Hitchin grafting representations I: Geometry

Pierre-Louis Blayac, Ursula Hamenstädt, and Théo Marty

July 9, 2025

Abstract

We give a geometric interpretation of Fock–Goncharov positivity and show that bending deformations of Fuchsian representations stabilize a uniform Finsler quasi-convex disk in the symmetric space $\mathrm{PSL}_d(\mathbb{R})/\mathrm{PSO}(d)$.

Contents

Introduction

The *Teichmüller space* $\mathcal{T}(S)$ of a closed oriented surface S of genus $g \geq 2$ is the space of *marked* hyperbolic structures on S, homeomorphic to $\mathbb{R}^{6g-6}$. Equivalently, it can be described as a distinguished component of the space of conjugacy classes of homomorphisms $\pi_1(S) \to \mathrm{PSL}_2(\mathbb{R})$, with target the group $\mathrm{PSL}_2(\mathbb{R})$ of orientation preserving isometries of the hyperbolic plane. It was discovered by Hitchin [Hit92] that an analog of the Teichmüller space also exists for conjugacy classes of representations of $\pi_1(S)$ into simple split real Lie groups of higher rank, which is also homeomorphic to a Euclidean space.

The so-called *Hitchin component* $\mathrm{Hit}(S)$ for the target group $\mathrm{PSL}_d(\mathbb{R})$ ($d \geq 3$) is the component of the *character variety* containing conjugacy classes of so-called *Fuchsian representations*, namely discrete representations which factor through an irreducible embedding $\mathrm{PSL}_2(\mathbb{R}) \to \mathrm{PSL}_d(\mathbb{R})$. Hitchin [Hit92] showed that the Hitchin component is homeomorphic to $\mathbb{R}^m$ for $m = (2g-2)\dim(\mathrm{PSL}_d(\mathbb{R}))$, and later Labourie [Lab06] and Fock–Goncharov [FG06] independently proved that all representations in the Hitchin component are faithful with discrete image. Thus each point in the Hitchin component defines a locally symmetric manifold with fundamental group isomorphic to $\pi_1(S)$. The mapping class group $\mathrm{Mod}(S)$ of S of isotopy classes of diffeomorphisms of S acts properly by precomposition on the Hitchin component.

In this article we study the geometry of a distinguished class of these locally symmetric manifolds, obtained by deforming Fuchsian representations via the famous *bending procedure* described below, which is in our opinion the simplest kind of deformation. We will see that the geometry of such manifolds is governed by the geometry of an explicit piecewise totally geodesic embedded subsurface. This subsurface will be constructed by grafting flat cylinders into a hyperbolic surface. In a sequel to this work [BHMM24], this will be used to investigate the so-called pressure metric on the Hitchin component defined by Bridgeman–Canary–Labourie–Sambarino [BCLS15].

Given a hyperbolic metric X on S and a simple closed geodesic $\gamma \subset X$, one can define two types of deformations of X. The first consists in shrinking the length of γ (this depends on the choice of a pair of pant in S on each side of γ). Another deformation consists in shearing the metric along γ, that is, rotating along γ the two components of $A \setminus \gamma$ where $A \subset S$ is an annulus containing γ as its core curve. This is the type of deformation we study here, adapted to the higher rank setting. Beyond their simplicity, shearing deformations have many interesting features: they are used to define the famous Fenchel–Nielsen coordinates on the Teichmüller space; they define a symplectic flow on the Teichmüller space; they can be generalised into *earthquakes* (shearing along measured geodesic laminations), which give another coordinate system for the Teichmüller space.

There is a more group theoretic interpretation of shearing. Namely, if γ is separating, then $\pi_1(S)$ splits into an amalgamated product and shearing can be thought of as a partial conjugation (conjugate only one factor of the product) of the representation of $\pi_1(S)$ into $\mathrm{PSL}_2(\mathbb{R})$ defining X, under an element in the infinite cyclic centralizer of γ in $\mathrm{PSL}_2(\mathbb{R})$. If γ is non-separating, then $\pi_1(S)$ is an HNN-extension and there is an analogous interpretation of shearing. This interpretation of shearing immediately extends to the Hitchin component, where now the (identity component of the) centralizer of $\gamma \in \mathrm{PSL}_d(\mathbb{R})$ is conjugate to the group of matrices $\exp(z)$ where z is a trace-free diagonal matrix. The resulting deformations of Fuchsian representations are called *bending* or *bulging* deformations, and we call *bending parameter* the matrix z used to perform the bending. Such deformations have frequently been looked at in the literature, see for example [Gol86; FK16; AZ23], and they correspond to the *rational* deformations of Fock and Goncharov [FG06]. Bending can also be carried out along a simple geodesic multicurve γ, and then the bending transformation consists of a k-tuple of transformations where k is the number of components of γ. Bending techniques have been used in other geometric contexts as well, for example to study Kleinian groups (see, e.g. [Thu97]), pseudohyperbolic geometry [Mes07], or projective geometry [JM87].

Our geometric study involves, instead of the usual Riemannian metric, a $\mathrm{PSL}_d(\mathbb{R})$-invariant Finsler metric $\mathfrak{F}$ on the symmetric space $\mathbb{X} = \mathrm{PSL}_d(\mathbb{R})/\mathrm{PO}(d)$. Consider the convex cone $\overline{\mathfrak{a}}^+$ of trace-free diagonal matrices with entries $x_1 \geq \cdots \geq x_d$ in descending order. It can be seen as a subset of the tangent bundle $T\mathbb{X}$ (in the tangent space of a fixed basepoint), where it is a fundamental set for the action of $\mathrm{PSL}_d(\mathbb{R})$. Then any positive linear functional α_0 on $\overline{\mathfrak{a}}^+$ satisfying symmetry and convexity assumptions (see Notation 1), for instance

$$\alpha_0(x) = (d-1)x_1 + (d-3)x_2 + \cdots + (1-d)x_d, \tag{1}$$

determines a $\mathrm{PSL}_d(\mathbb{R})$-invariant Finsler metric $\mathfrak{F}$ on $\mathbb{X}$. We denote by $d^{\mathfrak{F}}$ the distance function induced by $\mathfrak{F}$; this function is bi-Lipschitz equivalent to the distance function of the usual Riemannian metric.

Following Thurston (see [Tan97] for details), define the *abstract grafing* of a hyperbolic surface X along a simple geodesic multi-curve $\gamma^* = \gamma_1 \cup \cdots \cup \gamma_k$ with *grafting heights* $L_1, \ldots, L_k$ to be the surface obtained from X by cutting X open along the geodesics γ_i and inserting a flat cylinder of height L_i. Thus the abstract grafted surface can be thought of as a geometric structure on X which is piecewise flat or of constant curvature -1. It will be convenient to allow that the flat metric on the cylinders is defined by a non-eulidean norm on $\mathbb{R}^2$.

The following is our main result. It shows that in a very precise sense, abstract grafted surfaces serve as geometric models for Hitchin representations obtained from Fuchsian representations by bending. In its formulation, the *cylinder height* of a bending representation at a component γ_i of the bending locus γ is the Finsler distance between γ_i and its image under the bending transformation. The grafted surface mentioned below is defined in Definition 2.1, and illustrated in Figure 1.

Theorem A. *For every $\sigma > 0$, there exists $C_\sigma > 0$ such that the following holds.*

Consider a closed hyperbolic surface S, a multicurve $\gamma^* \subset S$ whose components have length at most σ, and a bending parameter z such that all cylinder heights of the resulting bending representation are bounded from below by some number $L > 0$.

Then there exists an abstract grafted surface S_z, with universal covering $\tilde{S}_z$, and a $\pi_1(S)$-equivariant embedding $\tilde{Q}_z : \tilde{S}_z \to (\mathbb{X}, d^{\mathfrak{F}})$ which is quasi-isometric with multiplicative *constant* $(1 + C_\sigma/(L+1))$ *and additive constant* C_σ. *Moreover, the image* $\tilde{S}^\iota_z = \tilde{Q}_z(\tilde{S}_z)$ *is* C_σ*-Finsler-quasiconvex in the sense that for all* $x, y \in \tilde{Q}_z(\tilde{S}_z)$, *there is* at least one *Finsler geodesic from* x *to* y *at distance at most* C_σ *from* $\tilde{S}^\iota_z$.

The result motivates us to mostly use the term *grafting representations* instead of talking about bending. Theorems 5.1 and 5.2 are more precise versions of this result.

In Section 12 of [KL18], Kapovich and Leeb proved that for any Hitchin representation, all orbit maps are Finsler quasi-convex. The main point of our result is that constants are independent of the representation, a result which cannot be obtained from Kapovich–Leeb's approach.

The results of Kapovich–Leeb rely on a version of a Morse lemma in higher rank symmetric spaces (Theorem 1.3 of [KLP18], see also [KL18]). The proof of Theorem A is independent of the results in [KLP18; KL18], but also embarks from a (different) Morse-type lemma for Finsler metrics.

Theorem B. *For every* $C > 0$ *there exists a number* $C' > 0$ *with the following property. Let* $c : [a, b] \to (\mathbb{X}, d^{\mathfrak{F}})$ *be a map such that*

$$d^{\mathfrak{F}}(c(s), c(u)) + d^{\mathfrak{F}}(c(u), c(t)) \leq d^{\mathfrak{F}}(c(s), s(t)) + C$$

for all $a \leq s \leq u \leq t \leq b$. *Then there exists a Finsler geodesic connecting* $c(a)$ *to* $c(b)$ *at Hausdorff distance at most* C' *to* c.

Theorem B does not hold for the Riemannian metric. We refer to Section 3 for more information.

A key ingredient in the proof of Theorem A is the construction of paths, with the properties stated in Theorem B for Hitchin representations obtained by bending a Fuchsian representation. This construction relies on positivity in the sense of Fock–Goncharov [FG06] and can be viewed as a geometric interpretation of positivity. It gives an alternative approach towards the understanding of Hitchin representations via uniformly Morse paths in the sense of [KLP18], with geometric control of a different nature.

Theorem A can be used to study the degeneration of a Hitchin grafting ray, letting the grafting parameter go to infinity. For instance, one can study the limit in the compactification of the Hitchin component introduced by Parreau [Par12; Par00]. A representation is identified with the projectivization of its $\bar{\mathfrak{a}}^+$-valued length function (the Jordan projection), living in a compact infinite-dimensional projective space. A limit point of this compactification can be seen, via a rescaling procedure, as the length function of an action of $\pi_1(S)$ on an affine building, which is an asymptotic cone of $\mathbb{X}$. In [BHMM24], we examine other bordifications of the Hitchin component such as Loftin's [Lof19].

The limit in Parreau's compactification of certain representations obtained via bending has already been studied by Parreau in her thesis (Section V.5 of [Par00]). More precisely, Parreau assumes $d = 3$ and deforms a Fuchsian representation by pinching a curve and bending along it. She obtains an explicit formula for the limit length function. Our results can be used to reprove this, and also to get formulas in other cases, for instance when no pinching is performed, and with $d \geq 3$. This is the content Theorem C stated below, after we mention a few other works.

Note that Parreau also defined in her context the grafted surface studied here (see Figure V.4 of [Par00]), although she did not prove quasi-convexity. In [Par22; LTW22; Rei23], limits in Parreau's compactification of other kinds of deformations are studied. In the first two papers, the authors construct an explicit invariant Finsler-convex subset for the limiting action on an affine building. (Finsler-convex means any two points of the subset can be connected by *at least one* Finsler geodesic; Parreau uses the terminology "weakly convex".) Our results also give such a natural invariant convex set.

For simplicity, suppose $\gamma^* \subset S$ is an oriented simple closed curve and S_0 is one of the component of $S - \gamma^*$. Following Parreau, one can define positive and negative intersection numbers $\iota^{\pm}(\gamma, \gamma^*)$ of a closed oriented curve $\gamma \subset S$ with γ^*.

Theorem C. *Let ρ_t be a Hitchin grafting ray with grafting locus γ^* and grafting parameter tz, where $z \in \mathfrak{a}$ with nonzero cylinder height. Then the following properties hold true.*

1. *$(\rho_t)_t$ converge, as t goes to infinity, to the point of Parreau's compactification given by the projectivization of the $\mathfrak{a}^+$-valued length function*
$$\gamma \mapsto \iota^+(\gamma, \gamma^*)\kappa(z) + \iota^-(\gamma, \gamma^*)\kappa(-z),$$
where κ is the Cartan projection.
2. *For the corresponding $\pi_1(S)$-action on an asymptotic cone B of $\mathbb{X}$, there is an invariant Finsler-convex subset $T \subset B$ isometric to the Bass–Serre tree of the graph of groups decomposition defined by γ^*, with edges decorated with $\kappa(z)$ or $\kappa(-z)$.*
3. *For a suitable choice of basepoint, the quotient manifolds $\rho_t(\pi_1(S))\backslash\mathbb{X}$ converge in the pointed Gromov–Hausdorff topology to the Fuchsian manifold defined by the bordered hyperbolic surface S_0.*

Compare the formula for the length function with Proposition V.5.9 of [Par00].

Organization of the article and outline of the proofs. All main results in this article build on the concept of *admissible paths* in the Lie group $\mathrm{PSL}_d(\mathbb{R})$. Such a path can be thought of as a (continuous) path which projects to a piecewise geodesic in the symmetric space $\mathbb{X}$. Each geodesic piece corresponds either to a geodesic arc in a totally geodesic embedded $\mathbb{H}^2$, or to an arc in a maximal flat, and these two types can be read off from the path in $\mathrm{PSL}_d(\mathbb{R})$. Furthermore, the construction is done in such a way that it encodes positivity properties of the path in the sense of [FG06]. This positivity then guarantees quantitative non-backtracking.

Admissible paths in $\mathrm{PSL}_d(\mathbb{R})$ and $\mathbb{X}$ and their analogs for abstract grafted surfaces are introduced in Section 2.5. Their natural appearance for grafting deformations of Fuchsian representations is via so-called *characteristic surfaces*, introduced in Section 2.4.

In Section 3, we introduce a very general class of invariant Finsler metrics on $\mathbb{X}$ defined by *polyhedral norms* in a Cartan subalgebra. We study geometric properties of these norms on $\mathbb{R}^{d-1}$ in Section 3.2. The remainder of Section 3 is devoted to the proof of Theorem B.

In Section 4, the relation of admissible paths in $\mathrm{PSL}_d(\mathbb{R})$ to positivity in the sense of [FG06] is established. These results are applied in Section 5 to prove Theorem A and Theorem C.

Acknowledgement: This project started as a working seminar in Fall 2021, during the pandemic, held in person at the Max Planck Institute for Mathematics in Bonn. We thank the MPI for the hospitality and financial support, and we thank Gianluca Faraco, Elia Fioravanti, Frieder Jäckel, Yannick Krifka, Andrea Egidio Monti, Laura Monk and Yongquan Zhang for many enjoyable discussions and good company. P.-L.B. is grateful to Dick Canary, Fanny Kassel and Ralph Spatzier for helpful discussions, and U.H. and P.-L.B. thank Andrés Sambarino for helpful discussions. All authors thank Beatrice Pozzetti for pointing out an error in an earlier version of this article.

1 Lie groups and symmetric spaces

This section collects some basic facts on Lie groups and symmetric spaces and introduces conventions and notations used later on.

Consider the simple Lie group $G = \mathrm{PSL}_d(\mathbb{R})$ and a representation $\tau : \mathrm{PSL}_2(\mathbb{R}) \to G$, that is, a locally injective Lie group homomorphism. Many (but not all) of our results work for other semisimple Lie groups, and their proofs are easier to write using the abstract language of semisimple Lie groups, which we recall below.

Recall that the Lie algebra $\mathfrak{sl}(2, \mathbb{R})$ of $\mathrm{PSL}_2(\mathbb{R})$ is the Lie algebra of trace free real $(2, 2)$-matrices. Denote by $\mathfrak{g} = T_{\mathrm{id}}G$ the Lie algebra of G, by $\mathfrak{a} \subset \mathfrak{g}$ a Cartan subalgebra (maximal abelian subalgebra when G is split) that contains $d\tau\left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$ and by $\mathfrak{a}^+ \subset \mathfrak{a}$ an open Weyl cone whose closure $\overline{\mathfrak{a}^+}$ contains $d\tau\left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$ (the definition of Weyl cone is recalled later in this section).

We require the representation τ to be *regular*, that is, $d\tau\left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$ belongs to the interior $\mathfrak{a}^+$ of the Weyl chamber, or equivalently, there is a unique Cartan subspace $\mathfrak{a}$ containing $d\tau\left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$.

Maximal compact subgroup and symmetric space

Let $K \subset G$ be a maximal compact subgroup which contains $\tau(\mathrm{PSO}(2))$ and whose Lie algebra $\mathfrak{k}$ is orthogonal (for the Killing form) to $\mathfrak{a}$.

The symmetric space of G is $\mathbb{X} = G/K$ with basepoint $\mathbf{x} = [\mathrm{id}] = K \in G/K$. Denote by $\pi_{\mathbb{X}} \colon G \to \mathbb{X}$ the projection. The space $\mathbb{X}$ is endowed with a nonpositively curved G-invariant Riemannian metric whose induced norm is denoted by $\|\cdot\|$, and whose distance function is denoted by $d_{\mathbb{X}}$. The metric is normalised so that $\|d\tau\left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)\| = 2$.

Maximal Flats

The subspace $\exp(\mathfrak{a}) \cdot \mathbf{x}$ is a totally geodesic embedded Euclidean subspace of $\mathbb{X}$ of maximal dimension. This flat will often be identified with the abelian subgroup $A = \exp(\mathfrak{a})$ which acts simply transitively on it. Each maximal Euclidean subspace of $\mathbb{X}$ can be represented as $g \cdot A$ for some $g \in G$. These maximal Euclidean subspaces are called *maximal flats.*

Root systems

Let $\mathcal{R} \subset \mathfrak{a}^*$ be the set of restricted roots of G, that is, the set of non-zero linear one-forms α on $\mathfrak{a}$ such that

$$\mathfrak{g}_\alpha := \{X \in \mathfrak{g} \mid [a, X] = \alpha(a)X \ \forall a \in \mathfrak{a}\} \neq 0.$$

Recall that G being split means that

$$\mathfrak{g} = \mathfrak{a} \oplus \bigoplus_{\alpha \in \mathcal{R}} \mathfrak{g}_\alpha.$$

The kernels of the restricted roots are the *walls* of $\mathfrak{a}$, and the *open Weyl cones* are the connected components of the complements of the walls. By our regularity assumption on τ there is a unique open Weyl cone containing $d\tau\left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$.

Let $\mathcal{R}^+ := \{\alpha \in \mathcal{R} : \alpha(d\tau\left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)) > 0\}$ be the set of positive roots, and $\mathcal{R}^- = -\mathcal{R}^+$ the set of negative roots. Let $\Delta \subset \mathcal{R}^+$ be the set of simple roots (the positive roots that are not sums of several other positive roots).

Minimal parabolic subgroups and flag variety

The normalizer $P := N_G(\bigoplus_{\alpha \in \mathcal{R}^+} \mathfrak{g}_\alpha) \subset G$ for the adjoint representation is a minimal parabolic subgroup. Its Lie algebra is $\mathfrak{p} := \mathfrak{a} \oplus \bigoplus_{\alpha \in \mathcal{R}^+} \mathfrak{g}_\alpha$. The *opposite* parabolic subgroup is $P^- := N_G(\bigoplus_{\alpha \in \mathcal{R}^-} \mathfrak{g}_\alpha)$.

The *flag variety* $\mathcal{F} := G/P$ is compact. In fact, K acts transitively on it, with finite point stabilizer.

A notable subgroup of P is $U = \exp\left(\bigoplus_{\alpha \in \mathcal{R}^+} \mathfrak{g}_\alpha\right)$. The dynamics coming from the geometry of the symmetric space and other homogeneous spaces of G has some contraction properties that are recorded in the following algebraic fact: for any $u = \exp(\sum_{\alpha \in \mathcal{R}^+} X_\alpha)$ in U, for any sequence $(a_n)_n \subset \mathfrak{a}^+$ that diverges from the walls (i.e. $\alpha(a_n) \to +\infty$ for any $\alpha \in \mathcal{R}_+$), we have

$$\exp(-a_n) \cdot u \cdot \exp(a_n) = \exp\left(\sum_{\alpha \in \mathcal{R}_+} e^{-\alpha(a_n)} X_\alpha\right) \underset{n \to \infty}{\longrightarrow} \mathrm{id}\,. \tag{2}$$

If $(a_n)_n$ does not diverge from all the walls but only some of them, then there is still a more complicated weaker contraction property.

Maps induced by τ

Let $T \subset \mathrm{PSL}_2(\mathbb{R})$ be the subgroup of upper triangular matrices. The *ideal boundary* $\partial_\infty \mathbb{H}^2 = \mathrm{PSL}_2(\mathbb{R})/T$ of the hyperbolic plane $\mathbb{H}^2 = \mathrm{PSL}_2(\mathbb{R})/\mathrm{PSO}(2)$ is naturally homeomorphic to the circle $\mathbb{R} \cup \{\infty\}$ under the map $t \mapsto [\left(\begin{smallmatrix} 1 & t \\ 0 & 1 \end{smallmatrix}\right) \cdot \left(\begin{smallmatrix} 1 & 0 \\ 1 & 1 \end{smallmatrix}\right)]$ and $\infty \mapsto [\mathrm{id}] = T \in \mathrm{PSL}_2(\mathbb{R})/T$.

One can check that $\tau(T) \subset P$ and that τ induces an embedding

$$\partial\tau : \partial_\infty \mathbb{H}^2 \hookrightarrow \mathcal{F} = G/P.$$

Transversality in the flag variety

Two flags $\xi, \eta \in \mathcal{F}$ are said to be *transverse* if there exists $g \in G$ such that $g\xi = \partial\tau(0)$ and $g\eta = \partial\tau(\infty)$; in this case we write $\xi \pitchfork \eta$. The set of transverse pairs of flags is an open dense subset of $\mathcal{F}^2$.

The set of flags transverse to $\partial\tau(0)$ is

$$\partial\tau(0)^{\pitchfork} := P^- \cdot \partial\tau(\infty) = \exp\left(\bigoplus_{\alpha \in \mathcal{R}^-} \mathfrak{g}_\alpha\right) \cdot \partial\tau(\infty). \tag{3}$$

Similarly, for any flag ξ denote by $\xi^{\pitchfork}$ the set of flags transverse to ξ. It is an open dense subset of $\mathcal{F}$. Note that by our convention, P^- equals the stabilizer of $\partial\tau(0)$ in $\mathcal{F}$.

Any two transverse flags are contained in the boundary of a unique maximal flat. The maximal flat asymptotic to the transverse flags $\partial\tau(0)$ and $\partial\tau(\infty)$ equals $A \cdot \mathbf{x} \subset \mathbb{X}$.

More generally, the flat between (that is, asymptotic to) any two transverse flags $(\xi, \eta) = g(\partial\tau(0), \partial\tau(\infty))$ is

$$F(\xi, \eta) := gA \cdot \mathbf{x} = gA \subset \mathbb{X}.$$

Jordan and Cartan projections, and loxodromic elements

For any $g \in G$, the *Cartan projection* is the unique element $\kappa(g) \in \overline{\mathfrak{a}^+}$ such that $g \in K\exp(\kappa(g))K$. Putting $A^+ = \exp(\overline{\mathfrak{a}^+})$, it is characterized by the fact that $\exp(\kappa(g))\mathbf{x}$ is the unique intersection point of $A^+\mathbf{x}$ with the K-orbit of $g\mathbf{x}$. Note that $d(\mathbf{x}, g\mathbf{x}) = ||\kappa(g)||$ (here d is the distance of the symmetric metric on $\mathbb{X}$).

Similarly, the G-orbit of any vector $v \in TX$ intersects $\overline{\mathfrak{a}^+} \subset \mathfrak{p}$ in precisely one point $\kappa(v)$ which is called the *Cartan projection* of v.

For the *Jordan projection* $\lambda(g) \in \overline{\mathfrak{a}^+}$ we choose the following unnatural but convenient definition (see Remark 5.31 of [BQ16])

$$\lambda(g) := \lim_{n\to\infty} \frac{1}{n}\kappa(g^n).$$

The element $g \in G$ is called *loxodromic* if $\lambda(g)$ is contained in the interior $\mathfrak{a}^+$ of $\overline{\mathfrak{a}^+}$, which is equivalent to saying that g has an attracting/repelling fixed pair of transverse flags (g^-, g^+). Then g acts as a translation on the flat $F(g^-, g^+)$.

That the representation τ is regular means that the image $\tau(g)$ of any loxodromic $g \in \mathrm{PSL}_2(\mathbb{R})$ is loxodromic in G.

Weyl Chambers and special directions

Since the symmetric space $\mathbb{X}$ is nonpositively curved, it admits a *visual boundary* $\partial_\infty\mathbb{X}$, which is naturally identified with the set of unit speed infinite geodesic rays starting at the basepoint $\mathbf{x}$.

By the normalization of the metric on $\mathbb{X}$, the representation τ induces an isometric embedding $\mathbb{H}^2 \hookrightarrow \mathbb{X}$. The isometric embeddings $A \hookrightarrow \mathbb{X}$, $\mathbb{H}^2 \hookrightarrow \mathbb{X}$ extend to embeddings of the visual boundaries $\partial_\infty A \hookrightarrow \partial_\infty\mathbb{X}$, $\partial_\infty\mathbb{H}^2 \hookrightarrow \partial_\infty\mathbb{X}$.

For any $g \in G$, we identify $\xi = g\partial\tau(\infty) \in \mathcal{F}$ with a compact subset of the visual boundary $\partial_\infty\mathbb{X}$, called a (closed) Weyl Chamber:

$$\xi = g\partial_\infty A^+ = g \cdot \{\lim_{t\to\infty} \exp(tv)\mathbf{x} : v \in \overline{\mathfrak{a}^+}\} \subset g\partial_\infty A \subset \partial_\infty\mathbb{X}.$$

It is the boundary at infinity of the *Weyl Cone* gA^+ based at $g\mathbf{x}$. The facets of $\xi = g\partial\tau(\infty) \in \mathcal{F}$ are the subsets of the form

$$g \cdot \partial_\infty \left(A^+ \cap \bigcap_{\alpha\in S} \ker\alpha \right) \subset \xi,$$

where S is a subset of Δ; they are boundaries at infinity of facets of the Weyl Cone gA^+.

Every G-orbit in $\partial_\infty\mathbb{X}$ intersects exactly *once* every Weyl Chamber. In particular, to every Weyl Chamber $\xi \in \mathcal{F}$ and every point p in the standard Weyl Chamber $\partial\tau(\infty)$ one can associate a point of ξ, which is the intersection point of ξ with $G \cdot p$. The embedding $\partial_\infty\mathbb{H}^2 \to \partial_\infty\mathbb{X}$ determines a special point of $\partial\tau(\infty)$. Its orbit under G determines a special point in every Weyl Chamber.

The Weyl group

Let us recall the definition of the Weyl group, denoted by Weyl. It is the intersection of the maximal compact subgroup K of G with the stabilizer of $\mathfrak{a}$, in restriction to $\mathfrak{a}$ (so quotiented out by the fixator of $\mathfrak{a}$ in G). It can also be described as the group of orthogonal transformations of $\mathfrak{a}$ generated by the reflections along the walls $(\ker\alpha)_{\alpha\in\Delta^+}$ of the closed Weyl Chamber $\overline{\mathfrak{a}^+}$. The Weyl group is finite, and any Weyl-orbit in $\mathfrak{a}$ intersects $\mathfrak{a}^+$ exactly once.

A Finsler metric coming from a linear functional on $\mathfrak{a}$

Notation 1. We fix a linear functional α_0 on $\mathfrak{a}$ which is positive on $\overline{\mathfrak{a}^+}$ and such that $\alpha_0(gv) < \alpha_0(v)$ for all $v \in \mathfrak{a}^+$ and $g \in$ Weyl*.*

We assume that α_0 is symmetric in the sense that if g is the transformation in the Weyl group that maps $\mathfrak{a}^+$ to its opposite $-\mathfrak{a}^+$ then $\alpha_0(gv) = -\alpha_0(v)$ for any $v \in \mathfrak{a}$.

Let us denote by $\alpha_0^\#$ the vector in $\mathfrak{a}$ such that $\alpha_0(v) = \langle v, \alpha_0^\#\rangle$ for any $v \in \mathfrak{a}$, where $\langle\cdot,\cdot\rangle$ is the inner product on $\mathfrak{a}$ defined by the Riemannian metric on $\mathbb{X}$. Then the assumption above on α_0 is equivalent to asking that $\alpha_0^\# \in \mathfrak{a}^+$. We also denote by $\alpha_0^\# \in \partial_\infty\mathfrak{a}^+$ the point at infinity to which the ray spanned by $\alpha_0^\#$ limits.

An example of a linear functional satisfying the above conditions is given in Equation 1.

Let $\mathfrak{k} \subset \mathfrak{g}$ be the Lie algebra of the maximal compact subgroup K of G. The orthogonal complement $\mathfrak{p}$ of $\mathfrak{k}$ with respect to the Killing form is naturally isomorphic to $T_\mathbf{x}\mathbb{X}$, and it contains $\mathfrak{a}$. For any vector $v \in T\mathbb{X}$ we set

$$\mathfrak{F}(v) = \alpha_0(\kappa(v)) \tag{4}$$

where as before, $\kappa(v)$ is the Cartan projection of v.

Proposition 1.1 (Lemmas 5.9-10 of [KL18]). *The following hold.*

1. *$\mathfrak{F}$ defines a G-invariant Finsler metric on $\mathbb{X}$.*

2. *The unparameterized Riemannian geodesics of $\mathbb{X}$ are also geodesics for $\mathfrak{F}$.*

3. *The translation length for $\mathfrak{F}$ of any element $g \in G$ acting on $\mathbb{X}$ is given by $\ell^{\mathfrak{F}}(g) := \alpha_0(\lambda(g))$ where $\lambda(g) \in \mathfrak{a}^+$ is the Jordan projection.*

In the sequel we always normalize the functional α_0 in such a way that the embedding $\mathbb{H}^2 \to \mathbb{X}$ which is isometric for the symmetric metric also is isometric for the Finsler metric $\mathfrak{F}$.

Busemann functions

The Busemann functions, or horofunctions, are generalizations of distance functions on $\mathbb{X}$: they record relative distances to a point at infinity. Geometrically, their level sets, called horospheres, are limits of spheres whose centers go to infinity. Since there are several kinds of metrics on $\mathbb{X}$, there are also several kinds of horofunctions.

Because of the contraction property of the subgroup $U \subset G$ explained in (2), horospheres centered at a point p in the interior of the Weyl chamber $\partial\tau(\infty)$ will always be U-invariant, since for any $u \in U$, for any $(x_n)_n$ converging to p we have $d(x_n, ux_n) \to 0$ for any G-invariant metric d. In fact, every distance sphere of $\mathbb{X}$ (for any G-invariant metric) is foliated by orbits of the stabilizer of the center, and U-orbits can be described as limits of these leaves when the center tends to a point in the interior of $\partial\tau(\infty)$.

The leaves foliating the spheres centered at a given point can be parameterized by vectors of $\mathfrak{a}^+$, which one may think of as vector-valued distances. Namely, the G-orbit of any pair $(x, y) \in \mathbb{X}^2$ intersects $\{\mathbf{x}\} \times A^+$ exactly once, and the intersection is denoted $(\mathbf{x}, \exp(\kappa(x, y)))$, where $\kappa(x, y) \in \overline{\mathfrak{a}^+}$ is thought of as a vector-valued distance from x to y. The orbits of the stabilizer of x are the level sets of $\kappa(x, \cdot)$. The Riemannian distance from x to y can be expressed as $||\kappa(x, y)||$, and the Finsler distance as defined in (4) can be expressed as

$$d^{\mathfrak{F}}(x, y) = \alpha_0(\kappa(x, y)). \tag{5}$$

Using the U-orbits one can define a vector-valued Busemann function centered at $\partial\tau(\infty)$: for any $x \in \mathbb{X}$, the U-orbit $U \cdot x$ intersects the standard flat $A \subset \mathbb{X}$ in exactly one point, and taking the logarithm we get a vector $b^{\mathfrak{a}}_{\partial\tau(\infty)}(\mathbf{x}, x) \in \mathfrak{a}$, that records the relative distance from x to $\partial\tau(\infty)$ compared to the basepoint $\mathbf{x}$. One can check using the contraction property of U that if $(y_n)_n \subset A^+$ converge to a point in the interior of the Weyl Chamber $\partial\tau(\infty)$ then the $\mathrm{Stab}(y_n)$-orbits converge to level sets of $b^{\mathfrak{a}}_{\partial\tau(\infty)}(\mathbf{x}, \cdot)$.

Using the action of K, one can extend these vector-valued Busemann functions to Weyl chambers other than $\partial\tau(\infty)$. For any Weyl Chamber $\xi = k\partial\tau(\infty) \in \mathcal{F} = G/P$, the *vector-valued Busemann function* or horofunction centered at ξ between $\mathbf{x}$ and x is $b^{\mathfrak{a}}_{\xi}(\mathbf{x}, x) = b^{\mathfrak{a}}_{\partial\tau(\infty)}(\mathbf{x}, k^{-1}x)$, and more generally for $x, y \in \mathbb{X}$ we have:

$$b^{\mathfrak{a}}_{\xi}(x, y) = -b^{\mathfrak{a}}_{\xi}(\mathbf{x}, x) + b^{\mathfrak{a}}_{\xi}(\mathbf{x}, y) = \lim_{n\to\infty} \kappa(z_n, x) - \kappa(z_n, y) \in \mathfrak{a},$$

where $(z_n)_n \subset \mathbb{X}$ is any sequence converging to a point of the visual boundary in the interior of ξ. One way to check the above formula is to find $k_n \in K$ such that $k_n z_n \in A^+$, use that $\kappa(z_n, x) = \kappa(k_n z_n, k_n x)$, and pass to a subsequence to make k_n converge to some $k \in K$.

A horosphere in $\mathbb{X}$ based at a point in the visual boundary $\partial_\infty A$ of $A = A \cdot \mathbf{x}$ equals the U-obit of a horosphere in A. Given a sequence $(z_n)_n$ in A going to a point p in the visual boundary of A, in the interior of the Weyl chamber $\partial_\infty A^+$, the sequence of Riemannian spheres in A which are centered at z_n and pass through the origin 0 converges to the hyperplane in A containing 0 and perpendicular to the ray from 0 to p. One can check that on the other hand, the sequence of Finsler spheres, that is, the level sets of $d^{\mathfrak{F}}$, converges to the kernel of α_0, which does not depend on p and which, in the notation 1, is the hyperplane perpendicular to $\alpha_0^\# \in \partial_\infty A^+$.

The $\mathbb{R}$-valued Busemann function $b_p(x, y)$ associated to the Riemannian metric on $\mathbb{X}$ and centered at a point p in the interior of some Weyl chamber $\xi \subset \partial_\infty \mathbb{X}$ is the limit $\lim_{n\to\infty} d(x, z_n) - d(y, z_n)$ where $z_n \to p$. If q is the intersection point of $\partial\tau(\infty)$ with the G-orbit of p then we have $b_p(x, y) = \langle b_\xi^{\mathfrak{a}}(x, y), v\rangle$ where $v \in \mathfrak{a}$ is the unit vector pointing at q and $\langle\cdot, \cdot\rangle$ is the inner product on $\mathfrak{a}$.

The *Busemann function* associated to our choice of Finsler metric is given by

$$b_\xi^{\mathfrak{F}}(x, y) = \alpha_0(b_\xi^{\mathfrak{a}}(x, y)) = \lim_{n\to\infty} d^{\mathfrak{F}}(x, z_n) - d^{\mathfrak{F}}(y, z_n) \in \mathbb{R}. \tag{6}$$

Here the last equality in the identity (6) is valid since $\mathfrak{F}$ is defined by a linear functional on $\overline{\mathfrak{a}^+}$. Moreover, for any loxodromic element $g \in G$ with attracting fixed point $\xi \in \mathcal{F}$, the translation length of g acting on $\mathbb{X}$ endowed with the Finsler metric $d^{\mathfrak{F}}$ equals the quantity $|b_\xi^{\mathfrak{F}}(x, gx)|$.

Note that using Notation 1 we have the following link between the Finsler and Riemannian Busemann functions: if $p \in \partial_\infty \mathbb{X}$ is the intersection point of $\xi \in \mathcal{F}$ with the G-orbit of the point in $\xi \subset \partial_\infty \mathbb{X}$ corresponding to $\alpha_0^\# \in \mathfrak{a}^+$ then

$$b_\xi^{\mathfrak{F}}(x, y) = \langle b_\xi^{\mathfrak{a}}(x, y), \alpha_0^\#\rangle = b_p(x, y), \tag{7}$$

in other words, Finsler horospheres are Riemannian horospheres.

When z_n converges to a point of the visual boundary which is not regular, that is, not in the interior of a Weyl Chamber, then the Riemannian Busemann functions are still well defined, the limit $\lim_{n\to\infty} d(x, z_n) - d(y, z_n)$ still exists. For the Finsler metric the situation is more complicated: up to passing to a subsequence of $(z_n)_n$, the limit $\lim_{n\to\infty} d^{\mathfrak{F}}(x, z_n) - d^{\mathfrak{F}}(y, z_n)$ is still well defined for all $x, y \in \mathbb{X}$ (we say that z_n converge in the horoboundary), but it will give a more complicated, less algebraic, function of x and y. This was described by Kapovich–Leeb in Lemma 5.18 of [KL18], and will be used in Section 3.3.

Before we state the result let us analyze geometrically the limits of spheres in the flat A. A Finsler ball is a convex polyhedron whose faces are contained in hyperplanes parallel to $\ker(\alpha_0)$ and their images under the action of the Weyl group by reflections. If the centers of a sequence of such convex polyhedra tend to infinity away from the walls of the Weyl chambers, then the convex polyhedra converge to a halfspace bounded by the image of $\ker(\alpha_0)$ under an element of the Weyl group. If the centers of such a sequence tend to infinity away from all walls *but one*, then the convex polyhedra converge to the intersection of two such halfspaces. For example, they could be $\{x : b_\xi^{\mathfrak{F}}(x,0) \leq 0\}$ and $\{x : b_\eta^{\mathfrak{F}}(x,0) \leq 0\}$ determined by two Weyl chambers ξ and η that share a codimension 1 face. So the associated Busemann function associated to this intersection should be $f(x) = \max(b_\xi^{\mathfrak{F}}(x,0), b_\eta^{\mathfrak{F}}(x,0))$.

The precise statement is as follows. Let $(z_n)_n \subset \mathbb{X}$ be a sequence converging to a point $p \in \partial_\infty \mathbb{X}$. The point p is contained in possibly infinitely many closed Weyl chambers, let $\mathcal{B}' \subset \mathcal{F}$ be the set of such Weyl Chambers. Let $\xi \in \mathcal{F}$ and let C be the Weyl cone based at $\mathbf{x}$ and asymptotic to ξ. If $\xi \notin \mathcal{B}$ then $d(z_n, C) \to \infty$. It could happen that $d(z_n, C) \to \infty$ even if $\xi \in \{\xi_1, \dots, \xi_\ell\}$. Let $\mathcal{B}' \subset \mathcal{B}$ be the set of ξ such that $d(z_n, C)$ remains bounded. In this case let p_ξ be the intersection point of ξ with $G \cdot \alpha_0^\#$ where we view $\alpha_0^\#$ as a point in $\partial_\infty A^+$. Up to passing to a subsequence, there exists $x_0 \in \mathbb{X}$ such that for any $x \in \mathbb{X}$ we have

$$d^{\mathfrak{F}}(x, z_n) - d^{\mathfrak{F}}(x_0, z_n) \underset{n\to\infty}{\to} \max_{\xi\in\mathcal{B}'} b_\xi^{\mathfrak{F}}(x, x_0) = \max_{\xi\in\mathcal{B}'} b_{p_\xi}(x, x_0). \tag{8}$$

In other words, the Finsler balls centered at z_n whose boundary contain x_0 converge to the intersections of the Riemannian horoballs centered at p_ξ for $\xi \in \mathcal{B}'$ whose boundary contains x_0.

Concrete description of the above objects for $G = \mathrm{PSL}_d(\mathbb{R})$

The Lie algebra $\mathfrak{g}$ is the algebra of trace free (d,d)-matrices. As Cartan subspace $\mathfrak{a}$ we choose the linear subspace of diagonal (d,d)-matrices with vanishing trace, and the open Weyl chamber $\mathfrak{a}^+$ is the open cone of diagonal matrices whose entries $(\lambda_1, \dots, \lambda_d)$ fulfill $\lambda_1 > \lambda_2 > \dots > \lambda_d$.

The subgroup $K \subset \mathrm{PSL}_d(\mathbb{R})$ is chosen as the group $\mathrm{PSO}_d(\mathbb{R})$, and $P \subset \mathrm{PSL}_d(\mathbb{R})$ is taken as the image in $\mathrm{PSL}_d(\mathbb{R})$ of the set of upper triangular matrices with positive entries on the diagonal and determinant one.

The flag variety $\mathcal{F}$ has the following explicit description. Namely, a *full flag* in $\mathbb{R}^d$ is a sequence

$$\xi = (\xi_1 \subset \xi_2 \subset \dots \subset \xi_d = \mathbb{R}^d)$$

where ξ_i is a linear subspace of $\mathbb{R}^d$ of dimension i for each $i \leq d$. Clearly $\mathrm{PSL}_d(\mathbb{R})$ acts transitively on the space of all full flags, with point stabilizer a minimal parabolic subgroup. Thus $\mathcal{F}$ is just the space of full flags in $\mathbb{R}^d$.

The Busemann functions also have a concrete description, using the identification between $\mathbb{X}$ and the set of inner products on $\mathbb{R}^d$ that induce the standard volume form. Namely, given $x, y \in \mathbb{X}$, let $||\cdot||_x$ and $||\cdot||_y$ denote the norms of the associated inner products on $\mathbb{R}^d$ and on the exterior products $\Lambda^k \mathbb{R}^d$ $(1 \leq k \leq d)$. Let $\xi = (\xi_1 \subset \xi_2 \subset \dots \subset \xi_d)$ be a full flag in $\mathbb{R}^d$. Let $v = (v_1, \dots, v_d) = b_\xi^{\mathfrak{a}}(x, y) \in \mathfrak{a}$ with $v_1 + \dots + v_d = 0$. Then for all $k \leq d$, we have

$$v_1 + \dots + v_k = \log \frac{||X||_x}{||X||_y}$$

where $X \in \Lambda^k \mathbb{R}^d$ is any representative of the k-plane ξ_k.

The homomorphism τ is obtained as follows. For $d \geq 3$ there exists up to conjugation a unique irreducible representation of $\mathrm{SL}_2(\mathbb{R})$ on $\mathbb{R}^d$. This representation determines the following embedding $\tau\colon \mathrm{PSL}_2(\mathbb{R}) \to \mathrm{PSL}_d(\mathbb{R})$. Let $\mathbb{R}^h_{d-1}[X,Y]$ be the set of degree $d-1$ homogeneous polynomials with real coefficients. A matrix $M = \left(\begin{smallmatrix} a & b \\ c & e \end{smallmatrix}\right) \in \mathrm{SL}_2(\mathbb{R})$ acts on the vector space $\mathbb{R}[X,Y]$ of polynomials in two variables by $M \cdot X = aX + cY$ and $M \cdot Y = bX + eY$. This action $\mathrm{SL}_2(\mathbb{R}) \curvearrowright \mathbb{R}[X,Y]$ preserves the d-dimensional linear subspace $\mathbb{R}^h_{d-1}[X,Y]$, with determinant one elements. So it induces an embedding $\mathrm{PSL}_2(\mathbb{R}) \to \mathrm{PSL}_d(\mathbb{R})$ which is just the representation τ.

Using a suitable basis we have $\tau(\mathrm{SO}(2)) \subset K$, the representation τ induces an isometric embedding $\mathbb{H}^2 = \mathrm{PSL}_2(\mathbb{R})/\mathrm{PSO}(2) \hookrightarrow \mathbb{X}$. Denote by $\hat{\mathbb{H}}^2 = \pi_{\mathbb{X}} \circ \tau(\mathrm{PSL}_2(\mathbb{R}))$ the image of $\mathbb{H}^2$ inside $\mathbb{X}$.

The following statement is a consequence of the fact that $d\tau(T^1\mathbb{H}^2)$ consists of regular vectors contained in a single G-orbit, and each such vector is tangent to a unique maximal flat. It is well known and immediate from the above discussion.

Fact 1.2. 1. *Every geodesic in $\hat{\mathbb{H}}^2$ lies in a unique maximal flat.*

2. *For any hyperbolic element $g \in \mathrm{PSL}_2(\mathbb{R})$, the centraliser of $\tau(g)$ in G acts by translations on the unique flat containing the image under $\pi_{\mathbb{X}}$ of the axis of g in $\mathbb{H}^2$.*

2 Hitchin grafting representations

The *Hitchin component* $\mathrm{Hit}(S)$ for conjugacy classes of representations $\pi_1(S) \to \mathrm{PSL}_d(\mathbb{R})$ is the connected component of the set of conjugacy classes of representations which factor through an irreducible representation $\mathrm{PSL}_2(\mathbb{R}) \to \mathrm{PSL}_d(\mathbb{R})$. In the sequel we always work with explicit representations rather than with conjugacy classes.

The Hitchin representations we are interested in are the familiar *bending* or *bulging* deformations of *Fuchsian* representations, that is, representations which factor through the embedding $\tau : \mathrm{PSL}_2(\mathbb{R}) \to \mathrm{PSL}_d(\mathbb{R})$. We refer to [Gol86; AZ23] for an account on the bending construction. The goal of this section is to introduce these representations as well as an abstract geometric model for them, and we establish some first geometric properties of the representations and the model. The precise relation between the geometry of bending representations and the geometry of the model will be established in Section 5 and constitutes the main result of this article.

The material in Subsections 2.1 – 2.3 is well known, and the purpose is to summarize the properties and the viewpoint we are going to pursue.

2.1 Abstract grafting

In this subsection we introduce abstract grafting of a hyperbolic surface as initiated by Thurston. We refer to [Tan97] for an early account on this construction. Contrary to the common definition in the literature, our grafting contains a twist which is needed for our purpose.

Consider a closed oriented surface S of genus $g \geq 2$ endowed with a hyperbolic metric. A *simple (geodesic) multi-curve* γ^* is the union of pairwise disjoint essential mutually not freely homotopic simple closed curves (geodesics) on S. We fix moreover an orientation on each component of γ^*.

Consider the special direction $u = d\tau\left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right) \in \mathfrak{a}$ given by τ. For any $z \in \mathfrak{a}$ and $\ell > 0$, let $\mathrm{Cyl}(\ell, z) \subset \mathfrak{a}/\ell u$ be the cylinder obtained by quotienting the strip $\{tu + sz : t \in \mathbb{R},\ s \in [0,1]\} \subset \mathfrak{a}$ under the translation by ℓu. The (Finsler) *height* of such cylinder is defined as

$$\text{height} = \min\{\mathfrak{F}(tu + z) : t \in \mathbb{R}\}. \tag{9}$$

We fix for every $\gamma \in \gamma^*$ a vector $z_\gamma \in \mathfrak{a}$; the collection $z = (z_\gamma)_{\gamma \in \gamma^*}$ is interpreted as a grafting parameter.

Definition 2.1. The *abstract grafting* of S along the geodesic multi-curve γ^* is the surface S_z obtained by cutting S open along each of the components γ of γ^*, inserting flat cylinders $C_\gamma = \mathrm{Cyl}(\ell_S(\gamma), z_\gamma)$ and gluing the surface back with the translation by z_γ.

If z_γ is not parallel to u for any $\gamma \in \gamma^*$, then this grafting comes with a natural homotopy equivalence $\pi_z : S_z \to S$ projecting the flat cylinders onto γ^*, which allow us to identify $\pi_1(S_z)$ and $\pi_1(S)$.

More precisely, for each $\gamma \in \gamma^*$, the metric completion of $S - \gamma$ is a surface whose boundary consists of two geodesics γ_1, γ_2 of the same length $\ell_S(\gamma)$. The choice of a parameterisation $\gamma(t)$ defines parameterisations $\gamma_1(t)$, $\gamma_2(t)$. Attach the flat cylinder C_γ to γ_1 and γ_2 by identifying $[tu] \in C_\gamma$ with $\gamma_1(t)$ and $[tu + z_\gamma]$ with $\gamma_2(t)$.

Let $C = \bigcup_{\gamma \in \gamma^*} C_\gamma \subset S_z$ and S' be the metric completion of $S - \gamma^*$, so that $S_z = (S' \sqcup C)/\sim$ where $\sim$ is the gluing procedure explained above. The projection map $\pi_z : S_z \to S$ satisfies the following. Its restriction $S' \to S$ is the continuous extension of the inclusion $S - \gamma^* \hookrightarrow S$. It projects each $[tu + sz_\gamma] \in C_\gamma$ to $\gamma(t) \in S$.

We call this operation *abstract grafting* to distinguish it from the *Hitchin grafting* that we will introduce for Hitchin representations. We shall refer to S_z as a *grafted surface*.

Note that if $z_\gamma = 0$ for every γ in γ^*, then the grafting is trivial and $S_z = S$. If all z_γ are parallel to u, then the grafted surface is hyperbolic and obtained from S by shearing along γ^* with shearing length given by the size of the parameters z_γ.

More generally, S_z is an orientable surface which admits a canonical piecewise smooth structure as well as a natural conformal structure which in turn induces a global C^1-structure on S. Any norm on $\mathfrak{a}$ which coincides with the norm induced by the hyperbolic metric on the distinguished direction u induces a Finsler metric on S_z which coincides with the hyperbolic metric on S' and whose restriction to the cylinders C_γ is flat.

In particular, the norm defined by the Riemannian metric of $\mathbb{X}$ can be used to endow the C^1-surface S_z with a C^0 Riemannian metric which is smooth everywhere except at the gluing locus, has constant curvature -1 in $S_z - \cup_\gamma C_\gamma$ and has constant curvature 0 in the interior of the cylinders C_γ. Since the curvature of this metric is non-positive whereever it is defined and the gluing is performed along geodesics, S_z is non-positively curved in the sense of Alexandrov and hence its universal covering $\tilde{S}_z$ is a CAT(0)-space.

Thus in this case every free homotopy class has a Riemannian geodesic representative whose length is minimal in the free homotopy class. Such a Riemannian geodesic is unique unless it is a core curve of a flat cylinder. If all the z_γ are orthogonal to the special direction u, then the natural homotopy equivalence $\pi_z : S_z \to S$ is 1-Lipschitz and hence in this case, free homotopy classes have bigger lengths in S_z than in S. Moreover, the unit tangent bundle $T^1 S_z$ of S_z is well defined, and there is a geodesic flow which is topologically mixing and admits a unique measure of maximal entropy [Kni98].

As we are interested in Finsler metrics on $\mathbb{X}$ using α_0 (see (4)) rather than the Riemannian one, we also endow S_z with a Finsler metric by equipping each cylinder C_γ with the quotient of the non-Euclidean norm $\mathfrak{F}$ on $\mathfrak{a}$. Observe that in general, for a given C^1-structure on S_z as constructed above, this metric is *discontinuous* at the gluing locus between the flat cylinders and the hyperbolic part. Additionally the metric on the flat part is sensitive in the direction of z, and does not depend only on the height of the grafting (contrarily to the Riemannian metric). Nevertheless it induces a well defined path metric on S_z.

The following observation will be useful later on when estimating lengths.

Lemma 2.2. *If all z_e are in $\ker(\alpha_0)$, then the natural projection $\pi_z : S_z \to S$ is 1-Lipschitz for the Finsler metric on S_z. In particular, all free homotopy classes of curves have bigger Finsler lengths in S_z than in S.*

Proof. By definition, the restriction of our projection map $\pi_z : S_z \to S$ to each flat cylinder $C_\gamma = \{tu + sz_\gamma\}/\ell_S(\gamma)u$ comes from the linear projection of $\mathfrak{a}$ onto the line spanned by u, parallel to the direction $z_\gamma \in \ker(\alpha_0)$. To conclude it suffices to note that this projection is 1-Lipschitz for the non-Euclidean norm on $\mathfrak{a}$, which was defined using α_0 (see (4)). $\square$

2.2 Particular case of an amalgamated product

In this section we explain briefly the construction of the following two Sections 2.3 and 2.4 in the special case where γ^* has only one component and is separating.

Let Σ be a closed orientable smooth surface of genus at least 2 and let $\gamma^* \subset \Sigma$ be a separating simple closed curve. Then γ^* splits Σ into two subsurfaces Σ_1 and Σ_2, and $\pi_1(\Sigma)$ can be written as an amalgamated product $\pi_1(\Sigma_1) \underset{\gamma^*}{*} \pi_2(\Sigma_2)$.

Consider a discrete and faithful representation $\rho : \pi_1(S) \to \mathrm{PSL}_2(\mathbb{R}) \xrightarrow{\tau} G$ such that $\rho(\gamma^*) = \exp(\ell_\rho(\gamma^*)u)$ where $u = d\tau\left(\begin{smallmatrix}1 & 0\\ 0 & -1\end{smallmatrix}\right)$ is the special direction of $\mathfrak{a}$. Let $z \in \mathfrak{a}$ be a grafting parameter. The *Hitchin grafting representation* $\rho_z : \pi_1(S) \to G$ is defined by requiring that $\rho_z(\gamma) = \rho(\gamma)$ for any $\gamma \in \pi_1(\Sigma_1)$ and $\rho_z(\gamma) = \exp(z) \cdot \rho(\gamma) \cdot \exp(-z)$ for any $\gamma \in \pi_1(\Sigma_2)$.

One can then define an immersion $Q_z : S_z \to \rho_z\backslash\mathbb{X}$ whose restriction to any of the hyperbolic pieces of S_z and to the flat cylinder is totally geodesic. Indeed, $\rho_z(\pi_1(\Sigma_1)) = \rho(\pi_1(\Sigma_1))$ preserves $\widetilde{\Sigma}_1 \subset \mathbb{H}^2 \subset \mathbb{X}$, inducing a totally geodesic embedding $\Sigma_1 = \rho_z(\pi_1(\Sigma_1))\backslash\widetilde{\Sigma}_1 \hookrightarrow \rho_z\backslash\mathbb{X}$. Similarly, if one identifies Σ_2 with $\rho(\pi_1(\Sigma_2))\backslash\widetilde{\Sigma}_2$ where $\widetilde{\Sigma}_2 \subset \mathbb{H}^2 \subset \mathbb{X}$, then $\rho_z(\pi_1(\Sigma_2)) = \exp(z) \cdot \rho(\pi_1(\Sigma_2)) \cdot \exp(-z)$ preserves $\exp(z)\widetilde{\Sigma}_2$ and hence induces a totally geodesic embedding $\Sigma_2 \hookrightarrow \rho_z \backslash \mathbb{X}$. In general the image of the boundary components $\partial\Sigma_1$ and $\partial\Sigma_2$ in $\rho_z \backslash \mathbb{X}$ are disjoint. However they can be connected by the natural totally geodesic embedding of the cylinder $C = \{tu + sz : t \in \mathbb{R},\ s \in [0,1]\}/\rho_z(\gamma^*)$ into $\rho_z \backslash \mathbb{X}$. Gluing these three embeddings yield a piecewise totally geodesic embedding of $S_z = (\Sigma_1 \cup C \cup \Sigma_2)/\sim$ into $\rho_z \backslash \mathbb{X}$.

2.3 Graphs of groups decomposition and bending

A classical reference for the theory of graph of groups is [Ser77]. We collect some facts we need. Let Σ be a closed orientable smooth surface and let $\gamma^* \subset \Sigma$ be a simple multi-curve. The multi-curve determines the following graph of groups decomposition of $\pi_1(\Sigma)$, which will be used to define a family of Hitchin representations.

Let G_{γ^*} be the oriented graph such that each vertex $v \in V$ corresponds to a component Σ_v of $\Sigma - \gamma^*$, and each edge $e \in E$ corresponds to an oriented component $\vec{\gamma}_e$ of γ^*. Given an edge $e \in E$, we denote by $\bar{e}$ the opposite edge of e, for which $\vec{\gamma}_{\bar{e}}$ corresponds to the curve $\vec{\gamma}_e$ with the reverse orientation. The oriented edge e is adjacent to the two (not necessarily distinct) components $\Sigma_{o(e)}, \Sigma_{t(e)}$ of $\Sigma - \gamma^*$ which contain $\vec{\gamma}_e$ in their boundary. One can embed G_{γ^*} into the surface Σ such that each vertex v lies in the interior of Σ_v and each edge e connects $o(e)$ to $t(e)$, crossing transversally $\vec{\gamma}_e$ once. Since we assume that Σ is oriented, choosing an orientation on γ^* is the same as choosing for each pair of opposite edges $e, \bar{e} \in E$ a preferred one by declaring that the ordered pair (u_1, u_2) consisting of the oriented tangent u_1 of the oriented edge e at x_e and the oriented tangent of the oriented geodesic γ^* defines the orientation of Σ.

The graph of groups decomposition of $\pi_1(S)$ defined by this datum associates to each vertex $v \in V$ the fundamental group $A_v := \pi_1(\Sigma_v, v)$ where v is seen as a point in the interior of Σ_v. To each edge e is associated the fundamental group $A_e := \pi_1(\vec{\gamma}_e, x_e) \simeq \mathbb{Z}$ of $\vec{\gamma}_e$, where x_e is the intersection point of $\vec{\gamma}_e$ with e (which is seen as an arc in Σ transverse to $\vec{\gamma}_e$). The inclusions $\vec{\gamma}_e \hookrightarrow \Sigma_{o(e)}$, $\vec{\gamma}_e \hookrightarrow \Sigma_{t(e)}$ determine the following monomorphisms, by connecting x_e to respectively $o(e)$ and $t(e)$ via e.

$$\alpha_{o(e)} : A_e = \pi_1(\vec{\gamma}_e, x_e) \hookrightarrow A_{o(e)} = \pi_1(\Sigma_{o(e)}, o(e)) \text{ and } \alpha_{t(e)} : A_e \hookrightarrow A_{t(e)}.$$

Note that $\alpha_{o(e)}(\vec{\gamma}_e) = \alpha_{t(\bar{e})}(\vec{\gamma}_{\bar{e}})^{-1}$.

That this construction indeed defines a decomposition of $\pi_1(\Sigma)$ as graph of groups with cyclic edge groups is well known. More precisely, choose a spanning tree $T \subset G_{\gamma*}$ of G_{γ^*}, with edge set $E_T \subset E$ invariant under the orientation reversing map $e \mapsto \bar{e}$. For a vertex $v \in V$ put A_v, and for an edge $e \in E$ put A_e. Denote by $\vec{\gamma}_e$ the oriented geodesic defined by the oriented edge e.

Let $\pi_1(G_{\gamma^*}, T)$ be the quotient group

$$\pi_1(G_{\gamma^*}, T) = {}^{(*_v A_v) * F_E}\!\diagup_{R}$$

where $*$ denote the free product, F_E is the free group generated by the edge set E, and R is the normal subgroup of $(*_v A_v) * F_E$ generated by the union of the sets

- $e \cdot \bar{e}$ for all $e \in E$,
- e for all $e \in E_T$,
- $e\alpha_{o(e)}(g)e^{-1}\alpha_{t(e)}(g)^{-1}$ for all $e \in E$ and $g \in A_e$, which we think of as $e\alpha_{o(e)}(g)e^{-1} \equiv \alpha_{t(e)}(g)$.

Thus $\pi_1(\Sigma)$ is obtained from simultaneous HNN-extension of the tree of groups defined by the spanning tree T.

Recall that the isomorphism between $\pi_1(G_{\gamma^*}, T)$ and $\pi_1(S)$ is constructed by choosing a basepoint $v_0 \in V$, embedding each vertex group A_v into $\pi_1(S, v_0)$ by connecting v to v_0 via the spanning tree T, and mapping F_E into $\pi_1(S, v_0)$ by connecting the endpoints $o(e)$ and $t(e)$ of each e to v_0 via the tree T.

Take a discrete and faithful representation $\rho \colon \pi_1(G_{\gamma^*}, T) \to \mathrm{PSL}_2(\mathbb{R}) \xrightarrow{\tau} G$ which factors through the embedding $\tau : \mathrm{PSL}_2(\mathbb{R}) \to \mathrm{PSL}_d(\mathbb{R})$. We use the graphs of groups decomposition of $\pi_1(\Sigma)$ to perform a bending of the representation in G with parameter $z = (z_\gamma)_{\gamma\in\gamma^*} \in \mathfrak{a}^{\gamma^*}$. This construction can be thought of as bending the surface S along the geodesic multicurve γ^* in the space of representations into G.

Let $\widetilde{\rho} : (*_v A_v) * F_E \to G$ be the composition of ρ with the projection

$$(*_v A_v) * F_E \to \pi_1(G_{\gamma^*}, T). \tag{10}$$

Fix an orientation on γ^*, so that for every $\gamma \in \gamma^*$, we get a preferred edge $e \in E$.

Then there exists $B \in \mathrm{PSL}_d(\mathbb{R})$ such that $\widetilde{\rho}(\alpha_{o(e)}(\vec{\gamma}_e)) = B\exp(\ell_\rho(\gamma)u)B^{-1}$, where $u = d\tau\left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$ is the special direction of $\mathfrak{a}$ (note that the ℓ_ρ-length of γ does not depend on the orientation of γ since α_0 was taken symmetric).

Set $\zeta_e = B\exp(z_\gamma)B^{-1}$, so that ζ_e commutes with $\widetilde{\rho}(\alpha_{o(e)}(\vec{\gamma}_e))$. Note that by definition of our relations R and $\alpha_{o(\bar{e})}(\vec{\gamma}_{\bar{e}}) = \alpha_{t(e)}(\vec{\gamma}_e)^{-1}$ we have

$$\widetilde{\rho}(\alpha_{o(\bar{e})}(\vec{\gamma}_{\bar{e}})) = \widetilde{\rho}(e)\widetilde{\rho}(\alpha_{o(e)}(\vec{\gamma}_e))^{-1}\widetilde{\rho}(e)^{-1} = \widetilde{\rho}(e)B\exp(-\ell_\rho(\gamma)u)B^{-1}\widetilde{\rho}(e)^{-1}.$$

Set $\zeta_{\bar{e}} = \widetilde{\rho}(e)B\exp(-z_\gamma)B^{-1}\widetilde{\rho}(e)^{-1} = \widetilde{\rho}(e)\zeta_e^{-1}\widetilde{\rho}(e)^{-1}$, that commutes with $\widetilde{\rho}(\alpha_{o(\bar{e})}(\vec{\gamma}_{\bar{e}}))$ and satisfies $\widetilde{\rho}(\bar{e})\zeta_{\bar{e}}\widetilde{\rho}(e)\zeta_e = 1$.

Geometrically, the group A_e acts on $\mathbb{H}^2$ as a translation on a geodesic $\widetilde{\gamma}$. By Fact 1.2, the image of $\widetilde{\gamma} \subset \mathbb{H}^2$ in $\mathbb{X}$ is contained in a unique maximal flat, and ζ_e preserves this flat and acts on it as a translation.

A Hitchin grafting representation is obtained by performing a partial conjugation of $\pi_1(G_{\gamma^*}, T)$ by the elements $\zeta = (\zeta_e)_{e\in E}$. Fix a basepoint $v_0 \in V$. For any $v \in V$, we denote by

$$\omega_v = \zeta_{e_1} \cdots \zeta_{e_n}$$

where $(e_1 \cdots e_n)$ is an oriented path in the tree T from v_0 to v. Since $\zeta_{\bar{e}} = \zeta_e^{-1}$ when e is in E_T and T is a tree, the value of ω_v does not depend on the chosen path.

Then define the representation $\widetilde{\rho}_z \colon (*_v A_v) * F_E \to G$ by

i) $\widetilde{\rho}_z(g) = \omega_v\widetilde{\rho}(g)\omega_v^{-1}$ for all $v \in V$ and $g \in A_v = \pi_1(\Sigma_v)$,

ii) $\widetilde{\rho}_z(e) = \omega_{o(e)}\widetilde{\rho}(e)\zeta_e\omega_{t(e)}^{-1}$ for all $e \in E$.

Lemma 2.3. *The representation $\widetilde{\rho}_z$ contains R in its kernel.*

Proof. For all $e \in E$, we have

$$\widetilde{\rho}_z(e\bar{e}) = \left(\omega_{o(e)}\widetilde{\rho}(e)\zeta_e\omega_{t(\bar{e})}^{-1}\right)\left(\omega_{o(e)}\widetilde{\rho}(\bar{e})\zeta_{\bar{e}}\omega_{t(\bar{e})}^{-1}\right) = 1$$

since $\omega_{t(\bar{e})} = \omega_{o(e)}$ and $\widetilde{\rho}(\bar{e})\zeta_{\bar{e}}\rho(e)\zeta_e = 1$.

For all $e \in E_T$, we have $\widetilde{\rho}(e) = 1$ and $\omega_{t(e)} = \omega_{o(e)}\zeta_e$, so

$$\widetilde{\rho}_z(e) = \omega_{o(e)}\widetilde{\rho}(e)\zeta_e\omega_{t(e)}^{-1} = 1$$

Take $e \in E$ and $g \in A_e$. Then

$$\begin{aligned}
\widetilde{\rho}_z(e\alpha_e(g)e^{-1}) &= \left(\omega_{o(e)}\widetilde{\rho}(e)\zeta_e\omega_{t(e)}^{-1}\right)\left(\omega_{t(e)}\widetilde{\rho}(\alpha_e(g))\omega_{t(e)}^{-1}\right)\left(\omega_{t(e)}\zeta_e^{-1}\widetilde{\rho}(e)^{-1}\omega_{o(e)}^{-1}\right)\\
&= \omega_{o(e)}\widetilde{\rho}(e)\zeta_e\widetilde{\rho}(\alpha_e(g))\zeta_e^{-1}\rho(e)^{-1}\omega_{o(e)}^{-1}\\
&= \omega_{o(e)}\widetilde{\rho}(e)\widetilde{\rho}(\alpha_e(g))\widetilde{\rho}(e)^{-1}\omega_{o(e)}^{-1} \qquad \text{since } \zeta_e \text{ and } \widetilde{\rho}(\alpha_e(g)) \text{ commute}\\
&= \omega_{o(e)}\widetilde{\rho}(e\alpha_e(g)\bar{e})\omega_{o(e)}^{-1}\\
&= \omega_{t(\bar{e})}\widetilde{\rho}(\alpha_{\bar{e}}(g))\omega_{t(\bar{e})}^{-1} = \widetilde{\rho}_z(\alpha_{\bar{e}}(g))
\end{aligned}$$

□

Definition 2.4. We denote by $\mathrm{Gr}_z^{\gamma^*}\rho\colon \pi_1(G_{\gamma^*}, T) \to G$ the representation induced by $\widetilde{\rho}_z$, and sometimes just ρ_z if there is only one hyperbolic structure involved. We call it the *Hitchin grafting representation* with data z along γ^*.

Up to conjugation, the representation ρ_z does not depend the choices made for the graph of group decomposition.

2.4 The characteristic surface for Hitchin grafting representations

Consider a Fuchsian representation $\rho : \pi_1(S) \to \mathrm{PSL}_2(\mathbb{R}) \to G$ and denote by S the hyperbolic surface defined by this representation. Choose some grafting datum z and let ρ_z be the Hithin grafted representation defined by ρ and z. As this representation is contained in the Hitchin component, it follows from Labourie [Lab06] and Fock–Goncharov [FG06] that ρ_z is faithful, with discrete image. In particular, the quotient manifold $\rho_z\backslash\mathbb{X}$ is homotopy equivalent to S; in fact ρ induces a natural homotopy class of homotopy equivalences between $\rho_z\backslash\mathbb{X}$ and S.

The goal of this subsection is to construct a geometrically controlled homotopy equivalence from an abstract grafted surface into $\rho_z\backslash\mathbb{X}$. The following proposition is the main result of this subsection.

Proposition 2.5. *Consider a Hitchin grafting representation ρ_z obtained from ρ and with grafting datum z. Recall that S_z denotes the abstract grafting of S from Definition 2.1, with universal covering $\widetilde{S}_z$. Then there exists a piecewise totally geodesic immersed surface $\widetilde{S}_z^\iota \subset \mathbb{X}$ and a ρ_z-equivariant immersion $\widetilde{Q}_z\colon \widetilde{S}_z \to \widetilde{S}_z^\iota \subset \mathbb{X}$.*

The map $\widetilde{Q}_z$ is a path isometry for the Riemannian (resp. Finsler) metric on $\widetilde{S}_z$ and the induced path metric on $\widetilde{S}_z^\iota$ from the Riemannian (resp. Finsler) metric on $\mathbb{X}$.

Before we prove the proposition, note that the surface $\widetilde{S}^{\iota}_z$ is ρ_z-invariant and hence descends to compact piecewise smooth immersed surface $S^{\iota}_z \subset \rho_z\backslash\mathbb{X}$. We call this surface *characteristic*. We shall show in Proposition 4.6 that the corresponding map $S_z \to S^{\iota}_z \subset \rho_z\backslash\mathbb{X}$ is actually an embedding, and hence that S^{ι}_z is not only an immersed surface but an embedded one.

Recall that the Riemannian (resp. Finsler) *cylinder height* of the Hitchin grafting representation ρ_z is the minimum of all $d(z_\gamma, \mathbb{R}u)$ (resp. $d^{\mathfrak{F}}(z_\gamma, \mathbb{R}u)$) for all $\gamma \in \gamma^*$, where $u = d\tau\left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$ is the special direction in $\mathfrak{a}$.

Proof of Proposition 2.5. Denote by $\widetilde{S}$ and $\widetilde{\gamma}^*$, respectively, the universal cover of S and the pre-image of γ^* in $\widetilde{S}$. Using the hyperbolic metric, we can fix an identification $S = \pi_1(S)\backslash\mathbb{H}^2$ so that $\widetilde{S} = \mathbb{H}^2$.

Let $\widetilde{S}_z$ be the universal cover of the abstract grafted surface S_z. This surface consists of a countable union S^{hyp}_z of simply connected hyperbolic surfaces with geodesic boundary, called hyperbolic pieces in the sequel, and a countable union of flat strips separating these hyperbolic pieces. Let $\mathcal{T} \subset \widetilde{S}$ be an embedded graph with one vertex v in the interior of each of the hyperbolic pieces $\widetilde{\Sigma}_v$ of $\widetilde{S} - \widetilde{\gamma}^*$ and where two such points are connected by an edge e if the pieces containing them are separated by a single component $\tilde{\gamma}_e$ of $\widetilde{\gamma}^*$. To each vertex v of $\mathcal{T}$ is also associated a hyperbolic piece $\widetilde{\Sigma}^z_v \subset \widetilde{S}_z$ which is naturally isometric to $\widetilde{\Sigma}_v$.

By construction, for any vertex v of $\mathcal{T}$ the stabilizer $A_v := \mathrm{Stab}_{\pi_1(S)}(\widetilde{\Sigma}_v)$ is mapped by ρ_z onto a conjugate $g_v\rho(A_v)g_v^{-1}$ of $\rho(A_v)$ in G and hence it stabilises a unique totally geodesic embedded bordered surface $\widehat{\Sigma}^z_v = g_v\widetilde{\Sigma}_v \subset \mathbb{X}$ which is naturally isometric to $\widetilde{\Sigma}_v$ and $\widetilde{\Sigma}^z_v$. Define $(\widetilde{Q}_z)_{|\widetilde{\Sigma}^s_v} : \widetilde{\Sigma}_v \to \widehat{\Sigma}^z_v$ to be this natural isometry. By the construction of ρ_z, the thus defined map $\widetilde{Q}_z : \widetilde{S}^{hyp}_z \to \mathbb{X}$ is equivariant with respect to the representation ρ_z.

Consider an edge e of $\mathcal{T}$ between two vertices $v = o(e)$ and $w = t(e)$ that projects onto a component $\gamma \subset \gamma^*$ matching the fixed orientation on γ^*. We also call $e \in \pi_1(S)$ the preferred generator of the stabilizer of $\widetilde{\gamma}_e = \widetilde{\Sigma}_v \cap \widetilde{\Sigma}_w$. Its holonomy $\rho_z(e)$ acts cocompactly by translation on boundary components $\widetilde{c}_1$ and $\widetilde{c}_2$ of $\widehat{\Sigma}^z_v$ and $\widehat{\Sigma}^z_w$, respectively.

Let as before $u = d\tau\left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$ be the special direction of $\mathfrak{a}$. By construction, there exists a unique $h \in \mathrm{PSL}_d(\mathbb{R})$ such that

- $h\rho_z(A_v)h^{-1} \subset \mathrm{PSL}_2(\mathbb{R}) \subset \mathrm{PSL}_d(\mathbb{R})$;
- $h\widehat{\Sigma}^z_v \subset \mathbb{H}^2 \subset \mathbb{X}$;
- $h\rho_z(e)h^{-1} = \exp(\ell_\rho(e)u)$;
- $h\widetilde{c}_1 \subset \mathbb{H}^2 \cap \mathfrak{a} \subset \mathbb{X}$ is an axis of $\exp(u)$, that is, it is invariant under $\exp(u)$ and $\exp(u)$ acts on it as a translation.

Recall that $h\rho_z(A_v)h^{-1} = hg_v\rho(A_v)g_v^{-1}h^{-1}$ and $hg_v \in \mathrm{PSL}_2(\mathbb{R})$. One can then check the following formula for the holonomy of the adjacent stabilizer A_w:

$$h\rho_z(A_w)h^{-1} = \exp(z_\gamma)hg_v\rho(A_w)g_v^{-1}h^{-1}\exp(-z_\gamma),$$

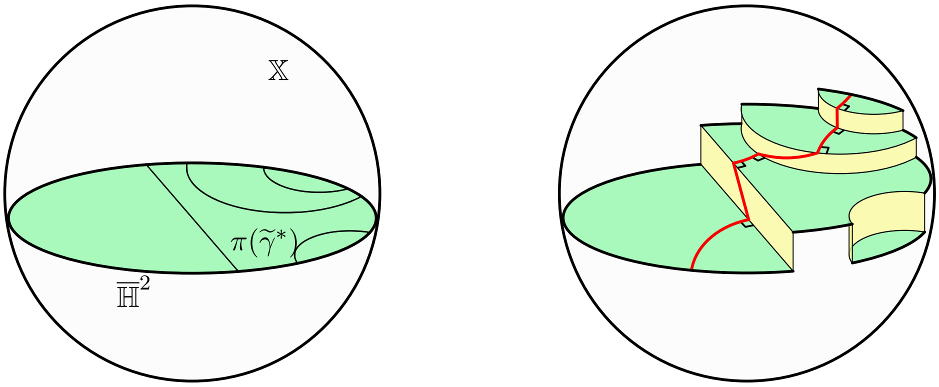

Figure 1: Geometric description of the Hitchin representation: the hyperbolic part in green, the flat part in yellow and an admissible path in red. The choices of directions of the flat parts (up or down) are arbitrary; recall that γ^* is a multicurve.

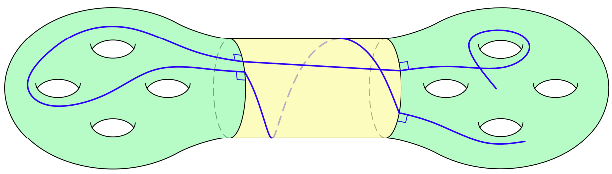

Figure 2: Admissible path in the closed surface obtained as an abstract grafting.

and hence $h\widehat{\Sigma}^z_w \subset \exp(z_\gamma)\mathbb{H}^2$ and $\widetilde{c}_2 = \exp(z_\gamma)\widetilde{c}_1 \subset \mathfrak{a}$ is another axis of $\exp(u)$. Thus the flat strip $h^{-1}\{tu + sz_\gamma\}$ is $\rho_z(e)$-invariant, connects $\widehat{\Sigma}^z_v$ to $\widehat{\Sigma}^z_w$, is the only such flat strip, and is naturally isometric to the flat strip between $\widetilde{\Sigma}^z_v$ and $\widetilde{\Sigma}^z_w$ in $\widetilde{S}_z$.

Doing this for all flat strips in $\widetilde{S}_z$ yields an extended map $\widetilde{Q}_z \colon \widetilde{S}_z \to \mathbb{X}$, which is an isometry on each hyperbolic and flat piece. Furthermore, by construction, the map $\widetilde{Q}_z$ is continuous and ρ_z-equivariant.

□

2.5 Admissible paths

In Section 5 we shall show that the characteristic surface not only is embedded, but it also can be used effectively to compare the large scale geometry of the locally symmetric manifold $\rho_z \backslash \mathbb{X}$ to the large scale geometry of the grafted surface. This comparison relies on the analysis of some specific paths which we introduce now.

2.5.1 Admissible paths in abstract grafted surfaces

We begin with introducing a family of paths in grafted surfaces, called *admissible paths*, which are from a technical point of view easier to handle than geodesics.

Let S_z be an abstract grafted surface with hyperbolic part S^{hyp} and cylinder part $\mathcal{C}$. In a nutshell, an admissible path c is a continuous path of S_z which is a geodesic everywhere except possibly at $S^{\text{hyp}} \cap \mathcal{C}$, where it might have a singularity. Moreover, we require that the "hyperbolic part" $c \cap S^{\text{hyp}}$ of the path c is orthogonal to $S^{\text{hyp}} \cap \mathcal{C}$ where it meets it.

It is clear that lifts of admissible paths to the universal cover are quasi-geodesics (although we will not need it). Our goal will be to show that the images of admissible paths under the map constructed in Proposition 2.5 are quasi-geodesics of the symmetric space, with control on the multiplicative constant.

Definition 2.6. Consider a closed hyperbolic surface S, a multicurve $\gamma^* \subset S$ and a grafting parameter z. Then S_z is the abstract grafted surface with hyperbolic part $S^{\text{hyp}} \subset S_z$ and flat (cylindrical) part $\mathcal{C} \subset S_z$. An *admissible path* in S_z is a continuous path $c \subset S_z$ such that

- c is geodesic outside $\mu = S^{\text{hyp}} \cap \mathcal{C}$;
- the hyperbolic part $c \cap S^{\text{hyp}}$ intersects γ^* orthogonally;
- a component of the flat part $c \cap \mathcal{C}$ connects the two distinct boundary components of the flat cylinder containing it.

Similarly one can define *admissible loops*.

Note that if z is trivial then $S_z = S$ and the above definition still makes sense. The flat part $\mathcal{C}$ is just γ^*, and the path is allowed to contain arcs in γ^* separating two geodesic arcs which emanate to the two distinct sides of γ^* in a tubular neighborhood of γ^*.

An *admissible path* in the universal cover $\tilde{S}_z$ is the lift of an admissible path in S_z.

Note that any two points of $\tilde{S}_z$ are connected by a unique admissible path; in other words, any path of S_z is homotopic (with fixed endpoints) to a unique admissible path.

Similarly, any loop in S_z not homotopic to a component of γ^* is freely homotopic to a unique admissible loop.

Observation 2.7. *The image under $\pi_z : S_z \to S$ (or the lift $\widetilde{S}_z \to \widetilde{S}$) of admissible paths in S_z are admissible paths in S.*

In fact, this induces a correspondence in the sense that any admissible path in S is the image under π_z of a unique admissible path in S_z.

2.5.2 Admissible paths in the symmetric space: geometric description

There are complete analogs of admissible paths in grafted surfaces for the symmetric space $\mathbb{X}$ of G, which are also called admissible paths. Such paths include the image of all admissible paths in $\tilde{S}_z$ under a path isometry $Q_\zeta : \tilde{S}_z \to \mathbb{X}$ constructed in Proposition 2.5.

Roughly speaking, admissible paths are piecewise geodesics that alternate between following a geodesic of the same type as the geodesics in the embedded $\mathbb{H}^2 \hookrightarrow \mathbb{X}$, and then following a geodesic in a flat, orthogonal to the previous geodesic, and then following a $\mathbb{H}^2$-type geodesic orthogonal to the previous flat... etc, see Figure 2.

The above description is not quite correct, in particular because it does not encapsulate the positivity assumption which is crucial in our proofs. There are several ways to define rigorously admissible paths. We are going to start with a geometric definition, which is easier to picture, and then in the next section we will give an algebraic definition. In the sequel the geometric definition will never be used, instead all the proofs will use the algebraic one, in particular because the positivity property of admissible paths is more naturally encoded in the algebraic definition.

Recall that $\mathbb{H}^2$ embeds isometrically into $\mathbb{X}$. In fact there are many isometric embeddings, and $\mathrm{PGL}_d(\mathbb{R})$ acts transitively on the set of all isometric embeddings. Let us call an $\mathbb{H}^2$-*frame* the datum of a point $x \in \mathbb{X}$ and a pair of orthogonal unit tangent vectors (v, w) which are tangent to a common embedded $\mathbb{H}^2$. Let Y be the space of $\mathbb{H}^2$-frames, on which $\mathrm{PGL}_d(\mathbb{R})$ also acts transitively. This action is even simply transitive since $\mathrm{PGL}_d(\mathbb{R})$ is real split.

On Y there is a natural geodesic flow $(\mathrm{geod}_t)_{t\in\mathbb{R}}$: given $(x, v, w) \in Y$ one can follow the geodesic ray spanned by v and parallel transport v and w along it. In other words, this action is the action of the one-parameter group of transvections on $\mathbb{X}$ along the geodesic ray spanned by v.

There is also a natural action of $\mathfrak{a}$, which we shall call the "orthogonal sliding action" and denote $(\mathrm{slide}_z)_{z\in\mathfrak{a}}$: given $(x, v, w) \in Y$ there is a unique maximal flat F containing w, and a unique identification of the tangent space of F with $\mathfrak{a}$ such that w is sent into $\mathfrak{a}^+$. Thus given $z \in \mathfrak{a}$ one can follow the geodesic ray spanned by the associated vector in F and parallel transport v and w along it. Note that the image of the vector v under this sliding action remains orthogonal to the flat F.

We define an (ω, L)-admissible path in $\mathbb{X}$ to be a path obtained by choosing an $\mathbb{H}^2$-frame and pushing it via the geodesic flow for some time at least ω, and then sliding orthogonally via $(\mathrm{silde}_{tz})_t$ for some time at least L using some direction $z \in \mathfrak{a}$, and then pushing along the geodesic flow again for time at least ω... etc.

In particular, an admissible path does not backtrack in any obvious way because it remembers directions along which the path can be continued. This is the property which can be thought of as a geometric interpretation of positivity in the sense of [FG06]. In fact we get *quantitative* positivity properties from the lower bound ω on the times we push along the geodesic flow.

2.5.3 Admissible paths in the symmetric space: algebraic definition

Let us now give an algebraic definition of admissible paths in $\mathbb{X}$. For this we will first define admissible paths in G. The description of these paths uses a basepoint for the action of G which is determined by the Fuchsian representation τ.

Notation 2. We set

- $a_t := \tau\begin{pmatrix} e^t & 0 \\ 0 & e^{-t} \end{pmatrix} \in G$;
- $r_\theta := \tau\begin{pmatrix} \cos(\theta/2) & \sin(\theta/2) \\ -\sin(\theta/2) & \cos(\theta/2) \end{pmatrix} \in G$;

- $a'_t := r_{\pi/2} \cdot a_t \cdot r_{\pi/2}^{-1} \in G$;
- *for every* $t \in \mathbb{R} \cup \{\infty\} = \partial\mathbb{H}^2$ *we write* $\xi_t = \partial_\infty \tau(t)$.

The group $G = \mathrm{PSL}_d(\mathbb{R})$ identifies with one component of the space of $\mathbb{H}^2$-frames Y introduced in the previous section via the orbit map $G \to Y$; $g \mapsto g \cdot F_o$, where $F_o = (o, v_o, w_o)$ is a fixed $\mathbb{H}^2$-frame, so that o is fixed by r_θ, and $v_o = \frac{d}{dt}_{|t=0} a'_t \cdot o$ and $w_o = \frac{d}{dt}_{|t=0} a_t \cdot o$ are tangent to the axes of a'_t and a_t, respectively.

Under this identification, the geodesic flow on Y corresponds to the multiplication on the right by a'_t: i.e. $\mathrm{geod}_t(gF_o) = (ga'_t)F_o$. On the other hand, the orthogonal sliding flow corresponds to the multiplication on the right by $\exp(z)$: that is, $\mathrm{slide}_z(gF_o) = (g \cdot \exp(z))F_o$ for any $z \in \mathfrak{a}$. This leads us to the following definition of admissible path.

Definition 2.8. A path $c : [0, T] \to G$ or $c : [0, \infty) \to G$ is said to be of

- *flat type* if $c(t) = g \cdot \exp(tz)$ for some $g \in G$ and $z \in \mathfrak{a}$ of norm 1 for the Finsler metric $\mathfrak{F}$;
- *hyperbolic type* if $c(t) = ga'_t$ for some $g \in G$.

An *admissible path* of G is a *continuous* (possibly infinite) concatenation of paths of flat and hyperbolic type.

It is moreover called (ω, L)-*admissible* for some parameters $\omega, L > 0$ if all hyperbolic (resp. flat) pieces, except maybe the first and last pieces, have length at least ω (resp. L).

A (ω, L)-*admissible path in* $\mathbb{X}$ is a path of the form $t \mapsto c(t) \cdot \mathbf{x}$ where c is a (ω, L)-admissible path of G; note that it is piecewise geodesic.

Remark . *Another way to describe admissible paths in G is the following: a path $c : [0, T] \to G$ is admissible, starting with a hyperbolic piece, if there exist $t_0 = 0 < t_1 < t_2 < \cdots < t_{n-1} < t_n = T$ and $z_1, z_3, \ldots, z_k \in \mathfrak{a}$ of norm 1 (where k is the biggest odd integer $< n$) such that for any $0 \le i < n$, for any $t \in [0, t_{i+1} - t_i]$,*

- *if i is even then $c(t_i + t) = c(t_i) \cdot a'_t$,*
- *if i is odd then $c(t_i + t) = c(t_i) \cdot \exp(tz_i)$.*

The following is fairly immediate from the definition of the construction of the characteristic surface of a Hitchin grafting representation ρ_z and the map $\tilde{Q}_z$ from Proposition 2.5. In its formulation, the *collar size* of a simple closed multi-geodesic $\gamma^* \subset S$ is the supremum of all numbers $r > 0$ such that the tubular neighborhood of radius r about γ^* is a union of annuli about the components of γ^*. By hyperbolic geometry, an upper bound on the length of the components of γ^* yields a lower bound on the collar size of γ^*.

Observation 2.9. *Consider a closed hyperbolic surface S, a multicurve $\gamma^* \subset S$ with collar size ω and a grafting parameter z such that all cylinder heights are at least L. Then the image under the grafting map $\tilde{Q}_z$ of any admissible path of $\tilde{S}_z$ is a (ω, L)-admissible path of $\mathbb{X}$.*

Remark 2.10. We define admissible loops *of a quotient of* $\mathbb{X}$ *as quotients of periodic infinite admissible paths.*

In Section 4 we recall the notion of positivity in G in the language of Lusztig [Lus94] and some basic results, and see that admissible paths have interesting positivity properties, coming from the fact that $(a'_t)_{t>0}$ are totally positive matrices, and $(\exp(z))_{z\in\mathfrak{a}}$ are totally nonnegative matrices.

3 A Morse-type lemma in the symmetric space

The goal of this section is to establish a Morse-type lemma in the symmetric space. Recall that in δ-hyperbolic geodesic metric spaces, the Morse lemma says that any (λ, C)-quasi-geodesic is at distance at most C' from a geodesic, where C' depends on δ, λ, C. In $\mathbb{R}^2$, equipped with any norm, this lemma does not hold true anymore for all quasi-geodesics, and the same can be said for the higher-rank symmetric spaces endowed with a Finsler or Riemannian metric, since they contain totally geodesic copies of $\mathbb{R}^2$. Recall that a path $c(t)$ in a metric space is a (λ, C)-quasi-geodesic if for all times t, s we have

$$\lambda^{-1}|t-s| - C \leq d(c(t), c(s)) \leq \lambda|t-s| + C.$$

Kapovich–Leeb–Porti [KLP18] (see also Section 12.1 of [KL18]) proved a Morse lemma for certain families of well-behaved quasi-geodesics in symmetric spaces of arbitrary rank, equipped with the standard Riemannian metric, and in Euclidean buildings, see also Section 7 of [BPS19]. There is also a version of the Morse Lemma for quasi-flats instead of quasi-geodesics, see [KL97; EF97]. We propose here a different approach to a generalization of the Morse lemma: we prove that nearby every *Finsler* $(1, C)$-quasi-geodesic, so with *multiplicative* error term of 1, there is at least one Finsler geodesic. Other Finsler geodesics could be far, as in $(\mathbb{X}, d^{\mathfrak{F}})$ there are Finsler geodesics with the same endpoints and arbitrarily large Hausdorff distance.

We first present our result using the notion of quasi-ruled paths as in [BHM11], and then translate it in terms of $(1, C)$-quasi-geodesics. A *C-quasi-ruled path* in a metric space (X, d_X) is a map $c : [0, T] \to X$ such that for any $0 \leq t \leq s \leq u \leq T$,

$$d_X(c(t), c(s)) + d_X(c(s), c(u)) \leq d_X(c(t), c(u)) + C.$$

Note that any reparameterization of a quasi-ruled path is quasi-ruled.

Theorem 3.1. *For any $C > 0$ there exists $C' > 0$ such that any Finsler C-quasi-ruled continuous path $c : [0, T] \to \mathbb{X}$ is at Hausdorff distance at most C' from a Finsler geodesic in $(\mathbb{X}, d^{\mathfrak{F}})$ connecting $c(0)$ to $c(T)$.*

One can translate the above theorem in terms of $(1, C)$-quasi-geodesic paths, using the following lemma, which is probably well-known to experts. We provide a proof in Subsection 3.1.

Lemma 3.2. *Let (X, d) be a geodesic metric space and $C \geq 0$. Any $(1, C)$-quasi-geodesic in X is $3C$-quasi-ruled and is at Hausdorff distance at most $1 + C$ from a continuous $(1, 2(1 + C))$-quasi-geodesic with the same endpoints.*

Conversely, any continuous C-quasi-ruled path is at Hausdorff distance at most $3C$ from a $(1, 3C)$-quasi-geodesic.

Note that for $C = 0$ this lemma says that any $(1, 0)$-quasi-geodesic is a (continuous, 0-quasi-ruled) geodesic, and that for any continuous 0-quasi-ruled path $c : [0, T] \to X$ from x to y there exists a geodesic $c' : [0, d(x, y)] \to X$ from x to y whose image is exactly the same as that of c.

Remark . *Theorem 3.1 is false for the Euclidean metric on $\mathbb{R}^d$ $(d \geq 2)$, and for the Riemannian metric on $\mathbb{X}$, as can be seen as follows.*

Let $\ell : \mathbb{R} \to \mathbb{R}^d$ be a line through $\ell(0) = 0$ parameterized by arc length for the Euclidean metric. For $n \geq 1$ put $x_n = \ell(-n), y_n = \ell(n)$. Consider the balls B_n^-, B_n^+ of radius n about x_n, y_n. As $n \to \infty$, the boundaries $\partial B_n^\pm$ of the balls $B_n^\pm$ converge in the pointed Gromov Hausdorff topology of $(\mathbb{R}^d, 0)$ to the hyperplane through 0 orthogonal to ℓ. Thus for any $m > 0$ and sufficiently large n, there are points $z_n^\pm$ on $\partial B_n^\pm$ of distance m to ℓ with $d(z_n^-, z_n^+) \leq 1$ (here d is the euclidean distance). Since the subsegment of ℓ connecting x_n to y_n is the unique euclidean geodesic between these points, the piecewise geodesics connecting x_n to y_n with breakpoints at z_n^-, z_n^+ violate the conclusion of Theorem 3.1.

This section is subdivided into four subsections. The first subsection is very short and provides a proof of Lemma 3.2 for the reader's convenience. In the second subsection, which is the longest, we establish Theorem 3.1 for polyhedral norms on $\mathbb{R}^d$, that is, norms whose norm one ball is a finite sided symmetric convex polyhedron. In the third subsection, we prove that any quasi-ruled continuous path in the symmetric space lies near a flat, using a description of Finsler horoballs of Kapovich–Leeb [KL18]. Finally the last section, which is very brief, contains the proof of Theorem 3.1.

3.1 Proof of Lemma 3.2

Let (X, d) be a geodesic metric space. Then for any $(1, C)$-quasi-geodesic $c : [0, T] \to X$ the following holds true.

- c is $3C$-quasi-ruled.
- c is at Hausdorff distance at most $1 + C$ from a continuous $(1, 2 + 2C)$-quasi-geodesic with the same endpoints.

The first property is an elementary computation which is left to the reader, and the second property is Lemma 1.11 in Chapter III.H of [BH99]. The two properties together yields the first part of the lemma.

Let us prove the second part of the lemma. For an arbitrary $C \geq 0$ consider a continuous C-quasi-ruled path $c : [0, T] \to X$ from x to y.

We now use an idea we found in the arXiv version of [BHM11], Lemma A.2. Observe that the map

$$f : c[0,T] \to [0, d(x,y)+C]; \quad z \mapsto d(x,z)$$

is continuous and is a $(1,C)$-quasi-isometry. Indeed, if $0 \leq t \leq s \leq T$ then

$$|f(c(t)) - f(c(s))| = |d(x,c(t)) - d(x,c(s))| \leq d(c(t),c(s)),$$

and

$$d(c(t),c(s)) \leq d(x,c(s)) - d(x,c(t)) + C \leq |f(c(t)) - f(c(s))| + C.$$

Moreover, since $c[0,T]$ is path-connected, and f is continuous and attains the values 0 and $d(x,y)$, by the Intermediate Value Theorem f attains all values in $[0,d(x,y)]$ and hence f is C-quasi-surjective. By a classical result from coarse geometry f admits a $(1,3C)$-quasi-inverse $g : [0,d(x,y)] \to c[0,T] \subset X$.

3.2 A Morse-type lemma for normed vector spaces

This subsection is entirely devoted to the study of the geometry of $\mathbb{R}^d$, equipped with a Finsler metric defined by a translation invariant norm on $T\mathbb{R}^d$. We begin with defining the Finsler metrics we are interested in. To this end call a cone in $\mathbb{R}^d$ *properly convex* if it is convex and its closure does not contain any affine subspace of $\mathbb{R}^d$ of dimension at least 1.

A (symmetric) polyhedral norm $|\cdot|$ on $\mathbb{R}^d$ is a norm of the form

$$|v| = \max\{\alpha(v) : \alpha \in \mathcal{A}\},$$

where $\mathcal{A}$ is a finite set of nonzero linear forms which spans $(\mathbb{R}^d)^*$, and which is symmetric in the sense that $-\mathcal{A} = \mathcal{A}$. This norm induces a metric $d(x,y) = |x-y|$ on $\mathbb{R}^d$ that is invariant under translations. The goal of this section is to show.

Proposition 3.3. *For any polyhedral norm $|\cdot|$ on $\mathbb{R}^d$, there exists $\mu > 0$ such that for any $C \geq 1$, any C-quasi-ruled continuous path $c : [0,T] \to \mathbb{R}^d$ is at Hausdorff distance at most μC from a geodesic in $(\mathbb{R}^d, |\cdot|)$ connecting $c(0)$ to $c(T)$.*

Note that this statement is false for a Euclidean norm on $\mathbb{R}^d$.

Proposition 3.3 has the following reformulation in terms of $(1,C)$-quasi-geodesics, thanks to Lemma 3.2.

Corollary 3.4. *For any polyhedral norm $|\cdot|$ on $\mathbb{R}^d$, there exists $\mu > 0$ such that for any $C \geq 1$, any $(1,C)$-quasi-geodesic $c : [0,T] \to \mathbb{R}^d$ is at Hausdorff distance at most μC from a geodesic in $(\mathbb{R}^d, |\cdot|)$ from $c(0)$ to $c(T)$.*

3.2.1 Diamonds

In this section we introduce several geometric objects relative to our polyhedral norm, including diamonds. We prove that $(1, C)$-quasi-geodesics stay at bounded distance from diamonds, which is the main technical step towards the proof of Proposition 3.3.

For any $\alpha \in \mathcal{A}$, the set $\mathcal{C}_\alpha = \{v \in \mathbb{R}^d : |v| = \alpha(v)\}$ is a polyhedral convex cone based at 0. Note that $\mathcal{C}_{-\alpha} = -\mathcal{C}_\alpha$. Up to removing unnecessary elements of $\mathcal{A}$, we may assume that $\mathcal{C}_\alpha$ has nonempty interior. We call *special cones* (based at $x \in \mathbb{R}^d$) the cones of $\mathbb{R}^d$ that are translates of a cone $\mathcal{C}_\alpha$ (by the translation $y \mapsto y + x$).

The unit closed ball $\bar{B}(0, 1)$, and more generally any closed ball $\bar{B}(x, r)$ for such a norm, is a polyhedral convex set, that is, a finite intersection of (affine) half-spaces of $\mathbb{R}^d$. More precisely,

$$\bar{B}(x, r) = \bigcap_{\alpha \in \mathcal{A}} \{y \in \mathbb{R}^d : \alpha(y - x) \leq r\} = x + r \cdot \bar{B}(0, 1).$$

The codimension-1 faces of $\bar{B}(x, r)$ are the intersections of its boundary $\partial B(x, r)$ with the special cones based at x.

Definition 3.5. Denote by $\mathcal{C}(x \to y)$ the intersection of all special cones based at x that contain y. We define the *diamond* of the pair x, y to be $D(x, y) = \mathcal{C}(x \to y) \cap \mathcal{C}(y \to x)$ (see Figures 3 and 4 for illustrations).

Note that $\mathcal{C}(y \to x) = y - x - \mathcal{C}(x \to y)$. This follows from the fact that for any $\alpha \in \mathcal{A}$, the special cone $x + \mathcal{C}_\alpha$ based at x contains y if and only if the special cone $y + \mathcal{C}_{-\alpha}$ based at y contains x.

Lemma 3.6. *For any $x, y \in \mathbb{R}^d$, we have*

$$\begin{aligned} D(x, y) &= \{z \in \mathbb{R}^d \mid d(x, z) + d(z, y) = d(x, y)\} \\ &= \cup\{\textit{geodesics from } x \textit{ to } y\} \end{aligned}$$

In particular, for any $z \in D(x, y)$, the concatenation of a geodesic from x to z with a geodesic from z to y is a geodesic from x to y.

Given a cone $\mathcal{C} = \mathcal{C}(0 \to y)$, set $\alpha_\mathcal{C}$ to be the mean of all $\alpha \in \mathcal{A}$ for which $\mathcal{C}_\alpha$ contains $\mathcal{C}$. It follows from the definition that $|z| \geq \alpha_\mathcal{C}(z)$ holds for all $z \in \mathbb{R}^d$, with equality exactly on $\mathcal{C}$.

Proof. For a cone $\mathcal{C} = \mathcal{C}(x \to y)$, consider the form $\alpha_\mathcal{C}$ defined above. The point z belongs to $D(x, y)$ if and only if $|z - x| = \alpha_{\mathcal{C}(x \to y)}(z - x)$ and $|z - y| = \alpha_{\mathcal{C}(y \to x)}(z - y) = -\alpha_{\mathcal{C}(x \to y)}(z - y)$. This implies

$$d(x, y) \leq d(x, z) + d(z, y) = \alpha_{\mathcal{C}(x \to y)}(z - x) + \alpha_{\mathcal{C}(x \to y)}(y - z) = \alpha_{\mathcal{C}(x \to y)}(y - x) = d(x, y)$$

and so $d(x, y) = d(z, x) + d(z, y)$.

Conversely if $|z - x| > \alpha_{\mathcal{C}(x \to y)}(z - x)$ or $|z - y| > -\alpha_{\mathcal{C}(x \to y)}(y - z)$, then the above inequality yields $d(z, x) + d(z, y) > d(x, y)$.

It is clear that if z lies on a geodesic from x to y, then $d(x,z)+d(z,y)=d(x,y)$. Reciprocally, if $d(x,z)+d(z,y)=d(x,y)$ then the concatenation of any geodesic from x to z and any geodesic from z to y is a geodesic from x to y. □

From the previous lemma and the triangle inequality we infer that for any point z not too far from a diamond $D(x,y)$ we almost have the triangle equality $d(x,z)+d(z,y)\simeq d(x,y)$. The following lemma is the key technical result towards the proof of Proposiion 3.3; it says that the converse also holds.

Lemma 3.7. *There is $\lambda_1>0$ such that for all $x,y,z\in\mathbb{R}^d$ it holds*

$$d(z,D(x,y))\leq\lambda_1(d(x,z)+d(z,y)-d(x,y)).$$

The two terms are equal to zero when z belongs to $D(x,y)$. One can think of the lemma in the following way: $z\mapsto f_{x,y}(z)=d(x,z)+d(z,y)-d(x,y)$ is convex, non-negative, and piecewise affine. Take a point $z\in\partial D(x,y)$ and follow a ray $\{z+tv,t\geq 0\}$ for a unit vector v for $|\cdot|$ at z whose euclidean angle (for some fixed euclidean inner product) to $D(x,y)$ is at least $\pi/2$. By this we mean the angle between v and any line segment in $D(x,y)$ starting at z. The restriction of $f_{x,y}$ to the ray is convex, piecewise affine and is equal to zero exactly at z. It follows that it grows at least linearly in t, the slope being given by the derivative at $t=0$. And so for $z'=z+tv$, one has $f_{x,y}(z')\geq t\cdot f'_{x,y}(0)\geq \mathrm{Cst} f'_{x,y}(0)\cdot d(z',D(x,y))$.

The issue is that the slope does not vary continuously in z, not even lower semi-continuously, so one can not hope to use a compactness argument to obtain a uniform bound on the union of the rays. One might study carefully the combinatorics of the map f to obtain a uniform bound on the slope. We instead take a slightly different approach, which requires one intermediate lemma.

Let us fix a Euclidean inner product $\langle,\rangle$ defining the Euclidean metric d_{eucl} on $\mathbb{R}^d$. By "orthogonal projection" to a closed convex set C we will mean closest-point projection for d_{eucl} to C, which is well defined by convexity of d_{eucl}.

Given a cone $\mathcal{C}\subset\mathbb{R}^d$ based at 0, define the dual cone of $\mathcal{C}$ to be the set

$$\mathcal{C}'=\{x\in\mathbb{R}^d,\langle x,\mathcal{C}\rangle\leq 0\}=\{x\in\mathbb{R}^d \text{ whose orthogonal projection to } \mathcal{C} \text{ is } 0\}.$$

Lemma 3.8. *Let $\mathcal{C}$ be a polyhedral convex cone of $\mathbb{R}^d$ based at 0, that is, the intersection of finitely many closed half-spaces $H_1,\dots,H_n$ containing 0 in their boundary. Let $\mathcal{C}'\subset\mathbb{R}^d$ be the polyhedral convex dual cone to $\mathcal{C}$. Then there exists $\lambda>0$ such that for any $x\in\mathcal{C}'$,*

$$d_{\mathrm{eucl}}(x,\mathcal{C})\leq\lambda\max\left(d_{\mathrm{eucl}}(x,H_1),\dots,d_{\mathrm{eucl}}(x,H_n)\right).$$

The results holds true for all $x\in\mathbb{R}^d$ (for a bigger constant λ), but the special case $x\in\mathcal{C}'$ is shorter to prove.

Proof. This is an immediate consequence of the fact that the function

$$f(x)=\max\left(d_{\mathrm{eucl}}(x,H_1),\dots,d_{\mathrm{eucl}}(x,H_n)\right)$$

is homogeneous, continuous, and positive on $\mathcal{C}'-\{0\}$. □

Note that in the previous lemma we allow $\mathcal{C}$ to have empty interior, or be reduced to $\{0\}$, or to be the entire space $\mathbb{R}^d$ (but this last case is not very interesting since then the dual $\mathcal{C}'$ is just $\{0\}$).

Proof of Lemma 3.7. Denote by $\{H_\alpha, \alpha \in \mathcal{A}\}$ the finite family of closed half spaces given by $H_\alpha = \{w \in \mathbb{R}^d, \alpha(w) \leq 0\}$.

For any subset S of $\mathcal{A}$, the intersection

$$\mathcal{C}_S = \cap_{\alpha \in S} H_\alpha$$

is a polyhedral convex cone. Let K_S be the dual cone to $\mathcal{C}_S$. We can apply Lemma 3.8 to K_S, and get a number $\lambda_S > 0$. Let $\lambda = \max\{\lambda_S \mid S \subset \mathcal{A}\}$.

We now prove the inequality for $x, y, z \in \mathbb{R}^d$ fixed. When z belongs to $D(x, y)$, we have $d(x, z) + d(z, y) - d(x, y) = 0 = d(z, D(x, y))$ by Lemma 3.6 and so the inequality holds.

Suppose that $z \notin D(x, y)$. If $d(x, z) \geq d(x, y)$ holds, then one has

$$d(x, z) + d(z, y) - d(x, y) \geq d(z, y) \geq d(z, D(x, y))$$

since $y \in D(x, y)$. So by symmetry in x, y we may assume that $r := d(x, z)$ is smaller than $R := d(x, y)$.

Let $B_x = \bar{B}(x, r)$ and $B_y = \bar{B}(y, R - r)$ be closed balls for the polyhedral norm $|\cdot|$, illustrated in Figure 3. They are polyhedral convex sets, i.e. finite intersections of affine half-spaces. More precisely

$$B_x = H_1 \cap \cdots \cap H_n \quad \text{and} \quad B_y = H'_1 \cap \cdots \cap H'_n$$

where $H_i = \{w \in \mathbb{R}^d, \alpha_i(w - x) \leq r\}$ and $H'_j = \{w \in \mathbb{R}^d, \alpha'_j(w - y) \leq R - r\}$ for some orderings $\mathcal{A} = \{\alpha_1, \ldots, \alpha_n\} = \{\alpha'_1, \ldots, \alpha'_n\}$ (it will be convenient later to have two different orderings).

The intersection

$$B_x \cap B_y = H_1 \cap \cdots \cap H_n \cap H'_1 \cap \cdots \cap H'_n$$

is a closed polyhedral convex subset of the diamond $D(x, y)$ by Lemma 3.6, with empty interior, and it is not empty (one can verify that $B_x \cap B_y$ contains the point $\frac{r}{R} y + \frac{R-r}{R} x$).

Let p be the Euclidean closest-point projection of z to $B_x \cap B_y$. Up to translation, we may assume that $p = 0$ to be able to use Lemma 3.8. Up to reordering we may also assume that the half-spaces containing p in their boundary are $H_1, \ldots, H_k$ and $H'_1, \ldots, H'_\ell$; note that k and ℓ are both positive since $p \in \partial B_x \cap \partial B_y$. Since $p = 0$ we have $H_i = \{w : \alpha_i(w) \leq 0\}$ for $i \leq k$ and $H'_j = \{w : \alpha'_j(w) \leq 0\}$ for $j \leq l$. Let $S = \{\alpha_i, i \leq k\} \cup \{\alpha'_j, j \leq \ell\} \subset \mathcal{A}$. Then using the notation introduced at the beginning of the proof we have

$$B_x \cap B_y \subset H_1 \cap \cdots \cap H_k \cap H'_1 \cap \cdots \cap H'_\ell = \mathcal{C}_S.$$

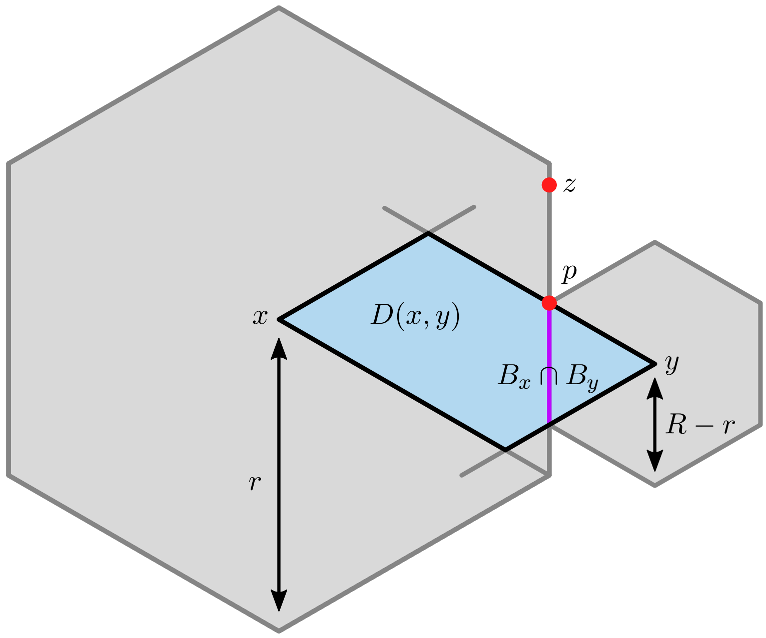

Figure 3: Illustration of the proof of Lemma 3.7. Three points x, y, z and a fourth point $p \in D(x,y)$ with $d(z,p) \leq \lambda(d(x,z) + d(y,z) - d(x,y))$.

Observe that p is also the Euclidean closest-point projection of z to $\mathcal{C}_S$. Indeed if by contradiction there was $p' \in \mathcal{C}_S$ closer to z, then by convexity of the euclidean distance function, any point of the line segment $(p,p']$ would be closer to z. But any point of $(p,p']$ close enough to p is contained in each $H_{k+1}, \dots, H_n$ and $H'_{\ell+1}, \dots, H'_n$ (since p is in their interior), and hence is in $B_x \cap B_y$, which contradicts that p is the Euclidean closest-point projection of z on $B_x \cap B_y$. In particular, z is contained in the dual cone K_S to $\mathcal{C}_S$.

By Lemma 3.8, the distance $d(z,p)$ is comparable to the distance between z and one of the half spaces $H_1, \dots, H_k, H'_1, \dots, H'_\ell$. But since $z \in B_x$, it must be contained in every H_i, so Lemma 3.8 implies that there exists a half space H'_j, $j \in \{1, \dots, \ell\}$, for which

$$\begin{aligned} d_{\mathrm{eucl}}(z, D(x,y)) &\leq d_{\mathrm{eucl}}(z,p) \\ &\leq \lambda_S d_{\mathrm{eucl}}(z, H'_j) \\ &\leq \lambda d_{\mathrm{eucl}}(z, B_y). \end{aligned}$$

Let us translate what this means for the polyhedral norm, using a constant ν such that $\nu^{-1} d \leq d_{\mathrm{eucl}} \leq \nu d$. The previous equation yields

$$d(z, D(x,y)) \leq \nu^2 \lambda d(z, B_y).$$

Let q be the intersection point of $[y,z]$ with ∂B_y, which satisfies $d(z,q) = d(z, B_y)$. Indeed for any $q' \in \partial B_y$ we have $d(q',y) = d(q,y)$, and so

$$d(z,q') = d(z,q') + d(q',y) - d(q',y) \geq d(z,y) - d(q,y) = d(z,q) = R - r$$

since $[y, z]$ is a geodesic for d. Then

$$\begin{aligned} d(x,z) + d(z,y) - d(x,y) &= r + d(z,y) - R \\ &= d(z,q) + d(q,y) + r - R \\ &= d(z,q) = d(z, B_y) \\ &\geq \nu^{-2}\lambda^{-1} d(z, D(x,y)). \end{aligned}$$ ☐

3.2.2 Proof of Proposition 3.3

In addition to diamonds we shall consider two other geometric objects: crowns and cores of diamonds. We refer to [KLP18] for closely related constructions and to Figure 4 for an illustration.

- $x, y \in \mathbb{R}^d$ are called *generic* if the diamond between them has nonempty interior, that is, if there is only one $\alpha \in \mathcal{A}$ such that $|x - y| = \alpha(x - y)$.
- The *Crown* $\mathrm{Cr}(x, y)$ is $\partial\mathcal{C}(x \to y) \cap \partial\mathcal{C}(y \to x)$ if x, y are generic, and otherwise it is just $D(x, y)$. Note that if $z \in \mathrm{Cr}(x, y)$ then the pair (x, z) is *not* generic.
- The *Core* $\mathrm{Co}(x, y)$ is the convex hull of the crown, which is just $D(x, y)$ if x and y are not generic.

Note that for generic x, y, the core $\mathrm{Co}(x, y)$ separates x from y in $D(x, y)$, in the sense that x and y are in different connected components of $D(x, y) \smallsetminus \mathrm{Co}(x, y)$.

Note also that for generic x, y the intersection of the core $\mathrm{Co}(x, y)$ with the boundary $\partial D(x, y)$ of the diamond is exactly the crown $\mathrm{Cr}(x, y)$.

We will need the following elementary result on properly convex cones, which implies that if a tip of a diamond is not too close to the crown then it is not too close to the core.

Given a set $K \subset \mathbb{R}^d$, denote by $\mathrm{Conv}(K)$ the convex hull of K.

Lemma 3.9. *Let $\mathcal{C} \subset \mathbb{R}^r$ be a closed properly convex cone with vertex 0. Then there exists $\lambda_2 > 0$ such that for any compact set $K \subset \mathcal{C}$, we have*

$$d(0, K) \leq \lambda_2 d(0, \mathrm{Conv}(K)).$$

Proof. Since $\mathcal{C}$ is properly convex, 0 is an extremal point of it and does not belong to the convex hull $\mathrm{Conv}(\mathcal{C} - B(0,1))$ where $B(0,1)$ is the open ball of radius one around 0 for the metric d. Let $\lambda^{-1} = d(0, \mathrm{Conv}(\mathcal{C} - B(0,1)))$ be the distance from 0 to $\mathrm{Conv}(\mathcal{C} - B(0,1))$.

Let $K \subset \mathcal{C}$ be a compact subset. If $0 \in K$ then $d(0, K) = 0 \leq \lambda d(0, \mathrm{Conv}(K))$.

Suppose $0 \notin K$ and and put $a = d(0, K)^{-1}$. Then the compact $a \cdot K$ is included in $\mathcal{C} - B(0,1)$ and so $\mathrm{Conv}(a \cdot K) \subset \mathrm{Conv}(\mathcal{C} - B(0,1))$, which yields $d(0, \mathrm{Conv}(a \cdot K)) \geq \lambda^{-1}$. As $\mathrm{Conv}(a \cdot K) = a \cdot \mathrm{Conv}(K)$, one has

$$d(0, \mathrm{Conv}(K)) = d(0, K) \cdot d(0, a \cdot \mathrm{Conv}(K)) \geq \lambda^{-1} d(0, K).$$ ☐

Finally we will need the following observation about quasi-ruled paths, whose proof is a simple calculation.

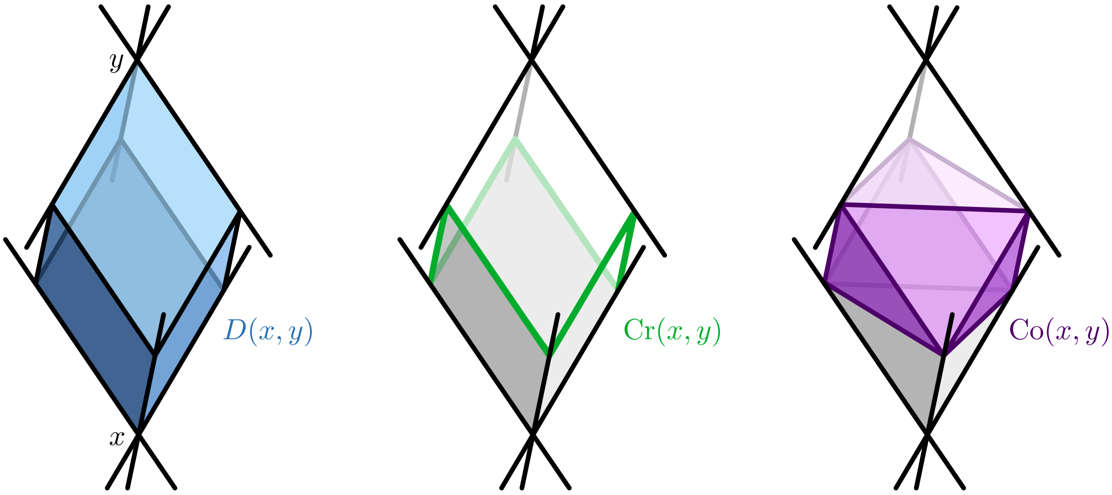

Figure 4: Illustration of the diamond, crown and core of two points $x, y \in \mathbb{R}^d$ in generic position.

Observation 3.10. *Let $a, b : [0, T] \to X$ and $c : [T, T'] \to X$ be paths in a metric space.*

1. *If a is C-quasi-ruled and $d(a(t), b(t)) \leq C'$ for all t then b is $C + 6C'$-quasi-ruled.*
2. *If a is C-quasi-ruled and if $d(c(t), a(T)) \leq C'$ for all t then the concatenation of the path a with c is $C + 2C'$-quasi-ruled.*

Proof of Proposition 3.3. We proceed by induction on the dimension d. In the case $d = 1$ there is nothing to show, so assume that for some $d \geq 2$, the claim holds true for all dimensions $< d$, with a constant depending on the dimension and on the polyhedral norm.

Let $|\cdot|$ be a polyhedral norm on $\mathbb{R}^d$. Note that the restriction of $|\cdot|$ to any linear subspace is polyhedral. Thus by the induction assumption, there exists a constant $\mu > 1$ so that the proposition is valid with this constant for paths contained in the linear subspaces $\{\alpha_{i_1} = \cdots = \alpha_{i_k} \mid \alpha_{i_j} \in \mathcal{A}\}$ which are the linear spans of the faces of the special cones $\mathcal{C}_\alpha$ for $\alpha \in \mathcal{A}$.

Let $\lambda_1 > 0$ be as in Lemma 3.7, $\nu > 1$ be such that $\nu^{-1} d \leq d_{\mathrm{eucl}} \leq \nu d$ where $d(x, y) = |x - y|$, and $\lambda_1' = \nu^2 \lambda_1$. Then for all $x, y \in \mathbb{R}^d$, if

$$\Pi_{xy} : \mathbb{R}^d \to D(x, y)$$

is the Euclidean closest-point projection onto $D(x, y)$ (which is well defined continuous since $D(x, y)$ is compact and convex, contrarily to the closest-point projection for d), then by Lemma 3.7 we get

$$d(z, \Pi_{xy}(z)) \leq \nu d_{\mathrm{eucl}}(z, \Pi_{xy}(z)) = \nu d_{\mathrm{eucl}}(z, D(x, y)) \leq \lambda_1'(d(x, z) + d(z, y) - d(x, y)). \tag{11}$$

Let $\lambda_2 > 0$ be the maximum of the constants from Lemma 3.9, applied to the special cones $\mathcal{C}_\alpha$, $\alpha \in \mathcal{A}$. Then for all $x, y \in \mathbb{R}^d$ we have

$$d(x, \mathrm{Cr}(x,y)) \leq \lambda_2 d(x, \mathrm{Co}(x,y)). \tag{12}$$

We claim that the statement of the proposition holds true for $(\mathbb{R}^d, |\cdot|)$ with the constant $\mu' = (1+2\lambda)(\mu(1+6\lambda_1') + \lambda_1')$, where $\lambda = \max(\lambda_1', \lambda_2)$.

To this end we proceed by induction on k where $d(x,y) \in (k-1, k]$.

Let $C \geq 1, k = 1$ and let c be a C-quasi-ruled continuous path from x to y such that $d(x,y) \leq 1$. Then the segment $[x,y]$ is at distance at most $C+1$ from c and hence the claim holds true in this case.

Let $k \geq 2$ and assume that the claim holds true for all continuous C-quasi-ruled paths from x to y such that $d(x,y) \leq k-1$. Let c be a C-quasi-ruled continuous path from x to y such that $d(x,y) \in (k-1,k]$.

There are two possible cases. In the first case, c stays $C\lambda$-far from the crown $\mathrm{Cr}(x,y)$.

Fix a point $z \in \mathrm{Co}(x,y)$ which is in the interior of the diamond $D(x,y)$. Let $Z \subset \mathrm{Co}(x,y)$ be the union of segments $[z,p]$ where $p \in \mathrm{Cr}(x,y)$. Then $Z \cap \partial D(x,y) = \mathrm{Cr}(x,y)$ and Z separates x from y in $D(x,y)$: if $a(t) \in D(x,y)$ is a continuous path from x to y then it must cross Z. If it crosses z there is nothing to prove. Otherwise, for each t the ray from z passing through $a(t)$ must cross $\partial D(x,y)$ at some point $b(t)$ that depends continuously on t, and at some time t we have $b(t) \in \mathrm{Cr}(x,y)$ hence $a(t) \in [b(t), z] \subset Z$.

The projection $\Pi_{xy} \circ c$ is a continuous path from x to y in $D(x,y)$, so it must cross Z at some time t.

Note that we have $\Pi_{xy}(c(t)) = c(t)$. Namely, otherwise $\Pi_{xy}(c(t))$ is contained in the boundary $\partial D(x,y)$, and hence contained in the crown. But since $\Pi_{xy}(c(t))$ is $C\lambda_1'$-close to $c(t)$ by Inequality (11), then $c(t)$ is $C\lambda_1'$-close and hence $C\lambda$-close to the crown, which contradicts our assumption.

Moreover, we must also have the inequalities

$$d(x, c(t)) \geq d(x, \mathrm{Co}(x,y)) \geq 1 \text{ and } d(y, c(t)) \geq 1.$$

Otherwise by Equation (12), it holds $d(x, \mathrm{Cr}(x,y)) \leq \lambda_2 \leq \lambda \leq C\lambda$, which contradicts our assumption that c stays $C\lambda$-away from the crown. Lemma 3.6 yields that

$$d(x, c(t)) = d(x,y) - d(y, c(t)) \leq k-1$$

and similarly $d(y, c(t)) \leq k-1$. We can apply the induction hypothesis (on k) to the path $c[0,t]$ and the path $c[t,T]$. As the concatenation of a geodesic connecting x to $c(t) \in D(x,y)$ and $c(t)$ to y is a geodesic, this suffices for the induction step.

In the second case, c passes at some time t at distance less than $C\lambda$ from the crown. Let $p \in \mathrm{Cr}(x,y)$ be such that $d(p, c(t)) \leq C\lambda$. Recall that this means the pairs (x,p) and (y,p) are not generic.

Concatenate $c_1 = c[0,t]$ with a geodesic from $c(t)$ to p, to get a continuous $(1+2\lambda)C$-quasi-ruled path c_1' from x to p by Observation 3.10. By Inequality (11), the projection

$\Pi_{xp} \circ c_1'$ is at distance at most $\lambda_1'(1+2\lambda)C$ from c_1. Using again Observation 3.10, it is a continuous $(1+2\lambda)C(1+6\lambda_1')$-quasi-ruled path in $D(x,p)$ from x to p.

As the pair (x,p) is not generic, we can apply our induction on the dimension and deduce that there is a geodesic $c_1'' \subset D(x,p)$ at distance at most $\mu C(1+2\lambda)(1+6\lambda_1')$ from $\Pi_{xp} \circ c_1'$, which is then at distance at most $C(1+2\lambda)(\mu(1+6\lambda_1')+\lambda_1')$ from the original path c.

With a similar construction for $c_2 = c[t,T]$, one obtains a geodesic c_2'' from p to y which is at distance at most $C(1+2\lambda)(\mu(1+6\lambda_1')+\lambda_1')$ from c.

The concatenation of c_1'' and c_2'' is by Lemma 3.6 a geodesic from x to y at distance at most $C(1+2\lambda)(\mu(1+6\lambda_1')+\lambda_1')$ from c, which concludes the proof. □

3.3 Projecting to a flat

In this subsection we extend Proposition 3.3 to the symmetric space $\mathbb{X} = \mathrm{PSL}_d(\mathbb{R})/\mathrm{PSO}(d)$ equipped with the Finsler metric $d^{\mathfrak{F}}$. We begin with extending the geometric notions from Section 3.2 to $\mathbb{X}$.

Recall, for instance from [KL18], that the *diamond* between two points $x, y \in \mathbb{X}$ is defined as follows. A *Weyl cone* of a flat of $\mathbb{X}$ is the translation under an element of $\mathrm{PSL}_d(\mathbb{R})$ of the standard Weyl cone $\exp \mathfrak{a}^+ \subset \exp \mathfrak{a}$ based at the basepoint $\mathbf{x}$ of the standard flat $\exp \mathfrak{a}$. Consider a flat F containing x and y, a Weyl cone $W \subset F$ based at x and containing y, and the opposite Weyl cone W' based at y (which automatically contains x). Then the *diamond* $D(x,y)$ is defined as

$$D(x,y) = W \cap W'.$$

It does not depend on choices, and as an intersection of convex subsets of $\mathbb{X}$, it is convex. The analog of Lemma 3.6 holds true.

Proposition 3.11 (Lemma 5.10 of [KL18]). *For all $x, y \in \mathbb{X}$, the diamond $D(x,y)$ is the set of points $z \in \mathbb{X}$ such that $d^{\mathfrak{F}}(x,z) + d^{\mathfrak{F}}(z,y) = d^{\mathfrak{F}}(x,y)$.*

The following non-uniform version of Lemma 3.7 for $(\mathbb{X}, d^{\mathfrak{F}})$ is used to reduce Theorem 3.1 to Proposition 3.3.

Proposition 3.12. *For any $C \geq 1$ there exists $C' > 0$ such that for all $x, y \in \mathbb{X}$, any $z \in \mathbb{X}$ such that $d^{\mathfrak{F}}(x,z) + d^{\mathfrak{F}}(z,y) \leq d^{\mathfrak{F}}(x,y) + C$ is at distance at most C' from $D(x,y)$.*

Proof. For the purpose of this proof, let us simplify notations by writing d instead of $d^{\mathfrak{F}}_{\mathbb{X}}$.

Suppose for a contradiction that there exists sequences $(x_n)_n$, $(y_n)_n$ and $(z_n)_n$ such that $d(x_n, z_n) + d(z_n, y_n) \leq d(x_n, y_n) + C$ for all n but the distance from z_n to $D(x_n, y_n)$ tend to infinity as $n \to \infty$.

First note that $d(x_n, z_n) \geq d(z_n, D(x_n, y_n)) \to \infty$, and similarly $d(z_n, y_n) \to \infty$. Since $d(x_n, y_n) \geq d(x_n, z_n) + d(z_n, y_n) - C$ we have $d(x_n, y_n) > d(x_n, z_n)$ for large enough n.

Let $w_n \in D(x_n, y_n)$ be at distance exactly $d(x_n, z_n)$ from x_n. Such a point exists by compactness of $D(x_n, y_n)$. Up to translating everything, we can assume that w_n equals the basepoint $\mathbf{x}$ of the symmetric space for all n, and that x_n, y_n are contained in the standard flat $A \subset \mathbb{X}$. It follows from the explicit construction of the diamond $D(x_n, y_n)$ that the points x_n, y_n are contained in antipodal Weyl cones with vertex $\mathbf{x}$, that is, we may assume that y_n is contained in the standard (closed) Weyl cone A^+, and x_n is contained in the opposite (closed) Weyl cone A^-.

Consider a Riemannian geodesic ray $(w_n(t))_{t\geq 0}$ (with unit *Finsler* speed) starting at $\mathbf{x}$ and passing through z_n, say at time $R_n = d(\mathbf{x}, z_n)$.

We claim that for any $0 \leq t \leq R_n$,

$$d(x_n, w_n(t)) + d(w_n(t), y_n) - d(x_n, y_n) \leq C \tag{13}$$

Namely, since by [KL18, Eq. (5.4)] Finsler balls of fixed radius are convex, the segment $t \to w_n(t)$ ($t \in [0, R]$) stays in the Finsler balls around x_n and y_n of respective radius $d(x_n, z_n)$ and $d(y_n, z_n)$.

Up to extracting a subsequence, we can assume that the sequence of rays $w_n(t)$ converge to a ray $w(t)$ starting at $\mathbf{x}$ and ending at some $w(\infty) \in \partial_\infty \mathbb{X}$. We now claim that

$$w(t) \text{ is contained in the flat } A. \tag{14}$$

Let us prove it. The ball of radius $d(y_n, z_n)$ around y_n contains $w_n(t)$ for any $t \leq R_n$, and this ball converges as $n \to \infty$ to a Finsler horoball, which hence contains $w(t)$ for all t. By the paragraph about Busemann functions in Section 1, this Finsler horoball is contained in a Riemannian horoball B_+ around the point $\xi_+ = \alpha_0^\# \in \partial_\infty A^+$ (see Notation 1).

Similarly, $w(t)$ is contained in a Riemannian horoball B_- around $\xi_- \in \partial_\infty A^- \cap G \cdot \alpha_0^\#$, which is antipodal to ξ_+ as α_0 was chosen symmetric with respect to the Cartan involution.

By Lemma 3.13 below, the *Tits distance* between $\xi_\pm$ and $w(\infty)$ is at most $\frac{\pi}{2}$. Recall that the Tits distance between $a, b \in \partial_\infty \mathbb{X}$ is the angle between two geodesic rays going to respectively a and b and that are contained in the same flat (see for instance [BH99, Part II, Ch. 9] for an account on the Tits metric).

Since ξ^- and ξ^+ are antipodal, their Tits distance is exactly π. It follows from Lemma 3.14 that the limit point $w(\infty)$ lies on the boundary at infinity of the unique flat which contains ξ^- and ξ^+ in its boundary, that is the standard flat A. Since $w(0)$ is also in that flat, the whole geodesic ray $(w(t))_t$ remains in the same flat A, and this ends the proof of the claim (14).

Now we combine (13), (14) and Lemma 3.7 to conclude. Fix $\lambda > 0$ such that for all $x, y, z \in A$ we have

$$d(z, D(x, y)) \leq \lambda(d(x, z) + d(z, y) - d(x, y)).$$

Then for any t, taking n large enough we have $d(x_n, w(t)) + d(w(t), y_n) - d(x_n, y_n) \leq C + 1$ and hence

$$d(w(t), D(x_n, y_n)) \leq \lambda(C + 1),$$

since $w(t) \in A$. Coming back to $w_n(t)$ we deduce, for n large enough, that

$$d(w_n(t), D(x_n, y_n)) \leq \lambda(C + 1) + 1.$$

But this should be a contradiction, as $w_n(t)$ is a ray going to z_n which is very far from $D(x_n, y_n)$. To get a true contradiction, one must be careful in the choice of w_n we made at the beginning: one must choose a point of the compact set $\partial B(x_n, d(x_n, z_n)) \cap D(x_n, y_n)$ which is closest to z_n for the Finsler metric. Then for any $t \leq R_n$ it holds

$$d(w_n(t), \partial B(x_n, d(x_n, z_n)) \cap D(x_n, y_n)) = t,$$

from which one easily deduces, using $d(x_n, z_n) - C \leq d(x_n, w_n(t)) \leq d(x_n, z_n)$, that

$$d(w_n(t), D(x_n, y_n)) \geq \frac{t}{2} - C,$$

whence a contradiction. $\square$

Lemma 3.13. *If a Riemannian geodesic ray $(w(t))_{t\geq 0}$ is contained in a Riemannian horoball centered at $\xi \in \partial_\infty \mathbb{X}$, then the Tits distance between ξ and the limit of $(w(t))_t$ inside $\partial_\infty \mathbb{X}$ is at most $\frac{\pi}{2}$.*

Proof. Denote by b_ξ the Riemannian Busemann function. The derivative at time t of $b_\xi(w(t), \mathbf{x})$ is $-\cos$ of the Riemannian angle at $w(t)$ between the ray w and the ray from $w(t)$ to p (see for instance [KL18, §3.1]). When $t \to \infty$, this angle converges to the Tits distance between the endpoint $w(\infty) \in \partial_\infty \mathbb{X}$ of the ray and ξ (see for instance [BH99, Part II, Prop. 9.8]). If we want $b_\xi(w(t), \mathbf{x})$ to remain bounded from above, its derivative cannot converge to a positive number. As a consequence, the Tits distance between $w(\infty)$ and ξ is at most $\frac{\pi}{2}$. $\square$

Denote by d^{Tits} the Tits distance on $\partial_\infty \mathbb{X}$.

Lemma 3.14. *Let $\xi_1, \xi_2, \zeta \in \partial_\infty \mathbb{X}$ be three points whose Tits distance satisfy the following*

$$d^{\mathrm{Tits}}(\xi_1, \zeta) + d^{\mathrm{Tits}}(\zeta, \xi_2) = d^{\mathrm{Tits}}(\xi_1, \xi_2) = \pi$$

Suppose that ξ_1 (and equivalently ξ_2) is a regular point and let F be the unique flat that contains ξ_1 and ξ_2 in its boundary at infinity $\partial_\infty F$.

Then ζ belongs to $\partial_\infty F$.

Proof. Assume that ζ is distinct from ξ_1 and ξ_2, otherwise it is trivial. The Tits metric is CAT(1) (see [BH99, Part II, Th. 9.13]). It implies that each pair of points at distance $d < \pi$ are joined by a unique geodesic of length d.

Let $\gamma_i : [0, \frac{\pi}{2}] \to \partial_\infty \mathbb{X}$ be the unique geodesic between ξ_i and ζ. Since for the triple ξ_1, ζ, ξ_2, equality holds in the triangle inequality, the concatenation of γ_1 and γ_2 (traveled backward) is a geodesic $\gamma \colon [0, \pi] \to \partial_\infty \mathbb{X}$ from ξ_1 to ξ_2. Since ξ_1 is regular, it lies in the interior of a Weyl chamber of $\partial_\infty F$. The definition of Tits distance forces $\gamma(t)$ to remains in the same Weyl chamber for small values of $t > 0$. In particular $\gamma(t)$ belongs to $\partial_\infty F$ for small t.

The points $\gamma(t)$ and ξ_2 are at distance $\pi - t < \pi$, and so there is a unique geodesic from $\gamma(t)$ to ξ_2. The boundary at infinity $\partial_\infty F$ is totally geodesic for the Tits metric. So by uniqueness, the geodesic $(\gamma([t, \pi]))$ remains in the boundary of $\partial_\infty F$. It follows that $\zeta = \gamma(d^{\mathrm{Tits}}(\xi_1, \zeta))$ lies in $\partial_\infty F$. $\square$

3.4 Proof of Theorem 3.1

We are now ready for the proof of Theorem 3.1.

Consider a continuous path $c : [0, T] \to \mathbb{X}$ such that for any $0 \leq t \leq s \leq u \leq T$, we have $d^{\mathfrak{F}}(c(t), c(s)) + d^{\mathfrak{F}}(c(s), c(u)) \leq d^{\mathfrak{F}}(c(t), c(u)) + C$.

Let F be a flat containing $c(0)$ and $c(T)$, so that it also contains the diamond $D(c(0), c(T))$. By Proposition 3.12 there exists $C' > 0$ only depending on C such that for any $0 \leq t \leq T$, there exists $a(t) \in D(c(0), c(T))$ at distance at most C' from $c(t)$, with $a(0) = c(0)$ and $a(T) = c(T)$. By the triangle inequality, the path $t \to a(t)$ is $C + 6C'$ quasi-ruled. By Lemma 3.2, up to enlarging $C + 6C'$ to a constant which also only depends on C, we may assume that $t \to a(t)$ is continuous.

Thus we can apply Proposition 3.3 to a, to find a geodesic $b : [0, T] \to F$ such that $d^{\mathfrak{F}}(a(t), b(t)) \leq C''$ for some $C'' > 0$ which depends only on C. Recall that geodesics for the restricted metric on F are also geodesics for the metric on $\mathbb{X}$.

We conclude that $d^{\mathfrak{F}}(c(t), b(t)) \leq C'' + C'$ for any t where $C'' + C'$ depends on C but not on the path c.

4 Fock–Goncharov positivity

This section is devoted to a geometric interpretation of positivity as introduced by Lusztig [Lus94] and imported into the context of Hitchin representations by Fock and Goncharov [FG06]. We collect the relevant algebraic results and relate them to admissible paths on the characteristic surface of a Hitchin grafting representation. Throughout this section, we put $G = \mathrm{SL}_d(\mathbb{R})$ although most of the discussion is valid for all split real simple Lie groups and although ultimately we are interested in $\mathrm{PSL}_d(\mathbb{R})$. For completeness, note that for $G = \mathrm{SL}_d(\mathbb{R})$, most of Lusztig's results and concepts were already known (see for instance [And87]), but we still use Lusztig's notation and formalism. In particular, we will use Lusztig's work to introduce the subsets

$$G_{>0} \subset G_{\geq 0} \subset G \text{ and } \mathcal{F}_{>0} \subset \mathcal{F}_{\geq 0} \subset \mathcal{F}$$

and some of their basic properties.

As $G = \mathrm{SL}_d(\mathbb{R})$, the subset $G_{>0} \subset G$ is the set of *totally positive matrices*, which are the matrices $A \in \mathrm{SL}_d(\mathbb{R})$ such that for any $1 \leq k \leq d-1$ the exterior product $\Lambda^k A \in \mathrm{SL}(\Lambda^k \mathbb{R}) = \mathrm{SL}_{d_k}(\mathbb{R})$ is positive, i.e. all its entries are positive. For general split Lie groups the definition of $G_{>0}$ is more complicated, see Sections 2.2, 2.12, 5.10 and 8.8 of [Lus94].

One can check that $\tau : \mathrm{PSL}_2(\mathbb{R}) \to \mathrm{PSL}_d(\mathbb{R})$ is positive in the sense that it maps projectivizations of positive matrices to projectivizations of totally positive matrices. Note also that Lusztig did not introduce the concept of positive representation from $\mathrm{PSL}_2(\mathbb{R})$ into G: this is due to Fock–Goncharov [FG06], who used it to prove among other things that all Hitchin representations are discrete. Many ideas in this section are inspired by the work of Fock–Goncharov.

Finally, the concept of positivity also exists in certain *nonsplit* real Lie group: see [BH12; GW18]. However in these settings the positive cone lies in a partial flag manifold instead of the full flag manifold, which is not enough for our purposes.

4.1 Reminders on positivity

The following summarizes the results from [Lus94] we are going to use. As before, $\mathcal{F}$ denotes the variety of full flags in $\mathbb{R}^d$.

Theorem 4.1 ([Lus94]). *There exist semigroups $G_{>0} \subset G_{\geq 0} \subset G$ and a subset $\mathcal{F}_{>0} \subset \mathcal{F}$ with the following properties.*

1. *[By definition, see §2.2] $\exp(\mathfrak{a}) \subset G_{\geq 0}$, in particular $1 \in G_{\geq 0}$.*
2. *[Th. 4.8] For the standard embedding $\tau : \mathrm{SL}_2(\mathbb{R}) \to \mathrm{SL}_d(\mathbb{R})$, the set $G_{>0}$ is an (open) connected component of*
$$\{g \in G : g\partial\tau(\infty) \pitchfork \partial\tau(\infty) \ \text{ and } \ g\partial\tau(0) \pitchfork \partial\tau(0)\}.$$
3. *[Th. 5.6] Every element of $G_{>0}$ is loxodromic, and in particular does not fix any point of $\mathbb{X}$.*
4. *[Th. 4.3 & Rem. 4.4] $G_{\geq 0}$ is the closure of $G_{>0}$.*
5. *[before Prop. 2.13] $G_{>0}G_{\geq 0} \subset G_{>0}$ and $G_{\geq 0}G_{>0} \subset G_{>0}$.*
6. *[Prop. 8.14] $\mathcal{F}_{>0}$ is an (open) connected component of $\partial\tau(\infty)^{\pitchfork} \cap \partial\tau(0)^{\pitchfork}$.*
7. *[Prop. 8.12] $G_{\geq 0}\mathcal{F}_{>0} \subset \mathcal{F}_{>0}$.*
8. *[By (7) above and by definition, see Th. 8.7] $\mathcal{F}_{>0} = G_{>0} \cdot \partial\tau(\infty)$.*

Example . *In the special case $G = \mathrm{SL}_3(\mathbb{R})$, one can visualise $\mathcal{F}_{>0}$: namely, recall that in this case, the flag variety $\mathcal{F}$ is identified with the set of pairs (p, ℓ) where p is a point of $\mathbb{RP}^2$ and ℓ is a line containing p.*

Let x, y and $z \in \mathbb{RP}^2$ be the image of the canonical basis of $\mathbb{R}^3$, let $[x, y]$, $[y, z]$ and $[z, x] \subset \mathbb{RP}^2$ be the image of the segments between the vectors of the canonical basis, and let $T \subset \mathbb{RP}^2$ be the triangle enclosed by these segments. Then $\mathcal{F}_{>0}$ is the set of pairs (p, ℓ) such that p is contained in the interior of T and ℓ intersects the (relative) interior of the segments $[x, y]$ and $[y, z]$. Then
$$G_{>0} = \{g \in \mathrm{SL}_3(\mathbb{R}) \mid g\overline{\mathcal{F}_{>0}} \subset \mathcal{F}_{>0}\}.$$
One can see that $g \in G_{>0}$ maps $\overline{T}$ into its interior (this corresponds to the fact that the entries of g, i.e. the minors of size 1, are positive), and hence g has an attracting fixed point in T. In general it is true that any totally positive matrix is diagonalisable with distinct positive eigenvalues.

We will also use the following, which should be well-known to the experts, but for which we did not find a reference.

Lemma 4.2. *9. $\mathcal{F}_{>0}$ is the interior of its closure, denoted by $\mathcal{F}_{\geq 0}$.*

10. $G_{>0}\mathcal{F}_{\geq 0} \subset \mathcal{F}_{>0}$.

11. Let us denote $G_{<0} := (G_{>0})^{-1}$ and $\mathcal{F}_{<0} = G_{<0} \cdot \partial\tau(\infty)$, and $G_{\leq 0}$ and $\mathcal{F}_{\leq 0}$ their respective closures. Then any pair in $\mathcal{F}_{<0} \times \mathcal{F}_{\geq 0}$ is transverse.

Remark . It is clear that $\mathcal{F}_{\geq 0}$ contains $G_{\geq 0} \cdot \partial\tau(\infty)$ but they are not equal in general. For instance consider the case $\mathrm{SL}_3(\mathbb{R})$. Denote the usual basis of $\mathbb{R}^3$ by (e_1, e_2, e_3). Consider the flags $\partial\tau(\infty) = (span(e_3), span(e_2, e_3))$ and $F = (span(e_2), span(e_2, e_3 - e_1))$.

Then F lies in $\mathcal{F}_{\geq 0}$ but not in $G_{\geq 0} \cdot \partial\tau(\infty)$. Indeed one can check $F \in \mathcal{F}_{\geq 0}$ by computing that $F = \lim_{\lambda \to 0^+} A_\lambda \partial\tau(\infty)$ where $A_\lambda = \left(\begin{smallmatrix} 1 & 3\lambda & \lambda \\ \lambda & 2 & 1 \\ \lambda^3 & \lambda & \lambda \end{smallmatrix}\right)$ is totally positive for small positive values of λ. (One can renormalize A_λ to make it determinant 1.)

But if a matrix $B \in G_{\geq 0}$ sends $\partial\tau(\infty)$ to the flag F, then it sends e_2 to a point $B \cdot e_2$ in $span(e_2, e_3 - e_1) \setminus span(e_2)$. Thus one can write $B \cdot e_2 = ae_2 + b(e_3 - e_1)$ with $b \neq 0$. Then one of the coefficients of B, in position $(3,1)$ or $(3,3)$, has a negative entry. It contradicts the fact that $B \in G_{\geq 0}$, and that matrices in $G_{\geq 0}$ have non-negative entries. The flag F may be thought as corresponding to a point at infinity of $G_{\geq 0}$ (for the compactification $\mathrm{SL}_3(\mathbb{R}) \hookrightarrow \mathrm{PGL}_3(\mathbb{R})$).

Proof of Lemma 4.2. Proof of (9). Put $\xi_\infty = \partial\tau(\infty)$ and $\xi_0 = \partial\tau(0)$. By definition, the sets $\xi_\infty^\pitchfork$ and $\xi_0^\pitchfork$ are open Bruhat cells in $\mathcal{F}$ and hence $Z = \xi_\infty^\pitchfork \cap \xi_0^\pitchfork$ is open. The complements of the Bruhat cells $\xi_\infty^\pitchfork$ and $\xi_0^\pitchfork$ are (real) connected projective varieties all of whose irreducible components are of codimension one. Then every point of the boundary of $\xi_\infty^\pitchfork \cap \xi_0^\pitchfork$ is contained in an irreducible subvariety of codimension one. Thus by property (6), the statement(9) is equivalent to the following. Let V be a component of Z. Denoting by $\overline{U}$ the closure of a set U, it holds

$$V = \mathcal{F} - \overline{\mathcal{F} - \overline{V}}. \tag{15}$$

As for any two open sets U_1, U_2 we have $\overline{U_1 \cup U_2} = \overline{U_1} \cup \overline{U_2}$, it suffices to write V as a finite intersection $V = \cap_{j=1}^{d-1} U_j$ where each of the sets U_j has property (15). To construct such sets note that if $\xi_\iota = (\xi_\iota^1 \subset \cdots \subset \xi_\iota^{d-1})$ $(\iota = \infty, 0)$ then the linear hyperplanes $\xi_\infty^{d-1}, \xi_0^{d-1}$ are transverse and hence they decomposes $\mathbb{R}^d$ into four connected components which are paired be the reflection $x \to -x$, say the components $A, -A, B, -B$. The closures of the components $A, -A$ and $B, -B$ intersect in a linear subspace of codimension 2. A component V of $Z = \xi_\infty^\pitchfork \cap \xi_0^\pitchfork$ consists of flags $\zeta = \zeta_1 \subset \cdots \subset \zeta_{d-1}$ with the property that ζ_1 is transverse to both $\xi_\infty^{d-1}, \xi_0^{d-1}$. But this means that for either A or B, say for A, any nonzero point on the line ζ_1 is contained in $A \cup -A$. As a consequence, up to exchanging A and B, if we define U_1 to be the set of all flags $\zeta = \zeta_1 \subset \cdots \subset \zeta_{d-1}$ such that $\zeta_1 - \{0\} \subset A \cup -A$ then $V \subset U_1$, and U_1 is an open set with property (15).

Now note that U_1 can also be described as follows. Choose a generator ω_ι of $\wedge^{d-1}\xi_\iota^{d-1}$ $(\iota = 0, \infty)$ and define U_1 to be the set of all flags ζ with the property that for some basis element e of ζ_1, the wedge products $e \wedge \omega_0, e \wedge \omega_\infty$ define the same (or the opposite) orientation of $\mathbb{R}^d$. Then U_1 is one of the sets described in the previous paragraph.

For $j \leq d-1$ define the set U_j as the set of all flags ζ so that for some generator e of $\Lambda^j \zeta_j$ and some generators $\omega_0^{d-j}, \omega_\infty^{d-j}$ of $\Lambda^{d-j}\xi_0^{d-j}, \Lambda^{d-j}\xi_\infty^j$ the orientations defined by $e \wedge \omega_0^{d-j}, e \wedge \omega_\infty^{d-j}$ coincide (or are opposite). For suitably choices of the sets U_j, we then have $V = \cap_j U_j$. Together with the first paragraph of this proof, (9) follows.

Proof of (10). Since $G_{>0}$ is open (by (2)), every $G_{>0}$-orbit in $\mathcal{F}$ is open. Being a union of such orbits, $G_{>0}\mathcal{F}_{\geq 0}$ is also open. Moreover it is contained in $\mathcal{F}_{\geq 0}$ (by (7)) and hence it is contained in its interior, which is precisely $\mathcal{F}_{>0}$ by (9).

Proof of (11). Consider $(\xi, \eta) \in \mathcal{F}_{<0} \times \mathcal{F}_{\geq 0}$. Since $\mathcal{F}_{>0}$ is open, $G_{>0}$ contains 1 in its closure, and $G_{>0}\mathcal{F}_{\geq 0} \subset \mathcal{F}_{>0}$, there exists $g \in G_{>0}$ such that $g\xi \in \mathcal{F}_{<0}$ and $g\eta \in \mathcal{F}_{>0}$.

By definition, there exists $h \in G_{>0}$ such that $g\xi = h^{-1}\xi_\infty$. Then $hg\xi = \xi_\infty$ and $hg\eta \in h\mathcal{F}_{>0} \subset \mathcal{F}_{>0} \subset \xi_\infty^{\pitchfork}$ by (6). Therefore $hg\xi$ and $hg\eta$ are transverse, and so are ξ and η. □

4.2 Positivity and injectivity of admissible paths

We now explain the assumption that the representation $\tau : \mathrm{SL}_2(\mathbb{R}) \to \mathrm{SL}_d(\mathbb{R})$ is positive: it means that τ maps every 2×2 matrix with positive entries into $G_{>0}$. It has the following consequences, which should be well-known to experts.

Lemma 4.3. *We have the following.*

1. $a'_t \in G_{>0}$ *for any* $t > 0$.
2. $G_{<0} = r_\pi G_{>0} r_\pi$ *(see Notation 2).*

Proof. Proof of (1). This is an immediate consequence of the fact that

$$\begin{pmatrix} \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} \\ -\frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} \end{pmatrix} \begin{pmatrix} e^t & 0 \\ 0 & e^{-t} \end{pmatrix} \begin{pmatrix} \frac{\sqrt{2}}{2} & -\frac{\sqrt{2}}{2} \\ \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} \end{pmatrix}$$

has positive entries.

Proof of (2). Let us prove that $r_\pi G_{>0}^{-1} r_\pi = G_{>0}$. The maps $g \mapsto g^{-1}$ and $g \mapsto r_\pi g r_\pi$ both preserve

$$\{g \in G : g\partial\tau(\infty) \pitchfork \partial\tau(\infty) \ \text{ and } \ g\partial\tau(0) \pitchfork \partial\tau(0)\},$$

and hence permute the connected components, and consequently so does $g \mapsto r_\pi g^{-1} r_\pi$. To prove that this map preserves the connected component $G_{>0}$ (see Theorem 4.1 (2)), it suffices to show that it fixes a point of $G_{>0}$. It is clear that it fixes for instance $a'_1 \in G_{>0}$. □

A first important consequence of all the facts about positivity that we have listed is the following.

Corollary 4.4. *For any admissible path* $c : [0, T] \to G$*, for all* $0 \leq s < t \leq T$ *we have* $c(s)^{-1}c(t) \in G_{\geq 0}$.

Proof. By definition, $c(s)^{-1}c(t)$ is a product of elements of G of the form a'_r for some $r > 0$ or $\exp(v)$ for some $v \in \mathfrak{a}$. All these elements belong to $G_{\geq 0}$ by Theorem 4.1 and Lemma 4.3, and so does their product since $G_{\geq 0}$ is a semigroup. $\square$

The previous result, combined with the fact that totally positive matrices are not the identity, tells us that admissible paths are injective, as explained below.

Proposition 4.5. *Any admissible path in $\mathbb{X}$ is injective.*

Proof. Let $c : [0, T] \to G$ be an admissible path and let $\mathbf{x} \in \mathbb{X}$ be an arbitrarily fixed point. Consider $0 \leq s < t \leq T$, and let us prove that $c(s)\mathbf{x} \neq c(t)\mathbf{x}$, i.e. that $c(s)^{-1}c(t)\mathbf{x} \neq \mathbf{x}$. By definition, $c(s)^{-1}c(t)$ is a product of elements of G of the form a'_r for some $r > 0$ or $\exp(v)$ for some $v \in \mathfrak{a}$. There are two cases.

Case 1: $c(s)^{-1}c(t) = \exp(v)$ for some nonzero $v \in \mathfrak{a}$, then it is clear that $\exp(v)\mathbf{x} \neq \mathbf{x}$.

Case 2: $c(s)^{-1}c(t)$ is a product of elements of G of the form a'_r for some $r > 0$ or $\exp(v)$ for some $v \in \mathfrak{a}$, with at least one element of the form a'_r. All the $\exp(v)$'s belong to $G_{\geq 0}$ by Theorem 4.1.1, and all the a'_r's belong to $G_{>0}$ by Lemma 4.3.1, and so $c(s)^{-1}c(t)$ belongs to $G_{>0}$ by Theorem 4.1.5. Therefore $c(s)^{-1}c(t)$ does not fix any point of $\mathbb{X}$ by Theorem 4.1.3, and $c(s)^{-1}c(t)\mathbf{x} \neq \mathbf{x}$. $\square$

This implies that the characteristic surface we have constructed in Section 2.4 is embedded.

Corollary 4.6. *The map $Q_z : S_z \to \rho_z \backslash \mathbb{X}$ constructed in Proposition 2.5 is injective.*

Proof. Consider $x, y \in S_z$ such that $Q_z(x) = Q_z(y)$, and let us prove that $x = y$.

Consider two lifts $\tilde{x}$ and $\tilde{y} \in \tilde{S}_z$ of respectively x and y. Then $\tilde{Q}_z(\tilde{x})$ and $\tilde{Q}_z(\tilde{y})$ have the same projection in $\rho_z \backslash \mathbb{X}$, which means that there exists $\gamma \in \pi_1(S)$ such that

$$\tilde{Q}_z(\tilde{x}) = \rho_z(\gamma)\tilde{Q}_z(\tilde{y}) = \tilde{Q}_z(\gamma\tilde{y}).$$

Consider an admissible path $c : [0, T] \to \tilde{S}_z$ from $\tilde{x}$ to $\gamma\tilde{y}$. Then by Observation 2.9 $\tilde{Q}_z \circ c : [0, T] \to \mathbb{X}$ is an admissible path from $\tilde{Q}_z(\tilde{x})$ to itself.

Since admissible paths of $\mathbb{X}$ are injective by Proposition 4.5, this means that $T = 0$ and $\tilde{x} = \gamma\tilde{y}$, and hence that $x = y$. $\square$

4.3 Positivity gives a control on default of the triangle inequality

The goal of this section is to prove the following result about totally positive transformations. We then use it to prove that admissible paths are (Finsler) quasi-ruled, quasi-convex and quasi-geodesics, which implies that the characteristic surface in the symmetric space associated to a Hitchin grafting representation is Finsler quasi-convex.

As before, we write $G = \mathrm{SL}_d(\mathbb{R})$, and we choose a basepoint $\mathbf{x} \in \mathbb{X} = G/K$, thought of as the projection of the identity in G.

Lemma 4.7. *For any $\omega > 0$, there exists $C_\omega > 0$ such that for all $g_+ \in G_{\geq 0}$ and $g_- \in a'_{-\omega}G_{\leq 0}$, we have*

$$d^{\mathfrak{F}}(g_-\mathbf{x}, \mathbf{x}) + d^{\mathfrak{F}}(\mathbf{x}, g_+\mathbf{x}) \leq d^{\mathfrak{F}}(g_-\mathbf{x}, g_+\mathbf{x}) + C_\omega.$$

To prove it we will need the following technical result. In its formulation, we use a K-invariant metric $d_{\mathcal{F}}$ on the flag variety $\mathcal{F}$. Furthermore, we denote by $\xi_0 \in \mathcal{F}$ the simplex $\partial\tau(\infty) = \exp(\mathfrak{a}^+) \in \mathcal{F}$. Distances are taken with respect to the distance function d defined by the symmetric Riemannian metric.

Lemma 4.8. *For every $\epsilon > 0$ there exists a number $C_\epsilon > 0$ only depending on ϵ with the following property. Consider $g \in G$ decomposed as $g = k\exp(\kappa(g))\ell$ with $k, \ell \in K$ (the maximal compact subgroup) and $\kappa(g) \in \mathfrak{a}^+$. Let $\xi \in \mathcal{F}$ be at $d_{\mathcal{F}}$-distance at least $\epsilon > 0$ from $k\xi_0^{\not\pitchfork}$. Then $\mathbf{x}$ is at distance at most C_ϵ from the Weyl cone with vertex at $g\mathbf{x}$ and boundary at infinity the simplex in $\partial_\infty\mathbb{X}$ corresponding to ξ.*

Proof. Since $d_{\mathcal{F}}(\xi, k\xi_0^{\not\pitchfork}) \geq \epsilon$, the simplices $\xi, k\xi_0$ are transverse and hence they are contained in a unique maximal flat F. We claim that there is $x \in F$ at distance at most $C_\epsilon > 0$ from $\mathbf{x}$, where C_ϵ only depends on ϵ. Indeed, this follows from the compactness of the set $\{(\xi, \eta) \in \mathcal{F}^2, d_{\mathcal{F}}(\xi, \eta) \geq \epsilon\}$ and continuity of the map which associates to two transverse flags the unique maximal flat whose visual boundary contains the Weyl chambers that corresponds to the two flags.

Note that $g\mathbf{x}$ is contained in the Weyl cone connecting $\mathbf{x}$ to $k\xi_0$ (because $k\mathbf{x} = \mathbf{x}, \ell\mathbf{x} = \mathbf{x}$ and $\exp(\kappa(g))\mathbf{x}$ is contained in the Weyl Cone connecting $\mathbf{x}$ to ξ_0). Thus as Weyl cones are convex cones, the endpoint $p \in \partial_\infty\mathbb{X}$ of the geodesic ray $[\mathbf{x}, p)$ starting at $\mathbf{x}$ and passing through $g\mathbf{x}$ is contained in the Weyl Chamber $k\xi_0$. The ray $[x, p)$ is contained in the Weyl Cone from x to $k\xi_0$.

We now apply the CAT(0)-property for the Riemannian symmetric metric to the asymptotic rays $[x, p)$ and $[\mathbf{x}, p)$. It yields that the point $y \in [x, p)$ at distance exactly $d(\mathbf{x}, g\mathbf{x})$ from x is of distance at most $d(\mathbf{x}, x) \leq C_\epsilon$ from $g\mathbf{x}$.

By construction, the geodesic ray $[x, p)$ is contained in the flat F, and its endpoint is contained in the Weyl chamber $k\xi_0$. The unique geodesic line η extending $[x, p)$ is contained in F and is backward asymptotic to a point q in the unique Weyl chamber ξ in the visual boundary of F which is transverse to $k\xi_0$. Recall that η passes through the C_ϵ-neighborhood of $\mathbf{x}$.

Using once more the CAT(0)-property, this time applied to the subray of η which connects y to q and the geodesic ray ζ connecting $g\mathbf{x}$ to q, we conclude that ζ passes through the C_ϵ-neighborhood of x and hence through the $2C_\epsilon$-neighborhood of $\mathbf{x}$. On the other hand, by construction, this ray is contained in the Weyl cone connecting $g\mathbf{x}$ to ξ. Together this is what we wanted to show. $\square$

Proof of Lemma 4.7. Decompose $g_\pm = k_\pm e^{\kappa(g_\pm)}\ell_\pm$ with $k_\pm, \ell_\pm \in K$ (the maximal compact subgroup) and $\kappa(g_\pm) \in \mathfrak{a}^+$. The plan is to use positivity and Lemma 4.8 to find Weyl Chambers $\xi_\pm$ such that their images $g_\pm\xi_\pm$ are transverse, the flat F through them passes near $\mathbf{x}$, and the Weyl Cone from $\mathbf{x}$ to $g_\pm\xi_\pm$ passes near $g_\pm\mathbf{x}$.

Recall the definition of the set $\mathcal{F}_{\geq 0}$, and for $\omega > 0$ the element a'_ω. Since $a'_\omega\mathcal{F}_{\geq 0}$ has nonempty interior and $\xi_0^{\not\pitchfork}$ is a closed set with empty interior, there exists $\epsilon = \epsilon_\omega > 0$ such that for every $k \in K$ there is $\xi \in a'_\omega\mathcal{F}_{\geq 0}$ at $d_{\mathcal{F}}$-distance at least ϵ from $k\xi_0^{\not\pitchfork}$. Similarly, for every $k \in K$ there is $\xi \in a'_{-\omega}\mathcal{F}_{\leq 0}$ at $d_{\mathcal{F}}$-distance at least ϵ from $k\xi_0^{\not\pitchfork}$.

Let $\xi_+ \in a'_\omega \mathcal{F}_{\geq 0}$ be at $d_{\mathcal{F}}$-distance at least ϵ from $\ell_+^{-1}\xi_0^{\not\pitchfork}$ and $\xi_- \in a'_{-\omega}\mathcal{F}_{\leq 0}$ be at $d_{\mathcal{F}}$-distance at least ϵ from $\ell_-^{-1}\xi_0^{\not\pitchfork}$.

We know $\xi_+ \in \mathcal{F}_{\geq 0}$ and $g_+ \in G_{\geq 0}$, so by Theorem 4.1.7 we have $g_+\xi_+ \in \mathcal{F}_{\geq 0}$, and similarly $g_-\xi_- \in a'_{-\omega}\mathcal{F}_{\leq 0}$. In particular, by Lemma 4.2.11 $g_+\xi_+$ and $g_-\xi_-$ are transverse. More precisely, if we denote as before by d the distance function of the symmetric metric, then the flat $F = F(g_-\xi_-, g_+\xi_+)$ through them contains a point x with $d(x, \mathbf{x}) \leq q_\omega$ for some $q_\omega > 0$ only depending on ω, because every pair in the set $a'_{-\omega}\mathcal{F}_{\leq 0} \times \mathcal{F}_{\geq 0}$ is transverse and this set is compact.

Since $d_{\mathcal{F}}(\xi_\pm, \ell_\pm^{-1}\xi_0^{\not\pitchfork}) \geq \epsilon$, Lemma 4.8 implies that $g_\pm \mathbf{x}$ is at Riemannian distance at most C_ϵ to the Weyl Cone connecting $\mathbf{x}$ to $g_\pm\xi_\pm$. By the CAT(0) property of $(\mathbb{X}, d)$, applied to the geodesics connecting $\mathbf{x}$ and x to all points in $g_\pm\xi_\pm$, the Hausdorff distance (for d) between this Weyl cone and the Weyl cone $W_\pm \subset F$ connecting x to $g_\pm\xi_\pm$ is at most $d(\mathbf{x}, x) \leq q_\omega$. As a consequence, there is $x_\pm \in W_\pm$ with

$$d(x_\pm, g_\pm \mathbf{x}) \leq q_\omega + C_\epsilon.$$

Since W_+ and W_- are two opposite Weyl Cones in F based at x, and since $x_\pm \in W_\pm$, we deduce from Proposition 3.11 that

$$d^{\mathfrak{F}}(x_-, x) + d^{\mathfrak{F}}(x, x_+) = d^{\mathfrak{F}}(x_-, x_+).$$

To conclude, recall that the Riemannian and Finsler metrics are comparable; that is, there exists $\lambda > 0$ such that $\lambda^{-1} d \leq d^{\mathfrak{F}} \leq \lambda d$. Then by the triangle inequality

$$\begin{aligned} d^{\mathfrak{F}}(g_-\mathbf{x}, \mathbf{x}) &+ d^{\mathfrak{F}}(\mathbf{x}, g_+\mathbf{x}) - d^{\mathfrak{F}}(g_-\mathbf{x}, g_+\mathbf{x}) \\ &\leq d^{\mathfrak{F}}(x_-, x) + d^{\mathfrak{F}}(x, x_+) - d^{\mathfrak{F}}(x_-, x_+) + 6\lambda(q_\omega + C_\epsilon) \\ &\leq 6\lambda(q_\omega + C_\epsilon). \end{aligned}$$ □

We now use the previous result to prove that admissible paths are quasi-ruled. For this we will need the following general fact about quasi-ruled paths.

Lemma 4.9. *Let (X, d) be a metric space and $x_1, \ldots, x_n \in X$ be such that for some constant $C > 0$ we have $d(x_i, x_j) + d(x_j, x_k) \leq d(x_i, x_k) + C$ for all $i < j < k$. Then any concatenation of geodesics $[x_1, x_2], [x_2, x_3], \ldots, [x_{n-1}, x_n]$ is $4C$-quasi-ruled.*

Proof. Let x, y, z be three points on such a concatenation, in this order, and let us check that $d(x, y) + d(y, z) \leq d(x, z) + 4C$. Let $[x_i, x_{i+1}]$, $[x_j, x_{j+1}]$ and $[x_k, x_{k+1}]$ be three geodesic pieces of the concatenation containing respectively x, y, z, such that $i \leq j \leq k$. Let us assume for the rest of the proof that $i < j < k$; if instead we have $i = j < k$ or $i < j = k$ then the proof is similar (in fact easier), and the case $i = j = k$ is obvious.

Using the triangle inequality and our assumption, and denoting $(a, b) = d(a, b)$ to lighten our estimates on distances, we have the following, which concludes the proof.

$$\begin{aligned}
(x, y) + (y, z) &\leq (x, x_{i+1}) + (x_{i+1}, x_j) + (x_j, y) + (y, x_{j+1}) + (x_{j+1}, x_k) + (x_k, z) \\
&\leq -(x_i, x) + (x_i, x_{i+1}) + (x_{i+1}, x_j) + (x_j, x_{j+1}) + (x_{j+1}, x_k) + (x_k, x_{k+1}) - (z, x_{k+1}) \\
&\leq -(x_i, x) + (x_i, x_j) + C + (x_j, x_k) + C + (x_k, x_{k+1}) - (z, x_{k+1}) \\
&\leq 2C - (x_i, x) + (x_i, x_k) + C + (x_k, x_{k+1}) - (z, x_{k+1}) \\
&\leq 3C - (x_i, x) + (x_i, x_{k+1}) + C - (z, x_{k+1}) \\
&\leq 4C + (x, z).
\end{aligned}$$ □

We now prove that admissible paths are quasi-ruled.

Proposition 4.10. *For any $\omega > 0$ there exists C_ω such that any $(\omega, 0)$-admissible path c in $\mathbb{X}$ is Finsler C_ω-quasi-ruled, and hence is at Hausdorff distance at most C'_ω from some Finsler geodesic by Theorem 3.1, where C'_ω only depends on C_ω.*

Proof. By definition $c(r) = a(r)\mathbf{x}$ for any r, where $a(r)$ is an admissible path in G.

Take $0 \leq t < s < u \leq T$, and let us show $d(c(t), c(s)) + d(c(s), c(u)) \leq d(c(t), c(u)) + C_\omega$ for a well-chosen C_ω. By Lemma 4.9 above we may assume that each of the points $a(t), a(s), a(u)$ is at the junctions of two pieces of the admissible path a, one of hyperbolic type, and the other of flat type (see Definition 2.8).

By Corollary 4.4, $a(s)^{-1}a(u) \in G_{\geq 0}$ and $a(s)^{-1}a(t) \in G_{\leq 0}$.

Note that $a(s)$ is adjacent to a hyperbolic-type piece of a, which has length at least ω, unless this piece is the first or last piece of a. If this hyperbolic-type piece is first or last and has length less than ω, then $c(s)$ is ω-close to either $c(t)$ or $c(u)$, and one conclude easily with the triangle inequality (taking $C_\omega \geq \omega$). Let us assume this hyperbolic-type piece has length at least ω.

If this piece is after $a(s)$ then $a(s)^{-1}a(u) \in a'_\omega G_{\geq 0}$. If on the contrary this piece is before $a(s)$ then $a(s)^{-1}a(t) \in a'_{-\omega} G_{\leq 0}$.

In any case, by Lemma 4.7, we can conclude:

$$d^{\mathfrak{F}}(c(t), c(s)) + d^{\mathfrak{F}}(c(s), c(u)) \leq d^{\mathfrak{F}}(c(t), c(u)) + C_\omega.$$ □

Finally, we deduce that the characteristic surface associated to a Hitchin grafting representation is quasi-convex.

Corollary 4.11. *Consider the map $\widetilde{Q}_z : \widetilde{S}_z \to \widetilde{S}^\iota_z \subset \mathbb{X}$ constructed in Proposition 2.5, such that the grafting locus $\gamma^* \subset S$ has collar size at least $\omega > 0$. Then $\widetilde{S}^\iota_z$ is Finsler C_ω-quasi-convex for some C_ω depending on ω, in the sense that any two points of $\widetilde{S}^\iota_z$ can be connected by a Finsler geodesic that stays at distance at most C_ω from $\widetilde{S}^\iota_z$.*

Proof. This is an immediate consequence of Proposition 4.10 and Observation 2.9. □

4.4 Estimates on eigenvalues of products of totally positive matrices

We will also need a quantitative version of the classical result that all elements of $G_{>0}$ are loxodromic. This quantitative result is probably well known to experts; since we did not find a precise reference for what we need, we give a proof (in the case $G = \mathrm{SL}_d(\mathbb{R})$). We will not need in the present paper the estimates on angles presented below, but they will be useful in the companion paper [BHMM24].

Recall that for any matrix g and $1 \leq k \leq d$, we denote by $\lambda_k(g)$ the logarithm of the norm of the k-th eigenvalue of g, such that $\lambda_1(g) \geq \lambda_2(g) \geq \cdots \geq \lambda_d(g)$. Recall that g is loxodromic if we have strict inequalities; in this case this gives a natural ordering on the eigenspaces of g (which are eigenlines).

Proposition 4.12. *For any $g \in G_{>0}$ there exists $\omega, \theta > 0$ such that the following holds.*

- *for any $h \in G_{\geq 0}$, for any $1 \leq k < d$, we have $\lambda_k(gh) \geq \lambda_{k+1}(gh) + \omega$ (so gh is loxodromic) and the angle between the sum of the k first eigenlines of gh and the sum of remaining eigenlines is at least θ.*
- *for all $h_1, \dots, h_n \in G_{\geq 0}$, denoting $h = gh_1gh_2 \cdots gh_n$ we have $\lambda_k(h) \geq \lambda_{k+1}(h) + n\omega$ for any $1 \leq k \leq d-1$.*

  In particular, $d^{\mathfrak{F}}(\mathbf{x}, h\mathbf{x}) \geq n\omega'$ for some constant ω' that only depends on ω.
- *for all $h, h' \in gG_{\geq 0}$, for any k, the angle between the sum of the k first eigenlines of h and the sum of last $d-k$ eigenlines of h' is at least θ.*

To prove Proposition 4.12, we first establish an intermediate result about positive matrices and use the fact that a matrix of size d is totally positive if and only if all its exterior products, seen as matrices of size $d_k = \dim(\Lambda^k \mathbb{R}^d)$ where $1 \leq k \leq d-1$, are *positive*, i.e. with positive entries.

Lemma 4.13. *For any positive matrix g of size d, there exists $\omega > 0$ such that the following holds.*

- *for any nonegative h, we have $\lambda_1(gh) \geq \lambda_2(gh) + \omega$ and the angle between the attracting eigenline of gh and the repelling hyperplane is at least θ.*
- *for all $A_1, \dots, A_n$ nonnegative, if $A = gA_1 \cdots gA_n$ then $\lambda_1(A) \geq \lambda_2(A) + n\omega$.*
- *for all h, h' nonnegative, for any k, the angle between the sum of the attracting eigenline of gh and the repelling hyperplane of gh' is at least θ.*

Proof. We will prove all three points at the same time. By density of positive matrices it suffices to prove the lemma for $A_1, \dots, A_n$ positive.

Let $\mathcal{C} \subset \mathbb{R}^d$ be the open convex cone of positive vectors, and $\Omega \subset \mathbb{RP}^{d-1}$ its projectivisation, which is a properly convex domain in the sense that there is an affine chart of $\mathbb{RP}^{d-1}$ containing Ω and in which Ω is bounded and convex.

Note that $A\overline{\mathcal{C}} \subset \mathcal{C} \cup \{0\}$ for any positive matrix A, so $A\overline{\Omega} \subset \Omega$.

Any properly convex domain $\Omega' \subset \mathbb{RP}^{d-1}$ can be endowed with a classical Finsler metric $d_{\Omega'}$ called the Hilbert metric, locally equivalent to the usual Riemannian metric of $\mathbb{RP}^{d-1}$, such that (see [Bir57], or see [PT14] for a broad introduction to Hilbert geometry)

1. it is projectively equivariant: $d_{h\Omega'} \circ h = d_{\Omega'}$ for any projective transformation h;

2. it is monotone with respect to inclusion: $d_{\Omega'} \leq d_{\Omega''}$ (on Ω'') for any $\Omega'' \subset \Omega'$;

3. if $\overline{\Omega}'' \subset \Omega'$ then there is $r < 1$ such that $d_{\Omega'} \leq r d_{\Omega''}$ (on Ω'').

Let g be a positive matrix: As $g\overline{\Omega} \subset \Omega$ there is $r < 1$ such that $d_\Omega \leq r d_{g\Omega}$. Then $g : \Omega \to \Omega$ is r-Lipschitz for d_Ω — hence it is a contraction — by equivariance of the Hilbert metric. In fact, $gA : \Omega \to \Omega$ also is an r-d_Ω-Lipschitz map for any positive matrix A since $gA\Omega \subset g\Omega$.

By the Banach fixed-point theorem, gA has a fixed point $p \in \Omega$ such that $(gA)^n x \to p$ for any $x \in \Omega$. In particular, gA is proximal with attracting line $p \in g\Omega$, and the repelling hyperplane does not intersect Ω.

This implies the existence of a positive lower bound θ on the angle (for the standard Euclidean metric) between the attracting line of gA and the repelling hyperplane of gA' for all positive matrices A, A'.

This settles our claims about angles in the statement of the lemma. It remains to prove the estimates on the gaps between the first two eigenvalues. This will be done by reinterpreting this gap as a contraction rate for the action of proximal transformations at their attracting eigenline and using our observation above that positive matrices contract the cone of positive vectors.

One can observe that for any proximal matrix h such that the angle between the attracting line and the repelling hyperplane is at least θ, the number $e^{\lambda_2(h)-\lambda_1(h)}$ is comparable to the contraction rate of h at its attracting fixed point in $\mathbb{RP}^{d-1}$ for the usual Riemannian metric, where the comparison error only depends on θ. In other words, if h is R-Lipschitz (for the Riemannian metric) at its attracting eigenline then

$$e^{\lambda_2(h)-\lambda_1(h)} \leq C_1 R$$

for some constant C_1 depending on θ.

Moreove, as we already mentioned, the restriction of the Hilbert metric d_Ω to the compact subset $\overline{g\Omega}$ is uniformly comparable to the standard Riemannian metric on $\mathbb{RP}^{d-1}$. Thus for any transformation $h : \Omega \to g\Omega$, if h is R-d_Ω-Lipschitz on $g\Omega$ then it is RC_2-Riemannian-Lipschitz for some constant C_2 that depends on $g\Omega$ and Ω. In particular it is RC_2-Lipschitz at its attracting eigenline, and hence

$$e^{\lambda_2(h)-\lambda_1(h)} \leq CR$$

where $C = C_1 C_2$.

In particular, for all $A_1, \ldots, A_n$ positive, $A = gA_1 \cdots gA_n$ is r^n-d_Ω-Lipschitz so

$$\lambda_1(A) - \lambda_2(A) \geq -\log(Cr^n).$$

In fact, for any k, the transformation A^k is r^{nk}-d_Ω-Lipschitz so

$$\lambda_1(A) - \lambda_2(A) = \frac{1}{k}(\lambda_1(A^k) - \lambda_2(A^k)) \geq \frac{-\log C}{k} + n \log\left(\frac{1}{r}\right),$$

and letting $k \to \infty$ we get

$$\lambda_1(A) - \lambda_2(A) \geq n \log\left(\frac{1}{r}\right).$$

Hence $\omega = \log \frac{1}{r}$ is the positive number we were looking for. □

Proof of Proposition 4.12. This is an immediate consequence of the previous lemma, the fact that a matrix $g \in \mathrm{GL}_d(\mathbb{R})$ is totally positive if and only if $\Lambda^k g \in \mathrm{GL}_{d_k}(\mathbb{R})$ is positive for any $1 \leq k \leq d$, and the fact that

$$\lambda_k(g) - \lambda_{k+1}(g) = \lambda_1(\Lambda^k g) - \lambda_2(\Lambda^k g).$$

We also used that the $\Lambda^{d-k}\mathbb{R}^d$ is naturally the dual to $\Lambda^k\mathbb{R}^d$, and that given a k-plane spanned by $v_1, \ldots, v_k \in \mathbb{R}^d$ and a transverse $d-k$-plane spanned by $w_1, \ldots, w_{d-k}$, the angle between these two subspaces is equal to the angle in $\Lambda^k\mathbb{R}^d$ between the vector $v_1 \wedge \cdots \wedge v_k$ and the hyperplane kernel of $w_1 \wedge \cdots \wedge w_{d-k} \in \Lambda^{d-k}\mathbb{R}^d$ seen as a linear form on $\Lambda^k\mathbb{R}^d$.

That it holds $d^{\mathfrak{F}}(\mathbf{x}, h\mathbf{x}) \geq \frac{n}{C'}$ comes from the fact that $d^{\mathfrak{F}}(\mathbf{x}, h\mathbf{x})$ is bounded from below by the Finsler translation length of h acting on $\mathbb{X}$, which is given by $\alpha_0(\lambda_1(h), \ldots, \lambda_d(h))$. Furthermore, the restriction of α_0 to the set of diagonal matrices with ordered diagonal entries $(v_1, \ldots, v_d)$ such that $v_k \geq v_{k+1}$ for each k is uniformly comparable to the maximums norm on the diagonal entries. □

Proposition 4.12 has the following consequence in terms of admissible paths.

Proposition 4.14. *For any $\omega > 0$ there exists $C_\omega > 0$ such that for any $(\omega, 0)$-admissible path $c : [0, T] \to G$, we have*

$$d^{\mathfrak{F}}(c(0) \cdot \mathbf{x}, c(T) \cdot \mathbf{x}) \geq \frac{k-2}{C_\omega},$$

where k is the number of singularities (i.e. $k+1$ is the number of geodesic pieces of c).

Observe that we need the -2 term in $(k-2)/C_\omega$ because we allow the first and last pieces of c to have length less than ω.

Proof. Without loss of generality we can assume $k \geq 3$, so that c contains at least one piece of hyperbolic type of length at least ω, and that $c(0) = 1$. Let r be the number of hyperbolic pieces of c of length at least ω; note that

$$\frac{k-2}{2} \leq r \leq \frac{k+2}{2}.$$

Then by definition of admissible (Definition 2.8) and Theorem 4.1, we can write $c(T) \in G_{>0}$ as the following product:

$$c(T) = g_0 a'_\omega g_1 a'_\omega g_2 \cdots g_{r-1} a'_\omega g_r,$$

where $g_0, \ldots, g_r \in G_{\geq 0}$.

By Proposition 4.12, there is $C > 0$ only depending on ω such that

$$d^{\mathfrak{F}}(c(0)\mathbf{x}, c(T)\mathbf{x}) \geq \frac{r}{C} \geq \frac{k-2}{2C}. \qquad \square$$

5 Geometric control: Uniform quasi-isometry

This section contains the main geometric results of this article. Recall that S is a hyperbolic closed surface, let $G = \mathrm{PSL}_d(\mathbb{R})$ and $\tau : \mathrm{PSL}_2(\mathbb{R}) \to G$ be the usual irreducible representation. In Section 2.3, given a collection of disjoint closed curves on S and an element of the Cartan subspace of G for each of these curves, we have recalled the definition of bending $\tau(\pi_1(S))$ inside G along these closed curves via the elements of the Cartan subspace. Moreover, in Section 2.4, we associated to such a bending an abstract grafting S_z of S (where z is the grafting parameter) and an equivariant, 1-Lipschitz, and piecewise totally geodesic map

$$\tilde{Q}_z : \tilde{S}_z \to \mathbb{X}$$

from its universal covering $\tilde{S}_z$ to the symmetric space $\mathbb{X}$ of G which projects to a map $Q_z : S_z \to \rho_z \backslash \mathbb{X}$.

Note that G is real split and τ is a regular and positive representation in the sense that it maps (projectivizations of) positive matrices in $\mathrm{PSL}_2(\mathbb{R})$ to (projectivizations of) totally positive matrices in G (see Section 4). Then the bent representation of $\pi_1(S)$ is Hitchin, which implies by independent (and different) work of Labourie [Lab06] and Fock–Goncharov [FG06] that our equivariant map $\tilde{S}_z \to \mathbb{X}$ is a quasi-isometric embedding.

In this section we give an upper bound for the multiplicative error of this quasi-isometric embedding and establish a more precise version of Theorem A. Our proof does not rely directly on the work of Labourie and Fock–Goncharov, but follows from the results on totally positive matrices proved in Section 4, which were greatly inspired by Fock–Goncharov's use of positivity.

Theorem 5.1. *For every $\sigma > 0$, there exists $C_\sigma > 0$ such that the following holds.*

Consider a closed hyperbolic surface S, a multicurve $\gamma^* \subset S$ whose components have length at most σ, and a grafting parameter z such that all cylinder heights of the abstract grafting S_z are bounded from below by some number $L > 0$.

Let us endow $\mathbb{X}$ with the G-invariant admissible Finsler metric $\mathfrak{F}$ and S_z with the pullback of this metric under Q_z, denoted by $d^{\mathfrak{F}}_{\tilde{S}_z}$. Then the grafting map $\tilde{Q}_z : \tilde{S}_z \to \mathbb{X}$ is an injective quasi-isometric embedding with multiplicative constant $(1 + C_\sigma/(L+1))$ and additive constant C_σ; more precisely, for all $x, y \in \tilde{S}_z$ we have

$$\left(1 + \frac{C_\sigma}{L+1}\right)^{-1} d^{\mathfrak{F}}_{\tilde{S}_z}(x,y) - C_\sigma \leq d^{\mathfrak{F}}(\tilde{Q}_z(x), \tilde{Q}_z(y)) \leq d^{\mathfrak{F}}_{\tilde{S}_z}(x,y).$$

Moreover, the image $\tilde{S}_z^\iota = \tilde{Q}_z(\tilde{S}_z)$ is C_σ-Finsler-quasiconvex in the sense that for all $x, y \in \tilde{Q}_z(\tilde{S}_z)$, there is a Finsler geodesic from x to y at distance at most C_σ from $\tilde{S}_z^\iota$.

The facts that $\tilde{Q}_z$ is injective and $\tilde{S}_z^\iota$ is quasi-convex have already been established in Corollaries 4.6 and 4.11.

Note that the upper bound for $d^{\mathfrak{F}}(\tilde{Q}_z(x), \tilde{Q}_z(y))$ is an immediate consequence of the definition of $d^{\mathfrak{F}}_{\tilde{S}_z}$ as the pullback of $d^{\mathfrak{F}}$.

The remaining estimate will be obtained as a consequence of Observation 2.9 and an intermediate proposition stating that the images of admissible paths in $\mathbb{X}$ are quasi-geodesics, with control on the multiplicative error term. This intermediate result will be proved using Proposition 4.10 (that admissible paths are quasi-ruled) and Proposition 4.14 (a lower bound on the displacement of admissible paths).

The collar lemma for hyperbolic surfaces states that for any $\sigma > 0$, if

$$\omega = \sinh^{-1}\left(\frac{1}{\sinh(\sigma/2)}\right)$$

then any simple closed geodesic $\gamma^* \subset S$ of length at most σ will have a *collar size* bounded from below by ω, in the sense that $N_\omega(\gamma^*) = \{z \mid d(z, \gamma^*) \leq \omega\}$ is an annulus. Then Observation 2.9 says that for any multicurve $\gamma^* \subset S$ with components of length at most σ, for any grafting parameter z, the image of an admissible path of $\widetilde{S}_z$ under $\widetilde{Q}_z : \widetilde{S}_z \to \mathbb{X}$ will be a $(\omega, 0)$-admissible path of $\mathbb{X}$.

We will also prove the following coarse estimates on lengths.

Theorem 5.2. *In the setting of Theorem 5.1, let $(\rho_z)_z$ be the associated family grafted Hitchin representations. Then there is C'_σ only depending on σ such that for any $\gamma \in \pi_1(S)$,*

$$\ell^{\mathfrak{F}}(\rho_z(\gamma)) \geq \frac{L+1}{C'_\sigma}\iota(\gamma, \gamma^*).$$

Moreover, recalling that z is the datum of a vector $z_e \in \mathfrak{a}$ for each component $e \subset \gamma^*$, then C'_σ may be chosen so that if $z_e \in \ker(\alpha_0)$ for any e then

$$\ell^{\mathfrak{F}}(\rho_z(\gamma)) \geq \left(1 + \frac{C'_\sigma}{L+1}\right)^{-1}\ell_S(\gamma),$$

where $\ell_S(\gamma)$ is the length of γ in S.

5.1 Elementary observations

The triangle inequality easily yields the following.

Observation 5.3. *For any piecewise geodesic curve $c : [0, T] \to (\mathbb{X}, d^{\mathfrak{F}})$ with $m \geq 1$ geodesic pieces, that is additionally C-quasi-ruled, we have*

$$d^{\mathfrak{F}}(c(0), c(T)) \geq \mathrm{Len}^{\mathfrak{F}}(c) - (m-1)C.$$

We will also need the following technical estimates.

Lemma 5.4. *Consider $a, b, k, L \geq 0$ and $C \geq 1$, such that*

$$a \geq b - kC \quad \textit{and} \quad b \geq (k-2)L \quad \textit{and} \quad a \geq \frac{k-2}{C}.$$

Then

$$a \geq \left(1 + \frac{4(C+1)^3}{L+1}\right)^{-1} b - 2C.$$

Proof. We can assume $L > 0$. Use the second equation to get

$$k \leq \tfrac{b}{L} + 2$$

Consider first the case that $L \geq 2C \geq 2$. We By the first equation, we have

$$a \geq b(1 - \tfrac{C}{L}) - 2C$$

with the additional inequality

$$1 - \tfrac{C}{L} \geq \left(1 + 2\tfrac{C}{L}\right)^{-1} \geq \left(1 + \tfrac{4(C+1)^3}{L+1}\right)^{-1}$$

(indeed one can check that $0 \leq x \leq \frac{1}{2}$ implies $(1-x)(1+2x) \geq 1$).

If $L < 2C$ then use the third equation to obtain

$$a \geq \frac{b - 2C}{1 + C^2} \geq \frac{b}{1 + C^2} - 2C$$

and insert

$$\left(1 + C^2\right)^{-1} \geq \left(1 + \tfrac{4(C+1)^3}{L+1}\right)^{-1}.$$

□

5.2 Admissible paths are uniform quasi-geodesics

To prove that the embedding $\widetilde{Q}_z : \widetilde{S}_z \to X$ is quasi-isometric, by Observation 2.9 it suffices to show that admissible paths in $\mathbb{X}$ are quasi-geodesics. We prove this now with uniform constants using the consequences of positivity established in Section 4 and the elementary observations of the previous section.

Proposition 5.5. *For all $\omega > 0$, there exist $C_\omega > 0$ such that for every $L \geq 0$, all (ω, L)-admissible paths are $\left(1 + \frac{C_\omega}{(L+1)}, C_\omega\right)$-quasi-Finsler-geodesics (where the first constant is the multiplicative constant).*

Proof. Let $c : [0, T] \to \mathbb{X}$ be an (ω, L)-admissible path. It is clear that $d^{\mathfrak{F}}(c(0), c(T)) \leq \mathrm{Len}^{\mathfrak{F}}(c)$, so we only need to obtain a converse inequality.

By Proposition 4.10, c is C_ω-quasi-ruled, for some constant C_ω depending on ω. By Observation 5.3 we get

$$d^{\mathfrak{F}}(c(0), c(T)) \geq \mathrm{Len}^{\mathfrak{F}}(c) - kC_\omega,$$

where k is the number of singularities of c (i.e. $k+1$ is the number of geodesic pieces).

Since c contains at least $(k-2)/2$ geodesic pieces of flat type and length at least L, we also know that

$$\mathrm{Len}^{\mathfrak{F}}(c) \geq (k-2)\tfrac{L}{2}.$$

Finally by Proposition 4.14 we also have

$$d^{\mathfrak{F}}(c(0), c(T)) \geq \frac{k-2}{C'_\omega}$$

for some constant C'_ω depending on ω.

We conclude thanks to Lemma 5.4. □

5.3 Proof of Theorem 5.1

It is an immediate consequence of Observation 2.9, Corollary 4.6, Corollary 4.11 and Proposition 5.5. More precisely, by compactness of S, Corollary 4.6 ensures that the natural map $S_z \xrightarrow{Q_z} \rho_z\backslash\mathbb{X}$ is an embedding.

Consider an equivariant lift $\tilde{Q}_z : \tilde{S}_z \to \mathbb{X}$ of Q_z. By Corollary 4.11, any two points in $\tilde{Q}_z(\tilde{S}_z)$ can be connected by a Finsler geodesic in $\mathbb{X}$ which remains at distance at most C_ω from $Q_z(S_z)$. This establishes the last statement in Theorem 5.1. Note that these Finsler geodesics then project to Finsler geodesics in the quotient $\rho_z\backslash\mathbb{X}$ in the C_ω-neighborhood of $Q_z(S_z)$.

To show the distance-length control, note that any admissible path in S_z is sent by Q_z to an (ω, L)-admissible path inside $\rho_z\backslash\mathbb{X}$ (Observation 2.9), which is therefore a quasi-geodesic (Proposition 5.5).

5.4 Proof of Theorem 5.2

Fix $[\gamma] \in [\pi_1(S)]$ transverse to γ^* and recall that it corresponds to a free homotopy class in the characteristic surface $S_z^\iota \subset \rho_z\backslash\mathbb{X}$. Let $c \subset S_z^\iota \subset \rho_z\backslash\mathbb{X}$ be the unique admissible loop in this free homotopy class. It has a $\rho_z(\gamma)$-invariant lift $\tilde{c} : \mathbb{R} \to \mathbb{X}$ such that 0 is a singularity and the geodesic piece of $\tilde{c}$ starting at time 0 is of hyperbolic type.

Denote by T the period of c, so that $\tilde{c}(t+T) = \rho_z(\gamma)\tilde{c}(t)$ for any t. By Proposition 5.5, for any $n \geq 1$ we have

$$d^{\mathfrak{F}}(\tilde{c}(0), \rho_z(\gamma)^n\tilde{c}(0)) = d^{\mathfrak{F}}(\tilde{c}(0), \tilde{c}(nT)) \geq \left(1 + \frac{C_\sigma}{L+1}\right)^{-1} n\,\mathrm{Len}^{\mathfrak{F}}(c) - C_\sigma.$$

Dividing by n and letting $n \to \infty$ we get

$$\ell^{\mathfrak{F}}(\rho_z(\gamma)) \geq \left(1 + \frac{C_\sigma}{L+1}\right)^{-1} \mathrm{Len}^{\mathfrak{F}}(c). \tag{16}$$

We may assume that $\iota(\gamma, \gamma^*) \geq 1$ (the case $\iota(\gamma, \gamma^*) = 0$ is trivial). The number $2\iota(\gamma, \gamma^*)$ of singularities of c is even and bounded from below by 2, and the same is true for the number of geodesic pieces, half of which are of flat type and have length at least L, and the other half are of hyperbolic type and have length at least $\omega = \sinh^{-1}(\sinh(\sigma/2)^{-1})$. Consequently,

$$\ell^{\mathfrak{F}}(\rho_z(\gamma)) \geq \left(1 + \frac{C_\sigma}{L+1}\right)^{-1} \frac{L+\omega}{2} \iota(\gamma, \gamma^*) \geq \frac{L+\omega}{2 + 2C'_\sigma} \iota(\gamma, \gamma^*).$$

Finally, if for each component $e \subset \gamma^*$ the vector z_e is taken in $\ker(\alpha_0)$ then we can apply Lemma 2.2, which says that $\mathrm{Len}^{\mathfrak{F}}(c)$ is bounded from below by the length of the image of c under the projection maps $S_z \to S$, which is itself greater than or equal to $\ell_S(\gamma)$, and this means Equation 16 implies the desired inequality.

5.5 Proof of Theorem C

We begin with the proof of the third part of Theorem C. Thus let S be a closed surface with hyperbolic metric h and let $S_0 \subset S$ be an essential subsurface, bounded by simple closed geodesics $\partial S_0 = \{\gamma_1, \ldots, \gamma_k\}$. Consider a one-parameter family ρ_t of Hitchin grafting representations with grafting datum tz for a tuple $z = (z_1, \ldots, z_k) \in \mathfrak{a}^k$ whose components are linearly independent from the direction of a tangent vector of $\mathbb{H} \subset \mathbb{X}$. Then for each t, the bordered hyperbolic surface S_0 is totally geodesic embedded in $\rho_t \backslash \mathbb{X}$.

Choose once and for all a basepoint $\mathbf{x} \in S_0$ and view this as a basepoint in $\rho_t \backslash \mathbb{X}$ for all t. By Theorem 5.2 and equivalence of the Riemannian and the Finsler metric, for each $R > 0$ there exists a number $t = t(R) > 0$ so that the shortest closed geodesic in $\rho_t \backslash \mathbb{X}$ which is not contained in S_0 intersects the complement of the R-ball about S_0. Thus for $t > t(R)$, the normal injectivity radius of S_0 in $\rho_t \backslash \mathbb{X}$ for the locally symmetric Riemannian metric is at least R, and the ball $B(S_0, R) \subset \rho_t \backslash \mathbb{X}$ of radius R about S_0 is homotopy equivalent to S_0.

By passing to a subsequence, we may assume that the pointed manifolds $(\rho_t \backslash \mathbb{X}, \mathbf{x})$ converge in the pointed Gromov Hausdorff topology to a locally symmetric pointed manifold $(N, \mathbf{x})$. This manifold contains S_0 as a totally geodesic embedded surface of infinite normal injectivity radius. But this just means that N equals the manifold defined by the Fuchsian representation $\rho | S_0$. This is precisely the statement of the last part of Theorem C.

To show the second part of Theorem C, choose a non-principal ultrafilter ω on $\mathbb{R}$ converging to $+\infty$ and a basepoint $\mathbf{x}$ in $\mathbb{X}$, viewed as the projection of the identity in $\mathrm{PSL}_d(\mathbb{R})$. Let ρ_t be a Hitchin grafting ray determined by a grafting parameter $z \in \mathfrak{a}$ and a simple closed geodesic $\gamma^* \subset S$. Consider the pointed metric spaces $\mathbb{X}_t = (\mathbb{X}, \mathbf{x}, \frac{1}{t} d^{\mathfrak{F}})$. We will see that for any $\gamma \in \pi_1(S)$, the distance in $\mathbb{X}_t$ between $\mathbf{x}$ and $\rho_t(\gamma) \cdot \mathbf{x}$ is bounded independently of t. Thus by passing to an ω-ultralimit, we obtain an action ρ_∞ of $\pi_1(S)$ on the ultralimit $\mathbb{X}_\infty$ of the metric spaces $\mathbb{X}_t$ defined by ω. This ultralimit is well known to be a Euclidean building whose apartments correspond to the maximal flats in $\mathbb{X}$, and it is equipped with a Finsler metric.

We first make some choice that will allow us to embed the Bass–Serre tree $\mathcal{T}$ of the graph of groups defined by $\gamma^* \subset S$ in $\mathbb{X}_t$. This embedding will not be ρ_t-equivariant, but its limit as $t \to \infty$ will be natural (independent of the choices made) and equivariant.

Let $\tilde{S} \simeq \mathbb{H}^2$ be the universal cover of S, and $\tilde{\gamma}^*$ the preimage of γ^*. Let Σ be the surface with boundaries obtained by cutting $\tilde{S}$ along $\tilde{\gamma}^*$. In each component $Y \subset \Sigma$ we pick a point $p(Y)$, and in each boundary component $B \subset \partial Y \subset \partial\Sigma$ we denote by $p(B) \in B$ the shortest distance projection of $p(Y)$.

For each t we have an abstract grafted surface $\tilde{S}_t$, in which Σ embeds isometrically, and we have a quasi-isometric embedding $Q_t : \tilde{S}_t \to \tilde{S}_t^\iota \subset \mathbb{X}$ that maps $p(Y_0)$ (where Y_0 is our preferred component of Σ) to the basepoint $\mathbf{x} \in \mathbb{X}$. For all t, $Y \subset \Sigma$ and $B \subset \partial Y$ we denote $p_t(Y) = Q_t(p(Y))$ and $p_t(B) = Q_t(p(B))$. Let $\mathcal{T}_t \subset \tilde{S}_t^\iota$ be the union of segments of the form $[p_t(Y), p_t(B)]$ if $B \subset \partial Y$ and $[p_t(B_1), p_t(B_2)]$ if $B_1, B_2 \subset \partial\Sigma$ project to the same component of $\tilde{\gamma}^*$ in $\tilde{S}$. Note that $\mathcal{T}_t$ is a tree isomorphic to the Basse–Serre tree $\mathcal{T}$.

The length of $[p_t(Y), p_t(B)]$ in $\mathbb{X}$ is independent of t, hence its length in $\mathbb{X}_t$ tends to zero as $t \to \infty$. The segment $[p_t(B_1), p_t(B_2)]$ is conjugate to a segment of the form $[0, v + tz] \subset \mathfrak{a}$ where v depends on B_1, B_2, so its length in $\mathbb{X}_t$ is bounded by a constant independent of t. As a consequence, every $p_t(Y)$, resp. $p_t(B)$, converge as $t \to \infty$ to a points $p_\infty(Y)$, resp. $p_\infty(B)$, in the building $\mathbb{X}_\infty$. In fact if $B \subset \partial Y$ then $p_\infty(B) = p_\infty(Y)$. Moreover, the whole tree $\mathcal{T}_t$ converges to a tree $\mathcal{T}_\infty \subset \mathbb{X}_\infty$ whose vertices are $(p_\infty(Y))_{Y \subset \Sigma}$, whose edges are straight segments conjugate to $[0, z] \subset \mathfrak{a}$, and which is also isometric to $\mathcal{T}$.

Now observe that $\mathcal{T}_\infty$ is invariant under the action of any $\gamma \in \pi_1(S)$: for any $Y \subset \Sigma$ and t, the distance in $\mathbb{X}$ between $\rho_t(\gamma)p_t(Y)$ and $p_t(\gamma Y)$ is independent of t, and hence the distance in $\mathbb{X}_t$ tends to zero, leading to $\rho_\infty(\gamma)p_\infty(Y) = p_\infty(\gamma Y)$ (and this also proves ρ_∞ is well defined).

Finally, note that $\mathcal{T}_\infty$ is Finsler-convex. Indeed for all $Y, Z \subset \Sigma$ and t, the path in $\mathcal{T}_t$ from $p_t(Y)$ to $p_t(Z)$ and the admissible path connecting the same two points are within bounded distance independent of t. By Proposition 4.10, the path in $\mathcal{T}_t$ is at distance $\leq K$ from an actual Finsler geodesic, where K does not depend on t. Hence in $\mathbb{X}_t$ it is at distance $\leq K/t$ from a Finsler geodesic, and at the limit the path in $\mathcal{T}_\infty$ is a Finsler geodesic itself. This yields the statement of the second part of Theorem C.

The first part of Theorem C follows immediately.

References

[And87] T. Ando. *Totally positive matrices.* Linear Algebra Appl. 90 (1987), pp. 165–219.

[AZ23] J. Audibert and M. Zshornack. *Rational approximation for Hitchin representations.* 2023. arXiv: 2310.15121.

[BCLS15] M. Bridgeman, R. Canary, F. Labourie, and A. Sambarino. *The pressure metric for Anosov representations.* Geometric and Functional Analysis 25.4 (June 2015), pp. 1089–1179.

[BH12] G. Ben Simon and T. Hartnick. *Invariant orders on Hermitian Lie groups.* J. Lie Theory 22.2 (2012), pp. 437–463.

[BH99] M. Bridson and A. Haefliger. *Metric Spaces of Non-Positive Curvature.* Springer Berlin Heidelberg, 1999.

[BHM11] S. Blachère, P. Haïssinsky, and P. Mathieu. *Harmonic measures versus quasiconformal measures for hyperbolic groups.* Ann. Sci. Éc. Norm. Supér. (4) 44.4 (2011), pp. 683–721.

[BHMM24] P.-L. Blayac, U. Hamenstädt, T. Marty, and A.E. Monti. *Hitchin grafting representations II: Dynamics.* 2024.

[Bir57] G. Birkhoff. *Extensions of Jentzsch's theorem.* Trans. Amer. Math. Soc. 85 (1957), pp. 219–227.

[BPS19] J. Bochi, R. Potrie, and A. Sambarino. *Anosov representations and dominated splittings.* J. Eur. Math. Soc. (JEMS) 21.11 (2019), pp. 3343–3414.

[BQ16] Y. Benoist and J. Quint. *Random Walks on Reductive Groups.* Springer International Publishing, 2016, pp. 153–167.

[EF97] A. Eskin and B. Farb. *Quasi-flats and rigidity in higher rank symmetric spaces.* J. Amer. Math. Soc. 10.3 (1997), pp. 653–692.

[FG06] V. Fock and A. Goncharov. *Moduli spaces of local systems and higher Teichmüller theory.* Publications mathématiques de l'IHÉS 103 (June 2006), pp. 1–211.

[FK16] P. Foulon and I. Kim. *Topological Entropy and bulging deformation of real projective structures on surface.* 2016. arXiv: `1608.06799`.

[Gol86] W. Goldman. *Invariant functions on Lie groups and Hamiltonian flows of surface group representations.* Inventiones Math. 85.2 (1986), pp. 263–302.

[GW18] O. Guichard and A. Wienhard. *Positivity and higher Teichmüller theory.* *European Congress of Mathematics.* Eur. Math. Soc., Zürich, 2018, pp. 289–310.

[Hit92] N. J. Hitchin. *Lie groups and Teichmüller space.* English. Topology 31.3 (1992), pp. 449–473.

[JM87] D. Johnson and J. J. Millson. *Deformation spaces associated to compact hyperbolic manifolds.* English. Discrete groups in geometry and analysis, Pap. Hon. G. D. Mostow 60th Birthday, Prog. Math. 67, 48-106 (1987). 1987.

[KL18] M. Kapovich and B. Leeb. *Finsler bordifications of symmetric and certain locally symmetric spaces.* Geom. Topol. 22.5 (2018), pp. 2533–2646.

[KL97] B. Kleiner and B. Leeb. *Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings.* Inst. Hautes Études Sci. Publ. Math. 86 (1997), pp. 115–197.

[KLP18] M. Kapovich, B. Leeb, and J. Porti. *A Morse lemma for quasigeodesics in symmetric spaces and Euclidean buildings*. Geom. Topol. 22.7 (2018), pp. 3827–3923.

[Kni98] G. Knieper. *The Uniqueness of the Measure of Maximal Entropy for Geodesic Flows on Rank 1 Manifolds*. The Annals of Mathematics 148.1 (1998), pp. 291–314.

[Lab06] F. Labourie. *Anosov flows, surface groups and curves in projective space*. Inventiones math. 165.1 (2006), pp. 51–114.

[Lof19] J. Loftin. *Convex $\mathbb{RP}^2$ structures and cubic differentials under neck separation*. English. J. Differ. Geom. 113.2 (2019), pp. 315–383.

[LTW22] J. Loftin, A. Tamburelli, and M. Wolf. *Limits of Cubic Differentials and Buildings*. 2022. arXiv: 2208.07532.

[Lus94] G. Lusztig. *Total positivity in reductive groups. Lie theory and geometry*. Vol. 123. Progr. Math. Birkhäuser Boston, Boston, MA, 1994, pp. 531–568.

[Mes07] G. Mess. *Lorentz spacetimes of constant curvature*. English. Geom. Dedicata 126 (2007), pp. 3–45.

[Par00] A. Parreau. *Dégenérescenses de sous-groupes discrets de groupes de Lie semisimples et actions de groupes sur des immeubles affines*. 2000.

[Par12] A. Parreau. *Compactification d'espaces de représentations de groupes de type fini*. Math. Z. 272 (2012), pp. 51–86.

[Par22] A. Parreau. *Invariant weakly convex cocompact subspaces for surface groups in A_2-buildings*. Transactions AMS. 375 (2022), pp. 2293–2339.

[PT14] A. Papadopoulos and M. Troyanov, eds. *Handbook of Hilbert geometry*. Vol. 22. IRMA Lectures in Mathematics and Theoretical Physics. European Mathematical Society (EMS), Zürich, 2014, pp. viii+452.

[Rei23] C. Reid. *Limits of Convex Projective Surfaces and Finsler Metrics*. 2023. arXiv: 2309.10290.

[Ser77] J. Serre. *Arbres, amalgames,* SL_2. Vol. No. 46. Société Mathématique de France, Paris, 1977, 189 pp. (1 plate).

[Tan97] H. Tanigawa. *Grafting, harmonic maps and projective structures on surfaces*. Journal of Differential Geometry 47.3 (1997), pp. 399–419.

[Thu97] W. P. Thurston. *Three-dimensional geometry and topology. Vol. 1. Ed. by Silvio Levy*. English. Vol. 35. Princeton Math. Ser. Princeton, NJ: Princeton University Press, 1997.

Pierre-Louis Blayac
Université de Strasbourg, IRMA, 7 rue René-Descartes, 67084 Strasbourg, France
e-mail: blayac@unistra.fr

Ursula Hamenstädt
Math. Institut der Univ. Bonn, Endenicher Allee 60, 53115 Bonn, Germany
e-mail: ursula@math.uni-bonn.de

Théo Marty
Institut de mathématiques de Bourgogne, 9 avenue Alain Savary, 21078 Dijon, France
e-mail: theo.marty@u-bourgogne.fr