserves horospheres at the vertices of the simplex σ . Busemann simplices are constructed so that the group G preserves horospheres *on the entire simplex*. One also has a good understanding of non-degenerate Busemann simplices in terms of the finite approximations $\beta_{R_i}(\sigma)$. These are the key features that lead to a bound

$$hdim(G) + k + 1 \leq n$$

for a non-degenerate Busemann k -simplex $\beta(\sigma)$. This dimension bound is our main technical result. If σ is a k -simplex in $\Delta_{[pAb]}$ then the group preserving horospheres at the vertices is *at least* \mathbb{Z}^{k+1} so if $\beta(\sigma)$ is non-degenerate we get from the dimension bound that $2(k+1) \leq n$ and therefore

$$k \leq \lfloor n/2 \rfloor - 1.$$

Since β is Lipschitz and its image $\text{Im}(\beta)$ is the union of the non-degenerate $\beta(\sigma)$, we get the *half-dimensional collapse* phenomenon

$$d \leq \lfloor n/2 \rfloor - 1.$$

9 Busemann Simplices (Mostly Metric Properties)

We now construct the Busemann simplices introduced at the end of the last section. These are limits of singular simplices on spheres centered at a fixed point x of finite radius R_i , i.e.

$$\lim_{R_i \rightarrow \infty} \sigma_{R_i, x} : \Delta^k \rightarrow (\partial_\infty, \angle_x).$$

Since Busemann simplices are defined as limits, we need to make sure that such limits exist and must also be Lipschitz and divergent. Therefore, we will first give some preliminary estimates on the finite radius approximations.

For reasons that will be clear by the end of this section, Busemann simplices can be defined whenever the vertices z_0, \dots, z_i are pairwise at an angle $\leq \pi/2$ from each other. Therefore, we will consider in this section only Busemann functions h_0, \dots, h_k whose gradients are pairwise at an angle $\leq \pi/2$ everywhere.

9.1 Preliminary estimates. Let $\Delta^k = \{(t_0, \dots, t_k) \mid t_0 + \dots + t_k = 1 \text{ and } t_i \geq 0 \text{ for each } i\}$. For each $t \in \Delta^k$, take the convex combination

$$f_t := t_0 h_0 + \dots + t_k h_k.$$

9.1.1 Infinitesimal Lipschitz estimate. Since $\nabla h_0, \dots, \nabla h_k$ are unit vectors that pairwise make an angle $\leq \pi/2$, the Hölder inequality gives

$$\frac{1}{\sqrt{k+1}} \leq \sum_{i=0}^k t_i^2 \leq |\nabla f_t| \leq 1.$$

Therefore the radial projection from the convex hull of the ∇h_i to the unit sphere is at most $\sqrt{k+1}$ -Lipschitz, so we get

$$\angle(\nabla f_t, \nabla f_{t'}) \leq \sqrt{k+1}|t - t'|_2.$$

9.1.2 Radius- R Lipschitz estimate. Fix a basepoint $x \in \widetilde{M}$. Since f_t is convex with no critical points as $|\nabla f_t| > 0$, the minima of f_t on the R -ball $B_x(R)$ form a convex subset of its boundary sphere $S_x(R)$, and strict convexity of the R -ball implies that this subset is a point. Denote it by $\sigma_R(t)$. This defines a map

$$\sigma_R : \Delta^k \rightarrow S_x(R).$$

LEMMA 22. *For each $R > 0$, the map*

$$\sigma_R : \Delta^k \rightarrow (S_x(R), \angle_x)$$

is $2\sqrt{k+1}$ -Lipschitz in the L^2 -metric on Δ^k .

Proof. The main idea of this proof is in Figure 5, which might be necessary for the purpose of following what is written next.

Let $p_1 = \sigma_R(t)$ and $p_2 = \sigma_R(t + \delta)$. Let $\alpha = \angle_x(p_1, p_2)$. Our goal is to bound α in terms of δ . For each $i = 1, 2$,

- let β_i be the angle at p_i between $-\nabla f_t$ and $-\nabla f_{t+\delta}$,
- let $\alpha_i = \angle_{p_i}(x, p_{i+1})$, where addition in i is taken mod 2,
- let μ_1 be the angle at p_1 between the $-\nabla f_{t+\delta}$ and the tangent to the geodesic from p_1 to p_2 , and
- let μ_2 be the angle at p_2 between the $-\nabla f_t$ and the tangent to the geodesic from p_2 to p_1 .

Then

$$\alpha_1 + \mu_1 + \beta_1 \geq \pi$$

since at p_1 the vector $-\nabla f_t$ is parallel to the tangent vector to the geodesic from x to p_1 because they are both orthogonal to the level set $f_t = f_t(p_1)$. Similarly,

$$\alpha_2 + \mu_2 + \beta_2 \geq \pi.$$

Therefore,

$$\beta_1 + \beta_2 \geq (\pi - \alpha_1 - \alpha_2) + (\pi - \mu_1 - \mu_2).$$

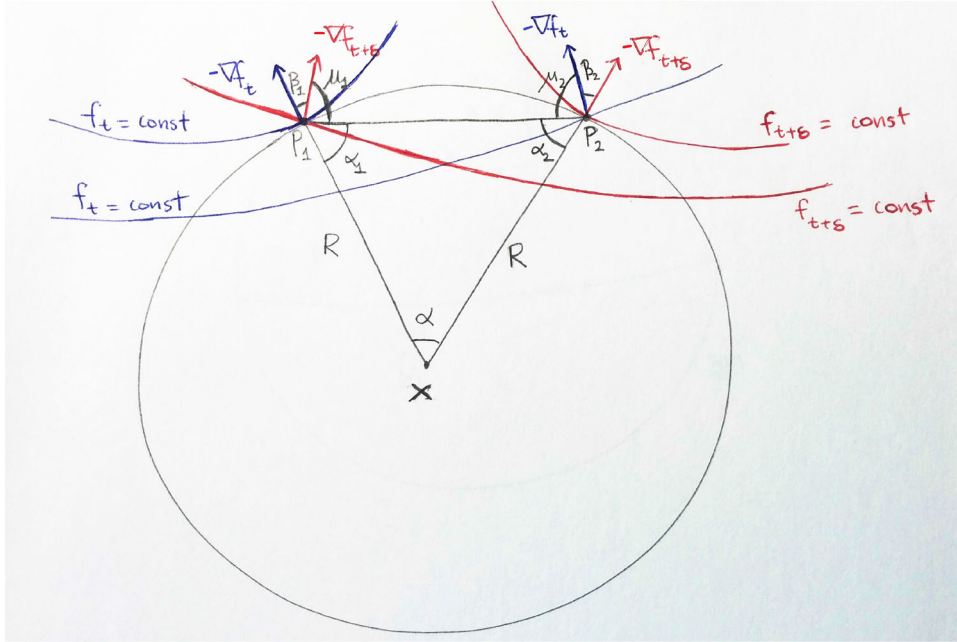


Figure 5: Lipschitz bound.

Now, α , α_1 and α_2 are the three angles of the triangle Δxp_1p_2 . Thus,

$$\alpha + \alpha_1 + \alpha_2 \leq \pi, \quad \text{so} \quad \alpha \leq \pi - \alpha_1 - \alpha_2.$$

Hence,

$$\beta_1 + \beta_2 \geq \alpha + (\pi - \mu_1 - \mu_2).$$

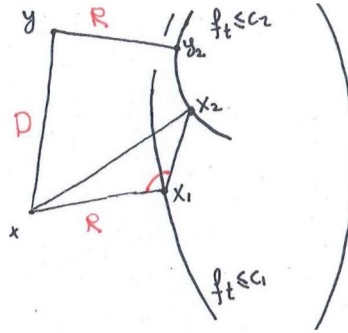
Next, we show that $\mu_1 \leq \pi/2$. Observe that p_1 does not belong to the sublevel set $f_{t+\delta} \leq f_{t+\delta}(p_2)$ since the sphere $S_x(R)$ lies on the other side of the level set $f_{t+\delta} = f_{t+\delta}(p_2)$. Therefore, $f_{t+\delta}(p_1) > f_{t+\delta}(p_2)$, so p_2 is contained in the sublevel set $f_{t+\delta} \leq f_{t+\delta}(p_1)$. Since at p_1 the vector $-\nabla f_{t+\delta}$ is orthogonal to the level set $f_{t+\delta} = f_{t+\delta}(p_1)$ it follows that $\mu_1 \leq \pi/2$. Similarly, we see that $\mu_2 \leq \pi/2$ and we obtain that

$$\alpha \leq \beta_1 + \beta_2.$$

Since each $\beta_i \leq \sqrt{k+1}|\delta|_2$ by 9.1.1, it follows that $\alpha \leq 2\sqrt{k+1}|\delta|_2$. □

9.2 Definition of Busemann simplices. The approximations $\sigma_R : \Delta^k \rightarrow S_x(R)$ depend on the choice of basepoint x . For the moment, let us emphasize this dependence and denote them by $\sigma_{R,x}$. Let us identify all the $(S_x(R), \angle_x)$ with $(\partial_\infty, \angle_x)$ via geodesic retraction. Then for each $R > 0$ we get a map

$$\sigma_{R_i,x} : \Delta^k \rightarrow (\partial_\infty, \angle_x),$$



and all these maps are Lipschitz with the same Lipschitz constant $2\sqrt{k+1}$. Therefore, by Arzela-Ascoli, there is a sequence of radii $R_i \rightarrow \infty$ so that $\sigma_{R_i, x}$ converge to a $2\sqrt{k+1}$ -Lipschitz map

$$\sigma := \lim_{R_i \rightarrow \infty} \sigma_{R_i, x} : \Delta^k \rightarrow (\partial_\infty, \angle_x).$$

We call any such map σ a *Busemann simplex*.

9.3 Properties of Busemann simplices

9.3.1 Independence of basepoint. Next we will show that σ does not depend on the choice of basepoint x but only on the sequence of radii R_i . The (easy) estimate we need for this is

$$d_{\widehat{M}}(\sigma_{R, x}(t), \sigma_{R, y}(t)) \leq D + \sqrt{2DR + D^2} \quad (18)$$

where $D := d(x, y)$. The key point is that R appears under the square root in the estimate (18). It follows from this (by Section 16.2) that $\{\sigma_{R_i, x}\}$ converges if and only if $\{\sigma_{R_i, y}\}$ converges, and that both converge to the same σ .

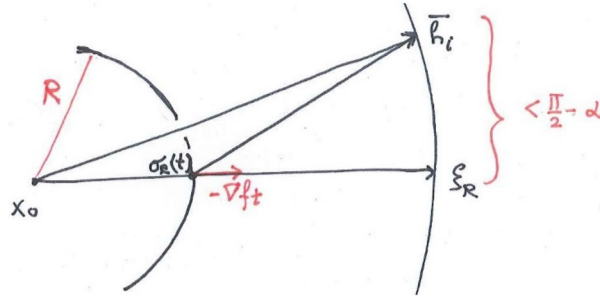
REMARK. Here is what we are using: If $x_i \rightarrow \xi$ and $\frac{d(y_i, x_i)}{d(x_0, x_i)} \rightarrow 0$ then $y_i \rightarrow \xi$.

Proof of estimate (18). We will use the following notation. Fix $c_1, c_2 \in \mathbb{R}$. Denote by x_i and y_i the closest point projections of x and y onto the sublevel set $\{f_t \leq c_i\}$, respectively. Suppose that $d(x, x_1) = R$ and $d(y, y_2) = R$. In other words, $x_1 = \sigma_{R, x}(t)$ and $y_2 = \sigma_{R, y}(t)$. Without loss of generality $c_2 \leq c_1$.

We need to bound $d(x_1, y_2)$ in terms of D and R . First, note that

$$\begin{aligned} d(x_1, y_2) &\leq d(x_1, x_2) + d(x_2, y_2) \\ &\leq d(x_1, x_2) + d(x, y) \\ &= d(x_1, x_2) + D, \end{aligned}$$

because closest point projection to $\{f_t \leq c_2\}$ is a semicontraction, i.e. a 1-Lipschitz map. It remains to bound $d(x_1, x_2)$. To do this, note that $c_2 \leq c_1$ implies $\angle_{x_1}(x, x_2) \geq \pi/2$, so



$$\begin{aligned}
 d(x_1, x_2)^2 + R^2 &= d(x_1, x_2)^2 + d(x, x_1)^2 \\
 &\leq d(x, x_2)^2 \\
 &\leq d(x, y_2)^2 \\
 &\leq (D + R)^2,
 \end{aligned}$$

where the first inequality is by triangle comparison with an obtuse triangle in Euclidean space, the second inequality is because x_2 is the closest point to x on $\{f_t \leq c_2\}$ while y_2 is just some point on this sublevel set, and the third is the triangle inequality. Simplifying, we get $d(x_1, x_2) \leq \sqrt{2DR + D^2}$. \square

9.3.2 Busemann simplices are Lipschitz in the Tits metric. A consequence of the basepoint independence is that $\sigma : \Delta^k \rightarrow \partial_\infty$ is $2\sqrt{k+1}$ -Lipschitz in the \angle_y metric for any point $y \in \widetilde{M}$, since we can use finite approximations $\{\sigma_{R_i, y}\}$ based at y to get the limit simplex σ , and for these Lemma 22 gives the $2\sqrt{k+1}$ -Lipschitz estimate. Therefore, Busemann simplices are $2\sqrt{k+1}$ -Lipschitz in the \angle -metric $\angle = \sup_{y \in \widetilde{M}} \angle_y$. Since for distances $< \pi$ the Tits metric agrees with the \angle -metric (in the sense that $\angle = \min(\text{Td}, \pi)$, see “Appendix B”) Busemann simplices are also locally $2\sqrt{k+1}$ -Lipschitz in the Tits metric. Since the Tits metric is a path metric, this implies they are $2\sqrt{k+1}$ -Lipschitz in the Tits metric.

9.3.3 Diameter bound. For each Busemann function h_i , denote by \bar{h}_i the center of the horosphere $h_i^{-1}(0)$. Busemann simplices are small in the following sense.

LEMMA 23. *If the set of vertices $\sigma^{(0)}$ has Td-diameter $< \pi/2 - \alpha$ then the entire Busemann simplex σ is contained in a ball of radius $\pi/2 - \alpha$ centered at any one of the vertices, i.e.*

$$\text{Td}(\bar{h}_i, \bar{h}_j) < \pi/2 - \alpha \quad \text{for all } i, j \implies \text{Td}(\bar{h}_i, \sigma(t)) \leq \pi/2 - \alpha.$$

Proof. Fix a basepoint x_0 and let $\sigma_R(t) := \sigma_{R, x_0}(t)$ be the finite approximation based at x_0 . Extend the geodesic segment $[x_0, \sigma_R(t)]$ to a geodesic ray $[x_0, \xi_R]$ with endpoint $\xi_R \in \partial_\infty$. Then, by looking at the ideal triangle with vertices $x_0, \sigma_R(t)$, and \bar{h}_i , we see that by nonpositive curvature we have

$$\angle_{x_0}(\sigma_R(t), \bar{h}_i) \leq \angle_{\sigma_R(t)}(\xi_R, \bar{h}_i). \tag{19}$$

Now, let $f_t = t_0 h_0 + \dots + t_k h_k$. At the point $\sigma_R(t)$ the negative of the gradient $-\nabla f_t$ points at ξ_R . The gradient ∇f_t is a convex combination of gradients of the Busemann functions h_i and in general is not a unit vector, but we can rescale it to be a unit vector $\nabla f_t / |\nabla f_t|$, which is a non-negative linear combination of the ∇h_i . Now, the visual boundary $(\partial_\infty, \angle_{\sigma_R(t)})$ with the angle metric $\angle_{\sigma_R(t)}$ can be identified with the unit sphere of the tangent space $T_{\sigma_R(t)} \widetilde{M}$. With respect to this identification, the normalized vector $\nabla f_t / |\nabla f_t|$ lies on the intersection of the unit sphere of $T_{\sigma_R(t)} \widetilde{M}$ and the “positive quadrant” of the vector subspace spanned by ∇h_i 's. It follows that in the $\angle_{\sigma_R(t)}$ -metric the point ξ_R is in the convex hull of the set $\{\bar{h}_0, \dots, \bar{h}_k\}$. Since this set has diameter $< \pi/2 - \alpha$ in the $\angle_{\sigma_R(t)}$ -metric, its convex hull is contained in an $\angle_{\sigma_R(t)}$ -metric $(\pi/2 - \alpha)$ -ball around each vertex \bar{h}_i , and therefore

$$\angle_{\sigma_R(t)}(\xi_R, \bar{h}_i) < \pi/2 - \alpha.$$

Using the earlier inequality (19) and taking the limit as $R_i \rightarrow \infty$ we get

$$\angle_{x_0}(\sigma(t), \bar{h}_i) \leq \pi/2 - \alpha.$$

Since this holds for every basepoint x_0 , we get the same bound in the Tits metric. \square

9.3.4 Invariance of horospheres on a Busemann simplex. Let

$$f = t_0 h_0 + \dots + t_k h_k$$

be a convex combination of Busemann functions with $\text{Td}(h_i, h_j) \leq \pi/2$. Then f is a convex function that does not attain its infimum. If we denote

$$s_i := f(\sigma_{R_i}(t)),$$

then, by definition, $\sigma_{R_i}(t)$ is the closest point projection of the basepoint x_0 to the sublevel set $\{f \leq s_i\}$. It follows from 3.9 of [BGS85] that the limit

$$\hat{f}(x) := \lim_{i \rightarrow \infty} d(x, \{f \leq s_i\}) - d(x_0, \{f \leq s_i\}) \quad (20)$$

exists and is equal to a Busemann function centered at $\sigma(t)$. If the Busemann functions h_i are G -invariant then f is also (obviously) and inspecting the formula (20) we see that \hat{f} is, as well. In summary, we get

LEMMA 24. *If G preserves horospheres at all the vertices \bar{h}_i of a Busemann simplex $\sigma : \Delta \rightarrow \partial_\infty$, then G preserves horospheres at every point $\sigma(t)$ of σ .*

REMARK. We do not know whether the same is true for geodesic or barycentric simplices. For those simplices, we only know that points on the simplex are fixed by G but we do not know that horospheres centered at those points are G -invariant.

10 The Busemann Map

In this section, we will construct a Lipschitz map β such that β satisfies the properties (1)–(3) of Subsection 6.1.3 and the restriction of β to each simplex σ in $\Delta_{[pAb]}$ is a Busemann simplex.

10.1 Construction. Define β on the vertices of $\Delta_{[pAb]}$ as in 5.3 by

$$\beta([A]) := \xi_{[A]}.$$

Fix a basepoint x . Doing the construction from 9.1.2 of the finite approximations σ_R on each simplex of $\Delta_{[pAb]}$ gives a map

$$\beta_R : \Delta_{[pAb]} \rightarrow S_x(R).$$

This map is $3k$ -Lipschitz in the \angle_x -metric: It is enough to check this on paths, where it follows from the fact (Lemma 22) that it is $3k$ -Lipschitz (note that $2\sqrt{k} + 1 \leq 3k$) on each simplex. Since $\Delta_{[pAb]}$ has countably many simplices, by Arzela-Ascoli we can take a sequence of radii $R_i \rightarrow \infty$ for which β_{R_i} converge to a $3k$ -Lipschitz map

$$\beta : \Delta_{[pAb]} \rightarrow (\partial_\infty, \angle_x). \quad (21)$$

Since the Lipschitz constant does not depend on x , this map is also Lipschitz in the Tits metric. We call it “the” *Busemann map*.

10.2 Equivariance. For any $\gamma \in \Gamma$ and simplex τ in $\Delta_{[pAb]}$, the Busemann simplices $\beta : \gamma\tau \rightarrow \partial_\infty$ and $\gamma\beta : \tau \rightarrow \partial_\infty$ have the same ordered set of vertices, so they are equal, i.e. $\gamma\beta = \beta\gamma$. Therefore β is Γ -equivariant.

10.3 Divergence. If σ is a simplex in $\Delta_{[pAb]}$ with vertices $[A_0] < \dots < [A_k]$, then the Busemann simplex $\beta(\sigma)$ is fixed pointwise by A_k because β is Γ -equivariant, and contained in a $(\pi/2 - \alpha)$ -neighborhood of a vertex $\beta([A_k])$ by Lemma 23. The group C_{A_k} preserves horospheres at this vertex by Proposition 9, so we get by Corollary 15 that the Busemann simplex $\beta(\sigma)$ diverges in M .

REMARK. Here is a slightly different way to see that the Busemann simplices diverge: Since β is continuous, the image $\beta(\sigma)$ is compact in the Tits metric. Cover it with finitely many $(\pi/2 - \alpha)$ -balls. Since C_{A_k} preserves horospheres on the *entire* Busemann simplex (Lemma 24) we can apply Proposition 14 to the centers of each of these balls and conclude that the portion of σ contained in each one of the finitely many $(\pi/2 - \alpha)$ -balls diverges. Therefore, the entire Busemann simplex σ diverges.

11 (Non)-degeneracy and Consequences

We saw earlier how it is important to have control over the dimension $d = \dim(\text{Im}(\beta \circ \mu), \angle_x)$. This d is bounded by the maximum of the dimensions of non-degenerate Busemann simplices in β . Thus, it would be good to know, especially if one wants to be efficient, when a Busemann simplex is degenerate and what we can do with non-degenerate ones. The goal of this section is to address these and to set things up for the next sections, where we will bound the number of vertices of non-degenerate simplices of β by $\lfloor n/2 \rfloor$. Since Busemann simplices are Lipschitz, this will bound the dimension d from above by $\lfloor n/2 \rfloor - 1$.

11.1 Non-degenerate simplices and linearly independent vectors. A simplex $\lambda : \Delta \rightarrow X$ is *non-degenerate* if the image $\lambda(\Delta)$ is not contained in the image of the boundary $\lambda(\partial\Delta)$. Any point $x \in \lambda(\Delta) \setminus \lambda(\partial\Delta)$ is called a *non-degenerate point* of λ . The meaning of non-degeneracy for the finite approximations σ_R is, partly, explained by the following simple lemma.

LEMMA 25. *Suppose σ is a Busemann simplex and let σ_R be a finite approximation of it. If $x \in \sigma_R(\Delta)$ is non-degenerate, then the gradient vectors $\{\nabla h_0, \dots, \nabla h_k\}$ are linearly independent at x .*

Proof. Note that $x = \sigma_R(t)$ for some $t = (t_0, \dots, t_k) \in \Delta$, and look at the convex combination $f_t = t_0 h_0 + \dots + t_k h_k$. At the point x , the gradient

$$\nabla f_t = t_0 \nabla h_0 + \dots + t_k \nabla h_k$$

is perpendicular to the sphere $S_{x_0}(R)$. Suppose the $\{\nabla h_i\}$ are linearly dependent at x , and write down a linear dependence relation as

$$\rho := s_0 \nabla h_0 + \dots + s_k \nabla h_k = 0,$$

where we can assume that at least one of the coefficients s_j is positive. So, there is a smallest $\varepsilon \geq 0$ such that $t_i - \varepsilon s_i = 0$ for some i . After reordering the indices, we may assume this happens for $i = 0$. Then $t_0 - \varepsilon s_0 = 0$ and $t_i - \varepsilon s_i \geq 0$ for all i , so at x we have

$$\begin{aligned} \nabla f_t &= \nabla f_t - \varepsilon \rho \\ &= (t_1 - \varepsilon s_1) \nabla h_1 + \dots + (t_k - \varepsilon s_k) \nabla h_k \\ &= a_1 \nabla h_1 + \dots + a_k \nabla h_k \end{aligned}$$

for some non-negative constants $\{a_i\}_{i=1}^k$. Setting

$$t' := \left(0, \frac{a_1}{|a_1|}, \dots, \frac{a_k}{|a_1|} \right) \in \partial\Delta,$$

the above equation can be rewritten as

$$\nabla f_t = |a_1| \nabla f_{t'}$$

at the point x , which implies that $x = \sigma_R(t')$. So, x is a degenerate point. \square

REMARK. The proof of this lemma is less delicate than it may appear at first glance. All we are doing is finding an intersection point a of the line $t+\mathbb{R}s$ with the boundary of the positive “octant” $\partial((\mathbb{R}_+)^{k+1})$ and observing that the boundary point $a/|a|_1 \in \partial\Delta$ defined by this is mapped to $\sigma_R(t)$.

The map $\sigma_R : \Delta \rightarrow \widetilde{M}$ may be very far from an embedding. However, we will see next that the situation is better if we restrict to the *preimages of non-degenerate points*. Putting these together for all R forms the set

$$(\mathbb{R}^+ \times \Delta)_\sigma^{nd} := \{(R, t) \in \mathbb{R}^+ \times \Delta \mid \sigma_R(t) \text{ is non-degenerate}\}. \tag{22}$$

This set is open because being non-degenerate is an open condition. Here, there is a slight abuse of notation since $(\mathbb{R}^+ \times \Delta)_\sigma^{nd}$ depends only on the vertices of σ and the basepoint x_0 , but not on the actual choice of σ . That is, one considers the vertex set $\{\bar{h}_0, \dots, \bar{h}_k\}$ of σ and associates to each \bar{h}_i a Busemann function h_i normalized so that $h_i(x_0) = 0$ and then, for every $t = (t_0, \dots, t_k)$, lets $f_t = \sum_{i=0}^k t_i h_i$ and lets $\sigma_R(t)$ be the minimum of f_t on the R -sphere about x_0 .

COROLLARY 26. *The map*

$$(\mathbb{R}^+ \times \Delta)_\sigma^{nd} \rightarrow \widetilde{M}, \tag{23}$$

$$(R, t) \mapsto \sigma_R(t) \tag{24}$$

is injective.

Proof. Note that $\sigma_R(t) = x = \sigma_R(t')$ gives the linear relation

$$t_0 \nabla h_0 + \dots + t_k \nabla h_k = c(t'_0 \nabla h_0 + \dots + t'_k \nabla h_k)$$

at the point x . If x is non-degenerate, the previous lemma implies this relation is trivial. So we must have $t = ct'$. Since $\sum t_i = 1 = \sum t'_i$, this implies $c = 1$ and $t = t'$. This proves the corollary. \square

11.2 Busemann cone. In section 12 it will often be useful to put all the (images of) Busemann simplices σ_R based at a single point x_0 together. The *Busemann cone* of σ (based at x_0) is the set

$$\sigma_{>0} := \bigcup_{R>0} \sigma_R(\Delta)$$

of all points in \widetilde{M} that lie on $\sigma_R = \sigma_{R,x_0}$ for some radius R . Sometimes it is convenient to also include the basepoint x_0 . The result is then a closed set

$$\sigma_{\geq 0} := \sigma_{>0} \cup \{x_0\}.$$

A point on the Busemann cone is *degenerate* if it is a degenerate point of σ_R for some R . As a matter of convention, we declare that x_0 is also a degenerate point. That is, a point on the Busemann cone is degenerate if it is not in the image of the map in Corollary 26.

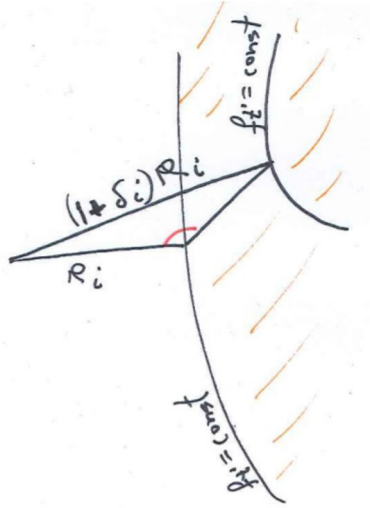


Figure 6: Comparison triangle.

11.3 A sequential criterion for degeneracy. Suppose $\sigma_{R_i} \rightarrow \sigma$ is a Busemann simplex. Since the sequence of radii $\{R_i\}$ is chosen somewhat non-canonically, sometimes we cannot avoid dealing with points on the Busemann cone that are not on σ_{R_i} . Therefore, it will be useful later to have the following lemma which gives a criterion for degeneracy of a point $\sigma(t)$ in terms of a sequence of points on the Busemann cone $\{q_i\}$ converging to $\sigma(t)$. Note that the q_i 's need not belong to σ_{R_i} .

LEMMA 27. *Let $\sigma : \Delta \rightarrow \partial_\infty$ be a Busemann simplex. Then, a point $\sigma(t)$ is degenerate if there is a sequence $\{q_i\}_{i=1}^\infty$ of degenerate points on the Busemann cone $q_i \in \sigma_{\geq 0}$ with*

$$\lim_{i \rightarrow \infty} \frac{d(q_i, \sigma_{R_i}(t))}{R_i} = 0. \quad (25)$$

Proof. The sublinearity (25) implies that $\angle_{x_0}(q_i, \sigma_{R_i}(t)) \rightarrow 0$, so $q_i \rightarrow \sigma(t)$, but this by itself does not yet mean that $\sigma(t)$ is degenerate. To prove degeneracy, we need to find a point $t' \in \partial\Delta$ satisfying $\sigma(t) = \sigma(t')$. We will now do this.

Since q_i are degenerate points on the Busemann cone, we have

$$q_i = \sigma_{R'_i}(t'_i)$$

for $t'_i \in \partial\Delta$ and radii $R'_i \rightarrow \infty$. Passing to a subsequence, we may assume

$$t'_i \rightarrow t' \in \partial\Delta.$$

Now, the sublinearity (25) and the triangle inequality imply that $\frac{|R'_i - R_i|}{R_i} \rightarrow 0$. Therefore we can write

$$R'_i = (1 + \delta_i)R_i \quad \text{for a sequence } \delta_i \rightarrow 0.$$

Since both $\sigma_{R_i}(t')$ and $\sigma_{R'_i}(t')$ are obtained as closest point projections of x_0 to (different) sublevel sets of the *same* convex function $f_{t'} := t'_0 h_0 + \dots + t'_k h_k$, comparison of the triangle with vertices $x_0, \sigma_{R_i}(t')$ and $\sigma_{R'_i}(t')$ (depicted in Figure 6) with an obtuse Euclidean triangle (see 15.2) gives

$$\frac{d(\sigma_{R'_i}(t'), \sigma_{R_i}(t'))}{R_i} \leq \sqrt{2\delta_i + \delta_i^2} \rightarrow 0.$$

Therefore the sequence $\{q_i\} = \{\sigma_{R'_i}(t')\}$ converges to the same point as the sequence $\{\sigma_{R_i}(t')\}$, and that point is

$$\sigma(t) = \lim q_i = \sigma(t') \in \partial\Delta. \quad \square$$

12 A Dimension Bound

Busemann simplices provide a way to connect points at infinity whose horospheres are preserved. Our goal in this section is to relate the dimension of a group G that acts on \widetilde{M} preserving a collection of horospheres and the dimension of the Busemann simplex with vertices at the centers of those horospheres. One expects these two dimensions to be complementary. Theorem 28 gives the expected bounds on these dimensions. This is responsible for the $n/2$ bound in the main theorems and is the climax of this paper. We will end this section by giving the proof of Theorem 7 in the introduction.

Recall that $\text{Fix}^0(G)$ is the set of points in ∂_∞ the horospheres centered at which are preserved by G .

Theorem 28. *If $\text{Fix}^0(\mathbb{Z}^r)$ contains the image of a non-degenerate Busemann k -simplex σ , then*

$$\dim \widetilde{M} \geq k + 1 + r. \tag{26}$$

REMARK. The method of proof of Theorem 28 applies to any subgroup $G < \Gamma$ that preserves horospheres on a non-degenerate Busemann k -simplex (see 12.6). But, for now it is helpful to focus on the case $G = \mathbb{Z}^r$. This is all we need for half-dimensional collapse.

The proof of Theorem 28 is not quite a natural argument, so before giving the technical proof, we will explain why it is the way it is in Section 12.1. Readers who did not like Section 2 should definitely skip this Subsection.

12.1 Problems and solutions. A natural approach to obtaining the bound $\dim \widetilde{M} \geq r + k + 1$ is to show that each intersection of horospheres has dimension $\geq r$ and that there is a $(k+1)$ -dimensional family of such intersections of horospheres. Consider the following *parameter space for intersections of horospheres*.

Pick representatives h_i for the vertices of the non degenerate Busemann simplex $\sigma : \Delta^k \rightarrow \text{Fix}^0(\mathbb{Z}^r)$, such that $h_i(x_0) = 0$, and look at the “horospherical coordinates” map

$$\vec{h}: \widetilde{M} \rightarrow \mathbb{R}^{k+1}, \quad (27)$$

$$x \mapsto (h_0(x), \dots, h_k(x)). \quad (28)$$

The map \vec{h} is (obviously) \mathbb{Z}^r -invariant, and its image $\vec{h}(\widetilde{M})$ in \mathbb{R}^{k+1} is the parameter space for all possible intersections of horospheres $\cap_{i=0}^k \{h_i = b_i\}$.

Let T_b be the intersection of horospheres $\vec{h}^{-1}(b)$. One can hope to show that each intersection of horospheres T_b has dimension $\geq r$ by letting \mathbb{Z}^r act on it. Ideally, T_b is an $(n - k - 1)$ -dimensional submanifold that is contractible. It will then follow that $r \leq (n - k - 1)$, so we obtain $\dim \widetilde{M} = n \geq r + 1 + k$.

However, in order for this intersection of horospheres to be an $(n - k - 1)$ -dimensional submanifold, the gradients of the corresponding Busemann functions need to be linearly independent everywhere on the submanifold. This leads us to the main problem: these gradients need not be linearly independent (even on just the intersection of the horospheres). So we cannot show that T_b has dimension $(n - k - 1)$. Neither can we bound its dimension from below by r . The map \vec{h} is \mathbb{Z}^r -invariant, so the group \mathbb{Z}^r acts on each T_b , but this by itself is not enough to bound the dimension of T_b from below by r . Pessimistically speaking, T_b could be discrete. So *we will not try to prove anything about T_b* . Instead, we will show that the complement $(\widetilde{M} - T_o)$ has dimension at least $(k + r + 1)$. Note that T_o is a closed subset of \widetilde{M} , so $(\widetilde{M} - T_o)$ is a manifold of dimension n and thus, we will obtain $n \geq k + r + 1$.

Recall that in algebraic topology, there are two ways to show that \mathbb{R}^m has dimension m . One way is to remove a point o from it and to show that a sphere centered at o is nontrivial in the $(m - 1)$ -homology of the complement $(\mathbb{R}^m - o)$. The second way is to let \mathbb{Z}^m act on \mathbb{R}^m the standard way and to show that the quotient is an m -torus, which has nontrivial m -homology. We will do something that is a combination of both ways in order to show that $(\widetilde{M} - T_o)$ is a manifold of dimension at least $(k + r + 1)$.

That is, we will try to bound the dimension of

$$(\widetilde{M} - T_o)/\mathbb{Z}^r,$$

for some fixed $b = o \in \mathbb{R}^{k+1}$.

So this problem has a topological solution if we can show $(\widetilde{M} - T_o)/\mathbb{Z}^r$ has a nontrivial homology class in dimension $(k + r)$ because it will imply that $n \geq (k + r + 1)$ since $(\widetilde{M} - T_o)/\mathbb{Z}^r$ is noncompact.

Finding a $(k + r)$ -homology class in $(\widetilde{M} - T_o)/\mathbb{Z}^r$. One can attempt to define a map $\mathbb{S}^k \times \mathbb{T}^r \rightarrow (\widetilde{M} - T_o)/\mathbb{Z}^r$ as follows. First, take a \mathbb{Z}^r -equivariant map $f: \mathbb{R}^r \rightarrow \widetilde{M}$ that takes $0 \mapsto x_0$, where x_0 is a fixed basepoint at which we will take the Busemann cone later. Then for each $T_b \neq \emptyset$, compose f with the closest point projection to T_b . If we can do this for b taking values in a k -sphere in $(\vec{h}(\widetilde{M}) - o)$, then we obtain a map $\mathbb{S}^k \times \mathbb{T}^r \rightarrow (\widetilde{M} - T_o)/\mathbb{Z}^r$. Specifically, if $\vec{h}(\widetilde{M})$ has nonempty interior (as a subset of \mathbb{R}^{k+1}), then we can take o to be an interior point and let b take value in a k -sphere surrounding o . A problem with this approach is that T_b needs not be convex so closest point projection is not well-defined. Nevertheless, intersections of horoballs are convex, so the solution is to project $f(\mathbb{R}^r)$ onto

$$\widehat{T}_b := \{x \in \widetilde{M} \mid h_i(x) \leq b_i, i = 0, 1, \dots, k\}$$

instead. However, the price we pay for this is that the projection of $f(\mathbb{R}^r)$ needs not land in T_b . Neither should it even be close to T_b . This is a problem.

But there is a solution, which is to use the magic of the *Busemann cone* $\sigma_{>0}$. We will show later in this Section 5 (in Lemma 30) that if $b \in \vec{h}(\sigma_{>0})$, then the closest-point projection $p(b, x_0)$ of $x_0 = f(0)$ onto the intersection \widehat{T}_b of horoballs is actually contained in the intersection of horospheres T_b , or in other words, $\vec{h}(p(b, x_0)) = b$. It follows that the projection of $f(\mathbb{R}^r)$ is contained in an L -neighborhood of T_b . This is because the action of \mathbb{Z}^r is cocompact on $f(\mathbb{R}^r)$ and closest-point projections are distance non-increasing. A result of this is that L is independent of b once we fix f , which suggests the following: if $\vec{h}(\sigma_{>0})$ contains a ball of radius larger than L , then the above attempt to define a map $\mathbb{S}^k \times \mathbb{T}^r \rightarrow (\widetilde{M} - T_o)/\mathbb{Z}^r$ can be realized, with o the center of the ball. Since L is arbitrarily large, it is not enough to show that $\vec{h}(\widetilde{M})$ has nonempty interior; we need to show that the subset $\vec{h}(\sigma_{>0})$ contains arbitrarily large balls. Before we explain this, we also need to check that the homology class obtained is nontrivial.

Showing the $(k + r)$ -homology class is nontrivial. Note that having such a map $\mathbb{S}^k \times \mathbb{T}^r \rightarrow (\widetilde{M} - T_o)/\mathbb{Z}^r$ is not enough, we also need this map to be nontrivial in homology, which holds if there is a map

$$(\widetilde{M} - T_o)/\mathbb{Z}^r \rightarrow \mathbb{S}^k \times \mathbb{T}^r$$

such that the composition

$$\mathbb{S}^k \times \mathbb{T}^r \rightarrow (\widetilde{M} - T_o)/\mathbb{Z}^r \rightarrow \mathbb{S}^k \times \mathbb{T}^r$$

has non-zero degree. There is a natural candidate for $(\widetilde{M} - T_o)/\mathbb{Z}^r \rightarrow \mathbb{T}^r$, which is the composition

$$(\widetilde{M} - T_o)/\mathbb{Z}^r \hookrightarrow \widetilde{M}/\mathbb{Z}^r \simeq \mathbb{T}^r.$$

In fact, this is a good candidate because the restriction $\mathbb{T}^r \rightarrow \mathbb{T}^r$ is a homotopy equivalence. A natural candidate for $(\widetilde{M} - T_o)/\mathbb{Z}^r \rightarrow \mathbb{S}^k$ is to take $\vec{h}: (\widetilde{M} - T_o) \rightarrow (\mathbb{R}^{k+1} - o) \simeq \mathbb{S}^k$ and then quotient out by the action of \mathbb{Z}^r using the fact that \vec{h} is \mathbb{Z}^r -invariant. This is also a good candidate.

In summary, what we are left to explain is how to show that the set of horospherical coordinates $\vec{h}(\sigma_{>0})$ of the Busemann cone $\sigma_{>0}$ contains arbitrarily large balls. This is where we need σ to be non-degenerate.

How to show that $\vec{h}(\sigma_{>0})$ contains arbitrarily large balls. First, let us comment on why $\vec{h}(\sigma_{>0})$ has nonempty interior as a subset of \mathbb{R}^{k+1} . The reason is because \vec{h} maps non-degenerate points in $\sigma_{>0}$ to interior points of the image $\vec{h}(\sigma_{>0})$. This is an Invariance-of-Domain argument. Recall that one obtains the Busemann cone $\sigma_{>0}$ by mapping in

$$\sigma_{>0}: \mathbb{R}^+ \times \Delta \rightarrow \widetilde{M}$$

as explained in Section 11. Since σ is non-degenerate, the approximation σ_{R_i} has non-degenerate points if i is large enough, and since being non-degenerate is an open condition, this implies that the set $(\mathbb{R}^+ \times \Delta)^{nd}$ of non-degenerate (R, t) -coordinates of $\sigma_{>0}$ is open in $\mathbb{R}^+ \times \Delta$ and therefore is a $(k+1)$ -dimensional manifold. Note that $\vec{h} \circ \sigma_{>0}$ maps $(\mathbb{R}^+ \times \Delta)^{nd}$ into \mathbb{R}^{k+1} , so if one can show that it is injective on $(\mathbb{R}^+ \times \Delta)^{nd}$, then one can use Invariance of Domain to show it is an open map and obtain that $\vec{h}(\sigma_{>0})$ has nonempty interior. To see that $\vec{h} \circ \sigma_{>0}$ is injective on $(\mathbb{R}^+ \times \Delta)^{nd}$, we need the maps \vec{h} and $\sigma_{>0}$ to be injective (when restricted to the relevant domains). The map $\sigma_{>0}$ restricted to $(\mathbb{R}^+ \times \Delta)^{nd}$ is injective by Corollary 26. The restriction of \vec{h} to the Busemann cone $\sigma_{>0}$ has an inverse $p(\cdot, x_0)$ since (see Corollary 31) $\vec{h}(p(x_0, b)) = b$ for all $b \in \sigma_{>0}$, and therefore is injective.

To see that $\vec{h}(\sigma_{>0})$ contains arbitrarily large balls, we suppose for contradiction that it does not, which implies that $\vec{h}(\sigma_{>0})$ has an L -net Q of boundary points (in the sense of point set topology) of $\vec{h}(\sigma_{>0})$. Since \vec{h} maps non-degenerate points in $\sigma_{>0}$ to interior points of the image $\vec{h}(\sigma_{>0})$, the points in Q belong to $\vec{h}(\sigma_{>0}(\mathbb{R}^+ \times \partial\Delta))$. Take $p(Q, x_0)$ to get a subset of degenerate points of $\sigma_{>0}$ and if this gives an L -net in $\sigma_{>0}$, then this will imply that all points of σ are degenerate, which contradicts the assumption that σ is non-degenerate. However, $p(Q, x_0)$ needs not give a net in $\sigma_{>0}$ because $p(\cdot, x_0)$ might stretch distance between points in a nonuniform way. Nevertheless, a quantitative estimate on how $p(\cdot, x_0)$ distorts distance, as in the proof of Lemma 35, implies that $p(Q, x_0)$ is a $L\sqrt{R}$ -net, by which we mean any point in $\sigma_{>0}$ that is a distance R from x_0 is $L\sqrt{R}$ -close to a point in $p(Q, x_0)$. Since \sqrt{R} is sublinear, this implies (by Lemma 27) all points in σ are degenerate, contradicting the hypothesis that σ is non-degenerate.

12.2 The setup. Let us now begin setting up the proof. Fix a basepoint $x_0 \in \widetilde{M}$ and let h_i be Busemann functions representing the vertices of the non-degenerate Busemann simplex $\sigma: \Delta^k \rightarrow \text{Fix}^0(\mathbb{Z}^r)$, normalized such that $h_i(x_0) = 0$. Putting

these Busemann functions together gives the *horospherical coordinates* map

$$\vec{h}: \widetilde{M} \rightarrow \mathbb{R}^{k+1}, \tag{29}$$

$$x \mapsto (h_0(x), \dots, h_k(x)). \tag{30}$$

The map \vec{h} is \mathbb{Z}^r -invariant because its coordinates are. It is a semicontraction in the sup norm $|\cdot|_\infty$ on \mathbb{R}^{k+1} because $|\nabla h_i| = 1$.

In the course of the argument, we will need to project various points to various (intersections of) horoballs. Define $\{\vec{h} \leq b\}$ as the \vec{h} -preimage of $(-\infty, b]$, so it is an intersection of horoballs. To keep track of all this, let $p(b, x)$ be the *closest point projection of the point x to the intersection of horoballs $\{\vec{h} \leq b\}$* . This defines a map

$$p: \mathbb{R}^{k+1} \times \widetilde{M} \rightarrow \widetilde{M}.$$

For a fixed b the map $p(b, \cdot)$ is a semicontraction, since it is the closest point projection to a convex set. It is \mathbb{Z}^r -equivariant since \mathbb{Z}^r preserves the Busemann functions h_i and thus also preserves the intersection of horoballs $\{\vec{h} \leq b\}$.

The “error”. A central role in the proof is played by the difference

$$\vec{h}(p(b, \cdot)) - b,$$

which measures the extent to which the closest point projection to the intersection of horoballs $\{\vec{h} \leq b\}$ fails to land in the intersection of horospheres $\{\vec{h} = b\}$. We will call it the “error”.

LEMMA 29. $\vec{h}(p(b, \cdot)) - b$ is a \mathbb{Z}^r -invariant semicontraction in the sup norm.

Proof. This follows directly from what we have said about $p(b, \cdot)$ and \vec{h} . □

12.3 Some key properties of the Busemann cone. Now let us turn to the Busemann cone $\sigma_{\geq 0}$. The key property of the Busemann cone is that the horospherical coordinates embed it in \mathbb{R}^{k+1} .

LEMMA 30. *For a Busemann simplex σ , the restriction*

$$\vec{h}|_{\sigma_{\geq 0}}: \sigma_{\geq 0} \rightarrow \mathbb{R}^{k+1}$$

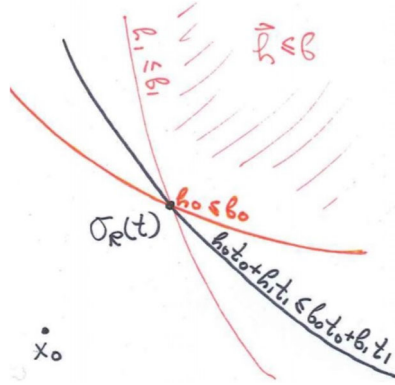
is a homeomorphism onto its image, with inverse $p(\cdot, x_0)$.

Proof. Let $b = (b_0, \dots, b_k) := \vec{h}(\sigma_R(t))$ be the image of a point on the Busemann cone. We need to show that

$$p(b, x_0) = \sigma_R(t).$$

The intersection of horoballs $\bigcap_i \{h_i \leq b_i\}$ contains $\sigma_R(t)$ and is contained in the sublevel set $\{h_0 t_0 + \dots + h_k t_k \leq b_0 t_0 + \dots + b_k t_k\}$. We claim that $\sigma_R(t)$ is the unique closest point to x_0 in this sublevel set. Therefore, it is also the unique closest point to x_0 on the intersection of horoballs. This proves that $p(b, x_0) = \sigma_R(t)$.

To see that $\sigma_R(t)$ is the unique closest point to x_0 on the sublevel set $\{h_0 t_0 + \dots + h_k t_k \leq b_0 t_0 + \dots + b_k t_k\}$, note that $b_i = h_i(\sigma_R(t))$ implies $\sigma_R(t)$ lies on the boundary of this sublevel set. Any given point of \widetilde{M} lies on the boundary of only one sublevel set of $h_0 t_0 + \dots + h_k t_k$, and by definition $\sigma_R(t)$ is on the boundary of the unique sublevel set of $h_0 t_0 + \dots + h_k t_k$ whose distance to x_0 is R . So the claim follows. \square



Denote the image of the Busemann cone in horospherical coordinates by

$$W := \vec{h}(\sigma_{\geq 0}).$$

Lemma 30 has several important consequences. The first is an estimate for the error in terms of the distance to the orbit of the basepoint $\mathbb{Z}^r x_0$. To get it, note first that the lemma immediately implies the error vanishes at x_0 .

COROLLARY 31. *For points $b \in W$ we have $\vec{h}(p(b, x_0)) = b$.*

It follows that the error is bounded by the distance to the orbit $\mathbb{Z}^r x_0$.

PROPOSITION 32. *For any $b \in W$ we have $|\vec{h}(p(b, x)) - b|_\infty \leq d(x, \mathbb{Z}^r x_0)$.*

Proof. We've shown that the error is a \mathbb{Z}^r -invariant semicontraction that vanishes at x_0 , so this is clear. \square

The second is a topological regular value theorem for Busemann simplices.

LEMMA 33. *If $x \in \sigma_{>0}^{nd}$ is a non-degenerate point on the Busemann cone, then $\vec{h}(x)$ is an interior point of W .*

Proof. By Corollary 26 and Lemma 30 the composition

$$(\mathbb{R}^+ \times \Delta)_\sigma^{nd} \hookrightarrow \sigma_{>0}^{nd} \hookrightarrow W \subset \mathbb{R}^{k+1}, \quad (31)$$

$$(R, t) \mapsto \sigma_R(t) \mapsto \vec{h}(\sigma_R(t)) \tag{32}$$

is injective. Since $(\mathbb{R}^+ \times \Delta)_\sigma^{nd}$ is an open subset of \mathbb{R}^{k+1} , Invariance of Domain implies that the image of this composition is also an open subset of \mathbb{R}^{k+1} and thus is contained in the interior of W . \square

The third says that boundary points of W come from degenerate points.

COROLLARY 34. *For every $b \in \partial W$ there is degenerate $q \in \sigma_{\geq 0}$ with $\vec{h}(q) = b$.*

Proof. First, we show that $b \in W$. To see this, note that if $b^i \in W$ converges to b then $p(b^i, x_0) \in \sigma_{\geq 0}$ converges to a point $q := p(b, x_0) \in \sigma_{\geq 0}$ because $p(\cdot, x_0)$ is a continuous map on all of \mathbb{R}^{k+1} and $\sigma_{\geq 0}$ is closed. Since $q \in \sigma_{\geq 0}$, Lemma 30 implies that $b = \vec{h}(q) \in W$. Since b is a boundary point, Lemma 33 implies that q is a degenerate point. \square

12.4 Finding large balls in W . In Lemma 33 we showed that if σ is non-degenerate then W contains open balls. In this subsection we will show that when σ is non-degenerate then W contains *arbitrarily large* balls. To do this, we will need to control how much the map $p(\cdot, x)$ distorts things.

LEMMA 35. *Suppose that $a, b \in \mathbb{R}^{k+1}$ with $b \leq a$, which means $b_i \leq a_i$ for all $0 \leq i \leq k$. Let $x \in \widetilde{M}$. Then*

$$d(p(a, x), p(b, x)) \leq \sqrt{2d(x, p(a, x))|a - b|_1 + |a - b|_1^2}.$$

Proof. We use the short hand notation $x_a := p(a, x)$. In other words, x_a is the closest point to x on the intersection of horoballs $\{\vec{h} \leq a\}$. In this notation, what we need to prove is

$$d(x_a, x_b) \leq \sqrt{2d(x, x_a)|a - b|_1 + |a - b|_1^2}. \tag{33}$$

First, note that since $b \leq a$ the point x_b lies in the intersection of horoballs $\{\vec{h} \leq a\}$ and therefore we have

$$\angle_{x_a}(x, x_b) \geq \pi/2.$$

Therefore

$$\begin{aligned} d(x, x_a)^2 + d(x_a, x_b)^2 &\leq d(x, x_b)^2 \\ &\leq d(x, x_{ab})^2 \\ &\leq (d(x, x_a) + d(x_a, x_{ab}))^2, \end{aligned}$$

where the first inequality follows by triangle comparison with an appropriate obtuse Euclidean triangle, the second inequality is because x_b is the closest point to x in the intersection of horoballs $\{\vec{h} \leq b\}$ while the projection x_{ab} of x_a to $\{\vec{h} \leq b\}$ is also in this intersection, and the third is the triangle inequality.

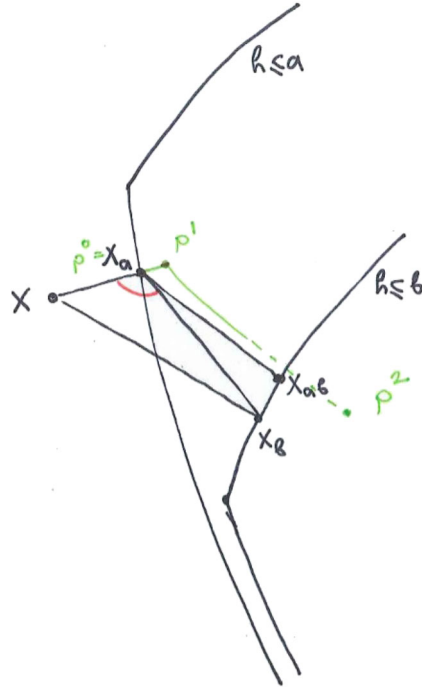
In order to get (33) from this, it is enough to prove

$$d(x_a, x_{ab}) \leq |a - b|_1. \quad (34)$$

Proof of (34). Starting at $x_a =: p^0$ go for a time $|a_0 - b_0|$ along $-\nabla h_0$ to arrive at a point which we call p^1 , then go for time $|a_1 - b_1|$ along $-\nabla h_1$ to arrive at a point denoted p^2 , and continue this way, going from p^i for time $|a_i - b_i|$ along $-\nabla h_i$ to arrive at p^{i+1} . We obtain a piecewise geodesic path

$$[p^0, p^1] \cup \dots \cup [p^k, p^{k+1}]$$

of length $|a - b|_1$ starting at x_a . Along this path all the Busemann functions h_i are monotone non-increasing since we assumed that $\angle(-\nabla h_i, -\nabla h_j) \leq \pi/2$ (see beginning of Section 9). Moreover, each h_i decreases by $|a_i - b_i|$ on the segment $[p^i, p^{i+1}]$ and therefore at the endpoint p^{k+1} we get $h_i(p^{k+1}) \leq b_i$ for all i . Thus, the endpoint p^{k+1} is contained in the intersection of horoballs $\{h_0 \leq b_0, \dots, h_k \leq b_k\}$. Since x_{ab} is the closest point to x_a in this intersection of horoballs, we get



$$\begin{aligned} d(x_a, x_{ab}) &\leq d(x_a, p^{k+1}) \\ &\leq |a - b|_1, \end{aligned}$$

which is what we wanted to show. \square

Now we are ready to prove the main result of this subsection.

PROPOSITION 36. *If $\sigma : \Delta \rightarrow \partial_\infty$ is a non-degenerate Busemann simplex, then W contains arbitrarily large $(k + 1)$ -balls.*

Proof. Fix $L > 0$. Let

$${}_a\Delta := \{b \in \mathbb{R}^{k+1} \mid b \leq a, |b - a|_1 \leq L\}.$$

We will show that W contains a set of the form ${}_a\Delta$ for some a , which will imply that W contains a ball of radius $L/(2k + 2)$ because ${}_a\Delta$ is a right angled triangle/simplex in \mathbb{R}^{k+1} with side length L . Since L is arbitrary, this implies that W contains arbitrarily large balls.

Suppose that W does not contain any set of the form ${}_a\Delta$. From the definition of Busemann simplices, we have a sequence $R_i \rightarrow \infty$ so that the maps σ_{R_i} converge to σ . In particular $\sigma_{R_i}(t) \rightarrow \sigma(t)$. Let $a^i := \vec{h}(\sigma_{R_i}(t))$. We are indexing the sequence $\{a^i\}$ with superscripts instead of subscripts to avoid possible confusion with the coordinates a_i and b_i of a and b . Since ${}_a\Delta$ is not contained in W , there is a point $b^i \in \partial W \cap {}_{a^i}\Delta$. By Corollary 34, we get a sequence of degenerate points $q_i \in \sigma_{\geq 0}$ on the Busemann cone for which $\vec{h}(q_i) = b^i$. We claim the distance $d(q_i, \sigma_{R_i}(t))$ grows sublinearly as a function of R_i . In fact, since $q_i = p(b^i, x_0)$ and $\sigma_{R_i}(t) = p(a^i, x_0)$, by Lemma 35, we get

$$d(q_i, \sigma_{R_i}(t)) \leq \sqrt{2R_iL + L^2}.$$

(By letting $a = a^i, b = b^i, x = x_0$ in Lemma 35 and noting that $|a^i - b^i|_1 \leq L$ and $d(x_0, p(a^i, x_0)) = R_i$.)

Thus, by Lemma 27, $\sigma(t)$ is a degenerate point. Since this is true for every $t \in \Delta$, we see that σ is a degenerate simplex. This contradicts the assumption that σ is non-degenerate. Thus W contains a set of the form ${}_a\Delta$. □

12.5 Proof of Theorem 28. We are now ready for the proof of the dimension bound. Given a non-degenerate Busemann k -simplex σ , we will use the projection family $p(\cdot, \cdot)$ to build a $(k + r)$ -dimensional homology cycle $S^k \times \mathbb{T}^r \rightarrow (\widetilde{M} - \vec{h}^{-1}(s))/\mathbb{Z}^r$ in the subset obtained by cutting out an appropriately chosen intersection of horospheres, and then show that the cycle is non-zero in the homology of the manifold $(\widetilde{M} - \vec{h}^{-1}(s))/\mathbb{Z}^r$. This will imply that the dimension of the manifold is greater than $k + r$ and will finish the proof of Theorem 28. The proof uses Propositions 32 and 36 that we established in the previous subsections.

12.5.1 Building the homology cycle. Since \widetilde{M} is contractible, the quotient $\widetilde{M}/\mathbb{Z}^r$ is homotopy equivalent to the r -torus $\mathbb{R}^r/\mathbb{Z}^r$. Any homotopy equivalence $g : \mathbb{T}^r \rightarrow \widetilde{M}/\mathbb{Z}^r$ is covered by a \mathbb{Z}^r -equivariant map $\tilde{g} : \mathbb{R}^r \rightarrow \widetilde{M}$. Form the Busemann cone $\sigma_{\geq 0}$ based at $\tilde{g}(0)$ and let $W = \vec{h}(\sigma_{\geq 0})$. Proposition 32 tells us that for all points $b \in W$ we have the uniform error bound

$$|\vec{h}(p(b, \tilde{g}(\cdot))) - b|_\infty \leq \text{diameter of } \tilde{g}([0, 1]^r). \tag{35}$$

Since σ is a non-degenerate Busemann k -simplex, the image of its Busemann cone in horospherical coordinates W contains arbitrarily large balls by Proposition 36. So W contains a closed L -ball about s whose radius L (in the sup metric) is greater than the diameter of $\tilde{g}([0, 1]^r)$. Let S_s^k be its boundary sphere. Then the map

$$S_s^k \times \mathbb{R}^r \rightarrow \widetilde{M}, \quad (36)$$

$$(b, z) \mapsto p(b, \tilde{g}(z)). \quad (37)$$

will miss the intersection of horospheres $\{\vec{h} = s\}$ because for every $b \in S_s^k$ the distance $|s - b|_\infty = L$ while the distance $|\vec{h}(p(b, \tilde{g}(\cdot))) - b|_\infty$ is strictly less than L . Thus, quotienting out by the \mathbb{Z}^r -action we get a homology cycle

$$S_s^k \times \mathbb{T}^r \rightarrow (\widetilde{M} - \vec{h}^{-1}(s))/\mathbb{Z}^r. \quad (38)$$

12.5.2 The homology cycle is non-trivial in the homology of $(\widetilde{M} \setminus \vec{h}^{-1}(s))/\mathbb{Z}^r$. Look at the composition

$$\begin{aligned} S_s^k \times \mathbb{R}^r &\rightarrow \widetilde{M} - \vec{h}^{-1}(s) \rightarrow (\mathbb{R}^{k+1} \setminus \{s\}) \times \widetilde{M}, \\ (b, z) &\mapsto p(b, \tilde{g}(z)) \mapsto (\vec{h}(p(b, \tilde{g}(z))), p(b, \tilde{g}(z))). \end{aligned} \quad (39)$$

We can homotope the second factor of this composition via the straight line homotopy in \widetilde{M} to the map $(b, z) \mapsto \tilde{g}(z)$. We can also homotope the first factor of this composition via the (Euclidean) straight line homotopy to the map $(b, z) \mapsto b$. As we travel along this straight line, the sup distance to b is monotone decreasing, so the inequality (35) is preserved, and therefore the straight line avoids $s \in \mathbb{R}^{k+1}$. Both of these homotopies are \mathbb{Z}^r -equivariant. In summary, we can \mathbb{Z}^r -equivariantly homotope the composition (39) to $(b, z) \mapsto (b, \tilde{g}(z))$ which covers the homotopy equivalence $id \times g$ of the \mathbb{Z}^r -quotients $S_s^k \times \mathbb{T}^r \rightarrow (\mathbb{R}^{k+1} \setminus \{s\}) \times \widetilde{M}/\mathbb{Z}^r$, so the original map (39) also defines a homotopy equivalence of the \mathbb{Z}^r -quotients. But, this original map factors through the map (38) representing the homology cycle, so that cycle is non-zero in homology.

12.5.3 Finishing the proof of Theorem 28. Since $(\widetilde{M} - \vec{h}^{-1}(s))/\mathbb{Z}^r$ is a non-compact n -manifold, it deformation retracts to an $(n - 1)$ -complex. We showed that this $(n - 1)$ -complex has non-zero $(r + k)$ -homology which implies that $r + k \leq n - 1$. This finishes the proof of Theorem 28.

12.6 Proof of Theorem 7 (replacing \mathbb{Z}^r by a general group G). In the proof of the dimension bound (Theorem 28), the Abelian nature of the group \mathbb{Z}^r only played a role in subsection 12.5, an even there the specifics of the group \mathbb{Z}^r played only a minor role. In this subsection we will phrase the argument given in 12.5 in a way that applies to any subgroup $G < \Gamma$. The topological conclusion is that we get the same result as before for any group G if we replace the rank of the Abelian group \mathbb{Z}^r by the homological dimension of the discrete group G . This is expressed in Corollary 38 below, which is a reformulation of Theorem 7 from the introduction. We stated it here in a slightly different, but equivalent, form. The main technical step is the following proposition which combines 12.5.1 and 12.5.2.

PROPOSITION 37. *Suppose that $\sigma : \Delta^k \rightarrow \text{Fix}^0(G)$ is non-degenerate Busemann k -simplex. Then the inclusion of any compact subset of $S^k \times \widetilde{M}/G$ factors up to homotopy through $(\widetilde{M} - \vec{h}^{-1}(s))/G$ for an appropriately chosen intersection of horospheres $\vec{h}^{-1}(s)$.*

Proof. It is clearly enough to prove the proposition for subsets of the form $S^k \times K$, where K is compact. Let $K \hookrightarrow \widetilde{M}/G$ be the inclusion of a subset of diameter $< L$, and let $\widetilde{K} \hookrightarrow \widetilde{M}$ be its inverse image in \widetilde{M} . It is a G -equivariant inclusion. Pick a point $x_0 \in \widetilde{K}$ and form the Busemann cone $\sigma_{\geq 0}$ based at this point. Since σ is a non-degenerate k -simplex, the image of the Busemann cone in horospherical coordinates contains a closed ball $\overline{B}_s(L) \subset W$ of radius L (in the sup metric) centered at some point $s \in W$. Let S_s^k be its boundary sphere and denote by

$$i : S_s^k \times \widetilde{K} \hookrightarrow \mathbb{R}^{k+1} \times \widetilde{M}$$

the standard inclusion. Let

$$f : S_s^k \times \widetilde{K} \rightarrow \mathbb{R}^{k+1} \times \widetilde{M}$$

be the map given by

$$f(b, z) := (\vec{h}(p(b, z)), p(b, z)),$$

and denote by f_t the straight line homotopy (along the Euclidean straight lines $(1 - t)b + t\vec{h}(p(b, z))$ in the first factor and geodesic lines from z to $p(b, z)$ in the second factor) starting at i and ending at f . Since points $b \in S_s^k \subset W$ are in the image of the Busemann cone, we have

$$\vec{h}(p(b, x_0)) = b \quad \text{for all } b \in S_s^k.$$

K has diameter $< L$ so every point z in \widetilde{K} is a distance $< L$ from a point of the orbit Gx_0 . Since $\vec{h}(p(b, \cdot))$ is a G -invariant semicontraction we get

$$\begin{aligned} |\vec{h}(p(b, z)) - b|_\infty &= |\vec{h}(p(b, z)) - \vec{h}(p(b, Gx_0))|_\infty \\ &\leq d(z, Gx_0) \\ &< L \end{aligned}$$

for all points $(b, z) \in S_s^k \times \widetilde{K}$. Therefore the straight line homotopy f_t between f and i avoids $s \times \widetilde{M}$, i.e. it stays in $(\mathbb{R}^{k+1} \setminus s) \times \widetilde{M}$. Finally, note that the map $f = f_1$ factors as

$$S_s^k \times \widetilde{K} \xrightarrow{p(\cdot, \cdot)} \widetilde{M} - \vec{h}^{-1}(s) \xrightarrow{\vec{h} \times id} (\mathbb{R}^{k+1} \setminus s) \times \widetilde{M}.$$

Quotienting out by the G -action, we see that $S_s^k \times K \hookrightarrow (\mathbb{R}^{k+1} \setminus s) \times \widetilde{M}/G$ can be homotoped to factor through the proper open subset $(\widetilde{M} - \vec{h}^{-1}(s))/G$ of \widetilde{M}/G . \square

COROLLARY 38. (Theorem 7) Suppose $\sigma : \Delta^k \rightarrow \text{Fix}^0(G)$ is a non-degenerate Busemann k -simplex. Then

$$hdim(G) + k \leq n - 1.$$

Proof. Let Σ be a non-zero r -dimensional homology class (with any system of coefficients) in \widetilde{M}/G . Then $S^k \times \Sigma$ is a non-trivial $(r+k)$ -dimensional homology class in $S^k \times \widetilde{M}/G$. It is supported on a compact subset $S^k \times K$, so the proposition implies that it defines a non-trivial homology class on the open manifold $(\widetilde{M} \setminus \vec{h}^{-1}(s))/G$. Since this manifold is homotopy equivalent to an $(n-1)$ -complex we conclude that $r+k \leq n-1$. \square

13 Finishing the Proofs of Theorems 1, 2, and 4

Our main application of Theorem 28 is the following corollary. It is the last ingredient we need in order to finish the proofs of Theorems 1, 2, and 4.

COROLLARY 39. Let $\beta : \Delta_{[pAb]} \rightarrow (\partial_\infty, \text{Td})$ be the Busemann map. Then

$$\dim(\text{Im } \beta) \leq \lfloor n/2 \rfloor - 1.$$

Proof. Let $\sigma = ([A_0] < \dots < [A_k])$ be a k -simplex in $\Delta_{[pAb]}$ and note that, because we are using virtual equivalence classes,

$$\text{rank}(A_k) \geq k + 1.$$

Note that the top Abelian group A_k preserves horospheres on the entire Busemann simplex $\beta(\sigma)$. So, if $\beta(\sigma)$ is a non-degenerate, Theorem 28 gives

$$n \geq k + 1 + \text{rank}(A_k).$$

The right hand side is $\geq 2(k+1)$ so we conclude

$$k \leq \lfloor n/2 \rfloor - 1.$$

Since β is Lipschitz, Lipschitz maps do not raise Hausdorff dimension, and topological dimension is less than or equal to Hausdorff dimension, we get for any non-degenerate Busemann simplex $\beta(\sigma)$ that

$$\dim(\beta(\sigma)) \leq \text{Hausdorff-dim}(\beta(\sigma)) \leq \dim(\sigma) \leq \lfloor n/2 \rfloor - 1.$$

The image $\text{Im}(\beta)$ is a countable union of non-degenerate Busemann simplices, so its topological dimension has the same upper bound. \square

Now, let β be the Busemann map. The above corollary implies that

$$d = \dim(\text{Im}(\beta \circ \mu)) \leq \dim(\text{Im } \beta) \leq \lfloor n/2 \rfloor - 1.$$

So, Theorem 18 applied to the Busemann map β proves Theorem 2. We saw in Section 7 and Proposition 21 that $d \leq \text{rank}_{Ab}(\pi_1 M) - 1$ and $d \leq \dim(\partial_\infty, \text{Td})$ so it also proves Theorem 4. Inspecting Theorem 18 we see that Theorem 1 (a) follows from the two bullets in Theorem 18 (in 3.4 we can pick $\delta < \varepsilon/2$) and that Theorem 1 (b) follows from $d \leq \lfloor n/2 \rfloor - 1$.

14 Low Dimensional Collapse: $\dim(\text{Im } \rho) \leq 1$

In this section we will explain the more specific consequences we get when our analogue of the rational Tits building has one dimensional image, i.e. when $d := \dim(\text{Im } \rho) \leq 1$. This happens if one of the following is true.

- The fundamental group $\pi_1 M$ does not contain \mathbb{Z}^3 , or
- $\dim(\partial_\infty, \text{Td}) \leq 1$, or
- $\dim M \leq 5$.

We have seen in Corollary 20 that if $d \leq 1$ then each component of $M_{\leq \varepsilon}$ is aspherical. Let ∂ be a component of $\partial M_{\leq \varepsilon}$. Next we ask “how bad can the fundamental group of ∂ be?” Notice it is an extension of a subgroup of $\pi_1 M$ by the fundamental group of a component $\hat{\partial}$ of $\partial \widetilde{M}_{\leq \varepsilon}$, i.e. it fits into an exact sequence

$$1 \rightarrow \pi_1(\hat{\partial}) \rightarrow \pi_1 \partial \rightarrow \pi_1 M. \tag{40}$$

The main theorem of this section is that if $d \leq 1$ then every finitely generated subgroup of $\pi_1(\hat{\partial})$ is free. Such a group is called a *locally free* group.

Theorem 40. *If $d \leq 1$ then $\pi_1(\hat{\partial})$ is locally free.*

The moral of the theorem is that the fundamental group of the aspherical manifold ∂ is built out of two ingredients that we understand, namely locally free groups and fundamental groups of nonpositively curved manifolds, so it is “not too exotic”.

14.1 Proof of Theorem 40 (Theorem 6 of the introduction). Let $G < \pi_1 \widetilde{M}_{\leq \varepsilon}$ be a finitely generated subgroup. Fix a presentation with finitely many generators for G , let $BG^{(2)}$ be its presentation 2-complex, (a single vertex, one edge for each generator, one 2-cell for each relation) and $BG^{(1)}$ the 1-skeleton of this 2-complex. Since $\widetilde{M}_{\leq \varepsilon} \cong \partial \widetilde{M}_{\leq \varepsilon} \times [0, \infty)$ is aspherical, the inclusion of groups is induced by a map

$$BG^{(2)} \xrightarrow{\varphi} \partial \widetilde{M}_{\leq \varepsilon}. \tag{41}$$

First, we will explain the proof in the simpler case when G is *finitely presented*. In this case, the complex $BG^{(2)}$ is finite, so φ can be homotoped in $\widetilde{M}_{\leq \varepsilon}$ to a map $\bar{\varphi}$ whose image $\bar{\varphi}(BG^{(2)})$ is contained in a graph (see the proof of Corollary 19). Thus the inclusion of $G < \pi_1(\hat{\partial})$ factors through a free group, implying G is free.

In general, we only know that G is finitely generated, i.e. that $BG^{(1)}$ is finite. So, if we try to homotope the map φ to factor through a graph then the finitely many generators $BG^{(1)}$ do not cause any problems but the infinitely many relations $BG^{(2)}$ might. They may “stick out” into the thick part in the course of the homotopy. The key idea is that this is not problematic because the images of the relations are one-dimensional at the end of the homotopy and we can simply attach them to the thin part as “strands” sticking out into the thick part without breaking the argument.

Let us now make precise sense of this idea. Think of φ as a map of pairs

$$(BG^{(2)}, BG^{(1)}) \xrightarrow{\varphi} (\widetilde{M}, \widetilde{M}_{\leq \varepsilon}). \tag{42}$$

CLAIM. We claim that the map of pairs φ can be homotoped inside $(\widetilde{M}, \widetilde{M}_{\leq \varepsilon})$ to a map of pairs $\overline{\varphi}$ whose image is in the 1-skeleton of a triangulation of $(\widetilde{M}, \widetilde{M}_{\leq \varepsilon})$, i.e.

$$(BG^{(2)}, BG^{(1)}) \xrightarrow{\overline{\varphi}} (\widetilde{M}^{(1)}, \widetilde{M}_{\leq \varepsilon}^{(1)}).$$

This is the main technical step in the proof of the theorem. It will follow from a relative version of the collapse argument given in the proof of Theorem 18. We will state it as a separate proposition and, since for this collapse argument it is not important that $d \leq 1$, we state it for general d .

PROPOSITION 41. Suppose K is a compact subset of $\partial \widetilde{M}_{\leq \varepsilon}$. Then the standard inclusion $i : (\partial \widetilde{M}_{\leq \varepsilon}, K) \hookrightarrow (\widetilde{M}, \widetilde{M}_{\leq \varepsilon})$ is homotopic as a map of pairs in $(\widetilde{M}, \widetilde{M}_{\leq \varepsilon})$ to a map

$$\overline{f} : (\partial \widetilde{M}_{\leq \varepsilon}, K) \rightarrow (\widetilde{M}^{(d)}, \widetilde{M}_{\leq \varepsilon}^{(d)})$$

whose image is contained in the d -skeleton of some triangulation of $(\widetilde{M}, \widetilde{M}_{\leq \varepsilon})$.

Proof. In the proof of Theorem 18 we homotoped the inclusion of a compact subset $K \hookrightarrow \partial \widetilde{M}_{\leq \varepsilon}$ inside $\widetilde{M}_{\leq \varepsilon}$ to a map with d -dimensional image. But, the proof really gave us a bit more. We constructed a homotopy between the inclusion $\partial \widetilde{M}_{\leq \varepsilon} \hookrightarrow \widetilde{M}_{\leq \varepsilon}$ and $c_t \circ \beta \circ \mu$ whose restriction to a compact subset K had, for large enough t , the additional property that the set K stayed inside $\widetilde{M}_{\leq \varepsilon}$. Since the image of $c_t \circ \beta \circ \mu$ is d -dimensional, one can deform it (not just its restriction to K , but the entire map $c_t \circ \beta \circ \mu$) by a δ -small deformation into a d -skeleton of a sufficiently fine triangulation of \widetilde{M} . This is done in two steps. First, one deforms the map by a δ -small deformation to a map f' whose image has Hausdorff dimension $\leq d$ (not just topological dimension $\leq d$). Second one deforms f' into a d -skeleton of a sufficiently fine triangulation. The triangulation can be taken to be a triangulation of pairs $(\widetilde{M}, \widetilde{M}_{\leq \varepsilon})$ and this second deformation can be done in a way that preserves simplices of the triangulation, i.e. a point inside a closed simplex stays inside that simplex (see ‘‘Appendix D’’ for more details on these two steps). \square

Now we return to our $d \leq 1$ situation and let $K = \varphi(BG^{(1)})$. Since the map of pairs φ factors through the inclusion i , we can apply Proposition 41 to get the desired homotopy of pairs $F : (BG^{(2)}, BG^{(1)}) \times I \rightarrow (\widetilde{M}, \widetilde{M}_{\leq \varepsilon})$ from $i \circ \varphi$ to $\overline{f} \circ \varphi$. This proves the claim. Next, note that the restriction of F to $BG^{(2)} \times \{0\} \cup BG^{(1)} \times I \cup BG^{(2)} \times \{1\}$ has image contained in $\widetilde{M}^{(1)} \cup \widetilde{M}_{\leq \varepsilon}$. Since this image is aspherical, we can extend this restriction across every $(2\text{-cell}) \times I$ in $BG^{(2)} \times I$ to get a map

$$\widehat{F} : BG^{(2)} \times I \rightarrow \widetilde{M}^{(1)} \cup \widetilde{M}_{\leq \varepsilon}.$$

This is a homotopy between the two maps

$$\varphi : BG^{(2)} \rightarrow \widetilde{M}^{(1)} \cup \widetilde{M}_{\leq \varepsilon} \tag{43}$$

and

$$\bar{f} \circ \varphi : BG^{(2)} \rightarrow \widetilde{M}^{(1)} \cup \widetilde{M}_{\leq \varepsilon}. \quad (44)$$

The first map (43) is π_1 -injective (the original map (41) is π_1 -injective by definition, and attaching the one dimensional “strands” $\widetilde{M}^{(1)}$ to the thin part $\widetilde{M}_{\leq \varepsilon}$ does not change this) while the image of the second map (44) is contained in a graph. Since these two maps are homotopic, we conclude that G is a free group. This proves that $\pi_1(\hat{\partial})$ is locally free.

Acknowledgments

Open access funding provided by Max Planck Society. We would like to thank Igor Belegradek for his interest and comments on earlier versions of this paper. We are very grateful to the two anonymous referees for their detailed comments and suggestions. The first author would like to thank the University of Muenster and the second author would like to thank the Max Planck Institute for Mathematics for their support and excellent working conditions.

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

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

15 Appendix A: A Useful Comparison

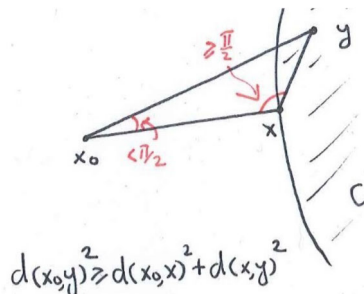
In this paper we repeatedly use the following standard comparison geometry argument.

15.1 Comparison with obtuse Euclidean triangle. Let $C \subset \widetilde{M}$ be a convex set and $p_C : \widetilde{M} \rightarrow C$ the closest point projection. Let $x_0 \notin C, y \in C$ and let $x = p_C(x_0)$ be the closest point to x_0 in C . Since C is convex and x is the closest point to x_0 in C , we observe (see figure above) that

$$\angle_x(x_0, y) \geq \pi/2. \quad (45)$$

Because of this, the Euclidean triangle with sides $[x_0, x]$ and $[x, y]$ meeting at an angle $\angle_x(x_0, y)$ has $d(x_0, y)^2 \geq d(x_0, x)^2 + d(x, y)^2$. Triangle comparison implies that the same is true for our triangle in \widetilde{M} . Consequently,

$$d(x, y) \leq \sqrt{d(x_0, y)^2 - d(x_0, x)^2}. \quad (46)$$



By comparison the sum of angles in any triangle is $\leq \pi$ so (45) also implies

$$\angle_{x_0}(x, y) < \pi/2. \quad (47)$$

Next is a typical example of how this triangle comparison is used.

15.2 Distance dependence of projections to sublevel sets. Suppose $f : \widetilde{M} \rightarrow \mathbb{R}$ is a convex function that does not attain its infimum. Pick a basepoint x_0 . Then for every $R \geq 0$ there is a unique constant $t(R)$ so that the sublevel set $\{f \leq t(R)\}$ meets the sphere $S_{x_0}(R)$ at a single point λ_R . Comparing with Euclidean triangles, we see how λ_R can vary with R .

LEMMA 42. $d(\lambda_{R(1+\delta)}, \lambda_R) \leq \sqrt{2\delta + \delta^2}R$.

Proof. The sublevel set $\{f \leq t(R)\}$ is convex, and λ_R is the point on it that is closest to x_0 . Since $t(R + \varepsilon R) \leq t(R)$, the point $\lambda_{R(1+\delta)}$ is also in the sublevel set $\{f \leq t(R)\}$. So, we can apply inequality (46) with $C = \{f \leq t(R)\}$, $x = \lambda_R$ and $y = \lambda_{R(1+\delta)}$ to get $d(\lambda_R, \lambda_{R(1+\delta)}) \leq \sqrt{2\delta + \delta^2}R$. A similar argument with $C = \{f \leq t(R - \delta R)\}$, $x = \lambda_{R(1-\delta)}$ and $y = \lambda_R$ shows that $d(\lambda_{R(1-\delta)}, \lambda_R) \leq \sqrt{2\delta - \delta^2}R$. \square

16 Appendix B: The Metrics \angle_x , \angle and Td on ∂_∞

In this ‘‘Appendix’’ we collect some results about metrics on ∂_∞ . Everything except for 16.7 can be found in Section 4 and Appendix 3 of [BGS85]. See also 3.5 of [Ebe96].

16.1 Angle metric. For every point $x \in \widetilde{M}$ we have the round metric $(\partial_\infty, \angle_x)$ which identifies the boundary at infinity with the unit sphere $T_x^1 \widetilde{M}$ in the tangent space at x . Taking the supremum over all $x \in \widetilde{M}$ we get the angle metric

$$\angle(\xi, \eta) := \sup_{x \in \widetilde{M}} \angle_x(\xi, \eta).$$

16.2 Description via distance in \widetilde{M} . This metric has the following alternative description. Fix a basepoint x_0 in \widetilde{M} and denote by $r_\xi = [x_0, \xi)$ the unit speed geodesic ray starting at x_0 and going to $\xi \in \partial_\infty$. Then

$$\sin\left(\frac{\angle(\xi, \eta)}{2}\right) = \lim_{t \rightarrow \infty} \frac{d(r_\xi(t), r_\eta(t))}{2t}.$$

16.3 Tits metric. By construction, the space $(\partial_\infty, \angle)$ has diameter π . So, in order to study large scale features of ∂_∞ one takes the induced path metric. This is called the Tits metric and denoted $(\partial_\infty, \text{Td})$. Then

$$\angle = \min(\text{Td}, \pi),$$

so for small scale purposes there is no difference between Td and \angle .

16.4 Relation between sphere and Td topology. The topology obtained from the metric $(\partial_\infty, \text{Td})$ is generally very different from the usual sphere topology $(\partial_\infty, \angle_x)$. But, we have the following relation, which is usually called *lower semicontinuity* of Td in the sphere topology: If $x_i \rightarrow x$ in the sphere topology then

$$\text{Td}(x, y) \leq \liminf_i \text{Td}(x_i, y).$$

Here is a typical application. Let C be a subset of ∂_∞ and \overline{C} its closure in the sphere topology. Then

$$\text{Td-diameter}(C) = \text{Td-diameter}(\overline{C}).$$

16.5 Curvature bounded above. A key feature of Td is that any points x, y with $\text{Td}(x, y) < \pi$ are connected by a unique Td -geodesic. Moreover, Td is $\text{CAT}(1)$, which means that for three points x, y, z mutually a distance $\leq \pi/2$ the geodesic triangle Δ_{xyz} is thinner than the corresponding comparison triangle with the same side lengths in the round sphere.

16.6 Sets of Td -diameter $\leq \pi/2$ have canonical centers. If K is a set in ∂_∞ that has $\text{Td-diameter} \leq \pi/2$ and is closed in the sphere topology, then there is a *unique* point $\xi \in \partial_\infty$ at which the function

$$\tau(\cdot) := \sup_{\eta \in K} \text{Td}(\eta, \cdot)$$

attains its infimum. This point $\xi = \xi_K$ is called *the center* of K . See 3.5.4–3.5.7 of [Ebe96].

16.7 Convex sets contain their centers. In general, the center ξ_K may not be in the set K . However, if K is *convex* in the Td -metric then $\xi_K \in K$. To see this, note that for any point ξ' with $\tau(\xi') \leq \pi/2$ we can take a sequence of points $x_i \in K$ with

$$\text{Td}(x_i, \xi') \rightarrow \inf_{y \in K} \text{Td}(y, \xi')$$

Passing to a subsequence, we can assume x_i converges in the sphere topology to some point $x \in K$, since by assumption K is closed in the sphere topology. Then, by lower semicontinuity of Td , the point x is a closest point to ξ' in the set K and comparison with the round sphere gives $\tau(x) \leq \tau(\xi')$. Since the center ξ_K is the unique infimum of τ , it must be contained in K .

REMARK. In Section 5 we show that if $K = \overline{C}$ is the closure (in the sphere topology) of a Td -convex set C , then we still get the conclusion $\xi_K \in K$.

16.8 Centers of mass for fixed sets of parabolics. One application of this is to finding canonical centers of mass in fixed sets of parabolic elements. Suppose $\gamma \in \Gamma$ is a parabolic element, and let $\text{Fix}(\gamma)$ in ∂_∞ be its fix set. This set is non-empty, and one can find a canonical center of mass inside of it by the following process.

- First, find a point $\xi \in \text{Fix}(\gamma)$ so that $\text{Td}(\xi, \eta) \leq \pi/2$ for all $\eta \in \text{Fix}(\gamma)$. Such a point exists by Appendix 3, Lemma 3 of [BGS85].
- Second take the set of *all such points*

$$B_\gamma = \{\xi \in \text{Fix}(\gamma) \mid \text{Td}(\gamma, \eta) \leq \pi/2 \text{ for all } \eta \in \text{Fix}(\gamma)\}.$$

This set is non-empty, has diameter $\leq \pi/2$ and is closed in the sphere topology. Therefore it has a unique center, which we denote ξ_γ .

- Since B_γ has diameter $\leq \pi/2$, any two points $\xi, \xi' \in B_\gamma$ are connected by a unique geodesic $[\xi, \xi']$. The endpoints of this geodesic are fixed by γ so uniqueness implies that the entire geodesic $[\xi, \xi']$ is fixed pointwise by γ . Moreover, CAT(1) comparison shows that for any point ξ'' on this geodesic $[\xi, \xi']$ we have $\text{Td}(\xi'', \eta) \leq \pi/2$ for all $\eta \in \text{Fix}(\gamma)$. Therefore B_γ is convex and hence

$$\xi_\gamma \in B_\gamma.$$

16.9 Centers of mass for fixed sets of Abelian groups. Since ξ_γ is the center of mass of $\text{Fix}(\gamma)$, it is fixed by anything that preserves this fixed set. So, it is fixed by the anything that commutes with γ . In particular, for any Abelian group A containing a parabolic element γ , the point ξ_γ is contained in the set

$$B_A := \{\xi \in \text{Fix}(A) \mid \text{Td}(\xi, \eta) \leq \pi/2 \text{ for all } \eta \in \text{Fix}(A)\}.$$

This set is Td-convex, has Td-diameter $\leq \pi/2$ and is closed in the sphere topology for the same reasons as B_γ . So, it has a unique center which we denote by $\xi_A \in B_A$. In this way we have constructed a center of mass $\xi_A \in \text{Fix}(A)$ for any Abelian group A containing a parabolic element.

REMARK. In Section 5 we take this construction one step further and obtain a unique center of mass for any *virtual equivalence class of Abelian groups* $[A]$ containing a *parabolic* by replacing the fix set $\text{Fix}(A)$ in this construction with the countable union of fixed sets

$$\bigcup_{n \in \mathbb{N}} \text{Fix}(n!A)$$

of finite index subgroups $n!A = \{\gamma^{n!} \mid \gamma \in A\}$. The added difficulties involved are that this union may not be fixed by any single element and also that it is, possibly, no longer closed in the sphere topology. These difficulties are dealt with in Section 5.

17 Appendix C: The Karlsson-Margulis Lemma (Isometries of Positive Infimum Displacement)

In this “Appendix” we describe a special case (Proposition 43 below) of Theorem 2.1 of [KM99] which is the “nonpositively curved geometry” part of their paper (as opposed to the “ergodic theory part”). There is nothing new here, but we found it comforting to know that the proof of this special case is elementary and does not resort to any ergodic theory.

17.1 Geodesic rays sublinearly tracking γ -orbits. Let γ be an isometry, pick a basepoint y , and let $y_n := \gamma^n y$. The infimum displacement of γ can be computed as the limit

$$A := \lim_{n \rightarrow \infty} \frac{d(y, y_n)}{n}. \tag{48}$$

PROPOSITION 43. *Suppose $A > 0$. Then there is a geodesic ray c for which*

$$\lim_{k \rightarrow \infty} \frac{d(y_k, c(Ak))}{k} = 0. \tag{49}$$

In other words, the Proposition says that if the infimum displacement of γ is positive, then there is a geodesic ray sublinearly tracking the positive γ -orbit $y, \gamma y, \gamma^2 y, \dots$. It follows that the positive γ -orbit converges to a unique limit point

$$\lim_{n \rightarrow \infty} \gamma^n y = c(\infty),$$

but the sublinear tracking is stronger than just this statement alone. Also note that there is another (different) geodesic ray c' that sublinearly tracks the negative γ -orbit $y, \gamma^{-1}y, \gamma^{-2}y, \dots$. In general the rays c and c' do not form a bi-infinite geodesic.

REMARK. When γ is a *hyperbolic* element, then the proposition is easy. If $y \in \widetilde{M}$ is a point at which γ has minimum displacement then the positive γ -orbit of y spans a geodesic ray c and its negative γ -orbit spans a geodesic ray c' . The union of these rays is a bi-infinite geodesic which is an axis of γ . The point of the proposition is that some aspects of this nice situation are still true for parabolic γ as long as the infimum displacement of γ is positive.

17.2 Finding good orbit points. Fix $\varepsilon > 0$. By the displacement formula (48), there is $K = K_\varepsilon$ such that for all $k \geq K$

$$(A - \varepsilon)k \leq d(y, y_k) \leq (A + \varepsilon)k. \tag{50}$$

Since the sequence $\{d(y, y_n) - (A - \varepsilon)n\}_n$ is *unbounded above*, there are arbitrarily large n such that the n th term of this sequence is larger than all its predecessors, i.e. for $0 \leq k \leq n$

$$d(y, y_{n-k}) - (A - \varepsilon)(n - k) \leq d(y, y_n) - (A - \varepsilon)n.$$

Since γ is an isometry $d(y, y_{n-k}) = d(y_k, y_n)$, so this can be rewritten as

$$(A - \varepsilon)k \leq d(y, y_n) - d(y_k, y_n). \quad (51)$$

The right-hand side is $\leq d(y, y_k)$ because of the triangle inequality. Thus, for $K \leq k \leq n$ we get

$$(A - \varepsilon)k \leq d(y, y_n) - d(y_k, y_n) \leq (A + \varepsilon)k \quad (52)$$

if n satisfies (51). Call y_n an ε -good orbit point if n satisfies (51). We emphasize that there are infinitely many ε -good orbit points, so we can always find a sequence $\{y_{n_i}\}$ of such points converging to a point at infinity.

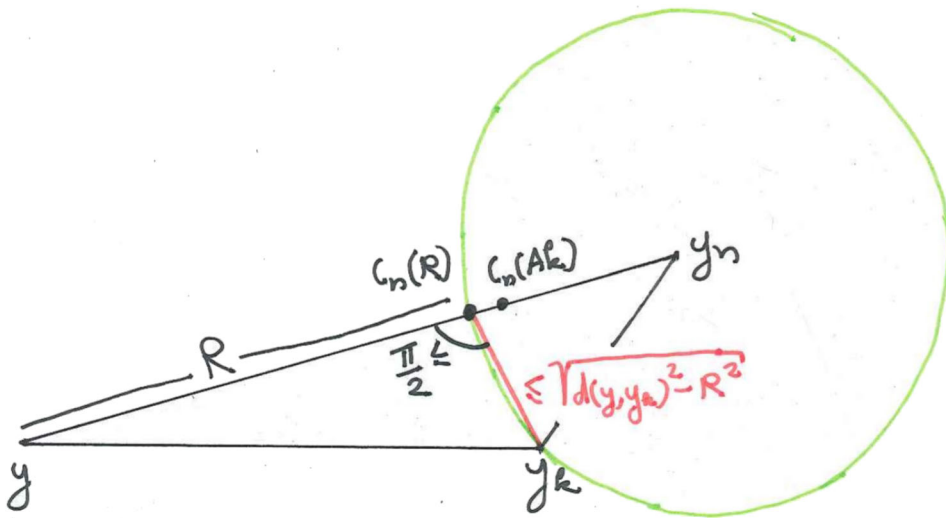
17.3 Finite segments. The heart of the proof is the following lemma. It says that for ε -good orbit points y_n the geodesic segments $[y, y_n]$ sublinearly track a large segment of the γ -orbit y_K, y_{K+1}, \dots, y_n up to an error determined by ε .

LEMMA 44. For $\varepsilon > 0$ pick K_ε satisfying (50) and $n = n_\varepsilon$ satisfying (51). Let $c_n = [y, y_n]$ be the geodesic segment from y to y_n and $c_n(Ak)$ the point obtained by going for a distance Ak along c_n . Then for any $K_\varepsilon \leq k \leq n$

$$\frac{d(y_k, c_n(Ak))}{k} \leq 2\sqrt{A\varepsilon} + \varepsilon. \quad (53)$$

Proof. Let $R := d(y, y_n) - d(y_k, y_n)$. Notice that

$$\angle_{c_n(R)}(y, y_k) \geq \pi/2$$



because the geodesic segment $[c_n(R), y_k]$ is contained in the convex ball $\{d(\cdot, y_n) \leq d(y_k, y_n)\}$ and the segment $[c_n(R), y]$ is perpendicular to this ball. So, comparing with the corresponding obtuse Euclidean triangle gives

$$\begin{aligned} d(y_k, c_n(R)) &\leq \sqrt{d(y, y_k)^2 - R^2} \\ &\leq \sqrt{(A + \varepsilon)^2 k^2 - (A - \varepsilon)^2 k^2} \\ &= 2\sqrt{A\varepsilon}k. \end{aligned}$$

Inequality (52) implies $|R - Ak| \leq \varepsilon k$ so we conclude that

$$\begin{aligned} d(y_k, c_n(Ak)) &\leq d(y_k, c_n(R)) + d(c_n(R), c_n(Ak)) \\ &\leq (2\sqrt{A\varepsilon} + \varepsilon)k. \end{aligned} \quad \square$$

17.4 Limits. What remains is to take limits as $n \rightarrow \infty$ and as $\varepsilon \rightarrow 0$ to get the desired geodesic ray c . We need to take some care in how we do this.

Fix $\varepsilon > 0$ and let $\{y_{n_i}\}$ be a convergent sequence of ε -good orbit points. Then the segments c_{n_i} converge to a geodesic ray

$$c^\varepsilon := \lim_{i \rightarrow \infty} c_{n_i}.$$

The lemma implies that inequality (53) holds for $n = n_i$ and all $K_\varepsilon \leq k \leq n_i$. Therefore, taking the limit of (53) _{$n=n_i$} as $i \rightarrow \infty$, we get for all $k \geq K_\varepsilon$ that

$$\frac{d(y_k, c^\varepsilon(Ak))}{k} \leq 2\sqrt{A\varepsilon} + \varepsilon. \tag{54}$$

Now we vary ε . By the triangle inequality we get from (54) that

$$\frac{d(c^\varepsilon(Ak), c^{\varepsilon'}(Ak))}{k} \leq (2\sqrt{A\varepsilon} + \varepsilon) + (2\sqrt{A\varepsilon'} + \varepsilon') \tag{55}$$

when $k \geq \max\{K_\varepsilon, K_{\varepsilon'}\}$. Therefore, any sequence $\{c^\varepsilon(\infty)\}_{\varepsilon \rightarrow 0}$ is Cauchy in the \angle_y -metric, which implies that c^ε converges to a geodesic ray

$$c := \lim_{\varepsilon \rightarrow 0} c^\varepsilon.$$

Next, notice that because the distance function $d(\cdot, \cdot)$ is convex, knowing (55) for arbitrarily large k implies it for all k . So

$$(55) \text{ holds for } \underline{all} \ k > 0.$$

Taking the limit as $\varepsilon' \rightarrow 0$ of inequality (55), we get that for all $k > 0$

$$\frac{d(c^\varepsilon(Ak), c(Ak))}{k} \leq 2\sqrt{A\varepsilon} + \varepsilon. \tag{56}$$

Finally, putting this together with (54) gives

$$\frac{d(y_k, c(Ak))}{k} \leq 2(2\sqrt{A\varepsilon} + \varepsilon).$$

for all $k \geq K_\varepsilon$. This proves the proposition.

18 Appendix D: δ -Deforming Subsets to Subcomplexes

Let $K \subset \widetilde{M}$ be a compact subset of Hausdorff dimension d . Let δ be greater than zero. A δ -deformation $K \times I \rightarrow \widetilde{M}$ is a homotopy of the inclusion $K \rightarrow \widetilde{M}$ during which each point of K at each point in time $t \in I$ is moved by at most δ from the original position. For any $\delta > 0$ we can δ -deform it to a d -dimensional subcomplex. In this ‘‘Appendix’’ we will recall this argument and also explain why it applies when instead of a compact set we have a countable union of compact sets $\cup_i K_i$. Then, we will explain how to get the same conclusion only assuming that K has topological dimension d (instead of Hausdorff dimension d).

The argument is by induction on skeleta of a fine triangulation of \widetilde{M} . In the course of the induction the inclusion $i : K \hookrightarrow \widetilde{M}$ will be δ -deformed to maps $f : K \rightarrow \widetilde{M}$ that are not necessarily inclusions. So, it is better to state the result like this.

Recall that P is a δ -fine triangulation of \widetilde{M} if every simplex of P has diameter $\leq \delta$.

PROPOSITION 45. *Let K be a compact set and suppose that $f : K \rightarrow \widetilde{M}$ is a map whose image $f(K)$ has Hausdorff dimension d . Then for every $\delta > 0$ the map f can be δ -deformed to a map \hat{f} whose image is contained in a d -dimensional subcomplex of \widetilde{M} .*

More generally, the same is true if $K = \cup_i K_i$ is a countable union of compact sets.

Let P be a smooth, δ -fine triangulation of \widetilde{M} . The main step is to show, for every $k > d$, that if $f(K)$ is contained in the k -skeleton $P^{(k)}$ then we can deform via a family $f_t : K \rightarrow P^{(k)}$ to a map $f_1 : K \rightarrow P^{(k-1)}$ to the $(k-1)$ -skeleton in such a way that during the deformation points of each closed simplex σ of $P^{(k)}$ remain in that simplex, i.e.

$$f_t(\sigma) \subset \sigma,$$

and

$$\dim(f_1(K)) \leq \dim(f(K)) = d.$$

Once one proves this main step, iterating it proves the Proposition.

Now we prove the main step. Note that since k is greater than the Hausdorff dimension of the image $f(K)$, for every k -simplex σ in $P^{(k)}$ there is a point $x \in \text{Interior}(\sigma) - f(K)$ of the interior of σ that is not contained in the image $f(K)$. First let's deal with the case when K is compact. In this case there is a ε_x -neighborhood of x that is disjoint from $f(K)$. Using the linear structure of σ , we can radially project away from x on the simplex σ (and not do anything outside σ) and get a map

$$r_1^\sigma : (P^{(k)} - \{x\}) \rightarrow (P^{(k)} - \text{Interior}(\sigma))$$

that is Lipschitz on $f(K)$ since $\varepsilon_x > 0$. Since this map is Lipschitz, it not increase Hausdorff dimension, i.e.

$$\dim(r_1^\sigma \circ f(K)) \leq \dim(f(K)) = d.$$

This map is part of an obvious radial projection homotopy r_t^σ that is the identity outside the interior of σ . Doing such a deformation on each k -simplex σ of $P^{(k)}$ gives the desired deformation of f_t of f into the $(k-1)$ -skeleton $P^{(k-1)}$. This deformation preserves simplices and does not increase Hausdorff dimension of $f(K)$ because this is true for all the individual deformations r_t^σ . This finishes the proof of the main step when K is compact.

Next, we explain why the same argument applies in general when K is a countable union of compact sets $\cup_i K_i$. We will use the fact that Hausdorff dimension of a countable union of compact sets is the supremum of the Hausdorff dimensions of the individual sets in union. This fact follows directly from the definition of Hausdorff dimension of a set as the infimum of numbers d' for which the d' -dimensional Hausdorff measure of the set is zero, since measures are countably additive. As before, we conclude that the radial projection r_1^σ is Lipschitz on each compact set $f(K_i)$, but the Lipschitz constant might depend on i . However, this doesn't matter because we still get for each individual compact set that

$$\dim(r_1^\sigma \circ f(K_i)) \leq \dim(f(K_i)) \leq d$$

and therefore we get the same bound

$$\dim(r_1^\sigma \circ f(\cup_i K_i)) \leq d$$

for the countable union. Therefore, the main step works for countable unions of compact sets $\cup_i K_i$.

A topological version. There is also a topological version of this, which says that if the topological dimension of K is $\leq d$ then we can δ -deform K to a d -dimensional subcomplex of \widetilde{M} . The simplest way to arrive at this version is to first δ -deform K to a map f' whose image has Hausdorff dimension $\leq d$ and then apply the proposition to f' . So, all we need is the following lemma.

LEMMA 46. *Suppose that $f : S \hookrightarrow \widetilde{M}$ is the inclusion of a subset of topological dimension $\leq d$. For any $\delta > 0$, we can δ -deform it to a map $f' : S \rightarrow \widetilde{M}$ whose image has Hausdorff dimension $\leq d$.*

Proof. Start with a cover of S by δ -balls in \widetilde{M} . Since S has topological dimension $\leq d$, this cover has a refinement whose nerve N has dimension $\leq d$ (by Theorem V.8 of [HW48]). Build a map $S \rightarrow N$ to this nerve using a partition of unity (as in Theorem V.9 of [HW48]) and a map from the nerve $N \rightarrow \widetilde{M}$ using geodesic simplices in \widetilde{M} .

In more detail, the second map is constructed as follows. First, take the barycentric subdivision of the nerve and call it N_1 . Second, send each vertex of the barycentric subdivision to a point in the corresponding intersection of elements of the cover. Third, join adjacent points by geodesics to extend the map to the 1-skeleton of N_1 . Fourth, use the standard order on the vertices of each 2-simplex in the barycentric subdivision to extend the map to the 2-skeleton of N_1 by forming geodesic simplices. Continue in this fashion to produce a map $N \rightarrow \widetilde{M}$. Note that since balls are convex, each simplex of N_1 is mapped into a δ -ball. It follows that the composition $S \rightarrow N \rightarrow \widetilde{M}$ is δ -close to the original map f . Moreover, its image has Hausdorff dimension $\leq d$ because N is d -dimensional and the map $N \rightarrow \widetilde{M}$ built using geodesic simplices is Lipschitz. \square

REMARK. One can try to prove the topological version directly by mimicking the proof of Proposition 45 without first deforming to a map with Hausdorff dimension $\leq d$. The difficulty with doing this direct argument is that it is not clear that the map r_1^σ does not raise topological dimensions of subsets of $P^{(k)} - \{x\}$.

19 Appendix E: Triangulating the Thick Thin Decomposition

In this ‘‘Appendix’’ we describe how to topologically isotope $\partial M_{\leq \varepsilon}$ to a smooth submanifold $S \subset M_{\leq \varepsilon}$. This implies that $\partial M_{\leq \varepsilon}$ has a triangulation. The procedure consists of two parts. We construct the smooth manifold S and specify a vector field V on $M_{\leq \varepsilon}$ which guides the isotopy.

Let us start with the vector field. For this, we need the following fact, which is established during the proof of Lemma 4 in Appendix 2 of [BGS85]. For every $x \in \widetilde{M}_{\leq \varepsilon}$ there is a vector v_x such that for every covering translation γ with $d_\gamma(x) \leq \varepsilon$ we have $\langle v_x, \nabla d_\gamma \rangle \leq 0$ and the inequality is strict for at least one such γ . Since this property is preserved under convex combinations of vectors, we can use partitions of unity to define a smooth vector field V on $M_{\leq \varepsilon}$ whose lift \widetilde{V} to $\widetilde{M}_{\leq \varepsilon}$ satisfies this property at every point.

To produce the smooth manifold, we pick a smooth function $g : \mathbb{R}^+ \rightarrow \mathbb{R}$ with $g(r) = 0$ for $r \geq \varepsilon$, $g'(r) < 0$ on $(0, \varepsilon)$ and $\lim_{r \rightarrow 0^+} g(r) = \infty$. The following function was used in [BGS85]:

$$\widetilde{F}(x) := \sum_{\gamma \neq 1} g \circ d_\gamma(x).$$

This function is Γ -invariant and thus descends to a function $F : M \rightarrow [0, \infty)$ on the quotient. This function is proper, smooth, and has no critical points outside a compact set (see Appendix 2 of [BGS85]). So, for a large enough constant c the level set $S := F^{-1}(c)$ is a smooth, separating submanifold in $M_{\leq \varepsilon}$.

The gradient of \widetilde{F} is

$$\nabla \widetilde{F} = \sum_{\gamma \neq 1} g' \nabla d_\gamma$$

and therefore by the definition of V we have

$$\langle V, \nabla F \rangle = \langle \tilde{V}, \nabla \tilde{F} \rangle > 0. \quad (57)$$

Therefore the integral curves of V meet S exactly once, so flowing along them from $\partial M_{\leq \varepsilon}$ to S (and stopping once one reaches S) defines a topological isotopy from $\partial M_{\leq \varepsilon}$ to S .

References

- [AN17] G. AVRAMIDI and T. T. NGUYEN PHAN. *Examples of noncompact nonpositively curved manifolds*. arXiv preprint [arXiv:1711.03629](https://arxiv.org/abs/1711.03629) (2017)
- [BGS85] W. BALLMANN, M. GROMOV, and V. SCHROEDER. *Manifolds of Nonpositive Curvature*. Birkhäuser **61** (1985)
- [Ebe96] P. EBERLEIN. *Geometry of Nonpositively Curved Manifolds*. University of Chicago Press, Chicago (1996)
- [Gro78] M. GROMOV. Manifolds of negative curvature. *JDG*, (2)13 (1978), 223–230.
- [HW48] W. HUREWICZ and H. WALLMAN. *Dimension Theory*, vol. 4. Princeton University Press, Princeton (1948)
- [KM99] A. KARLSSON and G. MARGULIS. A multiplicative ergodic theorem and nonpositively curved spaces. *Communications in Mathematical Physics*, (1)208 (1999), 107–123
- [Kle99] B. KLEINER. The local structure of length spaces with curvature bounded above, *Mathematische Zeitschrift*, (3)231 (1999), 409–456

GRIGORI AVRAMIDI · T. TÂM NGUYỄN-PHAN, Max Planck Institut für Mathematik, Bonn, Germany.

`gavramidi@mpim-bonn.mpg.de`
`tam@mpim-bonn.mpg.de`

Received: August 21, 2018

Revised: July 29, 2019

Accepted: September 5, 2019