

Geometry & Topology

Volume 29 (2025)

Morse actions of discrete groups on symmetric spaces: local-to-global principle

Michael Kapovich
Bernhard Leeb
Joan Porti

msp

msp

Geometry & Topology **29**:5 (2025) 2343–2390
DOI: 10.2140/gt.2025.29.2343
Published: 14 August 2025

Morse actions of discrete groups on symmetric spaces: local-to-global principle

MICHAEL KAPOVICH
BERNHARD LEEB
JOAN PORTI

Our main result is a local-to-global principle for Morse quasigeodesics, maps and actions. As an application of our techniques we show algorithmic recognizability of Morse actions and construct Morse "Schottky subgroups" of higher-rank semisimple Lie groups via arguments not based on Tits ping-pong. Our argument is purely geometric and proceeds by constructing equivariant Morse quasiisometric embeddings of trees into higher-rank symmetric spaces.

22E40, 53C35; 20F65

1. Introduction 2343
2. Preliminaries 2346
3. Morse maps 2350
4. Group-theoretic applications 2361
Appendix A. Further properties of Morse quasigeodesics 2371
Appendix B. Geometry of nonpositively curved symmetric spaces 2379
References 2388

1 Introduction

This is a sequel to our paper [Kapovich et al. 2017] and mostly consists of the material of Section 7 of our earlier paper [Kapovich et al. 2014] (the only additional material appears in Theorem 4.9 and the appendices). We recall that quasigeodesics in Gromov hyperbolic spaces can be recognized locally by looking at sufficiently large finite pieces; see [Coornaert et al. 1990]. In our earlier papers [Kapovich et al. 2017; 2018a; 2018b] as well as [Kapovich and Leeb 2018b; 2018a], for higher-rank symmetric spaces X (of noncompact type) we introduced an analogue of hyperbolic quasigeodesics, which we call *Morse quasigeodesics*. Morse quasigeodesics are defined relatively to a certain face τ_{mod} of the model spherical face σ_{mod} of X (which can be defined as a fundamental domain for the action on the Tits building of X of the identity component $\mathrm{Isom}_0(X)$ of the isometry group of X). In addition to the quasiisometry constants L and A, τ_{mod}-Morse quasigeodesics come equipped with two other parameters, a positive number D

© 2025 MSP (Mathematical Sciences Publishers). Distributed under the Creative Commons Attribution License 4.0 (CC BY). Open Access made possible by subscribing institutions via Subscribe to Open.

and a *Weyl-convex* subset Θ of the *open star* of τ_{mod} in the modal spherical chamber σ_{mod}. In [Kapovich et al. 2014; 2017; 2018b] we also defined τ_{mod}-Morse maps $Y \to X$ from Gromov-hyperbolic spaces to symmetric spaces. These maps are defined by the property that they send geodesics to *uniformly* τ_{mod}*-Morse quasigeodesics*, ie τ_{mod}-Morse quasigeodesics with a fixed set of parameters, (Θ, D, L, A).

The main result of this paper is a *local* characterization of Morse quasigeodesics in X:

Theorem 1.1 (local-to-global principle for Morse quasigeodesics) *For L, A, Θ, Θ', D there exist S, L', A', D' such that every S-local (Θ, D, L, A)-local Morse quasigeodesic in X is a (Θ', D', L', A')–Morse quasigeodesic.*

Here S-locality of a certain property of a map means that this property is satisfied for restrictions of this map to subintervals of length S. We refer to Theorem 3.34 and Theorem 3.34 for the details. Based on this principle, we prove in Section 3.7 a local-to-global principle for Morse maps from hyperbolic metric spaces to symmetric spaces.

Below, G is a semisimple Lie group acting isometrically and transitively on X, and K is a maximal compact subgroup of G, so that X is diffeomorphic to G/K. We will assume that G is commensurable with the isometry group $\mathrm{Isom}(X)$ in the sense that we allow finite kernel and cokernel for the natural map $G \to \mathrm{Isom}(X)$. However, we require G to act trivially on the model spherical chamber of X; see Section B for details.

We prove several consequences of the local-to-global principle:

1. (Theorems 4.4 and 4.7) The *structural stability* of Morse subgroups of G, generalizing Sullivan's structural stability theorem in rank one [1985]; see also [Kapovich et al. 2024] for a detailed proof. While structural stability for Anosov subgroups was known earlier (Labourie and Guichard–Wienhard), our method is more general and applies to a wider class of discrete subgroups; see [Kapovich and Leeb $\geq$ 2025].

Theorem 1.2 (openness of the space of Morse actions) *For a word hyperbolic group Γ, the subset of τ_{mod}-Morse actions of Γ on X is open in $\mathrm{Hom}(\Gamma, G)$.*

Theorem 1.3 (structural stability) *Let Γ be word hyperbolic. Then for τ_{mod}-Morse actions $\rho : \Gamma \curvearrowright X$, the boundary embedding $\alpha_\rho : \partial_\infty \Gamma \to \mathrm{Flag}(\tau_{\mathrm{mod}})$ depends continuously on the action ρ.*

In particular, actions sufficiently close to a faithful Morse action are again discrete and faithful. We supplement this structural stability theorem with a stability theorem on *domains of proper discontinuity*, Theorem 4.9.

2. (Section 4.3) The locality of the Morse property implies that Morse subgroups are algorithmically recognizable:

Theorem 1.4 (semidecidability of Morse property of group actions) *Let Γ be word hyperbolic. Then there exists an algorithm whose inputs are homomorphisms $\rho : \Gamma \to G$ (defined on generators of Γ) and which terminates if and only if ρ defines a τ_{mod}-Morse action $\Gamma \curvearrowright X$.*

In other words, the property of being a τ_{mod}-Morse subgroup of G is *semidecidable*. If an action of Γ on X is not Morse, the algorithm runs forever. Note that in view of [Kapovich 2016], there are no algorithms (in the sense of BSS computability) which would recognize if a representation $\Gamma \to \mathrm{Isom}(\mathbb{H}^3)$ is *not geometrically finite*.

3. (Section 4.2) We illustrate our techniques by constructing Morse–Schottky actions of free groups on higher-rank symmetric spaces. Unlike all previously known constructions, our proof does not rely on ping-pong arguments, but is purely geometric and proceeds by constructing equivariant quasiisometric embeddings of trees. The key step is the observation that a certain local *straightness* property for sufficiently spaced sequences of points in the symmetric space implies the global Morse property. This observation is also at the heart of the proof of the local-to-global principle for Morse actions.

Since [Kapovich et al. 2014] was originally posted, several improvements on the material of its Section 7 and, hence, of the present paper were made:

(a) Different forms of combination theorems for Anosov subgroups were proven in [Dey et al. 2019] and [Dey and Kapovich 2023; 2025], written in collaboration with Subhadip Dey. The first one was a generalization of the technique in Section 4.2 of the present paper, but the other two generalizations are based on a form of the ping-pong argument.

(b) Explicit estimates in the local-to-global principle for Morse quasigeodesics and, hence, Morse embeddings, were obtained by Max Riestenberg [2025]. Riestenberg's estimates are based on replacing certain limiting arguments used in the present paper with differential-geometric and Lie-theoretic arguments.

Organization of the paper The notions of Morse quasigeodesics and actions are discussed in detail in Section 3. In that section, among other things, we establish local-to-global principles for Morse quasigeodesics.

In Section 4 we apply local-to-global principles to discrete subgroups of Lie groups: We show that Morse actions are structurally stable and algorithmically recognizable. We also construct Morse–Schottky actions of free groups on symmetric spaces. In Appendices A and B we prove further properties of Morse quasigeodesics that we found to be useful in our work.

Acknowledgements Kapovich was supported by NSF grants DMS-12-05312 and DMS-16-04241, by KIAS (the Korea Institute for Advanced Study) through the KIAS scholar program, and by a Simons Foundation Fellowship, grant 391602. Porti was supported by Spanish State Research Agency, through grant PID2021-125625NB-100 and Maria de Maeztu Program CEX2020-001084-M.

The three authors are also grateful to the GEAR grant which partially supported the IHP trimester in Winter of 2012 through DMS 1107452, 1107263, 1107367 "*RNMS*: *Geometric structures and representation varieties*" (the GEAR Network), and to the Max Planck Institute for Mathematics in Bonn, where some of this work was done. We are grateful to the referee for carefully reading the paper and for making various useful suggestions.

2 Preliminaries

2.1 Basic notions of geometry of symmetric spaces

Throughout the paper we will be using definitions, notation and results of our earlier work. For the reader's convenience, we also review these notions in detail in Appendix B.

We refer the reader to our earlier papers for the various notions related to symmetric spaces such as *polyhedral Finsler metrics* on symmetric spaces [Kapovich and Leeb 2018b], the *opposition involution* ι of σ_{mod}, *model faces* τ_{mod} of σ_{mod} and the associated τ_{mod}-flag-manifolds $\mathrm{Flag}(\tau_{\mathrm{mod}})$ [Kapovich et al. 2017, Sections 2.2.2 and 2.2.3], *Weyl-convex* subsets $\Theta \subset \sigma_{\mathrm{mod}}$ [Kapovich et al. 2017, Definition 2.7], *type map* $\theta\colon \partial_\infty X \to \sigma_{\mathrm{mod}}$, open Schubert cells $C(\tau) \subset \mathrm{Flag}(\tau_{\mathrm{mod}})$ [Kapovich et al. 2017, Section 2.4], Δ-valued distances d_Δ on X [Kapovich et al. 2017, Section 2.6], Θ-regular geodesic segments (see [Kapovich et al. 2017, Section 2.5.3]), *parallel sets*, *stars*, *open stars* and Θ-stars, $\mathrm{st}(\tau)$, $\mathrm{ost}(\tau)$, and $\mathrm{st}_\Theta(\tau)$, *Weyl sectors* $V(x,\tau)$ [Kapovich et al. 2017, Section 2.4], *Weyl cones* $V(x,\mathrm{st}(\tau))$ and Θ-cones $V(x,\mathrm{st}_\Theta(\tau))$, *diamonds* $\diamondsuit_{\tau_{\mathrm{mod}}}(x,y)$ and Θ-diamonds $\diamondsuit_\Theta(x,y)$ [Kapovich et al. 2017, Section 2.5], τ_{mod}*-regular sequences and groups* [Kapovich et al. 2017, Section 4.2], τ_{mod}-convergence subgroups, flag-convergence, the Finsler interpretation of flag-convergence (see [Kapovich and Leeb 2018b, Sections 4.5 and 5.2] and [Kapovich et al. 2017]), τ_{mod}-limit sets $\Lambda_{\tau_{\mathrm{mod}}}(\Gamma) \subset \mathrm{Flag}(\tau_{\mathrm{mod}})$ [Kapovich et al. 2017, Section 4.5]), visual limit set [Kapovich et al. 2017, page 4], uniformly τ_{mod}-regular sequences and subgroups [Kapovich et al. 2017, Section 4.6]), Morse subgroups [Kapovich et al. 2017, Section 5.4] and, more generally, Morse quasigeodesics and Morse maps [Kapovich et al. 2018b, Definitions 5.31 and 5.33]), antipodal limit sets [Kapovich et al. 2017, Definition 5.1] and antipodal maps to flag-manifolds [Kapovich et al. 2018b, Definition 6.11].

In the paper we will be frequently using convexity of Θ-cones in X:

Proposition 2.1 [Kapovich et al. 2017, Proposition 2.10] *For every Weyl-convex subset $\Theta \subset \mathrm{st}(\tau_{\mathrm{mod}})$, for every $x \in X$ and $\tau \in \mathrm{Flag}(\tau_{\mathrm{mod}})$, the cone $V(x,\mathrm{st}_\Theta(\tau)) \subset X$ is convex.*

2.2 Standing notation and conventions

- We will use the notation X for a symmetric space of *noncompact type*, G for a semisimple Lie group acting isometrically and transitively on X, and K for a maximal compact subgroup of G, so that X is diffeomorphic to G/K. We will assume that G is commensurable with the isometry group $\mathrm{Isom}(X)$ in the sense that we allow finite kernel and cokernel for the natural map $G \to \mathrm{Isom}(X)$. In particular, the image of G in $\mathrm{Isom}(X)$ contains the identity component $\mathrm{Isom}(X)_o$.
- We let $\tau_{\mathrm{mod}} \subseteq \sigma_{\mathrm{mod}}$ be a fixed ι-invariant face type.
- We will use the notation $x_n \xrightarrow{f} \tau \in \mathrm{Flag}(\tau_{\mathrm{mod}})$ for the *flag-convergence* of a τ_{mod}-regular sequence $x_n \in X$ to a simplex $\tau \in \mathrm{Flag}(\tau_{\mathrm{mod}})$.

- We will be using the notation Θ, Θ' for an ι-invariant, compact, Weyl-convex [Kapovich et al. 2017, Definition 2.7] subset of the open star $\mathrm{ost}(\tau_{\mathrm{mod}}) \subset \sigma_{\mathrm{mod}}$.
- We will always assume that $\Theta < \Theta'$, meaning that $\Theta \subset \mathrm{int}(\Theta')$.
- Constants $L, A, D, \epsilon, \delta, l, a, s, S$ are meant to be always strictly positive and $L \geq 1$.

2.3 ζ-angles

We fix as auxiliary datum an ι-invariant type $\zeta = \zeta_{\mathrm{mod}} \in \mathrm{int}(\tau_{\mathrm{mod}})$. (We will omit the subscript in ζ_{mod} in order to avoid cumbersome notation for ζ-angles.) For a simplex $\tau \subset \partial_\infty X$ of type τ_{mod}, ie $\tau \in \mathrm{Flag}(\tau_{\mathrm{mod}})$, we define $\zeta(\tau) \in \tau$ as the ideal point of type ζ_{mod}. Given two such simplices $\tau_\pm \in \mathrm{Flag}(\tau_{\mathrm{mod}})$ and a point $x \in X$, define the *ζ-angles*

$$\angle_x^\zeta(\tau_-, \tau_+) := \angle_x^\zeta(\tau_-, \xi_+) := \angle_x(\xi_-, \xi_+), \tag{2.2}$$

where $\xi_\pm = \zeta(\tau_\pm)$.

Similarly, define the *ζ-Tits angle*

$$\angle_{\mathrm{Tits}}^\zeta(\tau_-, \tau_+) = \angle_{\mathrm{Tits}}^\zeta(\tau_-, \xi_+) := \angle_{\mathrm{Tits}}(\xi_-, \xi_+) = \angle_x(\xi_-, \xi_+), \tag{2.3}$$

where x belongs to a flat $F \subset X$ such that $\tau_-, \tau_+ \subset \partial_{\mathrm{Tits}} F$. Then simplices $\tau_\pm$ (of the same type) are antipodal if and only if

$$\angle_{\mathrm{Tits}}^\zeta(\tau_-, \tau_+) = \pi$$

for some, equivalently, every, choice of ζ as above.

Remark 2.4 We observe that the ideal points $\zeta_\pm$ are opposite, $\angle_{\mathrm{Tits}}(\zeta_-, \zeta_+) = \pi$, if and only if they can be seen under angle $\simeq \pi$ (ie close to π) from some point in X. More precisely, there exists $\epsilon(\zeta_{\mathrm{mod}})$ such that:

If $\angle_x(\zeta_-, \zeta_+) > \pi - \epsilon(\zeta_{\mathrm{mod}})$ for some point x, then $\zeta_\pm$ are opposite.

This follows from the angle comparison $\angle_x(\zeta_-, \zeta_+) \leq \angle_{\mathrm{Tits}}(\zeta_-, \zeta_+)$ and the fact that the Tits distance between ideal points of the fixed type ζ_{mod} takes only finitely many values.

For a τ_{mod}-regular unit tangent vector $v \in T_x X$ we denote by $\tau(v) \subset \partial_\infty X$ the unique simplex of type τ_{mod} such that ray ρ_v with the initial direction v represents an ideal point in $\mathrm{ost}(\tau(v))$. We put $\zeta(v) = \zeta(xy) = \zeta(\tau(v))$, where xy is a geodesic segment in X with the initial direction v. Note that $\zeta(v)$ depends continuously on $v \in TX$.

For a τ_{mod}-regular segment xy in X we let $\tau(xy) = \tau(v)$, where v is the unit vector tangent to xy.

For τ_{mod}-regular segments xy, xz and $\tau \in \mathrm{Flag}(\tau_{\mathrm{mod}})$, we define the ζ-angles

$$\angle_x^\zeta(\tau, y) = \angle_x^\zeta(y, \tau) = \angle_x^\zeta(\tau(xy), \tau), \quad \text{and} \quad \angle_x^\zeta(y, z) = \angle_x^\zeta(\tau(xy), \tau(xz)). \tag{2.5}$$

Note that the ζ-angle $\angle_x^\zeta(y,z)$ depends not on y,z but rather on the simplices $\tau(xy),\tau(xz)$. These ζ-angles will play the role of angles the between diamonds $\diamondsuit_{\tau_{\mathrm{mod}}}(x,y)$ and $\diamondsuit_{\tau_{\mathrm{mod}}}(x,z)$, meeting at x. Note that if X has rank 1, then the ζ-angles are just the ordinary Riemannian angles.

We observe that compactness of the set of Θ-regular unit vectors in T_xX and transitivity of the G-action on X imply that the map $v\mapsto\zeta(v)$ from the set of Θ-regular unit vectors to $G\zeta$ is uniformly continuous. This implies uniform continuity of the ζ-angle function $\angle_x^\zeta(\tau,y)$.

2.4 Distances to parallel sets versus angles

In this section we collect some geometric facts regarding parallel sets in symmetric spaces, primarily dealing with estimation of distances from points in X to parallel sets.

Remark 2.6 The constants and functions in this section are not explicit, and their existence is proven by compactness arguments. For explicit computations here and in Theorem 3.18, we refer the reader to the PhD thesis of Max Riestenberg [2025].

We first prove a lemma which strengthens Corollary 2.46 of [Kapovich et al. 2017].

Lemma 2.7 *Suppose that $\tau_\pm$ are antipodal simplices in $\partial_{\mathrm{Tits}}X$. Then every geodesic ray γ asymptotic to a point $\xi\in ost(\tau_+)$ is strongly asymptotic to a geodesic ray in $P(\tau_-,\tau_+)$.*

Proof If ξ belongs to the interior of the simplex τ_+, then the assertion follows from Corollary 2.46 of [Kapovich et al. 2017]:

> *Weyl sectors $V(x_1,\tau)$ and $V(x_2,\tau)$ are strongly asymptotic if and only if x_1 and x_2 lie in the same horocycle at τ.*

We now consider the general case. Suppose that ξ belongs to an open simplex $\mathrm{int}(\tau')$ such that τ is a face of τ'. Then there exists an apartment $a\subset\partial_{\mathrm{Tits}}X$ containing both ξ (and, hence, τ' as well as τ) and the simplex τ_-. Let $F\subset X$ be the maximal flat with $\partial_\infty F=a$. Then F contains a geodesic asymptotic to points in τ_- and τ_+. Therefore, F is contained in $P(\tau_-,\tau_+)$. On the other hand, by Corollary 2.46 of [Kapovich et al. 2017], applied to the simplex τ', we conclude that γ is strongly asymptotic to a geodesic ray in F. □

The following lemma provides a quantitative strengthening of the conclusion of Lemma 2.7:

Lemma 2.8 *Let Θ be a compact subset of $ost(\tau_+)$. Then those rays $x\xi$ with $\theta(\xi)\in\Theta$ are uniformly strongly asymptotic to $P(\tau_-,\tau_+)$, ie $d(\cdot,P(\tau_-,\tau_+))$ decays to zero along them uniformly in terms of $d(x,P(\tau_-,\tau_+))$ and Θ.*

Proof Suppose that the assertion of lemma is false, ie there exists $\epsilon>0$, a sequence $T_i\in\mathbb{R}_+$ diverging to infinity, and a sequence of rays $\rho_i=x_i\xi_i$ with $\xi_i\in\Theta$ and $d(x_i,P(\tau_-,\tau_+))\le d$, such that

$$d(y,P(\tau_-,\tau_+))\ge\epsilon \quad \text{for all } y\in\rho([0,T_i]). \tag{2.9}$$

Using the action of the stabilizer of $P(\tau_-, \tau_+)$, we can assume that the points x_i belong to a certain compact subset of X. Therefore, the sequence of rays $x_i\xi_i$ subconverges to a ray $x\xi$ with $d(x, P(\tau_-, \tau_+)) \le d$ and $\xi \in \Theta$. The inequality (2.9) then implies that the entire limit ray $x\xi$ is contained outside of the open ϵ-neighborhood of the parallel set $P(\tau_-, \tau_+)$. However, in view of Lemma 2.7, the ray $x\xi$ is strongly asymptotic to a geodesic in $P(\tau_-, \tau_+)$. Contradiction. □

We next relate distances from points $x \in X$ to parallel sets and the ζ-angles at x. Suppose that the simplices $\tau_\pm$, equivalently, the ideal points $\zeta_\pm = \zeta(\tau_\pm)$ (see Section 2.3), are opposite. Then

$$\angle_x^\zeta(\tau_-, \tau_+) = \angle_x(\zeta_-, \zeta_+) = \pi$$

if and only if x lies in the parallel set $P(\tau_-, \tau_+)$. Furthermore, $\angle_x^\zeta(\tau_-, \tau_+) \simeq \pi$ if and only if x is close to $P(\tau_-, \tau_+)$, and both quantities control each other near the parallel set. More precisely:

Lemma 2.10 (i) *If $d(x, P(\tau_-, \tau_+)) \le d$, then $\angle_x^\zeta(\tau_-, \tau_+) \ge \pi - \epsilon(d)$ with $\epsilon(d) \to 0$ as $d \to 0$.*

(ii) *For sufficiently small ϵ, $\epsilon \le \epsilon'(\zeta_{\mathrm{mod}})$, we have: The inequality $\angle_x^\zeta(\tau_-, \tau_+) \ge \pi - \epsilon$ implies that $d(x, P(\tau_-, \tau_+)) \le d(\epsilon)$ for some function $d(\epsilon)$ which converges to 0 as $\epsilon \to 0$.*

Proof The intersection of parabolic subgroups $P_{\tau_-} \cap P_{\tau_+}$ preserves the parallel set $P(\tau_-, \tau_+)$ and acts transitively on it. Compactness and the continuity of $\angle_\cdot(\zeta_-, \zeta_+)$ therefore imply that $\pi - \angle_\cdot(\zeta_-, \zeta_+)$ attains on the boundary of the tubular r-neighborhood of $P(\tau_-, \tau_+)$ a strictly positive maximum and minimum, which we denote by $\phi_1(r)$ and $\phi_2(r)$. Furthermore, $\phi_i(r) \to 0$ as $r \to 0$. We have the estimate

$$\pi - \phi_1(d(x, P(\tau_-, \tau_+))) \le \angle_x(\zeta_-, \zeta_+) \le \pi - \phi_2(d(x, P(\tau_-, \tau_+))).$$

The functions $\phi_i(r)$ are (weakly) monotonically increasing. This follows from the fact that, along rays asymptotic to ζ_- or ζ_+, the angle $\angle_\cdot(\zeta_-, \zeta_+)$ is monotonically increasing and the distance $d(\cdot, P(\tau_-, \tau_+))$ is monotonically decreasing. The estimate implies the assertions. □

The control of $d(\cdot, P(\tau_-, \tau_+))$ and $\angle_\cdot(\zeta_-, \zeta_+)$ "spreads" along the Weyl cone $V(x, \mathrm{st}(\tau_+))$, since the latter is asymptotic to the parallel set $P(\tau_-, \tau_+)$. Moreover, the control improves if one enters the cone far into a τ_{mod}-regular direction. More precisely:

Lemma 2.11 *Let $y \in V(x, \mathrm{st}_\Theta(\tau_+))$ be a point with $d(x, y) \ge l$.*

(i) *If $d(x, P(\tau_-, \tau_+)) \le d$, then*

$$d(y, P(\tau_-, \tau_+)) \le D'(d, \Theta, l) \le d,$$

with $D'(d, \Theta, l) \to 0$ as $l \to +\infty$.

(ii) *For sufficiently small ϵ, $\epsilon \le \epsilon'(\zeta_{\mathrm{mod}})$, we have: If $\angle_x(\zeta_-, \zeta_+) \ge \pi - \epsilon$, then*

$$\angle_y(\zeta_-, \zeta_+) \ge \pi - \epsilon'(\epsilon, \Theta, l) \ge \pi - \epsilon(d(\epsilon)),$$

with $\epsilon'(\epsilon, \Theta, l) \to 0$ as $l \to +\infty$.

Proof The distance from $P(\tau_-,\tau_+)$ takes its maximum at the tip x of the cone $V(x,\mathrm{st}(\tau_+))$, because it is monotonically decreasing along the rays $x\xi$ for $\xi\in\mathrm{st}(\tau_+)$. This yields the right-hand bounds d and, applying Lemma 2.10 twice, $\epsilon(d(\epsilon))$.

Those rays $x\xi$ with uniformly τ_{mod}-regular type $\theta(\xi)\in\Theta$ are uniformly strongly asymptotic to $P(\tau_-,\tau_+)$, ie $d(\cdot\,,P(\tau_-,\tau_+))$ decays to zero along them uniformly in terms of d and Θ; see Lemma 2.8. This yields the decay $D'(d,\Theta,l)\to 0$ as $l\to+\infty$. The decay of ϵ' follows by applying Lemma 2.10 again. □

3 Morse maps

In this section we investigate the Morse property of sequences and maps. The main aim of this section is to establish a local criterion for being Morse. To do so we introduce a local notion of *straightness* for sequences of points in X. Morse sequences are in general not straight, but they become straight after suitable modification, namely by sufficiently coarsifying them and then passing to the sequence of successive midpoints. Conversely, the key result is that sufficiently spaced straight sequences are Morse. We conclude that there is a local-to-global characterization of the Morse property.

3.1 Morse quasigeodesics

Definition 3.1 (Morse quasigeodesic) A (Θ,D,L,A)-*Morse quasigeodesic* in X is an (L,A)-quasigeodesic $p\colon I\to X$ (defined on an interval $I\subset\mathbb{R}$) such that for all $t_1,t_2\in I$, the subpath $p|_{[t_1,t_2]}$ is D-close to a Θ-diamond $\diamondsuit_\Theta(x_1,x_2)$ with $d(x_i,p(t_i))\le D$.

We will refer to a quadruple (Θ,D,L,A) as a *Morse datum* and abbreviate $M=(\Theta,D,L,A)$. Set $M+D'=(\Theta,D+D',L,A+2D')$. We say that M contains Θ if M has the form (Θ,D,L,A) for some $D\ge 0$, $L\ge 1$, $A\ge 0$.

The following lemma is immediate from the definition of an M-Morse quasigeodesic.

Lemma 3.2 (perturbation lemma) *If p, p' are paths in X such that p is M-Morse and $d(p,p')\le D'$, then p' is $M+D'$-Morse.*

A Morse quasigeodesic p is called a *Morse ray* if its domain is a half-line. If $I=\mathbb{R}$, then a Morse quasigeodesic is called a *Morse quasiline*.

Morse quasirays do not in general converge at infinity (in the visual compactification of X), but they τ_{mod}-converge at infinity. This is a consequence of:

Lemma 3.3 (conicality) *Every Morse quasiray $p\colon[0,\infty)\to X$ is uniformly Hausdorff close to a subset of a cone $V(p(0),\mathrm{st}_\Theta(\tau))$ for a unique simplex τ of type τ_{mod}.*

Proof The subpaths $p|_{[0,t_0]}$ are uniformly Hausdorff close to Θ-diamonds. These subconverge to a cone $V(x, \mathrm{st}_\Theta(\tau))$ x uniformly close to $p(0)$ and τ a simplex of type τ_{mod}. This establishes the existence. Since $p(n) \xrightarrow{f} \tau$, the uniqueness of τ follows from the uniqueness of τ_{mod}-limits; see [Kapovich et al. 2017, Lemma 4.23]. □

Definition 3.4 (end of Morse quasiray) We call the unique simplex given by the previous lemma the *end* of the Morse quasiray $p\colon [0,\infty) \to X$, and denote it by

$$p(+\infty) \in \mathrm{Flag}(\tau_{\mathrm{mod}}).$$

Hausdorff close Morse quasirays have the same end by Lemma 3.3. In Section 3.3 we will prove uniform continuity of ends of Morse quasirays with respect to the topology of *coarse convergence* of quasirays.

3.2 Morse maps

We now turn to *Morse maps* with more general domains (than just intervals).

Definition 3.5 Let Y be a Gromov-hyperbolic geodesic metric space. A map $f\colon Y \to X$ is called M-Morse if it sends geodesics in Y to M-Morse quasigeodesics.

Thus, every Morse map is a quasiisometric embedding. While this definition makes sense for general metric spaces, in [Kapovich et al. 2018b] we proved that the domain of a Morse map is necessarily hyperbolic.

More generally, one can define Morse maps on *quasigeodesic metric spaces*:

Definition 3.6 (quasigeodesic metric space) A metric space Z is called (l,a)-*quasigeodesic* if all pairs of points in Y can be connected by (l,a)-quasigeodesics. A space is called *quasigeodesic* if it is (l,a)-quasigeodesic for some pair of parameters l, a.

Every quasigeodesic space is quasiisometric to a geodesic metric space. Namely, if Z is a (λ,α)-quasigeodesic space then it is quasiisometric to the 1-skeleton of its $(\lambda+\alpha)$-Rips complex (one proves this by repeating the proof of [Druțu and Kapovich 2018, Theorem 8.52]). The quasigeodesic spaces considered in this paper are discrete groups equipped with word metrics, for which the claim is clear since we can use the Cayley graph as the geodesic space. The next definition is a slight generalization of the notion of Morse maps defined above.

Definition 3.7 (Morse embedding) Let (Θ, D, L, A) be a Morse datum. An (Θ, D, L, A, l, a)-*Morse embedding* from an (l,a)-quasigeodesic space Z into X is a map $f\colon Z \to X$ which sends (l,a)-quasigeodesics in Z to (Θ, D, L, A)-Morse quasigeodesics in X.

Of course, every (l,a)-quasigeodesic metric space is also (l',a')-quasigeodesic space for any $l' \geq l, a' \geq a$. The next lemma shows that this choice of quasigeodesic constants is essentially irrelevant.

Lemma 3.8 *Let $f: Z \to X$ be a map from a Gromov-hyperbolic (l,a)-quasigeodesic space Z. If f is $M = (\Theta, D, L, A, l, a)$-Morse then for any (l',a'), it sends (l',a')-quasigeodesics in Z to $M' = (\Theta, D', L', A')$-Morse quasigeodesics in X. Here the datum M' depends only on M, l', a' and the hyperbolicity constant δ of Z.*

Proof This follows from the definition of Morse quasigeodesics, and the Morse lemma applied to Z. □

Notice that the parameter Θ in the Morse datum M' is the same as in M. Hence, we arrive at:

Definition 3.9 A map $f: Z \to X$ of a quasigeodesic hyperbolic space Z is called Θ-*Morse* if it sends uniform quasigeodesics in Z to Θ-Morse uniform quasigeodesics in X.

This notion depends only on the quasiisometry class of Z, ie the precomposition of a Θ-Morse embedding with a quasiisometry is again Θ-Morse. For this to be true we have to require control on the images of quasigeodesics of arbitrarily bad (but uniform) quality.

Let Γ be a hyperbolic group with a fixed finite generating set S, and let Y be the Cayley graph of Γ with respect to S. For $x \in X$, an isometric action $\Gamma \curvearrowright X$ determines the *orbit map* $o_x: \Gamma \to \Gamma x \subset X$. Every such map extends to the Cayley graph Y of Γ, sending edges to geodesics in X.

Definition 3.10 An isometric action $\Gamma \curvearrowright X$, or a representation $\rho: \Gamma \to G$, is called M-Morse (with respect to a basepoint $x \in X$) if the (extended) orbit map $o_x: Y \to X$ is M-Morse. Similarly, a subgroup $\Gamma < G$ is Morse if the inclusion homomorphism $\Gamma \hookrightarrow G$ is Morse.

The Morse property of an action and the parameter Θ, of course, do not depend on the choice of a generating set of Γ and a basepoint x, but the triple (D, L, A) does. Thus, it makes sense to talk about Θ-Morse and τ_{mod}-Morse actions of hyperbolic groups, where $\Theta \subset \mathrm{ost}(\tau_{\mathrm{mod}})$. In [Kapovich et al. 2017; 2018b] and [Kapovich and Leeb 2018b] we gave many alternative definitions of Morse actions, including the equivalence of this definition to the notion of Anosov subgroups.

3.3 Continuity at infinity

Let X, Y be proper metric spaces. We fix a basepoint $y \in Y$.

Definition 3.11 A sequence of maps $f_n: Y \to X$ is said to *coarsely converge* to a map $f: Y \to X$ if there exists $C < \infty$ such that for every R there exists $N = N(C, R)$ for which

$$d(f_n|_B, f|_B) \leq C \quad \text{for all } n \geq N,$$

where $B = B(y, R)$.

Note the difference between this definition and the notion of uniform convergence on compacts: since we are working in the coarse setting, requiring the distance between maps to be less than ϵ close to zero is pointless.

In view of the Arzelà–Ascoli theorem, the space of (L, A)-coarse Lipschitz maps $Y \to X$ sending y to a fixed bounded subset of X is coarsely sequentially compact: every sequence contains a coarsely converging subsequence.

In the next lemma we assume that Y is a geodesic δ-hyperbolic space and X is a symmetric space of noncompact type. The lemma itself is an immediate consequence of the perturbation lemma, Lemma 3.2.

Lemma 3.12 *Suppose that $p_n\colon \mathbb{R}_+ \to X$ is a sequence of M-Morse rays which coarsely converges to a map $p\colon \mathbb{R}_+ \to X$. Then p is M'-Morse, where $M' = M + C$ and the constant C is the one appearing in the definition of coarse convergence.*

In particular, a coarse limit of a sequence of (uniformly) Morse quasigeodesics is again Morse.

For the next lemma, we equip the flag-manifold $\mathrm{F} = \mathrm{Flag}(\tau_{\mathrm{mod}})$ with some background metric d_{F}.

Lemma 3.13 *Suppose that $p_n\colon \mathbb{R}_+ \to X$ is a sequence of M-Morse rays coarsely converging to a M-Morse ray $p\colon \mathbb{R}_+ \to X$. Then the sequence $\tau_n := p_n(\infty)$ of ends of the quasirays p_n converges to $\tau = p(\infty)$. Moreover, the latter convergence is uniform in the following sense. For every $\epsilon > 0$ there exists n_0 depending only on M and C and $N(R, C)$ (appearing in Definition 3.11) such that for all $n \geq n_0$, $d_{\mathrm{F}}(\tau_n, \tau) \leq \epsilon$.*

Proof Suppose that the claim is false. Then in view of coarse compactness of the space of M-Morse maps sending y to a fixed compact subset of X, there exists a sequence (p_n) as in the lemma, coarsely converging to p, such that the sequence $p_n(\infty) = \tau_n$ converges to $\tau' \neq p(\infty) = \tau$. By the coarse convergence $p_n \to p$, there exists $C < \infty$ and a sequence $t_n \to \infty$ such that $d(p_n(t_n), p(t_n)) \leq C$. By the definition of Morse quasigeodesics, there exists a sequence of cones $V(x_n, \mathrm{st}(\tau_n))$ (with x_n in a bounded subset $B \subset X$) such that the image of p_n is contained in the D-neighborhood of $V(x_n, \mathrm{st}(\tau_n))$. Thus, the sequence $(p_n(t_n))$ flag-converges to τ', while $(p(t_n))$ flag-converges to τ. According to [Kapovich et al. 2017, Lemma 4.23], altering a sequence by a uniformly bounded amount does not change the flag-limit. Therefore, the sequence $(p(t_n))$ also flag-converges to τ'. Hence, $\tau = \tau'$. A contradiction. □

3.4 A Morse lemma for straight sequences

In order to motivate the results of this section we recall the following *sufficient condition* for a piecewise geodesic path in a Hadamard manifold Y of curvature ≤ -1 to be quasigeodesic; see eg [Kapovich and Liu 2019].

Proposition 3.14 *Suppose that c is a piecewise geodesic path in Y whose angles at the vertices are $\geq \alpha > 0$ and whose edges are longer than L, where α and L satisfy*

$$\cosh(L/2)\sin(\alpha/2) \geq \nu > 1. \tag{3.15}$$

Then c is an $(L(\nu), A(\nu))$-quasigeodesic.

By considering c with vertices on a horocycle in the hyperbolic plane, one sees that the inequality in this proposition is sharp.

Corollary 3.16 *If L is sufficiently large and α is sufficiently close to π, then c is (uniformly) quasigeodesic.*

In higher rank, we do not have an analogue of the inequality (3.15); instead, we will be generalizing the corollary. However, *angles* in the corollary will be replaced with *ζ-angles*. We will show (in a "string of diamonds" theorem, Theorem 3.30) that if a piecewise geodesic path c in X has sufficiently long edges, and ζ-angles between consecutive segments sufficiently close to π, then c is M-Morse for a suitable Morse datum.

In the following, we consider finite or infinite sequences (x_n) of points in X. For the next definition, we remind the reader of the definition of ζ-angles given in (2.5).

Definition 3.17 (straight and spaced sequence) We call a sequence (x_n) (Θ, ϵ)-*straight* if the segments $x_n x_{n+1}$ are Θ-regular and

$$\angle^{\zeta}_{x_n}(x_{n-1}, x_{n+1}) \geq \pi - \epsilon$$

for all n. We call it l-*spaced* if the segments $x_n x_{n+1}$ have length $\geq l$.

Note that every straight sequence can be extended to a biinfinite straight sequence.

Straightness is a local condition. The goal of this section is to prove the following local-to-global result asserting that sufficiently straight and spaced sequences satisfy a higher-rank version of the Morse lemma (for quasigeodesics in hyperbolic space).

Theorem 3.18 (Morse lemma for straight spaced sequences) *For Θ, Θ', δ there exist l, ϵ such that every (Θ, ϵ)-straight l-spaced sequence (x_n) is δ-close to a parallel set $P(\tau_-, \tau_+)$ with simplices $\tau_\pm$ of type τ_{mod}, and it moves from τ_- to τ_+ in the sense that its nearest-point projection $\bar{x}_n$ to $P(\tau_-, \tau_+)$ satisfies*

$$\bar{x}_{n\pm m} \in V(\bar{x}_n, \mathrm{st}_{\Theta'}(\tau_\pm)) \quad \text{for all } n \text{ and } m \geq 1. \tag{3.19}$$

Remark 3.20 (global spacing) (1) As a corollary of this theorem, we will show that straight spaced sequences are quasigeodesic:

$$d(x_n, x_{n+m}) \geq clm - 2\delta$$

with a constant $c = c(\Theta') > 0$. See Corollary 3.29. In particular, by interpolating the sequence (x_n) via geodesic segments we obtain a Morse quasigeodesic in X.

(2) Theorem 3.18 is a higher-rank generalization of two familiar facts from geometry of Gromov-hyperbolic geodesic metric spaces: The fact that local quasigeodesics (with suitable parameters) are global quasigeodesics and the Morse lemma stating that quasigeodesics stay uniformly close to geodesics. In the higher rank, quasigeodesics, of course, need not be close to geodesics, but, instead (under the straightness assumption), are close to diamonds/Weyl cones/parallel sets. A different version of the same phenomenon is established in our paper [Kapovich et al. 2018b, Theorem 1.3], where closeness of certain quasigeodesics to diamonds/Weyl cones/parallel sets is proven under the *uniform* τ_{mod}*-regularity* (rather than the straightness) assumption. As far as we know, neither result directly implies the other, ie neither uniform regularity directly implies straightness, nor vice-versa.

(3) One can obviously strengthen the Corollary 3.16 by stating that for each $\epsilon < \pi$ there exists $L_0(\epsilon)$ such that if $\alpha \geq \pi - \epsilon$ and $L \geq L_0(\epsilon)$, then c is a uniform quasigeodesic in X. A similar strengthening is false for symmetric spaces of rank ≥ 2. For instance, when $W \cong S_3$ and $\epsilon = 2\pi/3$, then no matter what Θ, Θ' and l are, the conclusion of Theorem 3.18 fails already for sequences contained in a single flat.

In order to prove the theorem, we start by considering half-infinite sequences and prove that they keep moving away from an ideal simplex of type τ_{mod} if they do so initially.

For the next definition we recall the definition of ζ-angles given in Section 2.3. Given a face $\tau \subset \partial_{\mathrm{Tits}} X$ of type τ_{mod} and distinct points $x, y \in X$, we have the angle

$$\angle_x^{\zeta}(\tau, y) := \angle_x^{\zeta}(z, y) = \angle_x(z, \zeta(xy)),$$

where z is a point (distinct from x) on the geodesic ray $x\xi$, where $\xi \in \tau$ is the point of type ζ. (See (2.5).)

Definition 3.21 (moving away from an ideal simplex) We say that a sequence (x_n) *moves* ϵ*-away* from a simplex τ of type τ_{mod} if

$$\angle_{x_n}^{\zeta}(\tau, x_{n+1}) \geq \pi - \epsilon \quad \text{for all } n.$$

Lemma 3.22 (moving away from ideal simplices) *For small* ϵ *and large* l, $\epsilon \leq \epsilon_0$ *and* $l \geq l(\epsilon, \Theta)$, *the following holds*:

If the sequence $(x_n)_{n\geq 0}$ *is* (Θ, ϵ)*-straight* l*-spaced and if*

$$\angle_{x_0}^{\zeta}(\tau, x_1) \geq \pi - 2\epsilon,$$

then (x_n) *moves* ϵ*-away from* τ.

Proof By Lemma 2.11(ii), the unit-speed geodesic segment $c\colon [0, t_1] \to X$ from $p(0)$ to $p(1)$ moves $\epsilon(d(2\epsilon))$-away from τ at all times, and $\epsilon'(2\epsilon, \Theta, l)$-away at times $\geq l$, which includes the final time t_1. For $l(\epsilon, \Theta)$ sufficiently large, we have $\epsilon'(2\epsilon, \Theta, l) \leq \epsilon$. Then c moves ϵ-away from τ at time t_1, which means that $\angle_{x_1}^{\zeta}(\tau, x_0) \leq \epsilon$. Straightness at x_1 and the triangle inequality yield that again $\angle_{x_1}^{\zeta}(\tau, x_2) \geq \pi - 2\epsilon$. One proceeds by induction. □

Note that there do exist simplices τ satisfying the hypothesis of the previous lemma. For instance, one can extend the initial segment x_0x_1 backwards to infinity and choose $\tau = \tau(x_1x_0)$.

Now we look at *biinfinite* sequences.

We assume in the following that $(x_n)_{n\in\mathbb{Z}}$ is (Θ,ϵ)-straight l-spaced for small ϵ and large l. As a first step, we study the asymptotics of such sequences and use the argument for Lemma 3.22 to find a pair of opposite ideal simplices $\tau_\pm$ such that (x_n) moves from τ_- towards τ_+.

Lemma 3.23 (moving towards ideal simplices) *For small ϵ and large l, $\epsilon \le \epsilon_0$ and $l \ge l(\epsilon,\Theta)$, the following holds*:

There exists a pair of opposite simplices $\tau_\pm$ of type τ_{mod} such that the inequality

$$\angle^{\xi}_{x_n}(\tau_\mp, x_{n\pm1}) \ge \pi - 2\epsilon \tag{3.24}$$

holds for all n.

Proof 1. For every n define a compact set $C_n^\mp \subset \mathrm{Flag}(\tau_{\mathrm{mod}})$

$$C_n^\pm = \{\tau_\pm : \angle^{\xi}_{x_n}(\tau_\pm, x_{n\mp1}) \ge \pi - 2\epsilon\}.$$

As in the proof of Lemma 3.22, straightness at x_{n+1} implies that $C_n^- \subset C_{n+1}^-$. Hence the family $\{C_n^-\}_{n\in\mathbb{Z}}$ form a nested sequence of nonempty compact subsets and therefore have nonempty intersection containing a simplex τ_-. Analogously, there exists a simplex τ_+ which belongs to C_n^+ for all n.

2. It remains to show that the simplices τ_-, τ_+ are antipodal. Using straightness and the triangle inequality, we see that

$$\angle^{\xi}_{x_n}(\tau_-, \tau_+) \ge \pi - 5\epsilon \tag{$\star$}$$

for all n. Hence, if $5\epsilon < \epsilon(\zeta)$, then the simplices τ_-, τ_+ are antipodal in view of Remark 2.4. □

The pair of opposite simplices (τ_-, τ_+) which we found determines a parallel set in X. The second step is to show that (x_n) is uniformly close to it.

Lemma 3.25 (close to parallel set) *For small ϵ and large l, $\epsilon \le \epsilon(\delta)$ and $l \ge l(\Theta,\delta)$, the sequence (x_n) is δ-close to $P(\tau_-,\tau_+)$.*

Proof The statement follows from the combination of the inequality ($\star$) in the second part of the proof of Lemma 3.23, and Lemma 2.10. □

The third and final step of the proof is to show that the nearest point projection $(\bar{x}_n)$ of (x_n) to $P(\tau_-,\tau_+)$ moves from τ_- towards τ_+.

Lemma 3.26 (projection moves towards ideal simplices) *For small ϵ and large l, with $\epsilon \le \epsilon_0$ and $l \ge l(\epsilon,\Theta,\Theta')$, the segments $\bar{x}_n\bar{x}_{n+1}$ are Θ'-regular and*

$$\angle^{\xi}_{\bar{x}_n}(\tau_-, \bar{x}_{n+1}) = \pi \quad \text{for all } n.$$

Proof By the previous lemma, (x_n) is δ_0-close to $P(\tau_-, \tau_+)$ if ϵ_0 is sufficiently small and l is sufficiently large. Since $x_n x_{n+1}$ is Θ-regular, the triangle inequality for Δ-lengths yields that the segment $\overline{x}_n \overline{x}_{n+1}$ is Θ'-regular, again if l is sufficiently large.

Let ξ_+ denote the ideal endpoint of the ray extending this segment, ie $\overline{x}_{n+1} \in \overline{x}_n \xi_+$. Then x_{n+1} is $2\delta_0$-close to the ray $x_n \xi_+$. We obtain that (in view of the uniform continuity of ζ-angles observed in Section 2.3)

$$\angle^{\zeta}_{\mathrm{Tits}}(\tau_-, \xi_+) \geq \angle^{\zeta}_{x_n}(\tau_-, \xi_+) \simeq \angle^{\zeta}_{x_n}(\tau_-, x_{n+1}) \simeq \pi,$$

where the last step follows from inequality (3.24). The discreteness of Tits distances between ideal points of fixed type ζ implies that in fact

$$\angle^{\zeta}_{\mathrm{Tits}}(\tau_-, \xi_+) = \pi,$$

ie the ideal points $\zeta(\tau_-)$ and $\zeta(\xi_+)$ are antipodal. But the only simplex opposite to τ_- in $\partial_\infty P(\tau_-, \tau_+)$ is τ_+, so $\tau(\xi_+) = \tau_+$ and

$$\angle^{\zeta}_{\overline{x}_n}(\tau_-, \overline{x}_{n+1}) = \angle^{\zeta}_{\overline{x}_n}(\tau_-, \xi_+) = \pi,$$

as claimed. □

Proof of Theorem 3.18 It suffices to consider biinfinite sequences.

The conclusion of Lemma 3.26 is equivalent to $\overline{x}_{n+1} \in V(\overline{x}_n, \mathrm{st}_{\Theta'}(\tau_+))$. Combining Lemmas 3.25 and 3.26, we thus obtain the theorem for $m = 1$.

The convexity of Θ'-cones, cf Proposition 2.1, implies that

$$V(\overline{x}_{n+1}, \mathrm{st}_{\Theta'}(\tau_+)) \subset V(\overline{x}_n, \mathrm{st}_{\Theta'}(\tau_+)),$$

and the assertion follows for all $m \geq 1$ by induction. □

Remark 3.27 The conclusion of the theorem implies flag-convergence $x_{\pm n} \to \tau_\pm$ as $n \to +\infty$. However, the sequences $(x_n)_{n \in \pm\mathbb{N}}$ do not in general converge at infinity, but accumulate at compact subsets of $\mathrm{st}_{\Theta'}(\tau_\pm)$.

3.5 Lipschitz retractions to straight paths

Consider a (possibly infinite) closed interval J in $\mathbb{R}$; we will assume that J has integer or infinite bounds. Suppose that $p\colon J \cap \mathbb{Z} \to P = P(\tau_-, \tau_+) \subset X$ is an l-separated, λ-Lipschitz, $(\Theta, 0)$-straight coarse sequence pointing away from τ_- and towards τ_+. We extend p to a piecewise geodesic map $p\colon J \to P$ by sending intervals $[n, n+1]$ to geodesic segments $p(n)\,p(n+1)$ via affine maps. We retain the name p for the extension.

Lemma 3.28 *There exists $L = L(l, \lambda, \Theta)$ and an L-Lipschitz retraction of X to p, ie an L-Lipschitz map $r\colon X \to J$ such that $r \circ p = \mathrm{Id}$. In particular, $p\colon J \cap \mathbb{Z} \to X$ is an $(\overline{L}, \overline{A})$-quasigeodesic, where $\overline{L}$ and $\overline{A}$ depend only on l, λ and Θ.*

Proof It suffices to prove existence of a retraction. Since P is convex in X, it suffices to construct a map $P \to J$. Pick a generic point $\xi = \xi_+ \in \tau_+$ and let $b_\xi \colon P \to \mathbb{R}$ denote the Busemann function normalized so that $b_\xi(p(z)) = 0$ for some $z \in J \cap \mathbb{Z}$. Then the Θ-regularity assumption on p implies that the slope of the piecewise linear function $b_\xi \circ p \colon J \to \mathbb{R}$ is strictly positive, bounded away from 0. The assumption that p is l-separated λ-Lipschitz implies that

$$l \le |p'(t)| \le \lambda$$

for each t (where the derivative exists). The straightness assumption on p implies that the function $h := b_\xi \circ p \colon J \to \mathbb{R}$ is strictly increasing. By combining these observations, we conclude that h is an L-bi-Lipschitz homeomorphism to $h(J)$ for some $L = L(l, \lambda, \Theta)$. Let $\rho \colon \mathbb{R} \to J$ denote the nearest-point projection to the interval J; it is a 1-Lipschitz map. Lastly, we define

$$r \colon P \to J \quad \text{by} \quad r = h^{-1} \circ \rho \circ b_\xi.$$

Since b_ξ is 1-Lipschitz, the map r is L-Lipschitz. By the construction, $r \circ p = \mathrm{Id}$. □

Corollary 3.29 *Fix $\delta > 0$ and Θ, Θ' such that $\Theta \subset \mathrm{int}(\Theta')$. Let $l = l(l, \Theta, \Theta', \delta)$ and $\epsilon = \epsilon(l, \Theta, \Theta', \delta)$ be the constants as in Theorem 3.18. Suppose that $p \colon J \cap \mathbb{Z} \to X$ is an l-spaced, λ-Lipschitz, (Θ, ϵ)-straight sequence. Then for $L = L(l - 2\delta, \lambda + 2\delta, \Theta')$ we have:*

(1) *There exists an $(L, 2\delta)$-coarse Lipschitz retraction $X \to J$.*

(2) *The map p is a (Θ', D', L', A')-quasigeodesic with D', L', A' depending only on $l, \lambda, \Theta, \Theta', \epsilon$.*

Proof The statement immediately follows the above lemma combined with Theorem 3.18. □

Reformulating in terms of piecewise geodesic paths, we obtain:

Theorem 3.30 (string of diamonds theorem) *For any pair of Weyl convex subsets $\Theta < \Theta'$ and a number $D \ge 0$, there exist positive numbers ϵ, S, L, A depending on the datum (Θ, Θ', D) such that the following holds.*

Suppose that c is an arc-length parametrized piecewise geodesic path (finite or infinite) in X obtained by concatenating geodesic segments $x_i x_{i+1}$ such that, for all i,

(1) *each segment $x_i x_{i+1}$ is Θ-regular and has length $\ge S$, and*

(2) *$\angle^\xi_{x_i}(x_{i-1}, x_{i+1}) \ge \pi - \epsilon$.*

Then the path c is (Θ', D, L, A)-Morse.

3.6 Local Morse quasigeodesics

According to Theorem 3.30, sufficiently straight and spaced straight piecewise geodesic paths are Morse. In this section we will now prove that, conversely, the Morse property implies straightness in a suitable sense, namely that for sufficiently spaced quadruples the associated midpoint triples are arbitrarily straight. (For the quadruples themselves this is in general not true.)

Definition 3.31 (quadruple condition) For points $x, y \in X$, we let $\mathrm{mid}(x, y)$ denote the midpoint of the geodesic segment xy. A map $p\colon I \to X$ satisfies the (Θ, ϵ, l, s)-*quadruple condition* if for all $t_1, t_2, t_3, t_4 \in I$ with $t_2 - t_1, t_3 - t_2, t_4 - t_3 \geq s$, the triple of midpoints

$$(\mathrm{mid}(t_1, t_2), \mathrm{mid}(t_2, t_3), \mathrm{mid}(t_3, t_4))$$

is (Θ, ϵ)-straight and l-spaced.

Proposition 3.32 (Morse implies quadruple condition) *For $L, A, \Theta, \Theta', D, \epsilon, l$ there exists a scale $s = s(L, A, \Theta, \Theta', D, \epsilon, l)$ such that every (Θ, D, L, A)-Morse quasigeodesic satisfies the $(\Theta', \epsilon, l, s')$-quadruple condition for every $s' \geq s$.*

Proof Let $p\colon I \to X$ be an (L, A, Θ, D)-Morse quasigeodesic, and let $t_1, \dots, t_4 \in I$ be such that $t_2 - t_1, t_3 - t_2, t_4 - t_3 \geq s$. We abbreviate $p_i := p(t_i)$ and $m_i = \mathrm{mid}(p_i, p_{i+1})$.

Regarding straightness, it suffices to show that the segment $m_2 m_1$ is Θ'-regular and that $\angle^{\xi}_{m_2}(p_2, m_1) \leq \frac{1}{2}\epsilon$ provided s is sufficiently large in terms of the given data.

By the Morse property, there exists a diamond $\diamondsuit_\Theta(x_1, x_3)$ such that

$$d(x_1, p_1), d(x_3, p_3) \leq D \quad \text{and} \quad p_2 \in N_D(\diamondsuit_\Theta(x_1, x_3)).$$

The diamond spans a unique parallel set $P(\tau_-, \tau_+)$. (Necessarily, we have $x_3 \in V(x_1, \mathrm{st}_\Theta(\tau_+))$ and $x_1 \in V(x_3, \mathrm{st}_\Theta(\tau_-))$.)

We denote by $\bar{p}_i$ and $\bar{m}_i$ the projections of p_i and m_i to the parallel set.

We first observe that m_2 (and m_3) is arbitrarily close to the parallel set if s is large enough. If this were not true, a limiting argument would produce a geodesic line at strictly positive finite Hausdorff distance $\in (0, D]$ from $P(\tau_-, \tau_+)$ and asymptotic to ideal points in $\mathrm{st}_\Theta(\tau_\pm)$. However, all lines asymptotic to ideal points in $\mathrm{st}_\Theta(\tau_\pm)$ are contained in $P(\tau_-, \tau_+)$.

Next, we look at the directions of the segments $\bar{m}_2\bar{m}_1$ and $\bar{m}_2\bar{p}_2$ and show that they have the same τ-direction. Since $\bar{p}_2$ is $2D$-close to $V(\bar{p}_1, \mathrm{st}_\Theta(\tau_+))$, the point $\bar{p}_1$ is $2D$-close to $V(\bar{p}_2, \mathrm{st}_\Theta(\tau_-))$, and hence also $\bar{m}_1$ is $2D$-close to $V(\bar{p}_2, \mathrm{st}_\Theta(\tau_-))$. Therefore, $\bar{p}_1, \bar{m}_1 \in V(\bar{p}_2, \mathrm{st}_{\Theta'}(\tau_-))$ if s is large enough. Similarly, $\bar{m}_2 \in V(\bar{p}_2, \mathrm{st}_{\Theta'}(\tau_+))$ and hence $\bar{p}_2 \in V(\bar{m}_2, \mathrm{st}_{\Theta'}(\tau_-))$. The convexity of Θ'-cones, see Proposition 2.1, implies that also $\bar{m}_1 \in V(\bar{m}_2, \mathrm{st}_{\Theta'}(\tau_-))$. In particular, $\angle^{\xi}_{\bar{m}_2}(\bar{p}_2, \bar{m}_1) = 0$ if s is sufficiently large.

Since m_2 is arbitrarily close to the parallel set if s is sufficiently large, it follows by another limiting argument that $\angle^{\xi}_{m_2}(p_2, m_1) \leq \epsilon/2$ if s is sufficiently large.

Regarding the spacing, we use that $\bar{m}_1 \in V(\bar{p}_2, \mathrm{st}_{\Theta'}(\tau_-))$ and $\bar{m}_2 \in V(\bar{p}_2, \mathrm{st}_{\Theta'}(\tau_+))$. It follows that

$$d(\bar{m}_1, \bar{m}_2) \geq c \cdot (d(\bar{m}_1, \bar{p}_2) + d(\bar{p}_2, \bar{m}_2))$$

with a constant $c = c(\Theta') > 0$, and hence that $d(m_1, m_2) \geq l$ if s is sufficiently large. □

Theorem 3.18 and Proposition 3.32 say that the Morse property for quasigeodesics is equivalent to straightness (of associated spaced sequences of points). Since straightness is a local condition, this leads to a local-to-global result for Morse quasigeodesics, namely that the Morse property holds globally if it holds locally up to a sufficiently large scale.

Definition 3.33 (local Morse quasigeodesic) An S-local (Θ, D, L, A)-*Morse quasigeodesic* in X is a map $p\colon I \to X$ such that for all t_0, the subpath $p|_{[t_0,t_0+S]}$ is a (Θ, D, L, A)-Morse quasigeodesic.

Note that local Morse quasigeodesics are uniformly coarse Lipschitz.

Theorem 3.34 (local-to-global principle for Morse quasigeodesics) *For L, A, Θ, Θ', D there exist S, L', A', D' such that every S-local (Θ, D, L, A)-local Morse quasigeodesic in X is an (Θ', D', L', A')-Morse quasigeodesic.*

Proof We choose an auxiliary Weyl convex subset Θ'' such that $\Theta < \Theta'' < \Theta'$.

Let $p\colon I \to X$ be an S-local (Θ, D, L, A)-local Morse quasigeodesic. We consider its coarsification on a (large) scale s and the associated midpoint sequence, ie we put $p_n^s = p(ns)$ and $m_n^s = \mathrm{mid}(p_n^s, p_{n+1}^s)$.

Whereas the coarsification itself does not in general become arbitrarily straight as the scale s increases, this is true for its midpoint sequence due to Proposition 3.32. We want it to be sufficiently straight and spaced so that we can apply to it the Morse lemma from Theorem 3.18. Therefore we first fix an auxiliary constant δ, and further auxiliary constants l, ϵ as determined by Theorem 3.18 in terms of Θ', Θ'' and δ. Then Proposition 3.32 applied to the (Θ, D, L, A)-Morse quasigeodesics $p|_{[t_0,t_0+S]}$ yields that (m_n^s) is (Θ'', ϵ)-straight and l-spaced if $S \geq 3s$ and the scale s is large enough depending on $L, A, \Theta, \Theta'', D, \epsilon, l$.

Now we can apply Theorem 3.18 to (m_n^s). It yields a nearby sequence $(\bar{m}_n^s)$ with $d(\bar{m}_n^s, m_n^s) \leq \delta$, which has the following property: For all $n_1 < n_2 < n_3$, the segments $\bar{m}_{n_1}^s \bar{m}_{n_3}^s$ are uniformly regular and the points $m_{n_2}^s$ are δ-close to the diamonds $\diamondsuit_{\Theta'}(\bar{m}_{n_1}^s, \bar{m}_{n_3}^s)$.

Since the subpaths $p|_{[ns,(n+1)s]}$ filling in (p_n^s) are (L, A)-quasigeodesics (because $S \geq s$), and it follows that for all $t_1, t_2 \in I$ the subpaths $p|_{[t_1,t_2]}$ are D'-close to Θ'-diamonds with D' depending on L, A, s.

The conclusion of Theorem 3.18 also implies a global spacing for the sequence (m_n^s), compare Remark 3.20, ie $d(m_n^s, m_{n'}^s) \geq c \cdot |n - n'|$ with a positive constant c depending on Θ', l. Hence p is a global (L', A')-quasigeodesic with L', A' depending on L, A, s, c.

Combining this information, we obtain that p is an (Θ', D', L', A')-Morse quasigeodesic for certain constants L', A' and D' depending on L, A, Θ, Θ' and D, provided that the scale S is sufficiently large in terms of the same data. □

3.7 Local-to-global principle for Morse maps

We now deduce from our local-to-global result for Morse quasigeodesics, Theorem 3.34, a local-to-global result for Morse embeddings.

We restrict to the setting of maps of Gromov-hyperbolic (l,a)-quasigeodesic metric spaces Z to symmetric spaces X.

Definition 3.35 (local Morse embedding) We call a map $f\colon Z\to X$ an S-local (Θ, D, L, A)-Morse embedding if for any (l,a)-quasigeodesic $q\colon I\to Z$ defined on an interval I of length $\le S$, the image path $f\circ q$ is a (Θ, D, L, A)-Morse quasigeodesic in X.

Theorem 3.36 (local-to-global principle for Morse embeddings of Gromov hyperbolic spaces) *For $l, a, L, A, \Theta, \Theta', D$ there exists a scale S and a datum (D', L', A') such that every S-local (Θ, D, L, A)-Morse embedding from an (l,a)-quasigeodesic Gromov hyperbolic space into X is a (Θ', D', L', A')-Morse embedding.*

Proof Let $f\colon Z\to X$ denote the local Morse embedding. It sends every (l,a)-quasigeodesic $q\colon I\to Z$ to an S-local (Θ, D, L, A)-Morse quasigeodesic $p=f\circ q$ in X. By Theorem 3.34, p is (L', A', Θ', D')-Morse if $S\ge S(l,a,L,A,\Theta,\Theta',D)$, where L', A', D' depend on the given data. □

Below is a reformulation of this theorem in the case of geodesic Gromov-hyperbolic spaces.

Let Z be a δ-hyperbolic geodesic space. An R-ball $B(z,R)$ in Z need not be convex, but it is δ-quasiconvex. In particular, the restriction of the metric from Z to $B(z,R)$ results in a $(1,\delta)$-quasigeodesic metric space.

Theorem 3.37 (local-to-global principle for Morse embeddings of geodesic spaces) *For L, A, Θ, Θ', D and δ there exists a scale R and a datum (D', L', A') such that if Z is a δ-hyperbolic geodesic metric space and the restriction of f to any R-ball is $(\Theta, D, L, A, 1, \delta)$-Morse, then $f\colon Z\to X$ is (Θ', D', L', A')-Morse.*

4 Group-theoretic applications

As a consequence of the local-to-global criterion for Morse maps, in this section we establish that the Morse property for isometric group actions is an open condition. Furthermore, for two nearby Morse actions, the actions on their τ_{mod}-limit sets are also close, ie conjugate by an equivariant homeomorphism close to identity. In view of the equivalence of Morse property with the asymptotic properties discussed earlier, this implies structural stability for asymptotically embedded groups. Another corollary of the local-to-global result is the algorithmic recognizability of Morse actions.

We conclude the section by illustrating our technique by constructing Morse–Schottky actions of free groups on higher-rank symmetric spaces.

4.1 Stability of Morse actions

We consider isometric actions $\Gamma \curvearrowright X$ of finitely generated groups.

Definition 4.1 (Morse action) We call an action $\Gamma \curvearrowright X$ Θ-*Morse* if one (any) orbit map $\Gamma \to \Gamma x \subset X$ is a Θ-Morse embedding with respect to a(ny) word metric on Γ. We call an action $\Gamma \curvearrowright X$ τ_{mod}-*Morse* if it is Θ-Morse for some τ_{mod}-Weyl convex compact subset $\Theta \subset \mathrm{ost}(\tau_{\mathrm{mod}})$.

Remark 4.2 (Morse actions are τ_{mod}-regular and undistorted) (i) It follows immediately from the definition of Morse quasigeodesics that Θ-Morse actions are τ_{mod}-regular for the simplex type τ_{mod} determined by Θ.

(ii) Morse subgroups of G are *undistorted* in the sense that the orbit maps are quasiisometric embeddings. In [Kapovich and Leeb 2018b] we prove that Morse subgroups of G satisfy a stronger property: they are *coarse Lipschitz retracts* of G. This retraction property is stronger than nondistortion: every finitely generated subgroup which is a coarse retract of G is undistorted in G, but there are examples of undistorted subgroups which are not coarse retracts. For instance, the group $\Phi := F_2 \times F_2$ admits an undistorted embedding in the isometry group of $X = \mathbb{H}^2 \times \mathbb{H}^2$. On the other hand, pick an epimorphism $\phi\colon F_2 \to \mathbb{Z}$ and define the subgroup $\Gamma < \Phi$ as the kernel of the homomorphism

$$(\gamma_1, \gamma_2) \mapsto \phi(\gamma_1) - \phi(\gamma_2).$$

Then Γ is a finitely generated undistorted subgroup of Φ (see eg [Ol'shanskii and Sapir 2001, Theorem 2]), but is not finitely presented (see eg [Baumslag and Roseblade 1984]). Hence, $\Gamma < G = \mathrm{Isom}(\mathbb{H}^2) \times \mathrm{Isom}(\mathbb{H}^2)$ is undistorted but is not a coarse Lipschitz retract.

We denote by $\mathrm{Hom}_{\tau_{\mathrm{mod}}}(\Gamma, G) \subset \mathrm{Hom}(\Gamma, G)$ the subset of τ_{mod}-Morse actions $\Gamma \curvearrowright X$.

By analogy with *local Morse quasigeodesics*, we define *local Morse group actions* $\rho\colon \Gamma \curvearrowright X$ of a hyperbolic group (with a fixed finite generating set):

Definition 4.3 An action ρ is called S-locally (Θ, D, L, A)-*locally Morse*, or (Θ, D, L, A)-*locally Morse on the scale* S, with respect to a basepoint $x \in X$, if the orbit map $\Gamma \to \Gamma \cdot x \subset X$ induces an S-local (Θ, D, L, A)-local Morse embedding of the Cayley graph of Γ.

According to our local-to-global result for Morse embeddings (see Theorem 3.37), an action of a word hyperbolic group is Morse if and only if it is local Morse on a sufficiently large scale. Since this is a finite condition, it follows that the Morse property is stable under perturbation of the action:

Theorem 4.4 (Morse is open for word hyperbolic groups) *For any word hyperbolic group* Γ *the subset* $\mathrm{Hom}_{\tau_{\mathrm{mod}}}(\Gamma, G)$ *is open in* $\mathrm{Hom}(\Gamma, G)$. *More precisely, if* $\rho \in \mathrm{Hom}_{\tau_{\mathrm{mod}}}(\Gamma, G)$ *is* M*-Morse with respect to a basepoint* $x \in X$ *then there exists a neighborhood of* ρ *in* $\mathrm{Hom}(\Gamma, G)$ *consisting entirely of* M'*-Morse representations with respect to* x, *where* M' *depends only on* M.

Proof Let $\rho\colon \Gamma \curvearrowright X$ be a Morse action. We fix a word metric on Γ and a basepoint $x \in X$. Then there exist data $M = (L, A, \Theta, D)$ such that the orbit map $\Gamma \to \Gamma x \subset X$ extends to a (Θ, D, L, A)-Morse map of the Cayley graph Y on Γ.

We relax the Morse parameters slightly, namely we consider (L, A, Θ, D)-Morse quasigeodesics as $(L, A+1, \Theta, D+1)$-Morse quasigeodesics satisfying strict inequalities. For every scale S, the orbit map $\Gamma \to \Gamma x \subset X$ defines an $(L, A+1, \Theta, D+1, S)$-*local* Morse embedding $Y \to X$. Due to Γ-equivariance, this is a finite condition in the sense that it is equivalent to a condition involving only finitely many orbit points. Since we relaxed the Morse parameters, the same condition is satisfied by all actions sufficiently close to ρ.

Theorem 3.37 provides a scale S such that all S-local $(\Theta, D+1, L, A+1)$-Morse embeddings $Y \to X$ are M'-Morse for some Morse datum M' depending only on $(L, A+1, \Theta, D+1, S)$. It follows that actions sufficiently close to ρ are τ_{mod}-Morse. □

Lemma 4.5 *Let Φ be a finite group and G a semisimple Lie group with finitely many connected components and finite center. Then* $\mathrm{Hom}(\Phi, G)/G$ *is finite, where $g \in G$ act on representations $\rho\colon \Phi \to G$ by postcompositions with the inner automorphism* Inn_g *of G defined by g,*

$$\rho \mapsto \mathrm{Inn}_g \circ \rho.$$

Proof As before, we fix a maximal compact subgroup $K < G$. Since every homomorphism $\rho\colon \Phi \to G$ has compact image, its image is conjugate to a subgroup of K. Thus, it suffices to prove finiteness of $\mathrm{Hom}(\Phi, K)/K$. The group K has the structure of a real-algebraic group; hence, $\mathrm{Hom}(\Phi, K)$ also has a natural structure of a compact (with respect to the Lie group topology of K) algebraic subset of $K \times K \times \cdots \times K$ (n times, where n is the number of generators of Φ). For $\rho \in \mathrm{Hom}(\Phi, K)$ the Zariski tangent space $T_\rho(\mathrm{Hom}(\Phi, K)/K)$ is isomorphic to the group of 1-cocycles

$$Z^1(\Phi, \mathfrak{k}_{\mathrm{Ad}\circ\rho}),$$

where we regard the Lie algebra $\mathfrak{k}$ as a $\mathbb{R}\Phi$-module via the composition of ρ and the adjoint representation of K. For every finite group Φ and an $\mathbb{R}\Phi$-module M, we have $H^i(\Phi, M) = 0$ for all $i > 0$; see eg [Brown 1982, Corollary 10.2]. Applying this to $M = \mathfrak{k}_{\mathrm{Ad}\circ\rho}$, we see that every point in $\mathrm{Hom}(\Phi, K)/K$ is isolated; see eg [Raghunathan 1972, Theorem 6.7]. Since $\mathrm{Hom}(\Phi, K)/K$ is compact, it follows that it is finite. □

Corollary 4.6 *For every hyperbolic group Γ the space of* ***faithful*** *Morse representations*

$$\mathrm{Hom}_{inj,\tau_{\mathrm{mod}}}(\Gamma, G)$$

is open in $\mathrm{Hom}_{\tau_{\mathrm{mod}}}(\Gamma, G)$.

Proof Every hyperbolic group Γ has the unique maximal finite normal subgroup $\Phi \lhd \Gamma$ (if Γ is not elementary then Φ is the kernel of the action of Γ on $\partial_\infty \Gamma$). Since Morse actions are properly discontinuous, the kernel of every Morse representation $\Gamma \to G$ is contained in Φ. Since $\mathrm{Hom}(\Phi, G)/G$ is finite (see Lemma 4.5), it follows that the set of faithful Morse representations is open in $\mathrm{Hom}_{\tau_{\mathrm{mod}}}(\Gamma, G)$. □

The result on the openness of the Morse condition for actions of word hyperbolic groups (see Theorem 4.4) can be strengthened in the sense that the asymptotics of Morse actions vary continuously:

Theorem 4.7 (Morse actions are structurally stable) *The boundary map at infinity of a Morse action depends continuously on the action.*

Proof According to Theorem 4.4 nearby actions are uniformly Morse. The assertion therefore follows from the fact that the ends of Morse quasirays vary uniformly continuously; cf Lemma 3.13. □

Remark 4.8 (i) Since the boundary maps at infinity are embeddings, the Γ-actions on the τ_{mod}-limit sets are topologically conjugate to each other and, for nearby actions, by a homeomorphism close to the identity.

(ii) In rank one, our argument yields a different proof for Sullivan's structural stability theorem [1985] for convex cocompact group actions on rank-one symmetric spaces. Other proofs can be found in [Labourie 2006; Guichard and Wienhard 2012] (for Anosov subgroups in higher rank), [Corlette 1990; Izeki 2000; Bowditch 1998] for rank-one symmetric spaces.

Our next goal is to extend the topological conjugation from the limit set to the domains of proper discontinuity. Recall that in [Kapovich et al. 2018a] we constructed domains of proper discontinuity and cocompactness for τ_{mod}-Morse group actions on flag-manifolds $\mathrm{Flag}(\nu_{\mathrm{mod}}) = G/P_{\nu_{\mathrm{mod}}}$. Such domains depend on a certain auxiliary datum, a *balanced thickening* $\mathrm{Th} \subset W$, which is a $W_{\tau_{\mathrm{mod}}}$-left invariant subset satisfying certain conditions; see [Kapovich et al. 2018a, Section 3.4]. Let $\nu_{\mathrm{mod}} \subset \sigma_{\mathrm{mod}}$ be an ι-invariant face such that Th is invariant under the action of $W_{\nu_{\mathrm{mod}}}$ via the *right* multiplication (this is automatic if $\nu_{\mathrm{mod}} = \sigma_{\mathrm{mod}}$ since $W_{\sigma_{\mathrm{mod}}} = \{e\}$). The thickening $\mathrm{Th} \subset W$ defines a thickening $\mathrm{Th}(\Lambda_{\tau_{\mathrm{mod}}}(\Gamma)) \subset \mathrm{Flag}(\nu_{\mathrm{mod}})$. One of the main results of [Kapovich et al. 2018a] (Theorem 1.7) is that each τ_{mod}-Morse subgroup $\Gamma < G$ acts properly discontinuously and cocompactly on

$$\Omega_{\mathrm{Th}}(\Gamma) := \mathrm{Flag}(\nu_{\mathrm{mod}}) - \mathrm{Th}(\Lambda_{\tau_{\mathrm{mod}}}(\Gamma)).$$

Theorem 4.9 (stability of Morse quotient spaces) *Suppose that $\rho_n \colon \Gamma \to \rho_n(\Gamma) = \Gamma_n < G$ is a sequence of faithful τ_{mod}-Morse representations converging to a τ_{mod}-Morse embedding $\rho \colon \Gamma \hookrightarrow G$. Then*:

(1) *The sequence of thickenings $\mathrm{Th}(\Lambda_{\tau_{\mathrm{mod}}}(\Gamma_n))$ Hausdorff-converges to $\mathrm{Th}(\Lambda_{\tau_{\mathrm{mod}}}(\Gamma))$.*

(2) *If $\gamma_n \in \Gamma$ is a divergent sequence, then, after extraction, the sequence $(\rho_n(\gamma_n))$ flag-converges to a point in $\Lambda_{\tau_{\mathrm{mod}}}(\Gamma)$.*

(3) *For all sufficiently large n, there exists an equivariant diffeomorphism*

$$h_n \colon \Omega_{\mathrm{Th}}(\Gamma) \to \Omega_{\mathrm{Th}}(\Gamma_n) \subset \mathrm{Flag}(\tau_{\mathrm{mod}})$$

such that the sequence of maps (h_n) converges to the inclusion map $\Omega_{\mathrm{Th}}(\Gamma) \to \mathrm{Flag}(\tau_{\mathrm{mod}})$ uniformly on compacts.

(4) *In particular, the quotient orbifolds $\Omega_{\mathrm{Th}}(\Gamma_n)/\Gamma_n$ are diffeomorphic to $\Omega_{\mathrm{Th}}(\Gamma)/\Gamma$ for all sufficiently large n.*

Proof (1) First of all, suppose that a sequence $\tau_n \in \mathrm{Flag}(\tau_{\mathrm{mod}})$ converges to $\tau \in \mathrm{Flag}(\tau_{\mathrm{mod}})$. Then, since $\mathrm{Flag}(\nu_{\mathrm{mod}}) = G/P_{\nu_{\mathrm{mod}}}$, there is a sequence $g_n \in G$, $g_n \to e$, such that $g_n(\tau) = \tau_n$. Since

$$g_n(\mathrm{Th}(\tau)) = \mathrm{Th}(g_n\tau) = \mathrm{Th}(\tau_n),$$

it follows that we have Hausdorff-convergence of subsets $\mathrm{Th}(\tau_n) \to \mathrm{Th}(\tau)$. Moreover, this convergence of subsets is uniform: There exists $n_0 = n(\delta)$ such that if $d(\tau_n, \tau) < \delta$ for all $n \geq n_0$, then $d(\mathrm{Th}(\tau_n), \mathrm{Th}(\tau)) < \epsilon = \epsilon(\delta)$ for all $n \geq n_0$. Here $\epsilon \to 0$ as $\delta \to 0$. Since the sequence of limit sets $\Lambda_{\tau_{\mathrm{mod}}}(\Gamma_n)$ Hausdorff-converges to $\Lambda_{\tau_{\mathrm{mod}}}(\Gamma)$, it follows that the sequence of thickenings $\mathrm{Th}(\Lambda_{\tau_{\mathrm{mod}}}(\Gamma_n))$ Hausdorff-converges to $\mathrm{Th}(\Lambda_{\tau_{\mathrm{mod}}}(\Gamma))$. This proves (1).

(2) Consider a sequence of geodesic rays $e\xi_n$ in the Cayley graph Y of Γ such that γ_n lies in an R-neighborhood of $e\xi_n$ for all n. Then, in view of the uniform M'-Morse property for the representations ρ_n, each point $\rho_n(\gamma_n)(x)$ belongs to the D'-neighborhood of the Weyl cone $V(x, \mathrm{st}(\tau_n))$, where $\tau_n = \alpha_n(\xi_n)$, and $\alpha_n \colon \partial_\infty\Gamma \to \Lambda_{\tau_{\mathrm{mod}}}(\Gamma_n)$ is the asymptotic embedding. Thus, by the definition of flag-convergence, the sequences $(\rho_n(\gamma_n))$ and (τ_n) have the same flag-limit in $\mathrm{Flag}(\tau_{\mathrm{mod}})$. By part (1), the sequence (τ_n) subconverges to a point in $\Lambda_{\tau_{\mathrm{mod}}}(\Gamma)$. Hence, the same holds for $(\rho_n(\gamma_n))$.

(3) The proof of this part is mostly standard; see [Izeki 2000] in the case when X is a hyperbolic space. The quotient orbifold $O = \Omega_{\mathrm{Th}}(\Gamma)/\Gamma$ has a natural (F, G)-structure where $\mathrm{F} = \mathrm{Flag}(\nu_{\mathrm{mod}})$. The orbifold O has finitely many components, let Z be one of them and let $\hat{Z} \subset \Omega_{\mathrm{Th}}(\Gamma)$ be a component projecting to Z. It suffices to construct maps h_n on each component $\hat{Z}$ and then extend these maps to maps h_n of $\Omega_{\mathrm{Th}}(\Gamma)$ by ρ_n-equivariance.

The covering map $\hat{Z} \to Z$ induces an epimorphism $\phi \colon \pi_1(Z) \to \Gamma_Z$, where Γ_Z is the Γ-stabilizer of $\hat{Z}$. Let $\mathrm{dev} \colon \tilde{Z} \to \hat{Z} \subset \Omega_{\mathrm{Th}}(\Gamma)$ be the developing map, where $\tilde{Z} \to Z$ is the universal covering. By the Ehresmann–Thurston holonomy theorem (see [Lok 1984], [Canary et al. 1987], [Goldman 1988] or [Kapovich 2001, Section 7.1]), for all sufficiently large n, the homomorphism $\phi_n := \rho_n \circ \phi$ is the holonomy of an (F, G)-structure on Z. Moreover, the developing maps $\mathrm{dev}_n \colon \tilde{Z} \to \mathrm{F}$ converge to dev uniformly on compacts in the C^∞-topology. Since $\pi_1(\hat{Z})$ is contained in the kernel of ϕ, it is also in the kernel of ϕ_n. Hence, the maps dev_n descend to maps $\widehat{\mathrm{dev}}_n \colon \hat{Z} \to \mathrm{F}$. The sequence $\widehat{\mathrm{dev}}_n$ still converges to the identity embedding $\hat{Z} \hookrightarrow \mathrm{F}$ uniformly on compacts. Pick a compact fundamental set $C \subset \hat{Z}$ for the Γ_Z-action, ie a compact subset whose Γ-orbit equals $\hat{Z}$. In view of part (1) of the theorem, $\widehat{\mathrm{dev}}_n(C) \subset \Omega_{\mathrm{Th}}(\Gamma_n)$ for all sufficiently large n. Therefore, we can assume that $\widehat{\mathrm{dev}}_n(\hat{Z})$ is contained in a component $\hat{Z}_n$ of $\Omega_{\mathrm{Th}}(\Gamma_n)$. By the compactness of the quotient orbifolds, $\widehat{\mathrm{dev}}_n$ projects to a finite-to-one (smooth) orbi-covering map $c_n \colon Z \to Z_n := \hat{Z}_n/\rho_n(\Gamma_Z)$. Hence, $\widehat{\mathrm{dev}}_n \colon \hat{Z} \to \hat{Z}_n$ is a covering map as well. If $\hat{Z}_n$ were simply connected, it would follow that $\widehat{\mathrm{dev}}_n$ is a diffeomorphism as required (and this is how Izeki concludes his proof in [Izeki 2000]). We will prove that $\widehat{\mathrm{dev}}_n$ is a diffeomorphism by a direct argument.

Suppose that each $\widehat{\mathrm{dev}}_n$ is not injective. Then, by the equivariance of these maps, after extraction, there exist convergent sequences $z_n \to z$ and $z'_n \to z'$ in $\hat{Z}$ and a sequence $\gamma_n \in \Gamma$ such that

$$\rho_n(\gamma_n)\,\widehat{\mathrm{dev}}_n(z_n) = \widehat{\mathrm{dev}}_n(z'_n), \quad \gamma_n(z_n) \neq z'_n.$$

If the sequence (γ_n) were contained in a finite subset of Γ we would obtain a contradiction with the uniform convergence on compacts $\widehat{\mathrm{dev}}_n \to \mathrm{id}$ on $\hat{Z}$. Hence, after extraction, we may assume that (γ_n) is a divergent sequence. We therefore obtain a dynamical relation between the points z, z' via the sequence $(\rho_n(\gamma_n))$. According to part (2), the sequence $(\rho_n(\gamma_n))$ flag-accumulates to $\Lambda_{\tau_{\mathrm{mod}}}(\Gamma)$. The dynamical relation then contradicts fatness of the balanced thickening Th; see [Kapovich et al. 2018a, Section 5.2] and the proof of Theorem 6.8 in [Kapovich et al. 2018a].

We conclude that the maps

$$\widehat{\mathrm{dev}}_n \colon \hat{Z} \to \hat{Z}_n$$

are diffeomorphisms for all sufficiently large n. Since $\rho_n \colon \Gamma \to \Gamma_n$ are isomorphisms, equivariance of the developing maps implies that the maps $h_n \colon \Omega_{\mathrm{Th}}(\Gamma) \to \Omega_{\mathrm{Th}}(\Gamma_n)$ are diffeomorphisms for sufficiently large n.

(4) This part is an immediate corollary of part (3). □

Remark 4.10 (i) When X is a hyperbolic space, the equivariant diffeomorphism $h_n \colon \Omega(\Gamma) \to \Omega(\Gamma_n)$ combined with the equivariant homeomorphism of the limit sets $\Lambda(\Gamma) \to \Lambda(\Gamma_n)$ yield an equivariant homeomorphism $\partial_\infty X \to \partial_\infty X$; see [Tukia 1985; Izeki 2000]. Such an extension does not exist in higher rank since, in general, there is no equivariant homeomorphism of thickened limit sets $\mathrm{Th}(\Lambda_{\tau_{\mathrm{mod}}}(\Gamma)) \to \mathrm{Th}(\Lambda_{\tau_{\mathrm{mod}}}(\Gamma_n))$. This can be already seen for group actions on products of hyperbolic planes.

(ii) An analogue of Theorem 4.9 holds when we replace the group actions on flag-manifolds with actions on Finsler compactifications of the symmetric space and replace flag-manifold thickenings $\mathrm{Th}(\Lambda_{\tau_{\mathrm{mod}}})$ with Finsler thickenings $\mathrm{Th}_{F\ddot{u}}(\Lambda_{\tau_{\mathrm{mod}}}) \subset \partial_{F\ddot{u}} X$. Proving this requires extending Ehresmann–Thurston holonomy theorem to the category of smooth manifolds with corners and we will not pursue it here.

4.2 Schottky actions

In this section we apply our local-to-global result for straight sequences (Theorem 3.18) to construct Morse actions of free groups, generalizing and sharpening[1] Tits' ping-pong construction.

We consider two oriented τ_{mod}-regular geodesic lines a, b in X. Let $\tau_{\pm a}, \tau_{\pm b} \in \mathrm{Flag}(\tau_{\mathrm{mod}})$ denote the simplices which they are τ-asymptotic to, and let $\theta_{\pm a}, \theta_{\pm b} \in \sigma_{\mathrm{mod}}$ denote the types of their forward/backward ideal endpoints in $\partial_\infty X$. (Note that $\theta_{-a} = \iota(\theta_a)$ and $\theta_{-b} = \iota(\theta_b)$.) Let Θ be a compact convex subset of $\mathrm{ost}(\tau_{\mathrm{mod}}) \subset \sigma_{\mathrm{mod}}$, which is invariant under ι.

Definition 4.11 (generic pair of geodesics) We call the pair of geodesics (a, b) *generic* if the four simplices $\tau_{\pm a}, \tau_{\pm b}$ are pairwise opposite.

[1] In the sense that we obtain free subgroups which are not only embedded, but also asymptotically embedded in G.

Let $\alpha, \beta \in G$ be axial isometries with axes a and b respectively and translating in the positive direction along these geodesics. Then $\tau_{\pm a}$ and $\tau_{\pm b}$ are the attractive/repulsive fixed points of α and β on $\mathrm{Flag}(\tau_{\mathrm{mod}})$.

For every pair of numbers $m, n \in \mathbb{N}$ we consider the representation of the free group in two generators

$$\rho_{m,n}\colon F_2 = \langle A, B\rangle \to G$$

sending the generator A to α^m and B to β^n. We regard it as an isometric action $\rho_{m,n}\colon F_2 \curvearrowright X$.

Definition 4.12 (Schottky subgroup) A τ_{mod}-*Schottky subgroup* of G is a free τ_{mod}-asymptotically embedded subgroup of G.

If G has rank one, this definition amounts to the requirement that Γ is convex cocompact and free. Equivalently, this is a discrete finitely generated subgroup of G which contains no nontrivial elliptic and parabolic elements and has totally disconnected limit set; see [Kapovich 2008]. We note that this definition essentially agrees with the standard definition of Schottky groups in rank-one Lie groups, provided one allows fundamental domains at infinity for such groups to be bounded by pairwise disjoint compact submanifolds which need not be topological spheres; see [Kapovich 2008] for the detailed discussion.

Theorem 4.13 (Morse–Schottky actions) *If the pair of geodesics (a, b) is generic and if $\theta_{\pm a}, \theta_{\pm b} \in \mathrm{int}(\Theta)$, then the action $\rho_{m,n}$ is Θ-Morse for sufficiently large m, n. Thus, such $\rho_{m,n}$ is injective and its image is a τ_{mod}-Schottky subgroup of G.*

Remark 4.14 In particular, these actions are faithful and undistorted; compare Remark 4.2.

Proof Let $S = \{A^{\pm 1}, B^{\pm 1}\}$ be the standard generating set. We consider the sequences (γ_k) in F_2 with the property that $\gamma_k^{-1}\gamma_{k+1} \in S$ and $\gamma_{k+1} \neq \gamma_{k-1}$ for all k. They correspond to the geodesic segments in the Cayley tree of F_2 associated to S which connect vertices.

Let $x \in X$ be a basepoint. In view of Lemma 3.8 we must show that the corresponding sequences $(\gamma_k x)$ in the orbit $F_2 \cdot x$ are uniformly Θ-Morse. (Meaning for instance that the maps $\mathbb{R} \to X$ sending the intervals $[k, k+1)$ to the points $\gamma_k x$ are uniform Θ-Morse quasigeodesics.) As in the proof of Theorem 3.34 we will obtain this by applying our local-to-global result for straight spaced sequences (Theorem 3.18) to the associated midpoint sequences. Note that the sequences $(\gamma_k x)$ themselves cannot be expected to be straight.

Taking into account the Γ-action, the uniform straightness of all midpoint sequences depends on the geometry of a finite configuration in the orbit. It is a consequence of the following fact. Consider the midpoints $y_{\pm m}$ of the segments $x\alpha^{\pm m}(x)$ and $z_{\pm n}$ of the segments $x\beta^{\pm n}(x)$.

Lemma 4.15 *For sufficiently large m and n, the quadruple $\{y_{\pm m}, z_{\pm n}\}$ is arbitrarily separated and Θ-regular. Moreover, for any of the four points, the segments connecting it to the other three points have arbitrarily small ζ-angles with the segment connecting it to x.*

Proof The four points are arbitrarily separated from each other and from x because the axes a and b diverge from each other due to our genericity assumption.

By symmetry, it suffices to verify the rest of the assertion for the point y_m, ie we show that the segments $y_m y_{-m}$ and $y_m z_n$ are Θ-regular for large m, n and that $\lim_{m\to\infty} \angle^{\xi}_{y_m}(x, y_{-m}) = 0$ and $\lim_{n,m\to\infty} \angle^{\xi}_{y_m}(x, z_n) = 0$.

The orbit points $\alpha^{\pm m}x$ and the midpoints $y_{\pm m}$ are contained in a tubular neighborhood of the axis a. Therefore, the segments $y_m x$ and $y_m y_{-m}$ are Θ-regular for large m and $\angle_{y_m}(x, y_{-m}) \to 0$. This implies that also $\angle^{\xi}_{y_m}(x, y_{-m}) \to 0$.

To verify the assertion for (y_m, z_n) we use that, due to genericity, the simplices τ_a and τ_b are opposite and we consider the parallel set $P = P(\tau_a, \tau_b)$. Since the geodesics a and b are forward asymptotic to P, it follows that the points x, y_m, z_n have uniformly bounded distance from P. We denote their projections to P by $\overline{x}, \overline{y}_m, \overline{z}_n$.

Let $\Theta'' \subset \mathrm{int}(\Theta)$ be an auxiliary Weyl convex subset such that $\theta_{\pm a}, \theta_{\pm b} \in \mathrm{int}(\Theta'')$. We have that $\overline{y}_m \in V(\overline{x}, \mathrm{st}_{\Theta''}(\tau_a))$ for large m because the points y_m lie in a tubular neighborhood of the ray with initial point $\overline{x}$ and asymptotic to a. Similarly, $\overline{z}_n \in V(\overline{x}, \mathrm{st}_{\Theta''}(\tau_b))$ for large n. It follows that $\overline{x} \in V(\overline{y}_m, \mathrm{st}_{\Theta''}(\tau_b))$ and, using the convexity of Θ-cones (Proposition 2.1), that $\overline{z}_n \in V(\overline{y}_m, \mathrm{st}_{\Theta''}(\tau_b))$.

The cone $V(y_m, \mathrm{st}_{\Theta''}(\tau_b))$ is uniformly Hausdorff close to the cone $V(\overline{y}_m, \mathrm{st}_{\Theta''}(\tau_b))$ because the Hausdorff distance of the cones is bounded by the distance $d(y_m, \overline{y}_m)$ of their tips. Hence there exist points $x', z'_n \in V(y_m, \mathrm{st}_{\Theta''}(\tau_b))$ uniformly close to x, z_n. Since $d(y_m, x'), d(y_m, z'_n) \to \infty$ as $m, n \to \infty$, it follows that the segments $y_m x$ and $y_m z_n$ are Θ-regular for large m, n. Furthermore, since $\angle^{\xi}_{y_m}(x', z'_n) = 0$ and both $\angle_{y_m}(x, x') \to 0$ and $\angle_{y_m}(z_n, z'_n) \to 0$, it follows that $\angle^{\xi}_{y_m}(x, z_n) \to 0$. □

To conclude the proof of Theorem 4.13, the lemma implies that for any given l, ϵ the midpoint triples of the four point sequences $(\gamma_k x)$ are (Θ, ϵ)-straight and l-spaced if m, n are sufficiently large; compare the quadruple condition (Definition 3.31). This means that the midpoint sequences of all sequences $(\gamma_k x)$ are (Θ, ϵ)-straight and l-spaced for large m, n. Theorem 3.18 then implies that the sequences $(\gamma_k x)$ are uniformly Θ-Morse. □

Remark 4.16 (1) Generalizing the above argument to free groups with finitely many generators, one can construct Morse Schottky subgroups for which the set $\theta(\Lambda) \subset \sigma_{\mathrm{mod}}$ of types of limit points is arbitrarily Hausdorff close to a given ι-invariant Weyl convex subset Θ. This provides an alternative approach to the second main theorem in [Benoist 1997] using coarse geometric arguments.

(2) Theorem 4.13 was generalized (by arguments similar to the proof of Theorem 4.13) in [Dey et al. 2019] to free products of Morse subgroups of G.

4.3 Algorithmic recognition of Morse actions

In this section, we describe an algorithm which has an isometric action $\rho\colon \Gamma \curvearrowright X$ and a point $x \in X$ as its input and terminates if and only if the action ρ is Morse (otherwise, the algorithm runs forever).

We begin by describing briefly *Riley's algorithm* [1983] accomplishing a similar task, namely, detecting geometrically finite actions on $X = \mathbb{H}^3$. Suppose that we are given a finite (symmetric) set of generators $g_1 = 1, \dots, g_m$ of a subgroup $\Gamma \subset \mathrm{PO}(3,1)$ and a basepoint $x \in X = \mathbb{H}^3$. The idea of the algorithm is to construct a finite sided Dirichlet fundamental domain D for Γ (with center at x): every geometrically finite subgroup of $\mathrm{PO}(3,1)$ admits such a domain. (The latter is false for geometrically finite subgroups of $\mathrm{PO}(n,1)$ with $n \geq 4$ but is, nevertheless, true for convex cocompact subgroups.) Given a finite sided convex fundamental domain, one concludes that Γ is geometrically finite. Here is how the algorithm works: for each k define the subset $S_k \subset \Gamma$ represented by words of length $\leq k$ in the letters $g_1, \dots, g_m$. For each $g \in S_k$ consider the half-space $\mathrm{Bis}(x, g(x)) \subset X$ bounded by the bisector of the segment $xg(x)$ and containing the point x. Then compute the intersection

$$D_k = \bigcap_{g \in S_k} \mathrm{Bis}(x, g(x)).$$

Check if D_k satisfies the conditions of *Poincaré's fundamental domain theorem*. If it does, then $D = D_k$ is a finite sided fundamental domain of Γ. If not, increase k by 1 and repeat the process. Clearly, this process terminates if and only if Γ is geometrically finite.

One can enhance the algorithm in order to detect if a geometrically finite group is convex cocompact. Namely, after a Dirichlet domain D is constructed, one checks for whether:

(1) The ideal boundary of a Dirichlet domain D has isolated ideal points (they would correspond to rank-two cusps which are not allowed in convex cocompact groups).

(2) The ideal boundary of D contains tangent circular arcs with points of tangency fixed by parabolic elements (coming from the "ideal vertex cycles"). Such points correspond to rank-one cusps, which again are not allowed in convex cocompact groups.

Checking (1) and (2) is a finite process; after its completion, one concludes that Γ is convex cocompact.

We refer the reader to [Gilman 1995; 1997], [Gilman and Maskit 1991], [Kapovich 2016] and [Kapovich and Leeb 2018a, Section 1.8] for more details concerning discreteness algorithms for groups acting on hyperbolic planes and hyperbolic 3-spaces.

We now consider group actions on general symmetric spaces. Let Γ be a hyperbolic group with a fixed finite (symmetric) generating set; we equip the group Γ with the word metric determined by this generating set.

For each n, let $\mathcal{L}_n$ denote the set of maps $q\colon [0, 3n] \cap \mathbb{Z} \to \Gamma$ which are restrictions of geodesics $\widetilde{q}\colon \mathbb{Z} \to \Gamma$, such that $q(0) = 1 \in \Gamma$. In view of the geodesic automatic structure on Γ (see eg [Epstein et al. 1992, Theorem 3.4.5]), the set $\mathcal{L}_n$ can be described via a finite state automaton.

Suppose that $\rho\colon \Gamma \curvearrowright X$ is an isometric action on a symmetric space X; we fix a basepoint $x \in X$ and the corresponding orbit map $f\colon \Gamma \to \Gamma x \subset X$. We also fix an ι-invariant face τ_{mod} of the model spherical simplex σ_{mod} of X. The algorithm that we are about to describe will detect that the action ρ is τ_{mod}-Morse.

Remark 4.17 If the face τ_{mod} is not fixed in advance, we would run algorithms for each face τ_{mod} in parallel.

For the algorithm we will be using a special (countable) increasing family of Weyl convex compact subsets $\Theta = \Theta_i \subset \mathrm{ost}(\tau_{\mathrm{mod}}) \subset \sigma_{\mathrm{mod}}$ which exhausts $\mathrm{ost}(\tau_{\mathrm{mod}})$; in particular, every compact ι-invariant convex subset of $\mathrm{ost}(\tau_{\mathrm{mod}}) \subset \sigma_{\mathrm{mod}}$ is contained in some Θ_i:

$$\Theta_i := \left\{ v \in \sigma : \min_{\alpha \in \Phi_{\tau_{\mathrm{mod}}}} \alpha(v) \geq \frac{1}{i} \right\}, \tag{4.18}$$

where $\Phi_{\tau_{\mathrm{mod}}}$ is the subset of the set of simple roots Φ (with respect to σ_{mod}) which vanish on the face τ_{mod}. Clearly, the sets Θ_i satisfy the required properties. Furthermore, we consider only those L and D which are natural numbers.

Next, consider the sequence

$$(L_i, \Theta_i, D_i) = (i, \Theta_i, D_i) \quad \text{for } i \in \mathbb{N}.$$

In order to detect τ_{mod}-Morse actions we will use the local characterization of Morse quasigeodesics given by Theorem 3.18 and Proposition 3.32. Due to the discrete nature of quasigeodesics that we will be considering, it suffices to assume that the additive quasiisometry constant A is zero.

Consider the functions

$$l(\Theta, \Theta', \delta), \epsilon(\Theta, \Theta', \delta)$$

as in Theorem 3.18. Using these functions, for the sets $\Theta = \Theta_i, \Theta' = \Theta_{i+1}$ and the constant $\delta = 1$ we define the numbers

$$l_i = l(\Theta, \Theta', \delta), \epsilon_i = \epsilon(\Theta, \Theta', \delta).$$

Next, for the numbers $L = L_i, D = D_i$ and the sets $\Theta = \Theta_i, \Theta' = \Theta_{i+1}$, consider the numbers

$$s_i = s(L_i, 0, \Theta_i, \Theta_{i+1}, D_i, \epsilon_{i+1}, l_{i+1})$$

as in Proposition 3.32. According to this proposition, every $(L_i, 0, \Theta_i, D_i)$-Morse quasigeodesic satisfies the $(\Theta_{i+1}, \epsilon_{i+1}, l_{i+1}, s)$-quadruple condition for all $s \geq s_i$. We note that, a priori, the sequence s_i need not be increasing. We set $S_1 = s_1$ and define a monotonic sequence S_i recursively by

$$S_{i+1} = \max(S_i, s_{i+1}).$$

Then every $(\Theta_i, D_i, L_i, 0)$-Morse quasigeodesic also satisfies the $(\Theta_{i+1}, \epsilon_{i+1}, l_{i+1}, S_{i+1})$-quadruple condition.

We are now ready to describe the algorithm. For each $i \in \mathbb{N}$ we compute the numbers l_i, ϵ_i and, then, S_i, as above. We then consider finite discrete paths in Γ, $q \in \mathscr{L}_{S_i}$, and the corresponding discrete paths in X, $p(t) = q(t)x$ with $t \in [0, 3S_i] \cap \mathbb{Z}$. The number of paths q (and, hence, p) for each i is finite, bounded by the growth function of the group Γ.

For each discrete path p we check the $(\Theta_i, \epsilon_i, l_i, S_i)$-quadruple condition. If for some $i = i_*$, all paths p satisfy this condition, the algorithm terminates: it follows from Theorem 3.18 that the map f sends all normalized discrete biinfinite geodesics in Γ to Morse quasigeodesics in X. Hence, the action $\Gamma \curvearrowright X$ is Morse in this case. Conversely, suppose that the action of Γ is $(\Theta, D, L, 0)$-Morse. Then f sends all isomeric embeddings $\widetilde{q}\colon \mathbb{Z} \to \Gamma$ to $(\Theta, D, L, 0)$-Morse quasigeodesics $\widetilde{p}$ in X. In view of the properties of the sequence

$$(L_i, \Theta_i, D_i),$$

it follows that for some i,

$$(L, \Theta, D) \le (L_i, \Theta_i, D_i),$$

ie $L \le L_i$, $\Theta \subset \Theta_i$ and $D \le D_i$; hence, all the biinfinite discrete paths $\widetilde{p}$ are $(\Theta_i, D_i, L_i, 0)$-Morse quasigeodesic. By the definition of the numbers l_i, ϵ_i, S_i, it then follows that all the discrete paths $p = f \circ q, q \in \mathscr{L}_{S_i}$ satisfy the $(\Theta_{i+1}, \epsilon_{i+1}, l_{i+1}, S_{i+1})$-quadruple condition. Thus, the algorithm will terminate at the step $i + 1$ in this case.

Therefore, the algorithm terminates if and only if the action is Morse (for some parameters). If the action is not Morse, the algorithm will run forever. □

Remark 4.19 Applied to a rank-one symmetric space X and a hyperbolic group Γ without a nontrivial normal finite subgroup, the above algorithm verifies if the given representation $\rho\colon \Gamma \to \mathrm{Isom}(X)$ is faithful with convex-cocompact image. We could not find this result in the existing literature; cf however [Gilman and Keen 2016].

Appendix A Further properties of Morse quasigeodesics

This is the only part of the paper not contained in [Kapovich et al. 2014]. Here we collect various properties of Morse quasigeodesics that we found to be useful elsewhere in our work.

A.1 Finsler geometry of symmetric spaces

In [Kapovich and Leeb 2018b] (see also [Kapovich et al. 2017]), we considered a certain class of G-invariant "polyhedral" Finsler metrics on X. Their geometric and asymptotic properties turned out to be well adapted to the study of geometric and dynamical properties of regular subgroups. They provide a Finsler geodesic *combing* of X which is, in many ways, more suitable for analyzing the asymptotic

geometry of X than the geodesic combing given by the standard Riemannian metric on X. These Finsler metrics also play a basic role in the present paper. We briefly recall their definition and some basic properties, and refer to [Kapovich and Leeb 2018b, Section 5.1] for more details.

Let $\overline{\theta} \in \mathrm{int}(\tau_{\mathrm{mod}})$ be a type spanning the face type τ_{mod}. The *$\overline{\theta}$-Finsler distance* $d^{\overline{\theta}}$ on X is the G-invariant pseudometric defined by

$$d^{\overline{\theta}}(x,y) := \max_{\theta(\xi)=\overline{\theta}} \big(b_\xi(x) - b_\xi(y)\big)$$

for $x, y \in X$, where the maximum is taken over all ideal points $\xi \in \partial_\infty X$ with type $\theta(\xi) = \overline{\theta}$. It is positive, ie a (nonsymmetric) metric, if and only if the radius of σ_{mod} with respect to $\overline{\theta}$ is $< \pi/2$. This is in turn equivalent to $\overline{\theta}$ not being contained in a factor of a nontrivial spherical join decomposition of σ_{mod}, and is always satisfied if for instance X is irreducible.

If $d^{\overline{\theta}}$ is positive, it is equivalent to the Riemannian metric. In general, if it is only a pseudometric, it is still equivalent to the Riemannian metric d on uniformly regular pairs of points. More precisely, if the pair of points x, y is Θ-regular, then

$$L^{-1} d(x,y) \le d^{\overline{\theta}}(x,y) \le L d(x,y)$$

with a constant $L = L(\Theta) \ge 1$.

Regarding symmetry of the Finsler distance, one has the identity

$$d^{\iota\overline{\theta}}(y,x) = d^{\overline{\theta}}(x,y),$$

and hence $d^{\overline{\theta}}$ is symmetric if and only if $\iota\overline{\theta} = \overline{\theta}$. We refer to $d^{\overline{\theta}}$ as a Finsler metric *of type* τ_{mod}.

The $d^{\overline{\theta}}$-balls in X are convex but not strictly convex. (Their intersections with flats through their centers are polyhedra.) Accordingly, $d^{\overline{\theta}}$-geodesics connecting two given points x, y are not unique. To simplify notation, xy will stand for *some* $d^{\overline{\theta}}$-geodesic connecting x and y. The union of all $d^{\overline{\theta}}$-geodesic xy equals the τ_{mod}-diamond $\diamondsuit_{\tau_{\mathrm{mod}}}(x,y)$, that is, a point lies on a $d^{\overline{\theta}}$-geodesic xy if and only if it is contained in $\diamondsuit_{\tau_{\mathrm{mod}}}(x,y)$; see [Kapovich et al. 2017]. Finsler geometry thus provides an alternative description of diamonds. Note that with this description, the diamond $\diamondsuit_{\tau_{\mathrm{mod}}}(x,y)$ is also defined when the segment xy is not τ_{mod}-regular. Such a *degenerate* τ_{mod}-diamond is contained in a smaller totally geodesic subspace, namely in the intersection of all τ_{mod}-parallel sets containing the points x, y. The description of geodesics and diamonds also implies that the unparametrized $d^{\overline{\theta}}$-geodesics depend only on the face type τ_{mod}, and not on $\overline{\theta}$. We will refer to $d^{\overline{\theta}}$-geodesics as *τ_{mod}-Finsler geodesics*. Note that Riemannian geodesics are Finsler geodesics.

We will call a Θ-regular τ_{mod}-Finsler geodesic a *Θ-Finsler geodesic*. If xy is a Θ-regular (Riemannian) segment, then the union of Θ-Finsler geodesics xy equals the Θ-diamond $\diamondsuit_\Theta(x,y)$.

Every τ_{mod}-Finsler ray in X is contained in a τ_{mod}-Weyl cone, and we will use the notation $x\tau$ for a τ_{mod}-Finsler ray contained $V(x, \mathrm{st}(\tau))$. Similarly, every τ_{mod}-Finsler line is contained in a τ_{mod}-parallel set, and we denote by $\tau_-\tau_+$ an oriented τ_{mod}-Finsler line forward/backward asymptotic to two antipodal simplices $\tau_\pm \in \mathrm{Flag}(\tau_{\mathrm{mod}})$ and contained in $P(\tau_-, \tau_+)$.

Examples of Θ-regular Finsler geodesics can be obtained as follows. Let (x_i) be a (finite or infinite) sequence contained in a parallel set $P(\tau_-, \tau_+)$ such that each Riemannian segment $x_i x_{i+1}$ is τ_+-longitudinal (see Definition B.2) and Θ'-regular. Then the concatenation of these geodesic segments is a Θ-regular Finsler geodesic; see Remark B.3.

Conversely, every Θ-regular Finsler geodesic $c\colon I \to X$ can be *approximated* by a piecewise Riemannian Finsler geodesic c': pick a number $s > 0$ and consider a maximal s-separated subset $J \subset I$. Then take c' to be the concatenation of Riemannian geodesic segments $c(i)c(j)$ for consecutive pairs $i, j \in J$. In view of this approximation procedure, the "string of diamonds" theorem (Theorem 3.30) holds if instead of Riemannian geodesic segments $x_i x_{i+1}$, we allow Θ-regular Finsler segments.

A.2 Stability of diamonds

Diamonds can be regarded as Finsler-geometric replacements of *geodesic segments* in nonpositively curved symmetric spaces of higher rank.

Riemannian geodesic segments in Hadamard manifolds (and, more generally, CAT(0) metric spaces) depend *uniformly continuously* on their tips: by convexity of the distance function we have,

$$d_{\mathrm{Haus}}(xy, x'y') \le \max(d(x, x'), d(y, y')).$$

In [Kapovich et al. 2018b, Proposition 3.70] we proved that diamonds $\diamondsuit_{\tau_{\mathrm{mod}}}$ depend *continuously* on their tips.

Below we establish uniform control on how much sufficiently large Θ-diamonds vary with their tips.

Lemma A.1 *For $d' > d > 0$ there exists $C = C(\Theta, \Theta', d, d')$ such that the following holds*:

If a segment $x_-x_+ \subset X$ is Θ-regular with length $\ge C$ and $y_\pm \in B(x_\pm, d)$, then

(i) *the segment y_-y_+ is Θ'-regular, and*

(ii) $\diamondsuit_\Theta(x_-, x_+) \subset N_{d'}(\diamondsuit_{\Theta'}(y_-, y_+))$.

Proof (i) The Θ'-regularity of y_-y_+ for sufficiently large C follows from the Δ-triangle inequality.

(ii) Suppose that there exists no constant C for which (ii) holds. Then there are sequences of points $x_n^\pm$ with $d(x_n^-, x_n^+) \to +\infty$, $y_n^\pm$ with $d(x_n^\pm, y_n^\pm) \le d$, $x_n \in \diamondsuit_\Theta(x_n^-, x_n^+)$ and $y_n \in \diamondsuit_{\Theta'}(y_n^-, y_n^+)$ with $d(x_n, \diamondsuit_{\Theta'}(y_n^-, y_n^+)) = d(x_n, y_n) = d'$. We may assume convergence $x_n \to x_\infty$ and $y_n \to y_\infty$ in X.

After extraction, at least one of the sequences $(x_n^\pm)$ diverges. There are two cases to consider.

Suppose first that both sequences $(x_n^\pm)$ diverge. Then they are uniformly τ_{mod}-regular and, after extraction, we have τ_{mod}-flag convergence $x_n^\pm, y_n^\pm \to \tau_\pm \in \mathrm{Flag}(\tau_{\mathrm{mod}})$. The limit simplices $\tau_\pm$ are antipodal (because $x_n \to x_\infty$). We observe that

$$d(x_n, \partial\diamondsuit_{\Theta'}(x_n^-, x_n^+)), d(y_n, \partial\diamondsuit_{\Theta'}(y_n^-, y_n^+)) \to +\infty.$$

It follows that the sequences of diamonds $\diamondsuit_{\Theta'}(x_n^-, x_n^+)$ and $\diamondsuit_{\Theta'}(y_n^-, y_n^+)$ both Hausdorff converge to the τ_{mod}-parallel set $P = P(\tau_-, \tau_+)$. It holds that $x_\infty \in P$ because $x_n \in \diamondsuit_\Theta(x_n^-, x_n^+)$. On the other hand, $d(x_\infty, P) = d'$ because $d(x_n, \diamondsuit_{\Theta'}(y_n^-, y_n^+)) = d'$, a contradiction.

Second, suppose that only one of the sequences $(x_n^\pm)$ diverges, say, after extraction, $x_n^- \to x_\infty^-$ and $y_n^- \to y_\infty^-$ in X to limit points with $d(x_\infty^-, y_\infty^-) \le d$, and $x_n^+ \to \tau_+ \in \mathrm{Flag}(\tau_{\mathrm{mod}})$. Now the distance of x_n from the boundary of the Θ'-Weyl cone with tip x_n^+ and containing x_n goes to infinity and it follows that $\diamondsuit_{\Theta'}(x_n^-, x_n^+) \to V(x_\infty^-, \mathrm{st}_{\Theta'}(\tau_+))$ and, similarly, $\diamondsuit_{\Theta'}(y_n^-, y_n^+) \to V(y_\infty^-, \mathrm{st}_{\Theta'}(\tau_+))$. The asymptotic limit Weyl cones have Hausdorff distance $d(x_\infty^-, y_\infty^-)$. On the other hand, $x_\infty \in V(x_\infty^-, \mathrm{st}_{\Theta'}(\tau_+))$ and $d(x_\infty, V(y_\infty^-, \mathrm{st}_{\Theta'}(\tau_+))) = d'$, again a contradiction.

This shows that (ii) holds for sufficiently large C. □

We reformulate this result in terms of Finsler geodesics:

Lemma A.2 *There exists $C = C(\Theta, \Theta', d, d')$ such that the following holds: if x_-x_+ is a Θ-Finsler geodesic in X with $d(x_-, x_+) \ge C$, and $y_\pm$ are points with $d(y_\pm, x_\pm) \le d$, then every point x on x_-x_+ lies within distance d' of a point y on a Θ'-Finsler geodesic y_-y_+.*

We do not claim here that one can take the same Finsler geodesic y_-y_+ for all points x on x_-x_+.

We now apply this stability result to Morse quasigeodesics. One, somewhat annoying, feature of the definition of Θ-Morse quasigeodesics $p\colon I \to X$ is that $p([t_1, t_2])$ is not required to be uniformly close to a Θ-diamond spanned by $p(t_1)$, $p(t_2)$. (One reason is because the segment $p(t_1)p(t_2)$ need not be Θ-regular.) Nevertheless, Lemma A.1 implies:

Lemma A.3 *For every Morse datum $M = (\Theta, B, L, A)$ and $\Theta' > \Theta$, there exists $C = C(M, \Theta')$ and D' such that the following holds for all (Θ, B, L, A)-Morse quasigeodesics $p\colon I \to X$:*

Whenever $d(x_1, x_2) \ge C$, the segment $x_1x_2 = p(t_1)p(t_2)$ is Θ'-regular and $p([t_1, t_2])$ lies in the D'-neighborhood of the Θ'-diamond $\diamondsuit_{\Theta'}(x_1, x_2)$.

A.3 Finsler approximation of Morse quasigeodesics

The next theorem establishes that every (sufficiently long) Morse quasigeodesic is uniformly close to a Finsler geodesic with the same endpoints. In this theorem, for convenience of notation, we will be allowing Morse quasigeodesics p to be defined on closed intervals I in the extended real line; this is just a shorthand for a map $p\colon I' = I \cap \mathbb{R} \to X$ such that, as $t \to \pm\infty$ in I', $p(t) \to p(\pm\infty) \in \mathrm{Flag}(\tau_{\mathrm{mod}})$. When we say that such maps p, c are within distance D' from each other, this simply means that their restrictions to I' are within distance $\le D'$.

Theorem A.4 (Finsler approximation theorem) *For every Morse datum $M = (\Theta, D, L, A)$, $\Theta' > \Theta$, and positive number S, there exist $C = C(M, \Theta', S)$ and $D' = D'(M, \Theta', S)$ satisfying the following.*

Let $p\colon I = [t_-, t_+] \to X \cup \mathrm{Flag}(\tau_{\mathrm{mod}})$ be an M-Morse quasigeodesic between the points $x_\pm = p(t_\pm)$ in $X \cup \mathrm{Flag}(\tau_{\mathrm{mod}})$ such that $d(x_-, x_+) \geq C$. Then there exists a Θ'-Finsler geodesic x_-x_+ equipped with a monotonic parametrization $c\colon I \to x_-x_+$ such that

(a) *the maps $p, c\colon I \to X$ are within distance $\leq D'$ from each other, and*

(b) *x_-x_+ is an S-spaced piecewise Riemannian geodesic, ie the Riemannian length of each Riemannian segments of x_-x_+ is $\geq S$.*

Proof We will prove this in the case when both $x_\pm$ are in X, since the proofs when one or both points $x_\pm$ are in $\mathrm{Flag}(\tau_{\mathrm{mod}})$ are similar: one replaces diamonds with Weyl cones or parallel sets.

By the definition of an M-Morse quasigeodesic, for all subintervals $[s_-, s_+] \subset [t_-, t_+]$ there exists a Θ-diamond

$$\diamondsuit_\Theta(y'_-, y'_+)$$

whose D-neighborhood contains $p([s_-, s_+])$, and for $y_\pm = p(s_\pm)$ we have

$$d(y_\pm, y'_\pm) \leq D.$$

Therefore, applying Lemma A.1(i), we conclude that the Riemannian segment y_-y_+ is Θ'-regular provided that

$$d(y_-, y_+) \geq C_1 = C_1(M, \Theta'). \tag{A.5}$$

In view of the quasigeodesic property of p, there exists $s = s(M, \Theta')$ such that the separation condition

$$s_+ - s_- \geq s = s(M, \Theta')$$

implies the inequality (A.5). This, of course, also applies to $[s_-, s_+] = [t_-, t_+]$ and hence, using Lemma A.1(ii), we obtain

$$p(I) \subset N_D(\diamondsuit_\Theta(x'_-, x'_+)) \subset N_{D+D_1}(\diamondsuit_{\Theta'}(x_-, x_+)),$$

where $D_1 = D_1(M, \Theta')$. We let

$$\overline{y}_\pm \in \diamondsuit' := \diamondsuit_{\Theta'}(x_-, x_+) = V(x_-, \mathrm{st}_{\Theta'}(\tau_+)) \cap V(x_+, \mathrm{st}_{\Theta'}(\tau_-))$$

denote the nearest-point projections of $y_\pm = p(s_\pm)$. There exists $s' = s'(M, \Theta')$ such that as long as $s_+ - s_- \geq s'$, the Riemannian segments $\overline{y}_-\overline{y}_+$ are also Θ'-regular and have length $\geq S$. Furthermore, as in the proof of Proposition 3.32, we can assume that $s' = s'(M, \Theta')$ is chosen so that, whenever $s_+ - s_- \geq s'$, each segment $\overline{y}_-\overline{y}_+$ is τ_+-longitudinal. We fix this $s' = s'(M, \Theta')$ from now on.

We also assume, in what follows, that $t_+ - t_- \geq s'(M, \Theta')$; this is achieved by assuming that

$$L^{-1}(d(x_-, x_+) - A) \geq s' = s'(M, \Theta').$$

Take a maximal s'-separated subset $J \subset I$ containing $t_\pm$. For each $j \in J$ define the point

$$z_j := \overline{p(j)} \in \diamondsuit'.$$

Then for all consecutive $i, j \in J$, with $s' \le |j-i| \le 2s'$, we have

$$\textbf{(A.6)} \qquad L^{-1}s' - (A + 2D + 2D_1) \le d(z_i, z_j) \le 2Ls' + (A + 2D + 2D_1).$$

We then let c denote the concatenation of Riemannian segments $z_i z_j$ for consecutive $i, j \in J$, where we use the affine parametrization of $[i, j] \to z_i z_j$. Thus, c is a Θ'-Finsler geodesic. We now take the smallest $s'' \ge s'(M, \Theta')$ satisfying

$$S \le L^{-1}s'' - (A + 2D + 2D_1).$$

The inequalities (A.6) imply that c satisfies both requirements of the approximation theorem with

$$D' = 2Ls'' + (A + 2D + 2D_1) + (D + D_1) + (2Ls'' + A). \qquad \square$$

Remark A.7 In the case when the domain of p is unbounded, one can prove a slightly sharper result; namely, one can take $\Theta' = \Theta$. Compare [Kapovich and Leeb 2023, Section 6].

A.4 Altering Morse quasigeodesics

Below we consider certain modifications of M-Morse quasigeodesics p in X represented as concatenations $p = p_- \star p_0 \star p_+$, where $x_\pm$ are the endpoints of p_0, and $y_\pm, x_\pm$ are the endpoints of $p_\pm$. (As in the previous section, we will be allowing $y_\pm$ to be in $X \cup \mathrm{Flag}(\tau_{\mathrm{mod}})$.) These modifications will have the form $p' = p'_- \star p'_0 \star p'_+$, where $p'_\pm$ and p'_0 are all Morse. We will see that, under certain assumptions, the entire p' is again Morse (for suitable Morse datum M').

We begin by analyzing extensions of p to biinfinite paths.

Lemma A.8 (extension lemma) *Suppose that $p = p_- \star p_0 \star p_+$ is an M-Morse quasigeodesic as above. Assume, furthermore, that*

$$p_\pm \subset V_\pm = V(x_\pm, \mathrm{st}(\tau_\pm)).$$

Whenever $y_\pm$ is in X, we let $c_\pm$ be Θ-regular Finsler rays contained in $V_\pm$ and connecting $y_\pm$ to $\tau_\pm$. Then, for every $\Theta' > \Theta$, there exists a Morse datum M' containing Θ' (and depending only on M and Θ') such that the concatenation

$$\widehat{p} = c_- \star p \star c_+$$

is M'-Morse, provided that $d(x_\pm, y_\pm) \ge C = C(M, \Theta')$.

Proof We fix an auxiliary subset Θ_1 satisfying $\Theta < \Theta_1 < \Theta'$. We let $S = S(\Theta_1, \Theta', 1), \epsilon = \epsilon(\Theta_1, \Theta', 1)$ be constants as in the string of diamonds theorem (Theorem 3.30).

According to Theorem A.4, there exists a Θ'-regular Finsler geodesic

$$\overline{c} = y_- \overline{x}_- \star \overline{x}_- \overline{x}_+ \star \overline{x}_+ y_+$$

within distance $D_1 = D_1(M, \Theta', S)$ from the path p, such that $\overline{c}$ is the concatenation of segments of length $\geq S$ and $d(x_\pm, \overline{x}_\pm) \leq D_1$. We let $z_\pm y_\pm$ denote the subsegments of $\overline{x}_\pm y_\pm$ containing $y_\pm$.

Since $d(x_\pm, \overline{x}_\pm) \leq D_1$, for each $\epsilon > 0$ and a sufficiently large $C_1 = C_1(D_1, \Theta')$, the inequality $d(x_\pm, y_\pm) \geq C_1$ implies

$$\angle^{\xi}_{y_\pm}(x_\pm, \overline{x}_\pm) \leq \epsilon.$$

Therefore,

$$\angle^{\xi}_{y_\pm}(z_\pm, \tau_\pm) \geq \pi - 2\epsilon$$

and, hence, the piecewise geodesic path

$$\widehat{c} = c_- \star \overline{c} \star c_+$$

is $(\Theta_1, 2\epsilon)$-straight and S-spaced. By Theorem 3.30, the concatenation $\widehat{c}$ is M_1-Morse, where $M_1 = (\Theta', 1, L, A)$. Since the path $\widehat{p}$ is within distance D_1 from $\widehat{c}$, it is M'-Morse, where $M' = M_1 + D_1$. □

The next lemma was proven in [Dey et al. 2019, Thm. 4.11] in the case when p, p' are finite paths. The proof in the case of (bi)infinite paths is the same and we omit it.

Lemma A.9 (replacement lemma) *Suppose that $p' = p'_- \star p'_0 \star p'_+$ is an M-Morse quasigeodesic in X, which is a concatenation of M-Morse quasigeodesics p'_-, p'_0, p'_+. Let p_0 be an M-Morse quasigeodesic connecting the initial and the terminal point of p'_0. Then for every $\Theta' > \Theta$ there exists a Morse datum M' containing Θ' (again, depending only on M and Θ') such that the path*

$$p' = p'_- \star p_0 \star p'_+$$

is M'-Morse.

In the following three lemmata we will modify an M-Morse quasigeodesic $p = p_- \star p_0 \star p_+$ by altering $p_\pm$ and keeping p_0 unchanged ("wiggling the head and the tail of p").

Lemma A.10 (wiggle lemma, I) *Suppose p is as above and the paths $p_\pm$ are both infinite and p_0 connects the points x_+, x_-. We let $p'_\pm$ be M-Morse quasigeodesic rays with finite terminal points $x_\pm$ and set $p' := p'_- \star p_0 \star p'_+$. Then, given $\Theta' > \Theta$ there exists $\epsilon = \epsilon(M, \Theta') > 0$ and a Morse datum M' (depending only on M and Θ') containing Θ' such that if*

$$\mu := \max(\angle^{\xi}_{x_\pm}(p'_\pm(\pm\infty), p_\pm(\pm\infty))) < \epsilon,$$

then p' is M'-Morse.

Proof We fix an auxiliary compact Weyl-convex subset $\Theta_1 \subset \mathrm{ost}(\tau_{\mathrm{mod}})$ such that $\Theta < \Theta_1 < \Theta'$. Set $\tau_\pm = p_\pm(\pm\infty)$, $\tau'_\pm = p'_\pm(\pm\infty)$.

According to Lemma A.9, there exists a Morse datum M_1 containing Θ_1 such that for any Θ_1-regular Finsler geodesic rays $c_\pm := x_\pm \tau_\pm$, the concatenation $c_- \star p_0 \star c_+$ is M_1-Morse.

Let $M_2 > M_1 + 1$ be a Morse datum containing Θ' and let $S > 0$ be such that if a path q in X is S-locally $M_1 + 1$-Morse then q is M_2-Morse; see Theorem 3.34. Let ϵ be such that for $x \in X$ and $\tau, \tau' \in \mathrm{Flag}(\tau_{\mathrm{mod}})$, if $\angle_x^\zeta(\tau, \tau') < \epsilon$ then each Θ_1-regular Finsler segment of length $\le S$ in $V(x, \mathrm{st}(\tau'))$ is within unit distance from a Θ_1-regular Finsler segment of length $\le S$ in $V(x, \mathrm{st}(\tau))$. We assume now that $\mu < \epsilon$.

Since $p'_\pm$ are M-Morse rays, they are within distance $D_1 = D_1(M, \Theta_1)$ from Θ_1-regular Finsler rays $c'_\pm = x_\pm \tau'_\pm$ connecting $x_\pm$ and $\tau'_\pm$. Define a new path $c' := c'_- \star p_0 \star c'_+$.

By our choice of ϵ, the Θ_1-regular Finsler subsegment $s'_\pm = x_\pm y'_\pm$ of $c'_\pm$ of length S is within unit distance from a Θ_1-regular Finsler subsegment $s_\pm = x_\pm y_\pm$ of $c_\pm$ of length S, where $c_\pm = x_\pm \tau_\pm$ is a Θ_1-Finsler geodesic connecting $x_\pm$ to $\tau_\pm$.

The concatenation

$$s_- \star p_0 \star s_+$$

is M_1-Morse and, since $c'_\pm$ are Θ_1-Finsler geodesic, the path c' is S-locally $M_1 + 1$-Morse. By our choice of S, the path c' is M_2-Morse. Since c' is within distance D_1 from p', the path p' is $M_2 + D_1$-Morse. Lastly, we set $M' := M_2 + D_1$. □

We generalize this lemma by allowing finite Morse quasigeodesics. We continue with the setting and notation of Lemma A.10; as before, M' in the next two lemmata will depend only on M and Θ'. In Lemma A.11 we now allow paths $p_\pm$ and $p'_\pm$ to be finite, connecting $y_\pm, x_\pm$ and $y'_\pm, x_\pm$, respectively. (Some of $y_\pm, y'_\pm$ might be in $\mathrm{Flag}(\tau_{\mathrm{mod}})$.) However, we will now assume that the distances $d(x_\pm, y_\pm), d(x'_\pm, y_\pm)$ are sufficiently large, namely $\ge C$.

Lemma A.11 (wiggle lemma, II) *Given $\Theta' > \Theta$ there exist $C \ge 0$, $\epsilon > 0$ and a Morse datum M' containing Θ' such that if*

$$\mu := \max(\angle^\zeta_{x_\pm}(y'_\pm, y_\pm)) < \epsilon \quad \text{and} \quad \nu := \min(d(x_\pm, y_\pm), d(x_\pm, y'_\pm)) \ge C,$$

then p' is M'-Morse.

Proof Pick an auxiliary compact Weyl-convex subset Θ_2, with $\Theta < \Theta_2 < \Theta'$.

We define biinfinite geodesic extensions $\hat{p}, \hat{p}'$ as in Lemma A.8, by extending (if necessary) the paths $p_\pm, p'_\pm$ via Θ-Finsler geodesics $y_\pm \tau_\pm$ and $y'_\pm \tau'_\pm$. According to Lemma A.8, there exists $C > 0$ and a Morse datum M_2 (containing Θ_2), both depending on M and Θ_2, such that the path $\hat{p}$ is M_2-Morse. The same lemma applied to the paths $\hat{p}'_\pm$ implies that they are also M_2-Morse.

By the construction,

$$\mu := \angle^\zeta_{x_\pm}(y'_\pm, y_\pm) = \angle^\zeta_{x_\pm}(\tau'_\pm, \tau_\pm).$$

Now, the claim follows from Lemma A.10. □

Lastly, we prove a general wiggle lemma where we allow perturbing of the entire path p. We consider concatenations

$$p = p_- \star p_0 \star p_+ \quad \text{and} \quad p' = p'_- \star p'_0 \star p'_+$$

of M-Morse quasigeodesics, where we assume that p_0 and p'_0 are within distance D_0 from each other. The paths $p_\pm$ connect $y_\pm, x_\pm$, and the paths $p'_\pm$ connect $y'_\pm, x'_\pm$.

In the next lemma, we continue with the setting and notation of Lemma A.10.

Lemma A.12 (wiggle lemma, III) *Given $\Theta' > \Theta$ there exist $C \geq 0$, $\epsilon > 0$ and a Morse datum M' containing Θ' such that if*

$$\mu := \max(\angle^{\xi}_{x_\pm}(y'_\pm, y_\pm)) < \epsilon \quad \textit{and} \quad \nu := \min(d(x_\pm, y_\pm), d(x'_\pm, y'_\pm)) \geq C,$$

then p' is M'-Morse.

Proof As before, we fix an auxiliary compact Weyl-convex subset Θ_3, $\Theta < \Theta_3 < \Theta'$. Then $p'_\pm$ are within distance $D_3 = D_3(M, \Theta_3)$ from Θ_3-regular Finsler geodesics $c_\pm := y'_\pm x_\pm$. We apply Lemma A.11 to the pair of paths

$$p,\ p'' := c_- \star p_0 \star c_+.$$

It follows that p'' is M_3-Morse for some Morse datum M_3 containing Θ' provided that

$$\mu \leq \epsilon = \epsilon(M, \Theta_3, \Theta') \quad \text{and} \quad \nu \geq C = C(M, \Theta_3, \Theta').$$

Since the paths p'' and p' are within distance $D' := \max(D_0, D_3)$ from each other, the path p' is $M' := (M_3 + D')$-Morse. □

Appendix B Geometry of nonpositively curved symmetric spaces and their ideal boundaries

In this appendix we collect various definitions and facts about symmetric spaces of noncompact type, their ideal boundaries and their isometry groups. The material of this section is taken from [Kapovich et al. 2017; 2018b] and [Kapovich and Leeb 2018b].

Let X be a symmetric space of noncompact type. Throughout the paper we will be using the *visual* compactification of X, namely $\overline{X} = X \cup \partial_\infty X$ (see [Ballmann et al. 1985]), where $\partial_\infty X$ is defined as the set of *asymptotic* equivalence classes of geodesic rays in X: two rays are equivalent if they are at finite Hausdorff distance from each other. We refer to [Ballmann et al. 1985] for the detailed definition of the topology on $\overline{X}$. We will also occasionally use the notion of *strongly asymptotic* geodesic rays: these are rays $\rho_i : \mathbb{R}_+ \to X$ for $i = 1, 2$ such that

$$\lim_{t\to\infty} d(\rho_1(t), \rho_2(t)) = 0.$$

The space of classes of strongly asymptotic rays is a topological space, obtained by taking the quotient space of the space of all geodesic rays in X. The latter is topologized by the topology of uniform convergence on compacts. Given a point $\xi \in \partial_\infty X$, one defines X_ξ, the *space of strong asymptote classes at* ξ as follows. Consider the set $\mathrm{Ray}(\xi)$ consisting of all geodesic rays in X asymptotic to ξ (and, as before, equipped with the topology of uniform convergence on compacts). Then take the quotient of $\mathrm{Ray}(\xi)$ by the equivalence relation where two rays are equivalent if they are strongly asymptotic. One can identify the resulting space X_ξ as follows. Pick a point $\widehat{\xi} \in \partial_\infty X$ opposite to ξ and consider the parallel set $P(\xi, \widehat{\xi})$, which is the union of all geodesic rays in X which are forward/backward asymptotic to $\xi, \widehat{\xi}$. Each point $x \in P(\xi, \widehat{\xi})$ defines the ray $x\xi$ asymptotic to ξ. The projection of the subset of such rays to X_ξ is a homeomorphism. Thus, we can identify X_ξ with the parallel set $P(\xi, \widehat{\xi})$. One metrizes X_ξ by

$$d(\rho_1, \rho_2) = \lim_{t\to\infty} d(\rho_1(t), \rho_2(t)).$$

Then the projection $P(\xi, \widehat{\xi}) \to X_\xi$ is an isometry. We refer to [Kapovich et al. 2017, Section 2.8] for details.

Given a subset $A \subset X$ we let $\partial_\infty A$ denote the accumulation set of A in $\partial_\infty X$.

Building notions The visual boundary $\partial_\infty X$ of a symmetric space X admits a structure as a thick spherical building (the Tits building of X), which is a certain *spherical* simplicial complex; see [Eberlein 1996; Ballmann et al. 1985]. This complex is either connected (if X has rank ≥ 2) or discrete (if X has rank one, or equivalently, X is negatively curved). In the connected case, this building is equipped with the path-metric $\angle_{\mathrm{Tits}}$ induced by the spherical metrics on simplices. This metric space has diameter 1. In the case when the building is discrete, the distance between distinct points is π.

The connected component $\mathrm{Isom}_0(X)$ of the isometry group of X acts isometrically on this spherical building, and transitively on facets (top-dimensional simplices). The quotient $\partial_\infty X/\mathrm{Isom}_0(X)$ is a single spherical simplex, denoted by σ_{mod}, the *model spherical chamber* of the building $\partial_\infty X$. One can identify σ_{mod} with a facet of the building $\partial_\infty X$, a *fundamental domain* for the action $\mathrm{Isom}_0(X)$ on the building. We will use the notation $\theta : \partial_\infty X \to \sigma_{\mathrm{mod}}$ for the *type projection*, the quotient map $\partial_\infty X \to \partial_\infty X/\mathrm{Isom}_0(X)$. The full isometry group $\mathrm{Isom}(X)$ acts on both $\partial_\infty X$ and σ_{mod} and the map θ is equivariant with respect to these actions. Note that the action on σ_{mod}, in general, is nontrivial. Throughout the paper, G denotes a Lie group with finitely many components, which acts isometrically on X with finite kernel, so that the action of G on σ_{mod} is trivial and the image of G in $\mathrm{Isom}(X)$ has finite index. In particular, we can identify X with G/K, where K is a maximal compact subgroup of G. The reader can think of G as being equal to the kernel of the action of $\mathrm{Isom}(X)$ on σ_{mod}. However, this would exclude such natural examples as $G = \mathrm{SL}(2, \mathbb{R})$, which acts on the associated symmetric space (the hyperbolic plane) with finite nontrivial kernel (equal to the center of G).

For an algebraically inclined reader, the spherical building above is defined as a simplicial complex whose vertices are conjugacy classes of maximal parabolic subgroups of G. If W is the Weyl group of X (equivalently, the relative Weyl group of G) and r is the rank of X (equivalently, the split real rank of G),

then W is a reflection group acting isometrically on the unit sphere a_{mod} of dimension $r-1$ and σ_{mod} is a fundamental chamber of this action. From the building viewpoint, a_{mod} is the *model spherical apartment* of the Tits building $\partial_\infty X$.

We return to the geometric discussion of $\partial_\infty X$. We let $\iota\colon \sigma_{\mathrm{mod}} \to \sigma_{\mathrm{mod}}$ denote the *opposition involution*: it is the projection to σ_{mod} of *Cartan involutions* of X. We will use the notation $\overline{\zeta} = \theta(\zeta)$ for elements of σ_{mod}. Similarly (cf [Kapovich et al. 2017, Section 2.2.2]), we let $\tau_{\mathrm{mod}} \subseteq \sigma_{\mathrm{mod}}$ denote faces of σ_{mod}: these are the *model simplices* in the Tits building; we will use the notation τ for simplices in $\partial_\infty X$ of type τ_{mod}, so that $\theta(\tau) = \tau_{\mathrm{mod}}$. Given a face τ_{mod} of σ_{mod}, we let $\partial_{\tau_{\mathrm{mod}}}\sigma_{\mathrm{mod}}$ denote the union of faces of σ_{mod} disjoint from τ_{mod}; we use the notation $\mathrm{ost}(\tau_{\mathrm{mod}})$ for the complement $\sigma_{\mathrm{mod}} \setminus \partial_{\tau_{\mathrm{mod}}}\sigma_{\mathrm{mod}}$, the *open star* of τ_{mod} in σ_{mod}.

Let $W_{\tau_{\mathrm{mod}}}$ denote the stabilizer of the face τ_{mod} in W. A (convex) subset $\Theta \subset \sigma_{\mathrm{mod}}$ is said to be *Weyl-convex* (more precisely, Weyl-convex with respect to a face τ_{mod}) if the $W_{\tau_{\mathrm{mod}}}$-orbit of Θ is a convex subset in the sphere a_{mod}; see [Kapovich et al. 2017, Definition 2.7]. We will *always* use the notation Θ for compact ι-invariant Weyl-convex subsets of the open star $\mathrm{ost}(\tau_{\mathrm{mod}})$.

Note that, when X has rank one, the data τ_{mod} and Θ are obsolete since σ_{mod} is a singleton. In this case, we also have that $\partial\sigma_{\mathrm{mod}} = \varnothing$ and $\Theta = \mathrm{int}(\sigma_{\mathrm{mod}}) = \sigma_{\mathrm{mod}}$ is clopen.

Two points in $\partial_\infty X$ are called *antipodal* if their distance in the Tits building $\partial_\infty X$ equals π. Equivalently, these points are connected by a geodesic in X. Equivalently, antipodal points are swapped by a Cartan involution of X. Similarly, two simplices in $\partial_\infty X$ are said to be antipodal if they are swapped by a Cartan involution. We will use the notation $\tau, \widehat{\tau}$ for pairs of antipodal simplices. Their types in σ_{mod} are swapped by the opposition involution ι. Two simplices τ, τ' in $\partial_\infty X$ are antipodal if and only if $\iota(\theta(\tau')) = \theta(\tau)$ and the open simplices corresponding to τ, τ' contain antipodal points.

Identify σ_{mod} with a simplex in $\partial_\infty X$. Then G-stabilizers of faces τ_{mod} of σ_{mod} are *standard parabolic subgroups* of G, and are denoted by $P_{\tau_{\mathrm{mod}}}$; these are closed subgroups of G. The set of simplices of type τ_{mod} in $\partial_\infty X$ is identified with the quotient $G/P_{\tau_{\mathrm{mod}}}$, which is a smooth compact manifold called the *flag-manifold* $\mathrm{Flag}(\tau_{\mathrm{mod}})$ of type τ_{mod}; see [Kapovich et al. 2017, Sections 2.2.2 and 2.2.3]. Given $\tau \in \mathrm{Flag}(\tau_{\mathrm{mod}})$, one defines the *open Schubert cell* $C(\tau) \subset \mathrm{Flag}(\tau_{\mathrm{mod}})$, which is an open subset of $\mathrm{Flag}(\tau_{\mathrm{mod}})$ consisting of elements opposite to τ; see [Kapovich et al. 2017, Section 2.4]. We will use the notation $\mathrm{int}\,\tau$ for the *open simplex* obtained by removing from τ all its proper faces. The notation $\mathrm{st}(\tau)$ is used to denote the *star* of τ in $\partial_\infty X$, the union of all faces of $\partial_\infty X$ containing τ. Similarly, $\mathrm{ost}(\tau)$, the *open star* of τ in $\partial_\infty X$, is obtained from $\mathrm{st}(\tau)$ by removing all faces disjoint from τ. Accordingly, $\partial_{\tau_{\mathrm{mod}}}\sigma_{\mathrm{mod}}$ is defined as $\sigma_{\mathrm{mod}} \setminus \mathrm{ost}(\tau_{\mathrm{mod}})$.

Given a Weyl-convex compact subset $\Theta \subset \mathrm{ost}\,\tau_{\mathrm{mod}} \subset \sigma_{\mathrm{mod}}$, we will define the Θ-*star* of a simplex τ in $\partial_\infty X$ of type τ_{mod} as the preimage of Θ under the restriction $\theta\colon \mathrm{st}(\tau) \subset \partial_\infty X \to \sigma_{\mathrm{mod}}$.

Symmetric space notions An isometry of X is called a *transvection* if it preserves a geodesic in X and acts trivially on its normal bundle. *Transvections* in X are precisely the compositions of pairs of Cartan involutions.

A *flat* in X is an isometrically embedded totally geodesic Euclidean subspace of X. The dimension of a *maximal flat* in X is the *rank* of X. We will use the notation F for maximal flats in X. The *parallel set* $P(l)$ of a geodesic l in X is the union of all maximal flats in X containing l. We let $\tau_\pm$ denote the smallest simplices in $\partial_\infty X$ containing the elements $\xi_\pm \in \partial_\infty l$; these simplices are antipodal. We will use the notation $P(\tau_-, \tau_+)$ for the parallel set $P(l)$, since $P(l)$ can be described as the union of geodesics in X that are forward/backward asymptotic to points in the open simplices $\operatorname{int}\tau_\pm$.

Given a maximal flat $F \subset X$, the stabilizer G_F of F in G acts transitively on F. For each $x \in F$ the intersection $G_x \cap G_F$ acts on F as a finite reflection group, isomorphic to the Weyl group W of X. One frequently fixes a basepoint $x = o \in X$ and the *model maximal flat* $F_{\mathrm{mod}} \subset X$ containing o; the stabilizer G_x is then denoted by K, it is a maximal compact subgroup of G. A fundamental domain for the W-action on F_{mod} is the *model Euclidean Weyl chamber*, denoted by Δ. The ideal boundary $\partial_\infty \Delta$ is then identified with σ_{mod}, the model chamber in $\partial_\infty X$.

The Euclidean Weyl chamber Δ is the Euclidean cone over the simplex σ_{mod}. Thus, we can also "cone-off" various objects from σ_{mod}. In particular, we define the τ_{mod}-*boundary* $\partial_{\tau_{\mathrm{mod}}}\Delta$ of Δ as the union of rays $o\zeta$ with $\zeta \in \partial_{\tau_{\mathrm{mod}}}\sigma_{\mathrm{mod}}$ (see [Kapovich et al. 2017, Section 2.5.2]), and define the Θ-cone Δ_Θ as the union of rays $o\zeta$ with $\zeta \in \Theta$, where $\Theta \subset \operatorname{ost}\tau_{\mathrm{mod}} \subset \sigma_{\mathrm{mod}}$.

The group of transvections along geodesics in F is usually denoted by A; then (in the case $G < \mathrm{Isom}(X)$) $G = KAK$ is the *Cartan decomposition* of G. The more refined form of this decomposition is $G = KA_+K$, where $A_+ \subset A$ is the subsemigroup consisting of transvections mapping Δ into itself. Geometrically speaking, the Cartan decomposition states that each K-orbit $Ky \subset X$ intersects Δ in exactly one point. The projection $c\colon y \mapsto Ky \cap \Delta$ is 1-Lipschitz since each orbit Ky meets F_{mod} orthogonally and transversally. This projection leads to the notion of Δ-distance on X (see [Kapovich et al. 2017, Section 2.6] and [Kapovich et al. 2009]): Given a pair of points $x, y \in X$, find $g \in G$ such that $g(x) = o$ and $z = g(y) \in \Delta$. Then

$$\overrightarrow{oz} := d_\Delta(x, y).$$

The vector $d_\Delta(x, y)$ is the complete G-congruence invariant of pairs $(x, y) \in X^2$, and

$$d(x, y) = \|d_\Delta(x, y)\|,$$

where $\|\cdot\|$ is the Euclidean norm on F_{mod} (regarded as a Euclidean vector space with o serving as zero). Since c is 1-Lipschitz, the Δ-distance satisfies the triangle inequality

$$\|d_\Delta(x, y) - d_\Delta(x, z)\| \le \|d_\Delta(y, z)\| = d(y, z).$$

We refer the reader to [Kapovich et al. 2009] for in-depth discussion of *generalized triangle inequalities* satisfied by the Δ-valued distance function on X.

We say that a nondegenerate segment $xy \subset X$ is τ_{mod}-regular if

$$d_\Delta(x, y) \notin \partial_{\tau_{\mathrm{mod}}} \Delta,$$

and that it is Θ-regular if

$$d_\Delta(x, y) \in \Delta_\Theta.$$

See [Kapovich et al. 2017, Section 2.5.3]. In what follows, we will always assume that τ_{mod} is ι-invariant and $\Theta \subset \mathrm{ost}(\tau_{\mathrm{mod}})$ is an ι-invariant Weyl-convex compact subset of $\mathrm{ost}(\tau_{\mathrm{mod}}) \subset \sigma_{\mathrm{mod}}$.

An injective map f from an interval I in $\mathbb{R}$ to X is called τ_{mod}-*regular* if for all $s < t$ in I the geodesic segment $f(s)f(t)$ is τ_{mod}-regular.

Weyl cones Fix a simplex τ in $\partial_\infty X$ and a point $x \in X$. We define the *Weyl cone* $V(x, \mathrm{st}(\tau))$ with the tip x over the star $\mathrm{st}(\tau) \subset \partial_\infty X$ as the union of geodesic rays emanating from x and asymptotic to points of $\mathrm{st}(\tau)$. Similarly, assuming that τ has the type τ_{mod} and $\Theta \subset \mathrm{ost}(\tau_{\mathrm{mod}}) \subset \sigma_{\mathrm{mod}}$ is a Weyl-convex compact subset, we define the Θ-cone $V(x, \mathrm{st}_\Theta(\tau))$ with the tip x over the Θ-star $\mathrm{st}_\Theta(\tau) \subset \partial_\infty X$ as the union of geodesic rays emanating from x and asymptotic to points of $\mathrm{st}_\Theta(\tau)$. It was proven in [Kapovich et al. 2017, Section 2.5] that Weyl cones $V(x, \mathrm{st}(\tau))$ and Θ-cones $V(x, \mathrm{st}_\Theta(\tau))$ are convex in X.

In particular, these cones satisfy the *nested cones property*:

(1) If $y \in V(x, \mathrm{st}(\tau))$, then $V(y, \mathrm{st}(\tau)) \subset V(x, \mathrm{st}(\tau))$.

(2) If $y \in V(x, \mathrm{st}_\Theta(\tau))$, then $V(y, \mathrm{st}_\Theta(\tau)) \subset V(x, \mathrm{st}_\Theta(\tau))$.

Finsler geodesics in X These are geodesics with respect to certain G-invariant Finsler metrics on X defined in [Kapovich and Leeb 2018b]. The precise description of all Finsler geodesics is given in [Kapovich and Leeb 2018b, Section 5.1.3]. We merely use the following description of Finsler geodesics in lieu of the definition:

Definition B.1 (Finsler geodesics) Fix an ι-invariant face τ_{mod} of σ_{mod}. Let $I \subset \mathbb{R}$ be an interval. A (continuous) path $c\colon I \to X$ is called a *Finsler geodesic* (more precisely, a τ_{mod}-Finsler geodesic) if there exists a pair of antipodal flags $\tau_\pm \in \mathrm{Flag}(\tau_{\mathrm{mod}})$ such that $c(I) \subset P(\tau_+, \tau_-)$ and

$$c(t_2) \in V(c(t_1), \mathrm{st}(\tau_+)) \quad \text{for all } t_1 < t_2.$$

Moreover, given an ι-invariant compact Weyl-convex subset $\Theta \subset \mathrm{ost}(\tau_{\mathrm{mod}})$, a Finsler geodesic $c\colon I \to X$ is called a Θ-*Finsler geodesic* if, in addition to the above, it satisfies the following stronger condition:

$$c(t_2) \in V(c(t_1), \mathrm{st}_\Theta(\tau_+)) \quad \text{for all } t_1 < t_2,$$

ie c is Θ-regular.

Note that each τ_{mod}-Finsler geodesic is automatically a τ_{mod}-regular path in X.

Diamonds Intersecting cones, we define *diamonds* in X; see [Kapovich et al. 2017, Section 2.5]. Take two antipodal simplices τ_+, τ_- of the type $\tau_{\mathrm{mod}} = \iota\tau_{\mathrm{mod}}$ and points $x_\pm \in X$ such that

$$x_+ \in V(x_-, \mathrm{st}(\tau_+)), x_- \in V(x_+, \mathrm{st}(\tau_-)).$$

Then the intersection

$$\diamondsuit_{\tau_{\mathrm{mod}}}(x_-, x_+) = V(x_-, \mathrm{st}(\tau_+)) \cap V(x_+, \mathrm{st}_\Theta(\tau_-))$$

is the τ_{mod}-diamond with tips $x_\pm$. Similarly, suppose that $\Theta \subset \mathrm{ost}(\tau_{\mathrm{mod}}) \subset \sigma_{\mathrm{mod}}$ is a Weyl-convex ι-invariant compact subset and

$$x_+ \in V(x_-, \mathrm{st}_\Theta(\tau_+)), \quad x_- \in V(x_+, \mathrm{st}_\Theta(\tau_-)).$$

Then the Θ-diamond is defined as the intersection

$$\diamondsuit_\Theta(x_-, x_+) = V(x_-, \mathrm{st}_\Theta(\tau_+)) \cap V(x_+, \mathrm{st}_\Theta(\tau_-)).$$

Thus, diamonds are also convex in X.

Diamonds have a nice interpretation in terms of Finsler geometry of X: it is proven in [Kapovich and Leeb 2018b] that $\diamondsuit_{\tau_{\mathrm{mod}}}(x_-, x_+)$ is the union of all Finsler geodesics in X connecting the points $x_\pm$. Similarly, $\diamondsuit_\Theta(x_-, x_+)$ is the union of all Θ-regular Finsler geodesics in X connecting the points $x_\pm$.

Longitudinal notions The following definition is taken from [Kapovich et al. 2018b, Section 3.2]. Let $P = P(\tau_-, \tau_+) \subset X$ be a parallel set, where $\tau_\pm$ are antipodal simplices of type τ_{mod} in $\partial_{\mathrm{Tits}} X$.

Definition B.2 (longitudinal directions and segments in parallel sets) A nondegenerate geodesic segment xy in P is said to be *longitudinal* or, more precisely, τ_+-longitudinal, if it is contained in a geodesic ray $x\xi \subset P$, where $\xi \in \mathrm{ost}(\tau_+)$.

Remark B.3 (1) Longitudinal directions and segments are, in particular, τ_{mod}-regular.

(2) A nondegenerate segment xy is longitudinal if and only if all nondegenerate subsegments xy are longitudinal.

(3) If xy and yz are τ_+-longitudinal Θ-regular segments in P, then so is the segment xz. This follows, for instance, from the fact that

$$V(y, \mathrm{st}_\Theta(\tau_+)) \subset V(x, \mathrm{st}_\Theta(\tau_+)).$$

Regular sequences and groups The reader should think of *regularity* conditions for subgroups $\Gamma < G$ as a way of strengthening the discreteness assumption: The discreteness condition means that the sequence of distances $d(x, \gamma_n(x))$ diverges to infinity for every sequence of distinct elements $g_n \in \Gamma$. Regularity conditions on Γ require certain forms of divergence to infinity of vector-valued distances $d_\Delta(x, g_n x)$.

We first consider sequences in the euclidean model Weyl chamber Δ. Recall that

$$\partial_{\tau_{\mathrm{mod}}}\Delta = V(0, \partial_{\tau_{\mathrm{mod}}}\sigma_{\mathrm{mod}}) \subset \Delta$$

is the union of faces of Δ which do not contain the sector $V(0, \tau_{\mathrm{mod}})$. Note that

$$\partial_{\tau_{\mathrm{mod}}}\Delta \cap V(0, \tau_{\mathrm{mod}}) = \partial V(0, \tau_{\mathrm{mod}}) = V(0, \partial\tau_{\mathrm{mod}}).$$

The following definitions are taken from [Kapovich et al. 2017, Section 4.2].

Definition B.4 (i) A sequence (δ_n) in Δ is called τ_{mod}-*regular* if it drifts away from $\partial_{\tau_{\mathrm{mod}}}\Delta$, ie

$$d(\delta_n, \partial_{\tau_{\mathrm{mod}}}\Delta) \to +\infty.$$

(ii) A sequence (x_n) in X is τ_{mod}-*regular* if for some (any) basepoint $o \in X$ the sequence of Δ-distances $d_\Delta(o, x_n)$ in Δ has this property.

(iii) A sequence (g_n) in G is τ_{mod}-*regular*, if for some (any) point $x \in X$ the orbit sequence $(g_n x)$ in X has this property.

(iv) A subgroup $\Gamma < G$ is τ_{mod}-*regular* if all sequences of distinct elements in Γ have this property.

Next, we describe a stronger form of regularity following [Kapovich et al. 2017, Section 4.6].

Definition B.5 (i) A sequence $\delta_n \to \infty$ in Δ is *uniformly* τ_{mod}-*regular* if it drifts away from $\partial_{\tau_{\mathrm{mod}}}\Delta$ at a linear rate with respect to its norm,

$$\liminf_{n\to+\infty} \frac{d(\delta_n, \partial_{\tau_{\mathrm{mod}}}\Delta)}{\|\delta_n\|} > 0.$$

(ii) A sequence (x_n) in X is *uniformly* τ_{mod}-*regular* if for some (any) basepoint $o \in X$ the sequence of Δ-distances $d_\Delta(o, x_n)$ in Δ has this property.

(iii) A sequence (g_n) in G is *uniformly* τ_{mod}-*regular* if for some (any) point $x \in X$ the orbit sequence $(g_n x)$ in X has this property.

(iv) A subgroup $\Gamma < G$ is *uniformly* τ_{mod}-*regular* if all sequences of distinct elements in Γ have this property.

Note that (uniform) regularity of a sequence in X is independent of the basepoint and stable under bounded perturbation of the sequence (due to the triangle inequality for Δ-distances). A sequence (x_n) is uniformly τ_{mod}-regular if and only if there exists a compact $\Theta \subset \mathrm{ost}(\tau_{\mathrm{mod}})$ such that for each $x \in X$ all but finitely many vectors $d_\Delta(x, x_n)$ belong to Δ_Θ.

Remark B.6 The definition of regularity of sequences in G has the following dynamical interpretation, in terms of dynamics on the flag-manifold $\mathrm{Flag}(\tau_{\mathrm{mod}})$, which generalizes the familiar convergence property for sequences of isometries of a Gromov-hyperbolic space Y acting on the visual boundary of Y.

A sequence (g_n) in G is said to be τ_{mod}-*contracting* if there exists a pair of elements $\tau_\pm \in \mathrm{Flag}(\tau_{\mathrm{mod}})$ such that the sequence (g_n) converges to τ_+ uniformly on compacts in $C(\tau_-) \subset \mathrm{Flag}(\tau_{\mathrm{mod}})$. Here, as elsewhere, $C(\tau_-)$ is the open Schubert cell in $\mathrm{Flag}(\tau_{\mathrm{mod}})$ consisting of simplices antipodal to τ_-. In this situation, the simplex τ_+ is the τ_{mod}-limit of (g_n). A sequence (g_n) is τ_{mod}-regular if and only if there exists a pair of bounded sequences $a_n, b_n \in G$ such that the sequence of compositions $c_n g_n b_n$ is τ_{mod}-contracting. Equivalently, a sequence (g_n) is τ_{mod}-regular if and only if every subsequence in (g_n) contains a further subsequence which is τ_{mod}-contracting. We refer the reader to [Kapovich et al. 2017] for details.

For a subgroup $\Gamma < G$, uniform τ_{mod}-regularity is equivalent to the *visual limit set* $\Lambda \subset \partial_\infty X$ being contained in the union of the open τ_{mod}-stars. Here Λ is the accumulation set in $\partial_\infty X$ of some (any) Γ-orbit in X.

We next discuss τ_{mod}-limit sets. In a nutshell, the τ_{mod}-limit set of a τ_{mod}-regular subgroup $\Gamma < G$ is the accumulation set in the flag-manifold $\mathrm{Flag}(\tau_{\mathrm{mod}})$ of one (equivalently, every) Γ-orbit in X. However, the way the accumulation is defined is more subtle than in the rank-one case. If G (equivalently, X) has rank one, one can define convergence of a divergent sequence $x_n \in X$ to a point $\lambda \in \partial_\infty X$ by taking a sequence of geodesic rays $o\xi_n$ through x_n (with the fixed initial point o). Then $x_n \to \lambda$ if and only if $\xi_n \to \lambda$ in $\mathrm{Flag}(\tau_{\mathrm{mod}}) = \partial_\infty X$. In higher rank, a geodesic $o\xi_n$ through x_n need not even terminate in a face τ_n of $\partial_{\mathrm{Tits}} X$ of type τ_{mod}. But, if it does, then $x_n \to \tau \in \mathrm{Flag}(\tau_{\mathrm{mod}})$ if and only if $\tau_n \to \tau$ in $\mathrm{Flag}(\tau_{\mathrm{mod}})$. In general, due to the τ_{mod}-regularity assumption, one finds (for all sufficiently large n) a unique face τ_n of type τ_{mod} in $\mathrm{Flag}(\tau_{\mathrm{mod}})$ such that ξ_n belongs to the open star $\mathrm{ost}(\tau_n)$ of τ_n. Then, by the definition, $x_n \to \tau$ if and only if $\tau_n \to \tau$ in $\mathrm{Flag}(\tau_{\mathrm{mod}})$. Lastly, if $\Gamma < G$ is a τ_{mod}-regular subgroup, then the τ_{mod}-limit set $\Lambda_{\tau_{\mathrm{mod}}}(\Gamma) \subset \mathrm{Flag}(\tau_{\mathrm{mod}})$ is the subset consisting of accumulation points in $\mathrm{Flag}(\tau_{\mathrm{mod}})$ of one (equivalently, every) Γ-orbit in X, where convergence is understood as above.

A subset Λ of $\mathrm{Flag}(\tau_{\mathrm{mod}})$ is said to be *antipodal* if any two distinct elements of Λ are antipodal to each other. A map $\beta\colon Z \to \mathrm{Flag}(\tau_{\mathrm{mod}})$ is said to be *antipodal* if it sends distinct elements of Z to antipodal elements of $\mathrm{Flag}(\tau_{\mathrm{mod}})$. We will apply these notions to the limit sets of τ_{mod}-regular subgroups of G.

Morse quasigeodesics, maps and subgroups Lastly, we come to the heart of the matter, the notion of *Morse quasigeodesics* in X (Definition 3.1), *Morse embeddings* (or *Morse maps*) to X (Definition 3.7) and *Morse subgroups* of G and *Morse actions* on X (Definition 4.1).

Below we recall an alternative characterization of Morse quasigeodesics given in [Kapovich et al. 2018b]. A natural way to *coarsify* the notion of regularity for geodesic segments in X is as follows.

Let $B \geq 0$. We say that a *pair* (x, y) of (not necessarily distinct) points is (Θ, B)-*regular* if it is oriented B-close to some Θ-regular pair of points (x', y'), ie $d(x, x') \leq B$ and $d(y, y') \leq B$. This is equivalent to the property that the segment xy is oriented B-Hausdorff close to the Θ-regular segment $x'y'$, and we say also that the *segment* xy is (Θ, B)-regular.

We say that a (not necessarily continuous) *path* $p\colon I \to X$ is (Θ, B)-regular if for every subinterval $[a', b'] \subset I$, the segment $p(a')p(b')$ is (Θ, B)-regular.

If τ_{mod} and Θ are ι-invariant, then we say that a *subset* of X is (Θ, B)-regular if every pair of points in the subset has this property, and more generally, that a *map* into X is (Θ, B)-regular if it sends any pair of points to a (Θ, B)-regular pair of points in X. Note that the images of (Θ, B)-regular maps are (Θ, B)-regular subsets. We say that the subset or map is (*coarsely*) Θ-*regular* if it is (Θ, B)-regular for some constant B.

In [Kapovich et al. 2018b] we prove:

Theorem B.7 (regular implies Morse for quasigeodesics) *Uniformly coarsely τ_{mod}-regular quasigeodesics in X are uniform τ_{mod}-Morse quasigeodesics and vice-versa.*

Note that τ_{mod}-Morse embeddings are coarsely uniformly τ_{mod}-regular quasiisometric embeddings. We also have the converse, proven in [Kapovich et al. 2018b]:

Theorem B.8 (regular implies Morse for quasiisometric embeddings) *Coarsely uniformly τ_{mod}-regular quasiisometric embeddings from quasigeodesic metric spaces into model spaces are uniform τ_{mod}-Morse embeddings.*

Morse actions on X are *undistorted* in the sense that the orbit maps are quasiisometric embeddings. In particular, they are properly discontinuous. Furthermore, Θ-Morse actions are (coarsely) Θ-regular. In [Kapovich et al. 2018b] we prove a converse:

Theorem B.9 (URU implies Morse) *Uniformly τ_{mod}-regular undistorted isometric actions by finitely generated groups on X are uniformly τ_{mod}-Morse.*

Below we collect some of the equivalent characterizations of τ_{mod}-Morse (equivalently, τ_{mod}-Anosov) subgroups of G given in [Kapovich et al. 2017; 2018b] and [Kapovich and Leeb 2018b]:

Theorem B.10 *The following are equivalent for a subgroup $\Gamma < G$:*

(1) *Γ is τ_{mod}-Morse.*

(2) *Γ is τ_{mod}-URU, ie it is τ_{mod}-uniformly regular, finitely generated and one (equivalently, every) orbit map $\Gamma \to X$ is a quasiisometric embedding.*

(3) *Γ is τ_{mod}-asymptotically embedded, ie Γ is τ_{mod}-regular, Gromov-hyperbolic with the Gromov-compactification $\overline{\Gamma} = \Gamma \cup \partial_\infty \Gamma$, and there exists an antipodal homeomorphism $\beta : \partial_\infty \Gamma \to \Lambda_{\tau_{\mathrm{mod}}}(\Gamma)$ which, combined with one (equivalently, every) orbit map $\Gamma \to X$ defines a continuous map $\overline{\Gamma} \to X \cup \mathrm{Flag}(\tau_{\mathrm{mod}})$, where the latter is topologized using the natural topologies of X and $\mathrm{Flag}(\tau_{\mathrm{mod}})$, and the topology of τ_{mod}-flag convergence for sequences in X to elements of $\mathrm{Flag}(\tau_{\mathrm{mod}})$.*

(4) *Γ is τ_{mod}-RCA (regular, conical and antipodal), ie Γ is τ_{mod}-regular, its τ_{mod}-limit set is antipodal and every element τ of $\Lambda_{\tau_{\mathrm{mod}}}(\Gamma)$ is conical. The latter means that for some (equivalently, every) $x \in X$ there exists $C < \infty$ and an infinite sequence γ_n in Γ such that the sequence $(\gamma_n(x))$ flag-converges to τ in the C-neighborhood of the Weyl cone $V(x, \mathrm{st}(\tau))$.*

The map β in (3) is called the *boundary map* of Γ. It is unique as long as Γ is a nonelementary hyperbolic group.

References

[Ballmann et al. 1985] **W Ballmann**, **M Gromov**, **V Schroeder**, *Manifolds of nonpositive curvature*, Progr. Math. 61, Birkhäuser, Boston, MA (1985) MR

[Baumslag and Roseblade 1984] **G Baumslag**, **J E Roseblade**, *Subgroups of direct products of free groups*, J. London Math. Soc. 30 (1984) 44–52 MR

[Benoist 1997] **Y Benoist**, *Propriétés asymptotiques des groupes linéaires*, Geom. Funct. Anal. 7 (1997) 1–47 MR

[Bowditch 1998] **B H Bowditch**, *Spaces of geometrically finite representations*, Ann. Acad. Sci. Fenn. Math. 23 (1998) 389–414 MR

[Brown 1982] **K S Brown**, *Cohomology of groups*, Graduate Texts in Math. 87, Springer (1982) MR

[Canary et al. 1987] **R D Canary**, **D B A Epstein**, **P Green**, *Notes on notes of Thurston*, from "Analytical and geometric aspects of hyperbolic space" (D B A Epstein, editor), London Math. Soc. Lecture Note Ser. 111, Cambridge Univ. Press (1987) 3–92 MR

[Coornaert et al. 1990] **M Coornaert**, **T Delzant**, **A Papadopoulos**, *Géométrie et théorie des groupes: les groupes hyperboliques de Gromov*, Lecture Notes in Math. 1441, Springer (1990) MR

[Corlette 1990] **K Corlette**, *Hausdorff dimensions of limit sets, I*, Invent. Math. 102 (1990) 521–541 MR

[Dey and Kapovich 2023] **S Dey**, **M Kapovich**, *Klein–Maskit combination theorem for Anosov subgroups: free products*, Math. Z. 305 (2023) art. id. 35 MR

[Dey and Kapovich 2025] **S Dey**, **M Kapovich**, *Klein–Maskit combination theorem for Anosov subgroups: amalgams*, J. Reine Angew. Math. 819 (2025) 1–43 MR

[Dey et al. 2019] **S Dey**, **M Kapovich**, **B Leeb**, *A combination theorem for Anosov subgroups*, Math. Z. 293 (2019) 551–578 MR

[Druțu and Kapovich 2018] **C Druțu**, **M Kapovich**, *Geometric group theory*, American Mathematical Society Colloquium Publications 63, Amer. Math. Soc., Providence, RI (2018) MR

[Eberlein 1996] **P B Eberlein**, *Geometry of nonpositively curved manifolds*, Univ. Chicago Press (1996) MR

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

[Gilman 1995] **J Gilman**, *Two-generator discrete subgroups of* $\mathrm{PSL}(2,\mathbb{R})$, Mem. Amer. Math. Soc. 561, Amer. Math. Soc., Providence, RI (1995) MR

[Gilman 1997] **J Gilman**, *Algorithms, complexity and discreteness criteria in* $\mathrm{PSL}(2,\mathbb{C})$, J. Anal. Math. 73 (1997) 91–114 MR

[Gilman and Keen 2016] **J Gilman**, **L Keen**, *Canonical hexagons and the* $\mathrm{PSL}(2,\mathbb{C})$ *discreteness problem*, preprint (2016) arXiv 1508.00257v3

[Gilman and Maskit 1991] **J Gilman**, **B Maskit**, *An algorithm for 2-generator Fuchsian groups*, Michigan Math. J. 38 (1991) 13–32 MR

[Goldman 1988] **W M Goldman**, *Geometric structures on manifolds and varieties of representations*, from "Geometry of group representations" (W M Goldman, A R Magid, editors), Contemp. Math. 74, Amer. Math. Soc., Providence, RI (1988) 169–198 MR

[Guichard and Wienhard 2012] **O Guichard**, **A Wienhard**, *Anosov representations: domains of discontinuity and applications*, Invent. Math. 190 (2012) 357–438 MR

[Izeki 2000] **H Izeki**, *Quasiconformal stability of Kleinian groups and an embedding of a space of flat conformal structures*, Conform. Geom. Dyn. 4 (2000) 108–119 MR

[Kapovich 2001] **M Kapovich**, *Hyperbolic manifolds and discrete groups*, Progr. Math. 183, Birkhäuser, Boston, MA (2001) MR

[Kapovich 2008] **M Kapovich**, *Kleinian groups in higher dimensions*, from "Geometry and dynamics of groups and spaces" (M Kapranov, S Kolyada, Y I Manin, P Moree, L Potyagailo, editors), Progr. Math. 265, Birkhäuser, Boston, MA (2008) 487–564 MR

[Kapovich 2016] **M Kapovich**, *Discreteness is undecidable*, Internat. J. Algebra Comput. 26 (2016) 467–472 MR

[Kapovich and Leeb 2018a] **M Kapovich**, **B Leeb**, *Discrete isometry groups of symmetric spaces*, from "Handbook of group actions, IV" (L Ji, A Papadopoulos, S-T Yau, editors), Adv. Lect. Math. 41, International, Somerville, MA (2018) 191–290 MR

[Kapovich and Leeb 2018b] **M Kapovich**, **B Leeb**, *Finsler bordifications of symmetric and certain locally symmetric spaces*, Geom. Topol. 22 (2018) 2533–2646 MR

[Kapovich and Leeb 2023] **M Kapovich**, **B Leeb**, *Relativizing characterizations of Anosov subgroups, I*, Groups Geom. Dyn. 17 (2023) 1005–1071 MR

[Kapovich and Leeb $\geq$ 2025] **M Kapovich**, **B Leeb**, *Relativizing characterizations of Anosov subgroups*, *II*, in preparation

[Kapovich and Liu 2019] **M Kapovich**, **B Liu**, *Geometric finiteness in negatively pinched Hadamard manifolds*, Ann. Acad. Sci. Fenn. Math. 44 (2019) 841–875 MR

[Kapovich et al. 2009] **M Kapovich**, **B Leeb**, **J Millson**, *Convex functions on symmetric spaces, side lengths of polygons and the stability inequalities for weighted configurations at infinity*, J. Differential Geom. 81 (2009) 297–354 MR

[Kapovich et al. 2014] **M Kapovich**, **B Leeb**, **J Porti**, *Morse actions of discrete groups on symmetric space*, preprint (2014) arXiv 1403.7671

[Kapovich et al. 2017] **M Kapovich**, **B Leeb**, **J Porti**, *Anosov subgroups: dynamical and geometric characterizations*, Eur. J. Math. 3 (2017) 808–898 MR

[Kapovich et al. 2018a] **M Kapovich**, **B Leeb**, **J Porti**, *Dynamics on flag manifolds: domains of proper discontinuity and cocompactness*, Geom. Topol. 22 (2018) 157–234 MR

[Kapovich et al. 2018b] **M Kapovich**, **B Leeb**, **J Porti**, *A Morse lemma for quasigeodesics in symmetric spaces and Euclidean buildings*, Geom. Topol. 22 (2018) 3827–3923 MR

[Kapovich et al. 2024] **M Kapovich**, **S Kim**, **J Lee**, *Structural stability of meandering-hyperbolic group actions*, J. Inst. Math. Jussieu 23 (2024) 753–810 MR

[Labourie 2006] **F Labourie**, *Anosov flows, surface groups and curves in projective space*, Invent. Math. 165 (2006) 51–114 MR

[Lok 1984] **W L Lok**, *Deformations of locally homogeneous spaces and Kleinian groups*, PhD thesis, Columbia Univ. (1984) Available at `https://www.proquest.com/docview/303285563`

[Ol'shanskii and Sapir 2001] **A Y Ol'shanskii**, **M V Sapir**, *Length and area functions on groups and quasi-isometric Higman embeddings*, Internat. J. Algebra Comput. 11 (2001) 137–170 MR

[Raghunathan 1972] **M S Raghunathan**, *Discrete subgroups of Lie groups*, Ergebnisse der Math. 68, Springer (1972) MR

[Riestenberg 2025] **J M Riestenberg**, *A quantified local-to-global principle for Morse quasigeodesics*, Groups Geom. Dyn. 19 (2025) 37–107 MR

[Riley 1983] **R Riley**, *Applications of a computer implementation of Poincaré's theorem on fundamental polyhedra*, Math. Comp. 40 (1983) 607–632 MR

[Sullivan 1985] **D Sullivan**, *Quasiconformal homeomorphisms and dynamics, II: Structural stability implies hyperbolicity for Kleinian groups*, Acta Math. 155 (1985) 243–260 MR

[Tukia 1985] **P Tukia**, *On isomorphisms of geometrically finite Möbius groups*, Inst. Hautes Études Sci. Publ. Math. 61 (1985) 171–214 MR

Department of Mathematics, University of California, Davis
Davis, CA, United States

Mathematisches Institut, Universität München
München, Germany

Departament de Matemàtiques and Centre de Recerca Matemàtica, Universitat Autònoma de Barcelona
Barcelona, Spain

kapovich@math.ucdavis.edu, b.l@lmu.de, joan.porti@uab.cat

Proposed: Bruce Kleiner
Seconded: Tobias H Colding, David Fisher
Received: 9 January 2023
Revised: 26 May 2024

GEOMETRY & TOPOLOGY

msp.org/gt

MANAGING EDITORS

Robert Lipshitz	University of Oregon lipshitz@uoregon.edu
András I Stipsicz	Alfréd Rényi Institute of Mathematics stipsicz@renyi.hu

BOARD OF EDITORS

Mohammed Abouzaid	Stanford University abouzaid@stanford.edu
Dan Abramovich	Brown University dan_abramovich@brown.edu
Ian Agol	University of California, Berkeley ianagol@math.berkeley.edu
Arend Bayer	University of Edinburgh arend.bayer@ed.ac.uk
Mark Behrens	University of Notre Dame mbehren1@nd.edu
Mladen Bestvina	University of Utah bestvina@math.utah.edu
Martin R Bridson	University of Oxford bridson@maths.ox.ac.uk
Jim Bryan	University of British Columbia jbryan@math.ubc.ca
Dmitri Burago	Pennsylvania State University burago@math.psu.edu
Tobias H Colding	Massachusetts Institute of Technology colding@math.mit.edu
Simon Donaldson	Imperial College, London s.donaldson@ic.ac.uk
Yasha Eliashberg	Stanford University eliash-gt@math.stanford.edu
Benson Farb	University of Chicago farb@math.uchicago.edu
David M Fisher	Rice University davidfisher@rice.edu
Mike Freedman	Microsoft Research michaelf@microsoft.com
David Gabai	Princeton University gabai@princeton.edu
Stavros Garoufalidis	Southern U. of Sci. and Tech., China stavros@mpim-bonn.mpg.de
Cameron Gordon	University of Texas gordon@math.utexas.edu
Jesper Grodal	University of Copenhagen jg@math.ku.dk
Misha Gromov	IHÉS and NYU, Courant Institute gromov@ihes.fr
Mark Gross	University of Cambridge mgross@dpmms.cam.ac.uk
Rob Kirby	University of California, Berkeley kirby@math.berkeley.edu
Bruce Kleiner	NYU, Courant Institute bkleiner@cims.nyu.edu
Sándor Kovács	University of Washington skovacs@uw.edu
Urs Lang	ETH Zürich urs.lang@math.ethz.ch
Marc Levine	Universität Duisburg-Essen marc.levine@uni-due.de
Ciprian Manolescu	University of California, Los Angeles cm@math.ucla.edu
Haynes Miller	Massachusetts Institute of Technology hrm@math.mit.edu
Tomasz Mrowka	Massachusetts Institute of Technology mrowka@math.mit.edu
Aaron Naber	Northwestern University anaber@math.northwestern.edu
Peter Ozsváth	Princeton University petero@math.princeton.edu
Leonid Polterovich	Tel Aviv University polterov@post.tau.ac.il
Colin Rourke	University of Warwick gt@maths.warwick.ac.uk
Roman Sauer	Karlsruhe Institute of Technology roman.sauer@kit.edu
Stefan Schwede	Universität Bonn schwede@math.uni-bonn.de
Natasa Sesum	Rutgers University natasas@math.rutgers.edu
Gang Tian	Massachusetts Institute of Technology tian@math.mit.edu
Ulrike Tillmann	Oxford University tillmann@maths.ox.ac.uk
Nathalie Wahl	University of Copenhagen wahl@math.ku.dk
Anna Wienhard	Universität Heidelberg wienhard@mathi.uni-heidelberg.de

See inside back cover or msp.org/gt for submission instructions.

The subscription price for 2025 is US $865/year for the electronic version, and $1210/year (+$75, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Geometry & Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Geometry & Topology (ISSN 1465-3060 printed, 1364-0380 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

GT peer review and production are managed by EditFLOW® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

http://msp.org/

© 2025 Mathematical Sciences Publishers

GEOMETRY & TOPOLOGY

Volume 29 Issue 5 (pages 2251–2782) 2025

Quantitative Thomas–Yau uniqueness 2251

YANG LI

Knot surgery formulae for instanton Floer homology, I: The main theorem 2269

ZHENKUN LI and FAN YE

Morse actions of discrete groups on symmetric spaces: local-to-global principle 2343

MICHAEL KAPOVICH, BERNHARD LEEB and JOAN PORTI

Nearly geodesic immersions and domains of discontinuity 2391

COLIN DAVALO

Graded character sheaves, HOMFLY-PT homology, and Hilbert schemes of points on $\mathbb{C}^2$ 2463

QUOC P HO and PENGHUI LI

Asymptotically Calabi metrics and weak Fano manifolds 2547

HANS-JOACHIM HEIN, SONG SUN, JEFFREY VIACLOVSKY and RUOBING ZHANG

The distribution of critical graphs of Jenkins–Strebel differentials 2571

FRANCISCO ARANA-HERRERA and AARON CALDERON

Tian's stabilization problem for toric Fanos 2609

CHENZI JIN and YANIR A RUBINSTEIN

On the logarithmic slice filtration 2653

FEDERICO BINDA, DOOSUNG PARK and PAUL ARNE ØSTVÆR

Joyce structures on spaces of quadratic differentials 2695

TOM BRIDGELAND

Big monodromy for higher Prym representations 2733

AARON LANDESMAN, DANIEL LITT and WILL SAWIN