Geometry & Topology 15 (2011) 1927–1981

Coarse differentiation and quasi-isometries of a class of solvable Lie groups II

Irine Peng

In this paper, we continue with the results in [12] and compute the group of quasi-isometries for a subclass of split solvable unimodular Lie groups. Consequently, we show that any finitely generated group quasi-isometric to a member of the subclass has to be polycyclic and is virtually a lattice in an abelian-by-abelian solvable Lie group. We also give an example of a unimodular solvable Lie group that is not quasi-isometric to any finitely generated group, as well deduce some quasi-isometric rigidity results.

51F99; 22E40

1 Introduction

A (κ, C) *quasi-isometry* f between metric spaces X and Y is a map $f\colon X \to Y$ satisfying

$$\frac{1}{\kappa}d(p,q) - C \leq d(f(p), f(q)) \leq \kappa d(p,q) + C$$

with the additional property that there is a number D such that Y is the D neighborhood of $f(X)$. Two quasi-isometries f, g are considered to be equivalent if there is a number $E > 0$ such that $d(f(p), g(p)) \leq E$ for all $p \in X$.

Let $G = \mathbb{R}^m \rtimes_\varphi \mathbb{R}^n$, $G' = \mathbb{R}^{m'} \rtimes_{\varphi'} \mathbb{R}^{n'}$ be connected, simply connected nondegenerate unimodular split solvable groups (see Section 2.1 for definitions). We say a map from G to G' is *standard* if it splits as a product map that respects φ and φ' (see Definition 2.1.3). The main result of this paper is the following statement.

Theorem 5.3.6 (abridged) *Let G, G' be nondegenerate, unimodular, split abelian-by-abelian solvable Lie groups, and $\phi\colon G \to G'$ a (κ, C) quasi-isometry. Then ϕ is bounded distance from a composition of a left translation and a standard map.*

1.0.1 Definition A homomorphism $\varphi\colon \mathbb{R}^m \to \mathrm{GL}_n$ is called *diagonalizable* if its image can be conjugated into the set of diagonal matrices.

Corollary 5.3.8 *If φ is diagonalizable and φ' isn't, then there is no quasi-isometry between them.*

Published: 17 October 2011 DOI: 10.2140/gt.2011.15.1927

1.0.2 Corollary

$$\mathrm{QI}(G) = \left(\prod_{[\alpha]} \mathrm{Bilip}(V_{[\alpha]})\right) \rtimes \mathrm{Sym}(G)$$

Here $[\alpha]$ is an equivalence class of roots and $\mathrm{Sym}(G)$ is a finite group, analogous to the Weyl group in reductive Lie groups, that reflects the symmetries of G. (See Section 2.1.)

As an application the work by Dymarz [2] on quasi-conformal maps on the boundary of G and its generalization Dymarz and Peng [3], we have:

1.0.3 Corollary *If Γ is a finitely generated group quasi-isometric to a G, then Γ is virtually polycyclic.*

1.1 Proof outline

Our starting point is Theorem 1.3.3 from Peng [12] (restated here in Section 3), which says that given a large enough box, when we express it as a tessellation of smaller boxes, the restriction of the map to most of those smaller boxes takes a product structure $f \times g$, where $f\colon \mathbf{A} \to \mathbf{A}'$ and $g\colon \mathbf{H} \to \mathbf{H}'$ satisfy some particular conditions. Our first task is to show that the $\mathbf{A}'$ part of the standard maps f_i (see Definition 2.1.3) are affine. This is done in Section 3, where will see that the linear part is a scalar multiple of a finite order element in $O(n)$ (where n is the rank of G). We also give interpretations of the linear and constant parts of f_i in terms of properties of G' and the measure of certain sets in the box where f_i was partially defined. In Section 4, we show that the linear part of the f_i's in different boxes have to be the same up to scalar multiple in the case that the rank of G is 2 or higher. The rank 1 case is the same as the content in Eskin, Fisher and Whyte [5]. The proof for higher rank case basically consists of as many rank 1 arguments as appeared in [5]. In the last section, we put all theses partially defined standard maps together to produce a splitting of the original quasi-isometry.

Acknowledgements I would like to thank Alex Eskin and David Fisher for their incredible patience and assistance. I also wish to thank Tullia Dymarz for helpful conversations and coordinating results of her paper [2] with this one and to thank Mikhail Ershov for finding me the reference of Hasegawa [9] containing the example of a solvable Lie group that admits no lattices.

2 Preliminaries

Here we recall the settings from Peng [12] and define new terms that will be used in this paper. Please see Sections 1.1 and 2.1 therein for further details.

2.1 Geometry of a certain class of solvable Lie groups

Nondegenerate, split abelian-by-abelian solvable Lie groups Let $\mathfrak{g}$ be a (real) solvable Lie algebra, and $\mathfrak{a}$ be a Cartan subalgebra, which is a self-normalizing subalgebra and exists in any Lie algebras as long as the underlying field is infinite. Then there are finitely many nonzero linear functionals $\alpha_i \colon \mathfrak{a} \to \mathbb{C}$ called *roots*, such that

$$\mathfrak{g} = \mathfrak{a} \oplus \bigoplus_{\alpha_i} \mathfrak{g}_{\alpha_i},$$

where $\mathfrak{g}_{\alpha_i} = \{x \in \mathfrak{g} : \forall t \in \mathfrak{a}, \exists n \text{ such that } (\mathrm{ad}(t) - \alpha_i(t)\,\mathrm{Id})^n(x) = 0\}$, Id is the identity map on $\mathfrak{g}$, and $ad\colon \mathfrak{g} \to \mathrm{Der}_{\mathbb{R}}(\mathfrak{g})$ is the adjoint representation. Let $\triangle$ denotes for the set of roots. Then $\mathrm{Aut}(\mathfrak{a})$ acts on $\triangle$ in a natural way. We define $\mathrm{Perm}(\mathfrak{g})$ to be the subgroup consisting $A \in \mathrm{Aut}(\mathfrak{a})$ such that

(i) it leaves the set of roots invariant, ie $A\triangle = \triangle$;

(ii) for every $\alpha \in \triangle$, $\dim \mathfrak{g}_{\alpha \circ A} = \dim \mathfrak{g}_\alpha$.

In this way, elements of $\mathrm{Perm}(\mathfrak{g})$ induces a permutation on the set $\triangle$, and we define $\mathrm{Sym}(\mathfrak{g})$ to be the image of $\mathrm{Perm}(\mathfrak{g})$ in the group of permutations of $\triangle$. For a generic $\mathfrak{g}$, its $\mathrm{Perm}(\mathfrak{g})$ is trivial.

We say $\mathfrak{g}$ is *split abelian-by-abelian* if $\mathfrak{g}$ is a semidirect product of $\mathfrak{a}$ and $\bigoplus_i \mathfrak{g}_{\alpha_i}$, and both are abelian Lie algebras; *unimodular* if the roots sum up to zero; and *nondegenerate* if the roots span $\mathfrak{a}^*$. In particular, nondegenerate means that each α_i is real-valued, and the number of roots is at least the dimension of $\mathfrak{a}$. Being unimodular is the same as saying that for every $t \in \mathfrak{a}$, the trace of $\mathrm{ad}(t)$ is zero. We extend these definitions to a Lie group if its Lie algebra has these properties, and write $\mathrm{Perm}(G)$, $\mathrm{Sym}(G)$ to mean $\mathrm{Perm}(\mathfrak{g})$ and $\mathrm{Sym}(\mathfrak{g})$ where $\mathfrak{g}$ is the Lie algebra of G.

In summary, a connected, simply connected solvable Lie group G that is nondegenerate, split abelian-by-abelian necessary takes the form $G = \mathbf{H} \rtimes_\varphi \mathbf{A}$ such that

(i) both $\mathbf{A}$ and $\mathbf{H}$ are abelian Lie groups;

(ii) $\varphi\colon \mathbf{A} \to \mathrm{Aut}(\mathbf{H})$ is injective;

(iii) there are finitely many $\alpha_i \in \mathbf{A}^* \backslash 0$ which together span $\mathbf{A}^*$, and a decomposition of $\mathbf{H} = \bigoplus_i V_{\alpha_i}$;

(iv) there is a basis $\mathcal{B}$ of $\mathbf{H}$ whose intersection with each of V_{α_i} constitute a basis of V_{α_i}, such that for each $\mathbf{t} \in \mathbf{A}$, $\varphi(\mathbf{t})$ with respect to $\mathcal{B}$ is a matrix consists of blocks, one for each V_{α_i}, of the form $e^{\alpha_i(\mathbf{t})} N(\alpha_i(t))$, where $N(\alpha_i(t))$ is an upper triangular with 1's on the diagonal and whose off-diagonal entries are polynomials of $\alpha_i(t)$.

If in addition, G is unimodular, then $\varphi(\mathbf{t})$ has determinant 1 for all $\mathbf{t} \in \mathbf{A}$.

The *rank* of a nondegenerate, split abelian-by-abelian group G is defined to be the dimension of $\mathbf{A}$, and by a result of de Cornulier [1], if two such groups are quasi-isometric, then they have the same rank.

We write $\triangle$ for the set of roots of G and coordinatize points in G as $((\mathbf{x}_\alpha)_{\alpha\in\triangle}, \mathbf{t})$, where $\mathbf{x}_\alpha = (x_{1,\alpha}, x_{2,\alpha}, \dots, x_{\dim(V_\alpha),\alpha}) \in V_\alpha$, $\mathbf{t} \in \mathbf{A}$. A left invariant Finsler metric that is quasi-isometric to a left invariant Riemannian is given by

$$d\mathbf{t} + \sum_{\alpha\in\triangle} e^{-\alpha(\mathbf{t})}\Big(d\mathbf{x}_\alpha + \sum_j P_{j,\alpha}(\alpha(\mathbf{t}))dx_{j,\alpha}\Big),$$

where $P_{j,\alpha}$ is a polynomial. The following consequence is immediate.

2.1.1 Lemma *If G is nondegenerate, split abelian-by-abelian, then it can be quasi-isometric embedded into $\prod_{\alpha\in\triangle} H_{\dim(V_\alpha)+1}$, where $H_{s+1} = \mathbb{R}^s \rtimes_\psi \mathbb{R}$ is a nonunimodular solvable Lie group determined by $\psi(t) = e^t N(t)$, where $N(t)$ is a nilpotent matrix (upper triangular with 1's on the diagonal) with polynomial entries, equipped with a left-invariant Finsler metric given by*

$$dt + e^{-t}\Big(d\mathbf{x} + \sum_j P_j(t)dx_j\Big)$$

where $P_j(t)$ is a polynomial.

2.1.2 Remark When $\psi(t)$ is diagonal, H_{s+1} is just the usual hyperbolic space.

To understand the geometry of H_{s+1} better, we first note that the metric is bilipschitz to one given by $dt + e^{-t}(1 + \max_j P_j(t))d\mathbf{x}$, which is quasi-isometric to one given by $dt + e^{-t}Q(t)d\mathbf{x}$ for some polynomial $Q(t)$. So a function quasi-isometric to the metric on H_{s+1} is the following

$$d((x_1,t_1),(x_2,t_2)) = \begin{cases} |t_1 - t_2| & \text{if } e^{-t_i}Q(t_i)|x_1 - x_2| \le 1 \text{ for } i = 1,2, \\ U_Q(|x_1 - x_2|) - (t_1 + t_2) & \text{otherwise,} \end{cases} \tag{1}$$

where $U_Q(|x_1 - x_2|) = t_0$ satisfies

$$e^{-t_0}Q(t_0)|x_1 - x_2| = 1.$$

Since exponential grows faster than any polynomials, the function U_Q has the following property:

$$\ln(x) - C_Q \le U_Q(x) \le 2\ln(x) + C_Q \tag{2}$$

for some constant C depending only on the polynomial Q.

Back to the description of G, we declare two roots equivalent if they are positive multiples of each other, and write $[\Xi]$ for the equivalence class containing $\Xi \in \triangle$. Moreover, for $\Xi_i, \Xi_2 \in [\Xi]$, we say Ξ_1 less than Ξ_2 if $\Xi_2/\Xi_1 > 1$. This makes sense because all roots in a root class are positive multiples of each other. A left translate of $V_{[\Xi]} = \bigoplus_{\sigma \in [\Xi]} V_\sigma$ will be called a *horocycle of root class* $[\Xi]$.

A left translate of $\mathbf{A}$, or a subset of it, is called a *flat*. For $p = (x_\alpha)_\alpha$, $q = (y_\alpha)_\alpha$ points in $\mathbf{H}$, we compute subsets of $p\mathbf{A}$ and $q\mathbf{A}$ that are within distance $O(1)$ of each other according to the embedded metric in Lemma 2.1.1, as the p and q translate of the subset

$$\bigcap_{\alpha \in \triangle : \ln(|x_\alpha - y_\alpha|) \geq 1} \alpha^{-1}[U_\alpha(|x_\alpha - y_\alpha|), \infty] \subset \mathbf{A}.$$

Since the roots sum up to zero in a nondegenerate, unimodular, split abelian-by-abelian group, the set where two flats come together can be empty, ie the two flats have no intersection. If it is not empty, then the equation above says that it is an unbounded convex subset of $\mathbf{A}$, bounded by hyperplanes parallel to root kernels. We will often say that $p\mathbf{A}$ and $q\mathbf{A}$ *come together* at this set.

2.1.3 Definition Let G, G' be nondegenerate, split abelian-by-abelian Lie groups. A map from G to G' or a subset of them, is called *standard map* if it takes the form $f \times g$, where $g\colon \mathbf{H} \to \mathbf{H}'$ sends foliation by root class horocycles of G to that of G', and $f\colon \mathbf{A} \to \mathbf{A}'$ sends foliations by root kernels of G to that of G'. We will often refer to f as the $\mathbf{A}$ part of a standard map.

2.1.4 Remark Note that when G has at least $\operatorname{rank}(G) + 1$ many root kernels, the condition on f means that f is affine, and when G is rank 1, the condition on f is empty.

2.2 Notation

2.2.1 General remarks about neighborhoods

Neighborhoods We write $B(p, r)$ for the ball centered at p of radius r, and $N_c(A)$ for the c neighborhood of the set A. We also write $d_H(A, B)$ for the Hausdorff distance between two sets A and B. If $\Omega \subset \mathbb{R}^k$ is a bounded compact set, and $r \in \mathbb{R}$, we write $r\Omega$ for the bounded compact set that is scaled from Ω with respect to the barycenter of Ω.

Given a subset $X \subset A$, $x_0 \in X$, the (η, C) *linear neighborhood of* X *with respect to* x_0 is the subset of A consisting of points $y \in A$ such that there is a $\hat{x} \in X$ with

$d(y,\hat{x}) = d(y,X) \le \eta d(\hat{x},x_0) + C$. If a quasi-geodesic λ is within (η, C) linear (or just η–linear) neighborhood of a geodesic segment γ, where $\eta \ll 1$ and $C \ll \eta|\lambda|$, then we say that λ *admits a geodesic approximation* by γ.

2.2.2 Notation used in split abelian-by-abelian groups Let $G = \mathbf{H} \rtimes \mathbf{A}$ stands for a nondegenerate, unimodular, split abelian-by-abelian group. Fix a point $p \in G$. We define the following.

- For a root class $[\alpha]$, let $l_{[\alpha]} = \sum_{\xi \in [\alpha]} \xi$ be the sum of all roots in the equivalence class. Let R_G be the set of all $l_{[\alpha]}$'s for $[\alpha]$ ranging over all root classes of G, and G_R for the group of linear maps that leaves R_G invariant. Since our group is nondegenerate, ie that the roots span $\mathrm{Hom}(\mathbb{A}, \mathbb{R})$, it follows that R_G is finite, and so G_R is a subgroup of $O(n)$.
- For $\alpha \in \triangle$ a root, we write $\alpha_0 \in \mathbf{A}^*$ for the positive multiple of $l_{[\alpha]}$ of unit norm with respect to the usual Euclidean inner product on $\mathbf{A}$ and $\vec{v}_{[\alpha]} \in \mathbf{A}$ for the dual of α_0.
- Given $\vec{v} \in \mathbf{A}$, we define

$$W_{\vec{v}}^{+} = \bigoplus\nolimits_{\Xi(\vec{v})>0} V_{\Xi}$$
$$W_{\vec{v}}^{-} = \bigoplus\nolimits_{\Xi(\vec{v})<0} V_{\Xi}$$
$$W_{\vec{v}}^{0} = \bigoplus\nolimits_{\Xi(\vec{v})=0} V_{\Xi}.$$

  We say a vector $\vec{v} \in \mathbf{A}$ is *regular* if $\Xi(\vec{v}) \neq 0$ for all roots Ξ, and a linear functional $\ell \in \mathbf{A}^*$ is regular if its dual $\vec{v}_\ell$ is regular. We extend this definitions to a linear function ℓ by defining W_ℓ^+, W_ℓ^-, W_ℓ^0, as $W_{\vec{v}_\ell}^+$, $W_{\vec{v}_\ell}^-$, $W_{\vec{v}_\ell}^0$ respectively, where $\vec{v}_\ell \in \mathbf{A}$ is the dual of ℓ.
- The root kernels partition the unit sphere in $\mathbf{A}$ into convex subsets called *chambers*. For vectors $\vec{u}, \vec{v}$ in the interior of the same chamber $\mathfrak{b}$, $W_{\vec{u}}^+ = W_{\vec{v}}^+$, and we define $W_{\mathfrak{b}}^+$ for this common subspace of $\mathbf{H}$ and $W_{\mathfrak{b}}^-$ for its complement in $\mathbf{H}$ according to the root space decomposition, so that $\mathbf{H} = W_{\mathfrak{b}}^+ \oplus W_{\mathfrak{b}}^-$.
- Let $\pi_A \colon G \longrightarrow \mathbf{A}$ be the projection onto the $\mathbf{A}$ factor as $(\mathbf{x}, \mathbf{t}) \mapsto \mathbf{t}$.
- For a regular liner functional $\ell \in \mathbf{A}^*$ with unit norm, we define $\Pi_\ell \colon G \to W_\ell^- \rtimes \mathbb{R}\vec{v}_\ell$ as $(\mathbf{x}, \vec{t}) \mapsto ([\mathbf{x}]_{W_\ell^-}, \ell(\mathbf{t})\vec{v}_\ell)$. We extend this to unit form vector $\vec{v}$ by defining $\Pi_{\vec{v}}$ as $\Pi_{\ell_{\vec{v}}}$ whereby $\ell_{\vec{v}}$ is the dual of $\vec{v}$.

Box associated to a compact convex set For $\alpha \in \triangle$, let $b(r) \subset V_\alpha$ be maximal product of intervals of size $2r$ centered at the origin, ie $[-r, r]^{\dim(V_\alpha)}$. Let $\Omega \subset \mathbf{A}$ be a convex compact set with nonempty interior whose barycenter is the identity of $\mathbf{A}$. We define the *box associated to* Ω, $\mathbf{B}(\Omega)$, as the set $\prod_{j=1}^{\sharp} b(e^{\max(\alpha_j(\Omega))})\Omega$. Fix a net $\mathfrak{n}$. We write $\mathcal{P}(\Omega)$ for the set of $\mathfrak{n}$ points in $\mathbf{B}(\Omega)$, so it is a finite set.

2.2.3 Remark We will work with a representative of a quasi-isometry class defined for particular choices of nets. Since any two nets in a space are bounded distance apart, all the coarse arguments remains valid for that entire equivalence. Let $\widehat{p}\colon G \to X$ assigns $x \in G$, a closest net point. In this way we tend to think of a set $K \subset G$ as a subset of the net X via the identification of K and $\widehat{p}(K)$. In particular, all the objects in a box defined by bullets points above are finite sets to us, and we equip a box with a counting measure. By abuse of notation, we denote this counting measure by the absolute value sign $|\cdot|$.

The following lemma shows that G is amenable.

2.2.4 Lemma *Let $\Omega \subset \mathbf{A}$ be compact convex with nonempty interior centered at the origin. Then, $\mathbf{B}(r\Omega)$, $r \to \infty$ is a Fölner sequence. The volume ratio between $N_\epsilon(\partial(\mathbf{B}(r\Omega)))$ and $\mathbf{B}(r\Omega)$ is $O(\epsilon/\operatorname{diam}(\mathbf{B}(r\Omega)))$.*

Proof See [12, Lemma 2.2.7]. □

3 Shadows, slabs and coarsening

We recall the following from Peng [12].

3.0.5 Theorem *Let G, G' be nondegenerate, unimodular, split abelian-by-abelian Lie groups, and $\phi\colon G \to G'$ be a (κ, C) quasi-isometry. Given $0 < \theta, \delta < 1$, there are numbers L_0, $0 < \rho < 1$ with the following properties.*

If $\Omega \subset \mathbf{A}$ is a product of intervals of equal size at least L_0, then a tiling of $\mathbf{B}(\Omega)$ by isometric copies of $\mathbf{B}(\rho\Omega)$

$$\mathbf{B}(\Omega) = \bigsqcup_{j \in \mathbf{J}} \mathbf{B}(\omega_j) \sqcup \Upsilon$$

contains a subset $\mathbf{J}_0 \subset \mathbf{J}$ of relative measure at least $1-\theta$, such that for all $j \in \mathbf{J}_0$, there is a subset $\mathcal{P}_0(\omega_j) \subset \mathcal{P}(\omega_j)$ with relative measure at least $1-\theta$ such that $\phi|_{\mathcal{P}_0(\omega_j)}$ is within $O(\delta \operatorname{diam}(\omega_j))$ of a standard map $g_j \times f_j$.

In this section, we focus on a particular standard map $g_i \times f_i$ supported on the subset U_i of a good box $\mathbf{B}(\omega_i)$, $i \in \mathbf{I}_0$. We will first show that the f_i is affine for all ranks. Then we will interpret its the translational and linear parts: the linear part has to come from a finite set related to the geometry of G', and the translational part depends on measure of certain subsets in $\mathbf{B}(\omega_i)$. We will drop the subscript i from now on.

3.1 Definitions

In this section we define a list of objects that will be used for the remainder of this section.

Root class half planes A set of the form $p\,\alpha_0^{-1}[-\infty, c]$ (resp. $p\,\alpha_0^{-1}[c,\infty]$), where $p \in \mathbf{H}$, $c \in \mathbb{R}$, is called a $[\alpha]$ negative (resp. positive) half plane. We write $\mathcal{H}^-_{[\alpha]}$ (resp. $\mathcal{H}^+_{[\alpha]}$) for the set of $[\alpha]$ negative (resp. positive) half planes. When we refer to a $[\alpha]$ half plane in a bounded set, we mean $p\,\alpha_0^{-1}([c,d])$, for some $p \in \mathbf{H}$, $c, d \in \mathbb{R}$. We will also say that the length of this $[\alpha]$ half plane is $|c-d|$ (remember here that the domain of a root is $\mathbf{A}$, not the entire group G).

Upper root boundary We define the *upper boundary of root class* $[\alpha]$, $\partial^+_{[\alpha]}$, as the quotient of $\mathcal{H}^+_{[\alpha]}$ under the equivalence relation of bounded Hausdorff distance.

If two positive $[\alpha]$ half-planes E_p, E_q where $p, q \in \mathbf{H}$, are bounded Hausdorff distance apart, then p, q can only differ by $V_{[\alpha]}$ coordinates. This means each equivalence class can be identified with $V_{[\alpha]}$, and the collection of all equivalence classes, $\partial^+_{[\alpha]}$, can be identified with $\bigoplus_{[\beta]\neq[\alpha]} V_{[\beta]}$.

Lower root boundary We say two $[\alpha]$ negative half planes H_p, H_q are equivalent if there is a sequence $H_i \in \mathcal{H}^-_{[\alpha]}$ such that $H_0 = H_p$, $H_q = H_n$ and any two successive H_i's intersect at an unbounded convex set. This is an equivalence relation because if H_p is equivalent to H_q, and H_q is equivalent to H_r, then concatenation of the sequences used to connect the two pairs is a sequence that connects H_p and H_r. We define the *lower boundary of* $[\alpha]$ $\partial^-_{[\alpha]}$ as the quotient of $\mathcal{H}^-_{[\alpha]}$ under this equivalence relation.

We see that if H_p, $H_q \in \mathcal{H}^-_{[\alpha]}$ based at $p, q \in \mathbf{H}$ have nonempty intersection, then p and q cannot differ by $V_{[\alpha]}$ coordinate. On the other hand, if p and q differ only in some $V_{[\beta]}$ coordinate, where $[\beta] \neq [\alpha]$, then $H_p \cap H_q \neq \varnothing$, so H_p is equivalent to H_q in this case. This way, we see that the equivalence class containing H_p, $p \in \mathbf{H}$ are all those $H_q \in \mathcal{H}_{[\alpha]}$, where $q \in \mathbf{H}$ differ from p by some elements of $\bigoplus_{[\beta]\neq[\alpha]} V_{[\beta]}$, and consequently, $\partial^-_{[\alpha]}$ an be identified with $V_{[\alpha]}$.

The shadow of a set on a lower root boundary If $p \in G$, we write $\pi^-_{[\alpha]}(p) \subset \partial^-_{[\alpha]}$ for the set of equivalence classes that contain minimal negative $[\alpha]$ half planes through points at most distance ρ away from p. Here ρ is the scale of discretization. Since $V_{[\alpha]}$ is the direct sum of V_Ξ, where $\Xi \in [\alpha]$, we will write $\pi^-_\sigma(p)$, where $\sigma \in [\alpha]$, for the projection of $\pi^-_{[\alpha]}(p)$ to V_σ. For $A \subset G$, we write $\pi^-_{[\alpha]}(A) = \bigcup_{p\in A} \pi^-_{[\alpha]}(p)$.

Measures on lower root boundaries Since $\partial^-_{[\alpha]}$ is a homogeneous space (the subgroup $V_{[\alpha]} \subset \mathbf{H}$ acts faithfully and transitively on it), it admits a Haar measure. We normalize this measure $|\,.\,|$ by requiring that for each $\sigma \in [\alpha]$,

$$|\pi^-_\sigma(p)|e^{-\sigma(p)} = 1 \quad \text{for all } p \in G \tag{3}$$

Half-planes of a linear functional Let ℓ be a linear functional that is not multiple of a root. A positive (resp. negative) half plane of ℓ refers to a set of the form $p\ell^{-1}[c, \infty)$ (resp. $p\ell^{-1}(-\infty, c]$), denoted by $\mathcal{H}^+_\ell$ (resp. $\mathcal{H}^-_\ell$).

Upper and lower boundaries of a linear functional We call the intersection of half planes corresponding to two perpendicular linear functionals, a quarter plane. We define an equivalence relation on $\mathcal{H}^+_\ell$ (resp. $\mathcal{H}^-_\ell$) as follows. Two positive (resp. negative) ℓ half planes H_p, H_q are equivalent if there is a sequence of $H_i \in \mathcal{H}^+_\ell$ such that $H_0 = H_p$, $H_n = H_q$ and the intersection between any two successive H_i's does not contain a quarter plane. This is an equivalence relation because if H_p is equivalent to H_q, and H_q is equivalent to H_r, then the concatenation of the sequences used to connect the two pairs is a sequence that establishes equivalence between H_p and H_r. We see that if the positive (resp. negative) ℓ half planes H_p, H_q based at $p, q \in \mathbf{H}$ are equivalent, then p, q differ by an element of W^-_ℓ (resp. W^+_ℓ).

We define the *upper boundary of* ℓ, ∂^+_ℓ, (resp. *lower boundary of* ℓ, ∂^-_ℓ) as the quotient of $\mathcal{H}^+_\ell$ (resp. $\mathcal{H}^-_\ell$) under this equivalence relation. In light of the forgoing discussion, we see that ∂^+_ℓ (resp. ∂^-_ℓ) can be identified with W^+_ℓ (resp. W^-_ℓ).

Measure on upper and lower boundaries of a linear functional Let ℓ be a generic linear functional. Since ∂^+_ℓ can be identified with W^+_ℓ which itself is a direct sum of $V_{[\Xi]}$ where $\ell(\vec{v}_\xi) > 0$, and each of $V_{[\Xi]}$ can be identified with $\partial_{[\Xi]}$, we can identify ∂^+_ℓ with $\prod_{[\Xi]:\ell(\vec{v}_\Xi)>0} \partial_{[\Xi]}$, and equip it with the product measures on the constituent root boundaries. The same procedure can be applied to ∂^-_ℓ to turn into a measure space.

The shadow of a set on the upper/lower boundaries of a linear functional If $p \in G$, we write $\pi^-_\ell(p) \subset \partial^-_\ell$ (resp. $\pi^+_\ell(p) \subset \partial^+_\ell$) for the set of equivalence classes that contain minimal negative (resp. positive) ℓ half planes through points at most distance ρ away from p, where ρ is the scale of discretization. For $A \subset G$ and $* \in \{+, -\}$, $\pi^*_\ell(A)$ is the union of $\pi^+_\ell(p)$, where p ranges over all points of A.

Branching constant The branching constant $\mathbf{b}_{[\alpha]}$ of root class $[\alpha]$ is the number such that $e^{\mathbf{b}_{[\alpha]}L}$ represents the number of negative $[\alpha]$ half planes of length L leaving a point. It equals $l_{[\alpha]}/\alpha_0$. Note that this is a number because α_0 divides every root in $[\alpha]$.

The branching constant $\mathbf{b}_\ell$, of a generic linear functional ℓ, is a number such that $e^{\mathbf{b}_\ell L}$ represents the number of ℓ half planes of length L leaving a point. Its value is given by

$$\mathbf{b}_\ell = \sum_{\sigma\,:\,\sigma(\vec{v}_\ell)>0} \sigma(\vec{v}_\ell) = \sum_{\sigma\,:\,\sigma(\vec{v}_\ell)<0} \sigma(\vec{v}_\ell). \tag{4}$$

We will now redefine $\mathbf{b}_\ell$ to take into account of the case when the linear functional ℓ might not have norm 1. So if $\widetilde{\ell}$ is a linear functional whose norm is not 1, we will write $\mathbf{b}_{\widetilde{\ell}}$ for $\mathbf{b}_{\widetilde{\ell}/\|\widetilde{\ell}\|}$.

Distance functions on lower root boundaries Given $p, q \in \partial^-_{[\alpha]} \sim V_{[\alpha]}$, let $t_{p,q}$ be the minimal $t \in \mathbb{R}$ such that there exists negative $[\alpha]$ half planes in the equivalence class of p and q that are distance 1 (or ρ if the scale of discretization is not 1) at sets whose π_A projection is $\ell^{-1}_{[\alpha]}(t)$.

Fix a positive number c, we define a *pseudodistance* $D_{[\alpha]}$ between p, q as follows.

$$D_{[\alpha]}(p,q) = e^{ct_{p,q}}$$

This can then be made into a metric by the usual procedure of defining the distance between two points as the infimum of $\sum_i D_{[\alpha]}(p_i, p_{i+1})$, taken over all finite chain of points $\{p_i\}$ connecting those two points. In this way, the space $(\partial^-_{[\alpha]}, D_{[\alpha]})$ becomes those whose quasi-conformal maps are studied by Dymarz [2] and Dymarz and Peng [3].

Shadows, slabs Using the same root class as before, we define a $[\alpha]$ *block* to be any left translate of $\bigoplus_{[\beta]\neq[\alpha]} V_{[\beta]} \rtimes \ker(\alpha_0)$. Given a $[\alpha]$ block H, the restriction of α_0 to H is a constant, which we denote by $\alpha_0(H)$.

For $\rho > 1$, we define the *ρ–shadow of H*, $\mathrm{Sh}(H,\rho)$, as the union of minimal negative $[\alpha]$ half planes containing a point in $N_\rho H$. By minimal we mean with respect to the order induced by coarse inclusion, where a set A is said to be *coarse included* in a set B if A is within $O(1)$ Hausdorff neighborhoods of B. For $h_2 < h_1 < \alpha_0(H)$, we define a *slab of H*, denoted by $\mathrm{Sl}^1_2(H)$ as the intersection between $\mathrm{Sh}(H,\rho)$ with the part of $\alpha_0^{-1}([h_2,h_1])$ in the box, ie it is the subset of $\mathrm{Sh}(H,\rho)$ in the box whose α_0 value lies in between h_2 and h_1.

Generalized slabs For $E_- \subset \partial^-_{[\alpha]}$, $E^+ \subset \partial^+_{[\alpha]}$, compact $K \subset \ker(\alpha_0)$, $h_2 < h_1$, we call a set

$$\begin{aligned} &S(E_-, E^+, K, h_2, h_1) \\ &\qquad = \left\{(\mathbf{x}_{[\alpha]}, (\mathbf{x}_{[\beta]})_{[\beta]\neq[\alpha]}, \mathbf{t}) \,:\, \mathbf{t} \in [h_2,h_1]K, \mathbf{x}_{[\alpha]} \in E_-, (\mathbf{x}_{[\beta]})_{[\beta]\neq[\alpha]} \in E^+\right\} \end{aligned}$$

a *generalized* $[\alpha]$ *slab*. This generalizes the definition of slabs defined earlier.

Coarsening We define a *coarsening* process as follows. For $h \in \mathbb{R}$, $E^+ \subset \partial^+_{[\alpha]}$, *the coarsening of* E^+ *by* h, $\mathcal{C}_h(E^+)$ is a set of equivalence classes of positive $[\alpha]$ half planes such that an equivalence class $\mathcal{H}$ belongs to $\mathcal{C}_h(E^+)$ if $\mathcal{H}$ has a representative that intersects with a representative of an element of E^+, and that the π_A projection of the intersection is a subset of $\alpha_0^{-1}[h, \infty]$.

Similarly, for $E_- \subset \partial^-_{[\alpha]}$ a subset of the lower boundary, the coarsening of E_- by h, $\mathcal{C}_h(E_-)$, is defined to be a set of equivalence classes of negative $[\alpha]$ half planes such that an equivalence class $\mathcal{H}$ belongs to $\mathcal{C}_h(E_+)$ if $\mathcal{H}$ has a representative that intersects with a representative of an element of E_-, and that the π_A projection of the intersection is a subset of $\alpha_0^{-1}[-\infty, h]$ (this is the same as the set of points in $\partial^-_{[\alpha]}$ that is distance e^h from a point in E_-).

Observe that as long as $h_3 \le h_2$, $h_4 \ge h_1$, we have

$$S(E_-, E^+, K, h_2, h_1) = S(\mathcal{C}_{h_3}(E_-), \mathcal{C}_{h_4}(E^+), K, h_2, h_1).$$

3.1.1 Lemma *The number of* $[\alpha]$ *planes in* $S = S(\mathcal{C}_{h_3}(E_-), \mathcal{C}_{h_4}(E^+), K, h_2, h_1)$ *is comparable to*

$$\frac{\mathrm{Vol}(S)}{|K|(h_1 - h_2)} e^{\mathbf{b}_{[\alpha]}(h_1 - h_2)}.$$

That is, it is compatible to the area of the cross-section times $e^{\mathbf{b}_{[\alpha]}(h_1-h_2)}$.

Proof The slab S is a product set. That is, it is the product of a subset in $\mathbf{H}$ with a subset in $\mathbf{A}$. Furthermore component in the $\mathbf{H}$ factor is product set $\Pi_{[\Xi]}\Omega_{[\Xi]}$ for subsets $\Omega_{[\Xi]}$ in root spaces $V_{[\Xi]}$; the component in the $\mathbf{A}$ factor is $K \times [h_2, h_1]$. Since the group is unimodular, the volume of S is $(\Pi_{[\Xi]}|\Omega_{[\Xi]}|)|K||h_1 - h_2|$.

The map $\pi_{[\alpha]}\colon G \to V_{[\alpha]} \ltimes (R)\vec{v}_{[\alpha]}$ (which is defined as $(\mathbf{x}_{[\alpha]}, \mathbf{x}_{[\beta]}, \dots)\mathbf{t} \mapsto (\mathbf{x}_{[\alpha]}, \alpha_0(\mathbf{t}))$) induces an onto map $H_{[\alpha]}$ between negative $[\alpha]$ half planes and negative half rays in the negatively curved space $V_{[\alpha]} \ltimes (R)\vec{v}_{[\alpha]}$. Therefore to count the number of negative $[\alpha]$ planes in a set Q, we count the number of negative half rays in $\pi_{[\alpha]}(Q)$, and multiply it by the multiplicity of $H_{[\alpha]}$ restricted to Q.

The number of negative half rays in $\pi_{[\alpha]}(S)$ is $|\Omega_{[\alpha]}|e^{\mathbf{b}_{[\alpha]}(h_1-h_2)}$, and the multiplicity of $H_{[\alpha]}|_S$ is $\Pi_{[\Xi]\neq[\alpha]}|\Omega_{[\Xi]}|$. □

3.2 Improving the almost product map

Let h_1, h_2 be functions on $[\alpha]$ blocks such that for any $[\alpha]$ block H,

$$h_2(H) < h_1(H) < \alpha_0(H).$$

Our setting from Theorem 3.0.5 says that the restriction of ϕ on $\mathbf{B}(\omega)$,

$$\phi\colon \mathbf{B}(\omega) \to G'$$

is a quasi-isometry such that on a subset $U_* \subset \mathbf{B}(\omega)$ of relative measure at least $1-\theta$, the restriction $\phi|_{U_*}$, is within $\epsilon\,\mathrm{diam}(\mathbf{B}(\omega))$ of a standard map $\widehat{\phi} = g \times f$. By the construction of U_*, there is a large subset of $[\alpha]$ planes in $\mathbf{B}(\omega)$ whose images under ϕ are within half planes corresponding to another root class that depends only on $[\alpha]$. Consequently, we obtain induced maps on $\partial^+_{[\alpha]}$ and $\partial^-_{[\alpha]}$ by examining those positive and negative $[\alpha]$ half planes that contain the one of those half planes whose images can be approximated by $\widehat{\phi}$. By abuse of notation we continue to denote those maps by g. Similarly, there is a large subset of left translates of $\mathbf{H}$ in $\mathbf{B}(\omega)$ whose images are close to another translates of $\mathbf{H}'$. This means that f is defined on a relative large set of ω. Since ω is a subset of a space with polynomial growth, we know the $O(\epsilon\,\mathrm{diam}(\omega))$ neighborhood of this subset contains ω, and we can extend f to all of ω by assigning $p \in \omega$ the f value of the point in the relative large subset closest to it.

Because ambient space has exponential growth and we only have measure control on the set U_*, describing the image of $\mathbf{B}(\omega)$ under ϕ is generally difficult because $\phi(\mathbf{B}(\omega))$ need not be sublinear away from $\phi(U_*)$. By restricting to certain subsets of $\mathbf{B}(\omega)$, it is possible to describe the image of those special types of subsets as demonstrated by the following two lemmas.

3.2.1 Lemma *Given $\epsilon < \beta \ll \beta' \ll 1$, there exist constants $c_1 \ll 1$ depending on θ, ϵ, $c_2 \ll 1$ depending on c_1, and a subset $E_{**} \subset \partial^-_{[\alpha]}(B)$ of relative measure at least $1-c_1$, such that whenever H is a $[\alpha]$ block at least $2\kappa\beta'\,\mathrm{diam}(\mathbf{B}(\omega))$ away from ∂B and $\pi^-_{[\alpha]}(H) \cap E_{**} \neq \varnothing$, then $|h_1(H) - h_2(H)||\pi_A(H)| \geq \beta|\omega|$ implies*

$$|\,\mathrm{Sl}^1_2(H) \cap U_*| \geq (1-c_2)|\,\mathrm{Sl}^1_2(H)|. \tag{5}$$

Proof Let c_2 be a constant to be chosen later. Let $E_1 \subset \partial^-_{[\alpha]}(B)$ be the subset such that for $x \in E_1$, there exists a $[\alpha]$ block H_x such that $x \in I_x = \pi^-_{[\alpha]}(H_x)$ and Equation (5) fails. Then we have a cover of E_1 by intervals I_x. By Vitali covering there is a subset of I_k's such that $\sum_k |I_k| \geq 1/5|E_1|$, and whose elements are strongly disjoint ie for $j \neq k$, $d(I_j, I_k) \geq 1/2\max(|I_j|)$, which means that the corresponding $[\alpha]$ block H_k's are also disjoint. By construction $|\,\mathrm{Sl}^1_2(H_k) \cap U^c_*| \geq c_2|\,\mathrm{Sl}^1_2(H_k)|$.

Summing over k yields

$$|B \cap U^c_*| \geq c_2 \sum_k |\,\mathrm{Sl}^1_2(H_k)| \geq \frac{c_2}{2} \sum_k |h_1(H_k) - h_2(H_k)||\pi_A(H_k)||I_k||\pi^+_{[\alpha]}(H_k)|.$$

As $|B \cap U_*^c| \le \theta|\omega||\pi^-_{[\alpha]}(B)||\pi^+_{[\alpha]}(B)|$ and $|\pi^+_{[\alpha]}(H_k)| = |\pi^+_{[\alpha]}(B)|$, we obtain

$$|E_1| \le 5\sum_k |I_k| \le \frac{10\theta}{\beta c_2}|\partial^-_{[\alpha]}(B)|.$$

Now choose c_2 appropriately so that $10\theta/(\beta c_2) \le c_1$. □

3.2.2 Definition Let H be a $[\alpha]$ block in $\mathbf{B}(\omega)$, and q is the map on the hyperplanes parallel to kernels of α_0 induced by f, where $\hat{\phi} = g \times f$ is the standard map defined on U_*. We set

$$\widetilde{\mathrm{Sl}}^1_2(H) := S\big(\mathcal{C}_{q(h_1)}(g(\pi^-_{[\alpha]}(H))), \mathcal{C}_{q(h_2)}(g(\pi^+_{[\alpha]}(H))), f(\pi_A(H) \times [h_2, h_1])\big)$$

3.2.3 Lemma *Given $\epsilon < \beta \ll \beta' \ll \beta'' \ll 1$, there exist constants c_3, c_4 depending θ and ϵ, and a subset $E_* \subset \pi^-_{[\alpha]}(B)$ of relative measure at least $1 - c_3$ with the following properties.*

Let H_0 be a $[\alpha]$ block in $\mathbf{B}(\omega)$ such that

(a) *the distance between H_0 and ∂B is at least $4\kappa\beta''\,\mathrm{diam}(\mathbf{B}(\omega))$;*

(b) *the intersection $\pi^-_{[\alpha]}(H_0) \cap E_*$ is not empty.*

Suppose H is a $[\alpha]$ block in

$$S = U_* \cap S\big(\pi^-_{[\alpha]}(H_0), \pi^+_{[\alpha]}(H_0), \pi_A(H_0), \ell_{[\alpha]}(H_0) - \beta''\,\mathrm{diam}(\mathbf{B}(\omega)), \ell_{[\alpha]}(H_0)\big)$$

such that $\pi^-_{[\alpha]}(H) \cap E \ne \varnothing$. Then, $\beta'|\omega| \ge |h_1(H) - h_2(H)||\pi_A(H)| \ge \beta|\omega|$ implies

$$|\,\widetilde{\mathrm{Sl}}^1_2(H) \cap N_{O(1)}\phi(U_* \cap \mathrm{Sl}^1_2(H))| \ge (1 - c_4)|\,\widetilde{\mathrm{Sl}}^1_2(H)|. \tag{6}$$

Here c_3, c_4 approach zero as ϵ and θ approach zero.

Proof Since we are interested only in the restriction to ϕ to $\mathbf{B}(\omega)$, first we need to make sure that $\widetilde{\mathrm{Sl}}^1_2$ lies in $\phi(\mathbf{B}(\omega))$. Recall that H_0 is more than $4\kappa^2\beta''\,\mathrm{diam}(\mathbf{B}(\omega))$ away from the boundary of $\mathbf{B}(\omega)$. This means that S is also more than $4\kappa^2\beta''\,\mathrm{diam}(\mathbf{B}(\omega))$ away from ∂B. By assumption, $|h_1(H) - h_2(H)| \le \beta'|\omega|/|\pi_A(H)|$, so $|h_1 - h_2| \le \beta'\,\mathrm{diam}(\mathbf{B}(\omega))$. Take $q_0 \in \widetilde{\mathrm{Sl}}^1_2(H)$. Then q_0 is no further than $\beta'\,\mathrm{diam}(\mathbf{B}(\omega))$ away from a point q in the generalized $[\alpha]$ slab $S(g(\pi^-_{[\alpha]}(H)), g(\pi^+_{[\alpha]}(H)), f(\pi_A(H) \times [h_2(H), h_1(H)]))$, and there is $p \in S \subset U_*$ such that $d(q, \hat{\phi}(p)) \le \beta'\,\mathrm{diam}(\mathbf{B}(\omega))$.

By definition, $\hat{\phi}(p)$ lies on a $[\Xi]$ half plane that is $\epsilon\,\mathrm{diam}(\mathbf{B}(\omega))$ away from the image of a $[\alpha]$ half-plane containing p, where $[\Xi]$ the image of $[\alpha]$ under the permutation on root spaces induced by the standard map $\hat{\phi}$. This means $d(q, \phi(S)) \le 2\beta'\,\mathrm{diam}(\mathbf{B}(\omega))$, and therefore $d(q_0,\phi(S)) \le 3\beta' \operatorname{diam}(\mathbf{B}(\omega))$. Since $d(S,\partial B) > 4\beta'' \operatorname{diam}(\mathbf{B}(\omega))$, it follows that $q_0 \in \phi(B)$.

Let c_3 be a constant to be chosen later. Let $E_2 \subset \partial^-_{[\alpha]}(B)\backslash E_1$ be such that for $x \in E_2$ there is a $[\alpha]$ block H_x such that $x \in I_x = \pi^-_{[\alpha]}(H_x)$ and Equation (6) fails. Thus we have a cover of E_2 by intervals I_x. By Vitali covering, we can find a subset of I_k's such that the inequality opposite to above holds for each H_k, such that $\sum_k |I_k| \ge (1/5)|E_2|$, and that I_k's are strongly disjoint. (That is, for $j \ne k$, $d(I_j, I_k) \ge 1/2 \max(|I_j|, |I_k|)$. In particular, this means $\widetilde{\mathrm{Sl}}^1_2(H_k)$'s are disjoint as well.

So we have for every H_k,

$$\big|\widetilde{\mathrm{Sl}}^1_2(H_k) \cap N_{O(1)}\big(\phi(U^c_* \cap \mathrm{Sl}^1_2(H_k))\big)\big| \ge c_4 |\,\widetilde{\mathrm{Sl}}^1_2(H_k)|.$$

Adding up all the H_k's and use the fact that the I_k's are strongly disjoint we have

$$|\phi(\mathbf{B}(\omega) \cap U^c_*)| \ge \sum_k \big|\widetilde{\mathrm{Sl}}^1_2(H_k) \cap N_{O(1)}\big(\phi(U^c_* \cap \mathrm{Sl}^1_2(H_k))\big)\big| \ge c_4 \sum_k |\,\widetilde{\mathrm{Sl}}^1_2(H_k)|. \tag{7}$$

We now proceed to show that to any $[\alpha]$ blocking set H satisfying the hypothesis, most measures in $\widetilde{\mathrm{Sl}}^1_2(H)$ comes from $\phi(\mathrm{Sl}^1_2(H) \cap U_*)$. More precisely, we show that

$$\phi\big(U_* \cap \big(N_{O(\epsilon \operatorname{diam}(\mathbf{B}(\omega)))}(\mathrm{Sl}^1_2(H))\big)^c\big) \cap \widetilde{\mathrm{Sl}}^1_2(H) = \varnothing. \tag{8}$$

Suppose this is not true. Then there is a $p \in U_*$ that lies outside of the $O(\epsilon \operatorname{diam}(\mathbf{B}(\omega)))$ neighborhood of $\mathrm{Sl}^1_2(H)$ such that $\phi(p)$ is in $\widetilde{\mathrm{Sl}}^1_2$. The fact $\phi(p)$ lies in $\widetilde{\mathrm{Sl}}^1_2(H)$ means that there is an element $q \in \mathrm{Sl}^1_2(H) \cap U_*$ such that the negative $[\Xi]$ half planes through $\phi(p)$ and $\hat{\phi}(q)$ meet at a height at most $q(h_1)$. Here $[\Xi]$ is the image of $[\alpha]$ under the permutation on root spaces induced by $\hat{\phi}$. So the negative $[\alpha]$ planes through q and p meets at a height no bigger than $h_1 + O(\epsilon \operatorname{diam}(\mathbf{B}(\omega)))$, but this contradicts the assumption that q lies outside of $O(\epsilon \operatorname{diam}(\mathbf{B}(\omega)))$ neighborhood of $\mathrm{Sl}^1_2(H)$.

Equation (8) implies that $|\,\widetilde{\mathrm{Sl}}^1_2(H)| \ge (1-c)|\,\mathrm{Sl}^1_2(H) \cap U_*|$ for some constant c that goes to zero as ϵ goes to zero. So now, picking up (7) we have

$$\begin{aligned}
|\phi(\mathbf{B}(\omega) \cap U^c_*)| &\ge c_4 \sum_k |\,\widetilde{\mathrm{Sl}}^1_2(H_k)| \\
&\ge c_4(1-c) \sum_k |\phi(\mathrm{Sl}^1_2(H) \cap U_*)| \\
&\ge c_4(1-c)(1-c_2) \sum_k |\,\mathrm{Sl}^1_2(H_k)|
\end{aligned}$$

$$\geq c_4(1-c)(1-c_2)\sum_k |\pi_A(H_k)||h_1(H_k)-h_2(H_k)||I_k||\pi^+_{[\alpha]}(H_k)|$$
$$\geq c_4(1-c)(1-c_2)\beta|\omega|\sum_k |I_k||\pi^+_{[\alpha]}(H_k)|.$$

Since $\pi^+_{[\alpha]}(H_k)=\pi^+_{[\alpha]}(B)$, and $|\phi(U^c_*)|\leq \theta|\omega||\pi^+_{[\alpha]}(B)||\pi^-_{[\alpha]}(B)|$, we obtain

$$|E_2|\leq 5\sum_k |I_k|\leq \frac{5\theta}{c_4\beta(1-c_2)(1-c)}|\pi^-_{[\alpha]}(B)|.$$

So now choose c_4 appropriately so that $5\theta/(c_4\beta(1-c_2)(1-c))\leq c_3$. □

3.2.4 Corollary *Let H be a $[\alpha]$ block in $\mathbf{B}(\omega)$ satisfying the hypothesis of Lemmas 3.2.1 and 3.2.3, and S as in Lemma 3.2.3. Suppose $w_1,w_2\in\mathbb{R}$ satisfy $\beta'|\omega|\geq |w_1-w_2||\pi_A(H)|\geq\beta|\omega|$. Then*

$$|\mathcal{C}_{w_1}(g(\pi^-_{[\alpha]}(H)))||\mathcal{C}_{w_2}(g(\pi^+_{[\alpha]}(H)\cap S))|\geq d|\pi^-_{[\alpha]}(H)||\pi^+_{[\alpha]}(H)|, \tag{9}$$
$$|\mathcal{C}_{w_1}(g(\pi^-_{[\alpha]}(H)))||\mathcal{C}_{w_2}(g(\pi^+_{[\alpha]}(H)\cap S))|\leq b|\pi^-_{[\alpha]}(H)||\pi^+_{[\alpha]}(H)|, \tag{10}$$

where d and b depend only on κ, C.

Proof Note that from the structure of U and the fact that ϕ is a quasi-isometry, it follows that for $z_1,z_2\in\pi_A(\mathbf{B}(\omega))$ we have

$$\frac{1}{2\kappa}|z_1-z_2|-\epsilon\operatorname{diam}(\mathbf{B}(\omega))\leq|q(z_1)-q(z_2)|\leq 2\kappa|z_1-z_2|+\epsilon\operatorname{diam}(\mathbf{B}(\Omega)).$$

This means q is essentially monotone, so there exist h_1,h_2 such that $q(h_1(H))=w_1$ and $q(h_2(H))=w_2$. We now apply Lemmas 3.2.1 and 3.2.3 to the resulting $\mathrm{Sl}^1_2(H)$ and $\widetilde{\mathrm{Sl}}^1_2(H)$. Lemma 3.2.1 gives

$$(1-c_2)|\mathrm{Sl}^1_2(H)|\leq|\mathrm{Sl}^1_2(H)\cap U_*|\leq 2\kappa|\widetilde{\mathrm{Sl}}^1_2(H)\cap\phi(U_*)|\leq 2\kappa|\widetilde{\mathrm{Sl}}^1_2(H)|.$$

On the other hand, by Equation (6) in Lemma 3.2.3, we know that $|\widetilde{\mathrm{Sl}}^1_2(H)\cap\phi(U_*)|\geq(1-c_4)|\widetilde{\mathrm{Sl}}^1_2(H)|$. The structure of a standard map means that the ratio of measures of $\widetilde{\mathrm{Sl}}^1_2(H)\cap\phi(U_*)$ to that of $\mathrm{Sl}^1_2(H)\cap U_*$ lies in $[1/2\kappa,2\kappa]$. These facts shows that

$$(1-c_4)|\widetilde{\mathrm{Sl}}^1_2(H)|\leq|\widetilde{\mathrm{Sl}}^1_2\cap\phi(U_*)|\leq 2\kappa|\mathrm{Sl}^1_2(H)\cap U_*|\leq 2\kappa|\mathrm{Sl}^1_2(H)|.$$

The claims now follow from the volume formula below.

$$|\mathrm{Sl}^1_2(H)|=|\pi^-_{[\alpha]}(H)||\pi^+_{[\alpha]}(H)||\pi_A(H)||h_2-h_1|,$$
$$|\widetilde{\mathrm{Sl}}^1_2(H)|=|\mathcal{C}_{q(h_1)}g(\pi^-_{[\alpha]}(H))||\mathcal{C}_{q(h_2)}g(\pi^+_{[\alpha]}(H))||f(\pi_A(H))||q(h_1)-q(h_2)|.$$ □

3.3 The constant part of f

Continuing with the same notation from the previous section, we show in this subsection, that $q\colon \mathbf{A}/\ker(\ell_{[\alpha]}) \to \mathbf{A}/\ker(\ell_{[\Xi]})$ is affine and compute its constant/translational term. To compute the constant term of this affine map, we make use of Corollary 3.2.4, and the property that the standard map $\widehat{\phi}$ roughly preserves the number of root class half planes.

3.3.1 Lemma *Let H be a $[\alpha]$ block in $\mathbf{B}(\omega)$. Let $\mathcal{F}$, $\widetilde{\mathcal{F}}$ denote the set of (maximal) $[\alpha]$, $[\Xi]$ half planes in $\mathrm{Sl}^1_2(H)$ and $\widetilde{\mathrm{Sl}}^1_2(H)$ respectively. Then*

$$\log|\mathcal{F}| = \log|\widetilde{\mathcal{F}}| + O(\epsilon\,\mathrm{diam}(B)).$$

Proof The claim is based on the fact that the map between $\mathrm{Sl}^1_2(H) \cap U_*$ and $\widetilde{\mathrm{Sl}}^1_2(H) \cap \phi(U_*)$ induces a map between the set of $[\alpha]$ half planes in $\mathrm{Sl}^1_2(H)$ and the set of $[\Xi]$ half planes up to an error of $e^{O(\epsilon\,\mathrm{diam}(B))}$. Explicitly, let $\mathcal{F}'$ be the set of $[\alpha]$ half planes in $\mathrm{Sl}^1_2(H)$ that are more than $O(\epsilon\,\mathrm{diam}(B))$ away from ∂B, and spend at least $1-\sqrt{c_2}$ fraction of their measure in U_*. Then $\mathcal{F}'$ has a relative large measure in $\mathcal{F}$ by Chebyshev. Now, for each $\gamma \in \mathcal{F}'$ there is a $[\Xi]$ half plane $\widehat{\gamma} \in \widehat{\mathcal{F}}$ such that $\phi(\gamma \cap U_*)$ is within $\epsilon\,\mathrm{diam}(B)$ of $\widehat{\gamma}$. We define $\psi(\gamma) = \widehat{\gamma}$. Note ψ is at most $e^{\epsilon\,\mathrm{diam}(B)+\sqrt{c_2}\,\mathrm{diam}(B)}$ to one since two $[\alpha]$ planes with the same $\widehat{\phi}$ image must be within $\epsilon\,\mathrm{diam}(B)$ of each other whenever they are in U_*. Inverse of ψ is defined similarly. □

3.3.2 Lemma *For all $h_1, h_2 \in [z_{\mathrm{bot}}, z_{\mathrm{top}}]$, where $z_{\mathrm{top}} = \max \alpha_0(B)$, $z_{\mathrm{bot}} = \min \alpha_0(B)$,*

$$q(h_1) - q(h_2) = \frac{\mathbf{b}_{[\alpha]}}{\mathbf{b}_{[\Xi]}}(h_1 - h_2) + O(\epsilon\,\mathrm{diam}(B)),$$

where $[\Xi] = f_*[\alpha]$.

Proof It is sufficient to check this for a $[\alpha]$ block H, and h_1, h_2 satisfying the hypothesis of Lemmas 3.2.1 and 3.2.3. Let $\mathcal{F}$, $\widetilde{\mathcal{F}}$ denote the set of $[\alpha]$, $[\Xi]$ half planes in $\mathrm{Sl}^1_2(H)$ and $\widetilde{\mathrm{Sl}}^1_2(H)$ respectively. The number of $[\alpha]$ half planes in $\mathrm{Sl}^1_2(H)$ is

$$|\pi^-_{[\alpha]}(H)|\,|\pi^+_{[\alpha]}(H)|\,e^{\mathbf{b}_{[\alpha]}(h_1-h_2)},$$

while the number of $[\Xi]$ half planes is

$$|\mathcal{C}_{q(h_1)}g(\pi^-_{[\alpha]}(H))|\,|\mathcal{C}_{q(h_2)}g(\pi^+_{[\alpha]}(H))|\,e^{\mathbf{b}_{[\Xi]}(q(h_1)-q(h_2))}.$$

By Lemma 3.3.1,

$$|\mathcal{C}_{q(h_1)} g(\pi^-_{[\alpha]}(H))|\,|\mathcal{C}_{q(h_2)} g(\pi^+_{[\alpha]}(H))|\, e^{\mathbf{b}_{[\Xi]}(q(h_1)-q(h_2))} = |\pi^-_{[\alpha]}(H)|\,|\pi^+_{[\alpha]}(H)| e^{\mathbf{b}_{[\alpha]}(h_1-h_2)}\, e^{O(\epsilon\operatorname{diam}(B))}.$$

Simplifying,

$$q(h_1)-q(h_2) = \frac{\mathbf{b}_{[\alpha]}}{\mathbf{b}_{[\Xi]}}(h_1-h_2) + \frac{1}{\mathbf{b}_{[\Xi]}}\log\frac{|\pi^-_{[\alpha]}(H)||\pi^+_{[\alpha]}(H)|}{|\mathcal{C}_{q(h_1)} g(\pi^-_{[\alpha]}(H))||\mathcal{C}_{q(h_2)} g(\pi^+_{[\alpha]}(H))|} + O(\epsilon\operatorname{diam}(B)).$$

The claim now follows because by Corollary 3.2.4,

$$\frac{|\pi^-_{[\alpha]}(H)|\,|\pi^+_{[\alpha]}(H)|}{|\mathcal{C}_{q(h_1)} g(\pi^-_{[\alpha]}(H))|\,|\mathcal{C}_{q(h_2)} g(\pi^+_{[\alpha]}(H))|} = O(1).$$ □

3.3.3 Lemma *The $\mathbf{A}'$ part of a standard map is affine. Its natural action sends R_G to $R_{G'}$, hence can only take on one of finitely many possibilities.*

Proof Lemma 3.3.2 already shows this in the case that G has rank 1. We now consider G whose rank is at least 2. Every point of $\mathbf{A}$ and $\mathbf{A}'$ is uniquely determined by the intersection of $\operatorname{rank}(G)$ many translates of root kernels, so Lemma 3.3.2 shows that f is affine and takes foliations by root kernels of G to that of G'.

Let A denote for the linear part of f, and σ be the permutation that f induces on the root classes. The existence of a standard map $O(\epsilon\operatorname{diam}(B))$ away from a quasi-isometry means that for $\vec{u}$ ranges over a large subset of $\mathbb{S}^{n-1}$, $[\Xi]$ a root class, then $l_{[\Xi]}(\vec{u}) > 0$ if and only if $l_{\sigma([\Xi])}(A(\vec{u})) > 0$, and $l_{[\Xi]}(\vec{u}) = 0$ if and only if $l_{\sigma([\Xi])}(A(\vec{u})) = 0$. So $l_{\sigma([\Xi])} \circ A = c_{[\Xi]} l_{[\Xi]}$ for some $c_{[\Xi]} > 0$.

Since f is affine, the pushforward of $\ell \in \mathbf{A}^*$, $f_*(\ell) = \ell \circ A^{-1}$, is an element of $\mathbf{A}'^*$. Take ℓ a regular linear functional of unit norm, and H an $[\alpha]$ block. Then, inside of $\mathrm{Sl}^1_2(H)$, the number of maximal sets of the form $p\,\ell^{-1}[c,d]$, $p \in \mathbf{H}$, $d-c = L$ is

$$|\pi^-_{[\alpha]}(H)|\,|\pi^+_{[\alpha]}(H)| \exp\Bigg(L \sum_{[\Xi]:\,\Xi(\vec{v}_\ell)>0} l_{[\Xi]}(\vec{v}_\ell)\Bigg),$$

while the number of $f_*\ell$ half planes in $\widetilde{\mathrm{Sl}}^1_2(H)$ is

$$|\mathcal{C}_{q(h_1)} g(\pi^-_{[\alpha]}(H))|\,|\mathcal{C}_{q(h_2)} g(\pi^+_{[\alpha]}(H) \cap S)| \exp\Bigg(L \sum_{[\Xi]:\,\Xi(\vec{v}_\ell)>0} l_{\sigma([\Xi])}(A\vec{v}_\ell)\Bigg).$$

Simplifying using Lemma 3.3.1 and Corollary 3.2.4 yields

$$\sum_{[\Xi]:\,\Xi(\vec{v}_\ell)>0} l_{\sigma([\Xi])}(A\vec{v}_\ell) = \sum_{[\Xi]:\,\Xi(\vec{v}_\ell)>0} l_{[\Xi]}(\vec{v}_\ell) + O(\epsilon\,\mathrm{diam}(B)).$$

Since the map from S^{n-1} to disjoint union of root classes defined by sending $\vec{v}$ to $\{[\Xi] : \Xi(\vec{v}) > 0\} \sqcup \{[\beta] : \beta(\vec{v}) < 0\}$ is constant on chambers, and each chamber contains a basis of $\mathbf{A}$, we conclude that up to an error of $O(\epsilon\,\mathrm{diam}(\mathbf{B}(\omega)))$,

$$\sum_{[\Xi]:\,\Xi(\vec{v}_\ell)>0} l_{[\Xi]} = \sum_{[\Xi]:\,\Xi(\vec{v}_\ell)>0} l_{\sigma([\Xi])} \circ A = \sum_{[\Xi]:\,\Xi(\vec{v}_\ell)>0} c_{[\Xi]} l_{[\Xi]}.$$

In other words, we have two equations

$$\sum_{[\alpha]:\,\alpha(\vec{v})>0} (1 - c_{[\alpha]}) l_{[\alpha]} = 0,$$

$$\sum_{[\beta]:\,\beta(\vec{v})<0} (1 - c_{[\beta]}) l_{[\beta]} = 0,$$

therefore $\widetilde{R} = \{(1 - c_{[\alpha]}) l_{[\alpha]}, [\alpha]\text{an equivalence class of roots}\}$ is a finite set of linear functionals whose sum is zero and such that for any codimension 1 hyperplane, the sum of those elements in $\widetilde{R}$ lying entirely on a half plane is zero. Therefore $\widetilde{R}$ consists of zero linear functionals, so $c_{[\alpha]} = 1$ for all root equivalence classes $[\alpha]$. But this means $l_{\sigma([\alpha])} \circ A = l_{[\alpha]}$. So A is a linear map that sends R_G to $R_{G'}$. □

The next proposition gives an interpretation to the translational part of f.

3.3.4 Proposition *Let $\epsilon < \beta \ll \beta' \ll \beta'' \ll 1$ be as in Lemma 3.2.1. Take a generalized $[\alpha]$ slab $S(E_-, E^+, K, h_{\mathrm{bot}}, h_{\mathrm{top}})$ in $\mathbf{B}(\omega)$. Suppose $h_{\mathrm{bot}} < z_{\mathrm{bot}} < z_{\mathrm{top}} < h_{\mathrm{top}}$ with $4\beta(h_{\mathrm{top}} - h_{\mathrm{bot}}) \le (z_{\mathrm{top}} - z_{\mathrm{bot}}) \le \beta'(h_{\mathrm{top}} - h_{\mathrm{bot}})$, and $|h_{\mathrm{top}} - z_{\mathrm{top}}|, |z_{\mathrm{bot}} - h_{\mathrm{bot}}| > 4\kappa^2\beta''(h_{\mathrm{top}} - h_{\mathrm{bot}})$. Then for $z \in [z_{\mathrm{bot}}, z_{\mathrm{top}}]$,*

$$q(z) = \frac{\mathbf{b}_{[\alpha]}}{\mathbf{b}_{f_*[\alpha]}} z + \frac{1}{\mathbf{b}_{f_*[\alpha]}} \log \frac{|\mathcal{C}_{q(z_{\mathrm{top}})} g(\pi^-_{[\alpha]}(H))|}{|\pi^-_{[\alpha]}(H)|} + O(\epsilon\,\mathrm{diam}(B)), \tag{11}$$

where H is a $[\alpha]$ block such that $\alpha_0(H_{\mathrm{top}}) = z_{\mathrm{top}}$.

Proof We know from Lemma 3.3.2 that

$$q(z) = \frac{\mathbf{b}_{[\alpha]}}{\mathbf{b}_{f_*[\alpha]}}(z) + \Big(q(z_{\mathrm{top}}) - \frac{\mathbf{b}_{[\alpha]}}{\mathbf{b}_{f_*[\alpha]}}(z_{\mathrm{top}})\Big) + O(\epsilon\,\mathrm{diam}(B)).$$

We now find an alternative expression for the term in the parenthesis.

Let $H_{\mathrm{top}} \subset B$ be a $[\alpha]$ block such that $\alpha_0(H_{\mathrm{top}}) = z_{\mathrm{top}}$. Then, by (3) and Lemma 3.3.2,

$$\log \frac{|\pi_{[\alpha]}^{-}(H)| e^{-\mathbf{b}_{[\alpha]} z_{\mathrm{top}}}}{|\mathcal{C}_{q(z_{\mathrm{top}})} g(\pi_{[\alpha]}^{-}(H))| e^{-\mathbf{b}_{f_*[\alpha]} q(z_{\mathrm{top}})}} = O(\epsilon \operatorname{diam}(B)).$$

Simplifying gives

$$\frac{1}{\mathbf{b}_{f_*[\alpha]}} \log \frac{|\pi_{[\alpha]}^{-}(H)|}{|\mathcal{C}_{q(z_{\mathrm{top}})} g(\pi_{[\alpha]}^{-}(H))|} = \frac{\mathbf{b}_{[\alpha]}}{\mathbf{b}_{f_*[\alpha]}} z_{\mathrm{top}} - q(z_{\mathrm{top}}) + O(\epsilon \operatorname{diam}(B)). \qquad \square$$

3.3.5 Corollary *There is a linear map $T\colon \mathbf{A} \to \mathbf{A}'$ and a vector $\vec{c} \in \mathbf{A}'$ such that for $z \in \mathbf{A}$,*

$$f(z) = T(z) + \vec{c} + O(\epsilon \operatorname{diam}(B)). \tag{12}$$

Proof By Proposition 3.3.4 and Lemma 3.3.3 we have

$$\log \frac{|\mathcal{C}_{q(h_1)} g(\pi_{[\alpha]}^{-}(H))|}{|\pi_{[\alpha]}^{-}(H)|} = f_*\alpha_0(z) - \alpha_0(z) + O(\epsilon \operatorname{diam}(B)), \tag{13}$$

for every root class $[\alpha]$, and $z \in \mathbf{A}$. Since $|\pi_{[\alpha]}(H)| = e^{\alpha_0(H)}$ and $|\mathcal{C}_{q(h_1)} g(\pi_{[\alpha]}^{-}(H))| = e^{f_*\alpha_0(f \times g(H))}$, where $f \times g(H)$ is a $f_*[\alpha]$ block. Since the left hand side of (13) is linear in α_0 it follows that the constant terms in (11) in different one dimensional directions are compatible, and this implies the existence of $\vec{c} \in \mathbf{A}'$ that gives rise to (11) along each root class. $\square$

3.3.6 Remark In view of Corollary 3.3.5, by modifying f by $O(\epsilon \operatorname{diam}(B))$, we can assume that Equations (12) and (11) holding without error terms.

4 Aligning the linear part of standard maps

Again, we refer to Theorem 3.0.5. By Lemma 3.3.3, the linear part of f_i's that appeared in the conclusion of Theorem 3.0.5 sends R_G of G, to $R_{G'}$ of G'. A priori, the linear parts of the f_i's do not have to be the same. For a generic G and G', the group of permutations between R_G and $R_{G'}$ is trivial, in which case the linear parts of f_i are all the same.

In this section, we show that this is true in general. That is, when rank of G is 2 or higher, the linear parts of the f_i's have to be the same. When G is rank 1, there are two root classes and the argument follows exactly the same proof as in Eskin, Fisher and Whyte [5], with the modification of replacing x and y horocycles by left translates of horocycles corresponding to those two root classes.

We aim to prove the following by the end of this section.

4.0.7 Theorem *Let G, G' be nondegenerate, unimodular, split abelian-by-abelian Lie groups. Let $\phi\colon G \to G'$ be a (κ, C) quasi-isometry. Given $0 < \epsilon, \theta < 1$, there is a number L_0 such that if $\Sigma \subset \mathbf{A}'$ is a product of intervals of equal size at least L_0, then there is a subset $U \subset \phi^{-1}(\mathbf{B}(\Sigma))$ of relative measure at least $1-\theta$ such that for some permutation σ between root classes of G and root classes of G', any root class horocycle $p\,V_{[\alpha]} \subset G$ that spends at least $\sqrt{\theta}$ proportion of its measure in U satisfies*

$$d(\phi(pV_{[\alpha]} \cap U), p'V_{\sigma([\alpha])} \cap \phi(U)) = O(\epsilon\,\mathrm{diam}(\Sigma)).$$

Choose $\delta, \eta \ll 1$. By Lemma 2.2.4, G' is amenable and boxes have small boundary compared to its volume, therefore the same is true of its image under ϕ^{-1}. So we can take a large box $\mathbf{B}$ in G', tile $\phi^{-1}(\mathbf{B})$ by large boxes and apply Theorem 3.0.5 to each of the tiling boxes in $\phi^{-1}(\mathbf{B})$ to obtain a further tiling by smaller boxes, where most of them support standard maps. In other words, we have a tiling of $\phi^{-1}(\mathbf{B})$ by smaller boxes $\mathbf{B}_i = \mathbf{B}(\Omega_i)$, where $\Omega_i \subset \mathbf{A}$ is compact convex, and all of Ω_i's are translates of each other.

$$\mathbf{B} = \bigsqcup_{i\in\mathbf{I}} \phi(\mathbf{B}_i) \sqcup \Upsilon \tag{14}$$

where $|\Upsilon| \le O(1/\mathrm{diam}(\mathbf{B}))|\mathbf{B}|$, and there is a subset $\mathbf{I}_0 \subset \mathbf{I}$ of relative measure at least $1-\theta$ so that for each $i \in \mathbf{I}_0$, there is a subset $U_i \subset \mathbf{B}_i$, also of relative measure at least $1-\theta$, such that $\phi|_{U_i}$ is $O(\epsilon\,\mathrm{diam}(\mathbf{B}_i))$ away from a standard map $g_i \times f_i$, where f_i is affine. Furthermore, the linear part of f_i sends R_G to $R_{G'}$, and is of finite order.

Since each U_i has large measure, $U_* = \bigcup_{i\in\mathbf{I}_0} U_i$ has relative measure at least $1-O(\theta)$ in $\phi^{-1}(\mathbf{B})$. For a subset $V \subset G'$, we write $\mathbf{I}(V)$ for those $i \in \mathbf{I}$ such that $\mathbf{B}_i \cap V \neq \varnothing$.

Theorem 4.0.7 finishes the alignment step because of the following consequences.

4.0.8 Corollary *In the conclusion of Theorem 3.0.5, the linear part of f_i's, $i \in \mathbf{I}_0$, are all the same.*

Proof By Lemma 3.3.3 we know that linear part can only be one of finitely many possibilities. Furthermore, they are uniquely determined by the permutations they induce on the root classes, which is also reflected by the permutations on the different root class horocycles. Therefore Theorem 4.0.7 implies that this permutation the same for all good boxes, and the claim follows. □

4.0.9 Corollary *Given $0 < \epsilon, \theta < 1$, there is a number L_0 such that if $\Omega \subset \mathbf{A}$ is a product of intervals of equal size at least L_0, then there is a subset $U \subset \phi^{-1}(\mathbf{B}(\Omega))$ of relative measure at least $1-\theta$, and a standard map $\widehat{\phi} = g \times f$ where f is affine defined on it such that*

$$d(\phi|_U, \widehat{\phi}) = O(\epsilon \operatorname{diam}(\mathbf{B}(\Omega))).$$

Proof Repeat the proof of Theorem 3.0.5 to $\mathbf{B}(\Omega)$. □

4.1 S graph, $\widehat{S}$ graph and the H_ℓ graphs

We continue with the setting from (14) and discretize $\mathbf{B}$ in this section so that it reflects the structures of the standard maps $\widehat{\phi}_i$, for $i \in \mathbf{I}_0$. Recall that each box $\mathbf{B}_i$ tiling the set $\phi^{-1}(\mathbf{B})$ is isometric to $\mathbf{B}(\omega)$, the box associated to some convex compact set $\omega \subset \mathbf{A}$. Let $\rho_i, i = 1, 2, 3, 4, 5$, be numbers such that $\rho_1 \gg C$, and $\rho_i \ll \rho_{i+1} \ll \epsilon \operatorname{diam}(\omega)$.

The S graph Take a ρ_1 net in $\mathbf{B}(R)$ and connect two net points by an edge if 1) their π_A images are no further than ρ_1 from each other and 2) that they are no further than $10\rho_1$ apart. We metrize this graph by letting lengths of edges be the distance between the corresponding points in G', so all edges have length $O(\rho_1)$. We refer to the discretization restricted to $\mathbf{B}$ as the S graph.

The sets U' and U Let U' be the union of the U_i's, for $i \in \mathbf{I}_g$ as in (14). We will now produce another set U with the property that $U \subset N_{\rho_5}(U')$. First let $U'' = N_{\rho_1}(U')$, and then let $U^c = N_{\rho_5+\rho_1}(U''^c)$.

The sets $\widetilde{U}$ and $\widetilde{U}^0$ We define $\widetilde{U}, \widetilde{U}^0 \subset U'$ using the following criteria. Let $p \in U'$, and $\widehat{\phi}_p = f_p \times g_p$ be local approximation of ϕ around p. We say $p \in \widetilde{U}$ (resp. $\widetilde{U}^0$) if for every chamber $\mathfrak{b}$, and its image chamber $(f_p)_*(\mathfrak{b})$, for at least some positive (small lower bound) measure of linear functionals ℓ with $\vec{v}_\ell \in \mathfrak{b}$, and positive measures of $\vec{u} \in (f_p)_*(\mathfrak{b})$,

- $p(\ker(\ell) \ltimes W_{\mathfrak{b}^+}) \cap \phi^{-1}(\mathbf{B})$ spends at least $1-\sqrt{\theta}$ proportion of its measure in U' (resp. U);
- the rank 1 subspace $\phi(p)G'_{\vec{u}}$ in $\mathbf{B}$ spends at least $1-\sqrt{\theta}$ proportion of its measure in $\phi(U')$ (resp. $\phi(U)$).

That $\widetilde{U}$ and $\widetilde{U}^0$ exist and of relative large measure follow from the same properties of U' and U.

Favourable chamber horocycle $W_{\mathfrak{b}}^{\pm}$ We say a chamber horocycle $W_{\mathfrak{b}}^{+}$ or $W_{\mathfrak{b}}^{-}$ is *favourable* (resp. *very favourable*) if it has nonempty intersection with $\widetilde{U}$ (resp. $\widetilde{U}^0$).

From now on we fix a favourable chamber horocycle $H_0 = qW_{\mathfrak{b}}^{+}$ and a very favourable chamber horocycle $H = pW_{\mathfrak{b}}^{+}$ distance $3\rho_5/5$ away, such that for a positive measure of linear functional ℓ in $\mathfrak{b}$, $\ker(\ell) \ltimes H$ lies in the shadow of $\ker(\ell) \ltimes H_0$. Let $\mathcal{S}_0 \subset \mathfrak{b}$ be this subset of positive measure. Also fix constants α, β, β', β'' so that $\epsilon \ll \alpha \ll \beta \ll \beta' \ll \beta'' \ll 1$. Let $h_1 = \ell(H) - (\alpha+\beta)R$ and $h_2 = \ell(H) - (\alpha+\beta+\beta'/2)R$.

The sets $\mathbf{I}(H)$ and $\mathbf{I}_g(H)$ Let $\mathbf{I}(H) \subset \mathbf{I}$ be the index of those B_i that have nonempty intersection with H. Inside of $\mathbf{I}(H)$, we define $\mathbf{I}_g(W)$ to be those i such that $|H \cap U' \cap B_i| \geq (1-\theta^{1/3})|H \cap B_i| > 0$.

The sets $W(H)_\ell$ and $\widehat{W}(H)_\ell$ For each $\ell \in \mathcal{S}_0$, we will construct $W(H)_\ell$ as the union of $W(H)^{i,\ell}$ for i ranges over $\mathbf{I}(H)$.

- First, suppose $i \in \mathbf{I}_g(H)$. Since H is very favourable, it follows that the set $H_\ell := \ker(\ell) \ltimes H$ has the property that

  $$|\operatorname{Sl}_2^1(H_\ell) \cap U_i| \geq (1-c_2)|\operatorname{Sl}_2^1(H_\ell)|.$$

  For each j between h_2 and h_1, if we write $\rho(j)$ for the relative proportion of $\operatorname{Sl}_2^1(H_\ell) \cap \ell^{-1}(j) \cap U_i^c$ in $\operatorname{Sl}_2^1(H_\ell) \cap \ell^{-1}(j)$, the above condition means that

  $$\sum_{j=h_2}^{h_1} \rho(j) \leq 2c_2$$

  which means that for some height $h_\ell^i \in [h_2, h_1]$, $\rho(h_\ell^i) \leq 2\sqrt{c_2}$. Let

  $$W(H)_{i,\ell} = \operatorname{Sh}(\ker(\ell) \ltimes H_0, \rho_1) \cap \ell^{-1}(h_\ell^i) \cap B_i.$$

- Now for $i \in \mathbf{I}(H) \backslash \mathbf{I}_g(H)$, we define

  $$W(H)_{i,\ell} = \operatorname{Sh}(\ker(\ell) \ltimes H_0, \rho_1) \cap \ell^{-1}(h_1) \cap B_i.$$

The set $\widehat{W}(H)_\ell$ will be constructed as a union of $\widehat{W}(H)_{i,\ell}$'s for i ranges over $\mathbf{I}_g$. Define

$$\widehat{W}(H)_{i,\ell} := S(\pi_-(H_\ell), \pi_+(H_\ell), \pi_A(H_\ell), (f_i)_*(\ell)(h_\ell^i)).$$

ℓ–Shadow vertices For each $\ell \in \mathcal{S}_0$, the set of ℓ–shadow vertices will be defined as the union of good ℓ–shadow vertices and bad ℓ–shadow vertices. The bad ℓ–shadow vertices are define to be the S vertices that lie in ρ_1 neighborhoods of $\phi(U'^c \cap W(H)_\ell)$, and those in $\widehat{W}(H)_\ell$ that lie within $10\epsilon R$ of boundary of $\mathbf{B}(R)$, as well as those that are more than $\beta'' R$ away from $\phi(U')$. The S vertices in $\widehat{W}(H)_\ell$ that are not bad shadow vertices are defined to be good ℓ–shadow vertices.

4.1.1 Lemma *For each $\ell \in \mathcal{S}_0$, there is a constant c_5 depending ϵ, θ such that the proportion of bad shadow vertices is at most c_5, and c_5 approaches zero with ϵ, θ.*

Proof Bad shadow vertices are defined in two stages. First we have the set S_1 of vertices in $\widehat{W}(H_\ell)$ that are either within $10\kappa\beta''\operatorname{diam}(\omega)$ of $\partial\mathbf{B}$ or outside of $\beta'\operatorname{diam}(\omega)$ neighborhood of a point in $\phi(U_*)$ whose ℓ value is h_0^i smaller than $\ell(H_\ell)$. That this set has small measure in $\widehat{W}(H_\ell)$ follows from two facts. First, the subset that are close to $\partial\mathbf{B}$ has relative small measure by Lemma 2.2.4.

Second, if the proportion of S_1 in $\widehat{W}(H_\ell)$ is θ, then the set of points in $\widetilde{\mathrm{Sl}}^1_2(H_\ell) \cap \phi(U_*^c)$ contained in a ℓ half plane through a point of S_1 has measure at most θ relative to $\widetilde{\mathrm{Sl}}^1_2(H_\ell)$. However by Lemma 3.2.3, the proportion of $\widetilde{\mathrm{Sl}}^1_2(H_\ell) \cap \phi(U_*^c)$ in $\widetilde{\mathrm{Sl}}^1_2(H_\ell)$ is at most c_4. Therefore $\theta \le c_4$. In the second stage, we enlarge the set of bad vertices in $\widehat{W}(H)$ by adding the set $N_{\rho_1}\phi(U_*^c \cap \widehat{W}(H_\ell))$. That this set has small measure follows from our choice of h_0^i. □

The $\widehat{S}_\ell$—graph For each $\ell \in \mathcal{S}_0$, we modify the S–graph near $\phi(H_\ell)$ to produce a $\widehat{S}_\ell$ graph which reflects divergence property dictated by the standard maps.

For $x \in W_\ell^+$, $y \in W_\ell^-$ and $t \in \mathbb{R}$, we write $\gamma_{x,y}(t)$ for the preimage of $(x, y, t) \in \mathbf{H} \rtimes \mathbb{R}_{\vec{v}_\ell}$ under the projection $\pi_\ell\colon G \to \mathbf{H} \rtimes \mathbb{R}_{\vec{v}_\ell}$; similarly we write $\gamma_{x,y}([c,d])$ for the $\bigcup_{t\in[c,d]} \gamma_{x,y}(t)$, which is a ℓ half plane of length $|c-d|$.

For each $\ell \in \mathcal{S}_0$, $i \in \mathbf{I}_g(H)$, let

$$K_{\ell,i} = \bigcup_{x\in W_\ell^+, y\in W_\ell^-} \gamma_{x,y}([(f_i)_*\ell(H_0) - \rho_5/5, (f_i)_*\ell(h_\ell^i)]) \cap \widehat{W}_{\ell,i}.$$

We begin by replacing $K_{\ell,i}$ as a subset of the S graph by disjoint union of $\gamma_{x,y}$'s, then define the $\widehat{S}$ graph by declaring new sets of vertices and incidence relations. For each $t_j \in (j/\rho_1)(f_i)_*\ell(H_0) - \rho_5/5 - (f_i)_*\ell(h_\ell^i)$ call the S vertices of $\gamma_{x,y}(t_j) \in K_{\ell,i}$ *pre-vertices.*

In the range, we tile left cosets of $W^-_{(f_i)_*\ell}$ in $((f_i)_*\ell)^{-1}(t_j)$ by rectangles T_-'s of diameter $10\rho_1$; in the domain, we tile left cosets of W_ℓ^+ in $\ell^{-1}(q_i^{-1}(t_j))$ by rectangles T_+'s of diameter $10\kappa^2\rho_1$. We identify two vertices p, q if

(i) p, q are in the same T_-;

(ii) $\hat{\phi}_i^{-1}(p)$ and $\hat{\phi}_i^{-1}(q)$ are in the same T_+ which has the property that

$$|\partial_{\ell}^-(T_+) \cap \widetilde{E}_{**}| \geq 1/2|\partial_{\ell}^-(T_+)|,$$

where $\widetilde{E}_{**}$ is the union of E_{**} in each root lower boundaries that appear in ∂_{ℓ}^- from each $\mathbf{B}_i$, $i \in \mathbf{I}_0$ as given by Lemma 3.2.1.

We also remove any edges in K_i that ends at a bad shadow vertex. A $\widehat{S}_{\ell}$ vertex is called *irregular* if it arise from the procedure above. Otherwise it is called *regular*.

In our original S graph, every point has the same valence provided the vertex is not close to the boundary. However, upon the changes made for the $\widehat{S}_{\ell}$ graph, the homogeneity of valence is not so clear. That this change in valency is bounded is given by the lemma below.

4.1.2 Lemma *There are M_l, $M_u \in \mathbb{R}$ depending only on κ, C such that for any two $\widehat{S}$ vertices, the ratio of numbers of $f_*\ell$ half planes through them is bounded between M_l and M_u.*

Proof Let v be an irregular vertex, and let $z_{\text{top}} = \max(f_*\ell)(\mathbf{B})$. First, the number of $f_*\ell$ half planes of length $z_{\text{top}} - (f_*\ell)(v)$ containing v and $(f_*\ell)^{-1}(z_{\text{top}})$ is $e^{\mathbf{b}_{f_*\ell}(z_{\text{top}} - f_*\ell(v))}$.

As v is an irregular vertex, there exists a ℓ block H' in $\mathbf{B}_i$ such that $v \in \phi_i(H')$ and $\partial_{\ell}^-(H')$ contains a point of E_{**}. The number of $f_*\ell$ half planes containing v and $\widehat{W}_i(H_\ell)$ (whose $f_*\ell$ and q_i value is h_ℓ^i) is

$$\begin{aligned}
&\approx |\mathcal{C}_{q_i(h_\ell^i)} g_i(\pi_\ell^-(H'))| e^{-\mathbf{b}_{f_*\ell} h_\ell^i} && \text{by (3) and (4)} \\
&\approx |\mathcal{C}_{q_i(h_\ell^i)} g_i(\pi_\ell^-(H'))| \frac{|\mathcal{C}_{q_i(h_2)} g(\pi_\ell^+(H') \cap S)|}{|\pi_\ell^+(H')|} e^{-\mathbf{b}_\ell q_i^{-1}(h_\ell^i)} && \text{by Remark 3.3.6} \\
&\approx |\pi_\ell^-(H')| e^{-\mathbf{b}_\ell q_i^{-1}(h_\ell^i)} && \text{by Corollary 3.2.4} \\
&\approx e^{\mathbf{b}_\ell \left(\ell(H') - q_i^{-1}(h_\ell^i)\right)} && \text{by (3) and (4)} \\
&\approx e^{\mathbf{b}_{f_*\ell} \left(q_i(\ell(H')) - h_\ell^i\right)} && \text{by Remark 3.3.6,}
\end{aligned}$$

where $\approx$ means up to a multiple constant that depends on the discretization. □

The H_ℓ graphs We now define the H_ℓ graph as a subgraph of the $\widehat{S}_\ell$ graph consisting of the irregular ℓ–vertices, as well as the bad ℓ–shadow vertices. The irregular ℓ–vertices are called good H_ℓ vertices and those bad ℓ–shadow vertices are called bad H_ℓ vertices. An edge of the H_ℓ graph is concatenation of edges in the $\widehat{S}_\ell$ graph, all of the same direction, that connects two H_ℓ vertices, or connect a good H_ℓ vertex with $\partial\mathbf{B}$. Because the valence of the $\widehat{S}_\ell$ graph is bounded, the same is true for the H_ℓ graph.

4.2 Averaging over the H_ℓ graph

For each $\ell \in \mathcal{S}_0$, let $\mathcal{V}_\ell$ and $\mathcal{E}_\ell$ denote for the sets of vertices and edges in the H_ℓ graph. Those in $\mathcal{E}_\ell$ with one vertex in $\partial\mathbf{B}$ is called an *leaf* edge. Because we would need to consider concatenations of edges that avoid the bad ℓ shadow vertices, we will need to extract subsets of $\mathcal{V}_\ell$ and $\mathcal{E}_\ell$ such that paths in the induced subgraph have the desired property, and we will achieve this by a series of averaging processes.

In the following, the δ_i's, ϵ_i's θ_i's and μ_i's are all numbers less than 1 and approach zero as ϵ and θ in Theorem 3.0.5 approach zero.

4.2.1 Definition

$\mathcal{V}_{1,\ell}$: The set of good vertices as in the definition of the H_ℓ graph.

$\mathcal{E}_{1,\ell}$: Either connects two vertices in $\mathcal{V}_1$ or is a leaf edge based on a vertex of $\mathcal{V}_1$.

$\mathcal{V}_{2,\ell}$: The subset of $\mathcal{V}_1$ where at least $1-\nu_2$ proportion of the edges are in $\mathcal{E}_1$.

$\mathcal{E}_{3,\ell}$: An $\mathcal{E}_1$ edge e such that for all $\widehat{S}$ vertices $x \in e$, $1-\theta_3$ fraction of the forward branching edges (ie edges in the same direction as e) branching at x are in $\mathcal{E}_1$.

$\mathcal{E}_{4,\ell}$: An $\mathcal{E}_3$ such that for any $\widehat{S}$ vertex $x \in e$, $1-\theta_4$ fraction of the reverse edges (ie edges in opposite direction as e) branching at x are in $\mathcal{E}_1$.

$\mathcal{V}_{3,\ell}$: The subset of $\mathcal{V}_2$ where at least $1-\nu_3$ proportion of edges belong to $\mathcal{E}_4$.

$\mathcal{V}_{4,\ell}$: The subset of $\mathcal{V}_3$ that is not a strange vertex (see Definition 4.2.2 below).

4.2.2 Definition An S–vertex w is said to be *marked* by a $\mathcal{V}_{1,\ell}$ vertex v if w comes within $O(\rho_2)$ of the cloud of v. A vertex $v \in \mathcal{V}_\ell$ is called *strange* if there is a nonsingular direction $\vec{u}$ such that at least $1-\nu_4$ proportion of points that can be reached from v along edges in direction $\vec{u}$ are marked by vertices in $\mathcal{V}_{1,\ell} \backslash \mathcal{V}_{3,\ell}$.

4.2.3 Lemma *At least $1-\delta_1$ fraction of $\mathcal{V}_\ell$ are in $\mathcal{V}_{1,\ell}$, and at least $1-\epsilon_1$ fraction of $\mathcal{E}_\ell$ are in $\mathcal{E}_{1,\ell}$.*

Proof The relative large fraction of $\mathcal{V}_{1,\ell}$ in $\mathcal{V}_\ell$ follows from the fact that H was chosen to be a very favourable chamber horocycle. To see the second claim, recall that Lemma 4.1.2 says that the valence of any S–vertex is bounded by two numbers M_u and M_l. This implies that

$$\frac{1}{M_u}\frac{|\mathcal{V}_\ell|}{|\mathcal{E}_\ell|} \le \frac{1}{M_l}.$$

Now since
$$|\mathcal{E}^c_{1,\ell}| \le 2M_u|\mathcal{V}^c_{1,\ell}|,$$
it follows that
$$\frac{|\mathcal{E}^c_{1,\ell}|}{|\mathcal{E}_\ell|} \le 2\frac{M_u}{M_l}\frac{|\mathcal{V}^c_{1,\ell}|}{|\mathcal{V}_\ell|},$$
so ϵ_1 goes to zero because $\epsilon_1 \le \delta_1$ and δ_1 goes to zero. □

4.2.4 Corollary *At least* $1-\theta_2$ *fraction of* $\mathcal{V}_\ell$ *are in* $\mathcal{V}_{2,\ell}$.

We now proceed to show the existence of those other edges and vertices.

4.2.5 Lemma *At least* $1-\epsilon_3$ *proportion of edges in the* H_ℓ *graph are in* $\mathcal{E}_{3,\ell}$.

Proof Suffice to show that for $v \in \mathcal{V}_{2,\ell}$, almost all edges leaving v belongs to $\mathcal{E}_{3,\ell}$. Take $v \in \mathcal{V}_{2,\ell}$ and let $\mathcal{E}(v)$ denotes all edges incident to v whose directions are nondegenerate, and so we can identify $\mathcal{E}(v)$ with $\prod_{\vec{v}\in\widetilde{S}} W^+_{\vec{v}}$, where $\widetilde{S}$ is the dense subset of nondegenerate vectors in the unit sphere of $\mathbf{A}'$. Note that under this identification $\mathcal{E}(v)$ is a subset of doubling space and Vitaly's covering lemma applies.

We know by definition of $\mathcal{V}_{2,\ell}$, most edges in $\mathcal{E}(v)$ belong to $\mathcal{E}_{2,\ell}$, ie

$$|\mathcal{E}^c_{2,\ell} \cap \mathcal{E}(v)| \le \delta_2|\mathcal{E}(v)|.$$

Let $A_v = \mathcal{E}(v) \cap \mathcal{E}^c_3$ denotes the subset of $\mathcal{E}_v$ not in $\mathcal{E}_3$. We know that for any $e \in A_v$, there is a $x \in e$ such that at least θ_3 of the edges (in the same direction as e) branching from e at x is not in $\mathcal{E}_1$, so there is a neighborhood U of e such that

$$|\mathcal{E}^c_2 \cap U \cap \mathcal{E}(v)| \ge \theta_3|U \cap \mathcal{E}(v)|.$$

We thus get a cover of A_v by the U's. By Vitaly covering lemma, there exists disjoint U_j's such that

$$\sum_j |U_j| \ge \frac{1}{2}|A_v|.$$

So $$|A_v| \le 2\sum|U_j| \le \frac{2}{\theta_3}\sum|U_j \cap \mathcal{E}^c_2 \cap \mathcal{E}(v)| \le \frac{2}{\theta_3}|\mathcal{E}^c_2 \cap \mathcal{E}(v)| \le \frac{2}{\theta_3}\delta_2|\mathcal{E}(v)|,$$

and we see that we can make A_v relatively small by choosing appropriate θ_3. □

4.2.6 Lemma *At least $1-\epsilon_4$ fraction of the edges in the H_ℓ graph are in $\mathcal{E}_{4,\ell}$.*

Proof We already know that the nonleaf edges have the desired reverse branching property. So suffice to assume that the proportion of leaf edges is at least $\epsilon_1^{1/6} = \alpha$.

Let Y be the set of paths (concatenation of edges in the same direction) joining $\partial\mathbf{B}$, and $Y' \subset Y$ be those consisting of paths which pass through a point not in $\mathcal{V}_1$. Let $D(\gamma) = 1$ if $\gamma \in Y'$ and $D(\gamma) = 0$ otherwise. By Lemma 4.2.3 and our assumption on the proportion of leaf edges, we have

$$\sum_{\gamma\in Y} D(\gamma) \le \epsilon_1 |\mathcal{E}_\ell| \le \frac{\epsilon_1}{\alpha}|\mathcal{E}_{\mathrm{leaf}}|.$$

For a point $v \in \partial\mathbf{B}$, let Y_v denote the set of geodesics emanating from v. We have

$$\sum_{v\in\partial\mathbf{B}} \sum_{\gamma\in Y_v} D(\gamma) \le \sum_{v\in\partial\mathbf{B}} \frac{\epsilon_1}{\alpha}|\mathcal{E}_{\mathrm{leaf}}(v)|,$$

where $\mathcal{E}_{\mathrm{leaf}}(v)$ denotes the set of leaf edges emanating from v. Let θ' to be chosen later, and define

$$P = \left\{ v \in \partial\mathbf{B} : \sum_{\gamma\in Y_v} D(\gamma) \ge \theta' |\mathcal{E}_{\mathrm{leaf}}(v)| \right\}.$$

By Chebyshev, the set P has relative small measure, so it suffices to show that for $v \notin P$, most of the leaf edges leaving v are in $\mathcal{E}_4$.

Take $v \in \partial\mathbf{B}\backslash P$, then

$$\sum_{\gamma\in Y_v} D(\gamma) \le \theta' |\mathcal{E}_{\mathrm{leaf}}(v)|$$

and let A_v be those leaf edges leaving v that are not in $\mathcal{E}_4$. As before, we identify $\mathcal{E}_{\mathrm{leaf}}(v)$ with $\prod_{\vec{u}} W^+\vec{u}$, where $\vec{u}$ ranges over nonsingular unit vectors. The definition of A_v means that for any $e \in A_v$, there is a $x \in e$ such that at least θ_4 fraction of the edges branching from e at x are not in $\mathcal{E}_2$. So then there is a neighborhood $U \subset Y_v$ with $e \in U$ such that $|\mathcal{E}_1^c \cap U| \ge \theta_4 |U|$, and so

$$\sum_{\gamma\in U} D(\gamma) = |\mathcal{E}_1^c \cap U| \ge \theta_4 |U|.$$

By extracting from the U's, a disjoint subcollection that covers at least half of the measure of A_v we have

$$|A_v| \le 2\sum_j |U_j| \le \frac{2}{\theta_4}\sum_j \sum_{\gamma\in U_j} D(\gamma) \le \frac{2}{\theta_4}\sum_{\gamma\in Y_v} D(\gamma) \le \frac{2\theta'}{\theta_4}|\mathcal{E}_{\mathrm{leaf}}(v)|.$$

Choose θ' appropriately. □

4.2.7 Lemma *At least $1-\delta_6$ fraction of $\mathcal{V}_\ell$ are in $\mathcal{V}_{4,\ell}$.*

Proof The claim follows from the observations that the number of H_ℓ vertices that can mark a given S–vertex is $O(\rho_1)$ (because they must come from good boxes), and that $\mathcal{V}_{1,\ell}\backslash\mathcal{V}_{3,\ell}$ has relative small measure in $\mathcal{V}_{1,\ell}$. □

So now to each H_ℓ graph, we have a relative large subset $\mathcal{V}_\ell^0 \subset \mathcal{V}_\ell$ of H_ℓ vertices for which most geodesics leaving them are very good in quality.

4.2.8 Lemma *There is an S–vertex $\phi(x)$, where $x \in U_*$ with the following properties. There is a relative large subset $\mathcal{S}_0 \subset \mathfrak{b}$ and a positive subset $\mathcal{S}_0'$ of $\mathfrak{b}'$, where $\mathfrak{b}'$ is the image of $\mathfrak{b}$ under the standard map supported around x, such that*

- *for every $\ell \in \mathcal{S}_0$, $\mathcal{V}_\ell^0 \cap B(\phi(x), \epsilon R) \neq \varnothing$;*
- *for every $p \in \mathcal{V}_\ell^0 \cap B(\phi(x), \epsilon R) \neq \varnothing$, $\vec{u} \in \mathcal{S}_0'$, most geodesics leaving p in direction $\vec{u}$ belong to $\mathcal{E}_4$.*

Proof To each V_ℓ we further break it up according to its **H** coordinate, so that $V_\ell = \bigcup_{x\in H} V_{x,\ell}$ where we recall that H is the very favourable chamber horocycle. Let χ be the characteristic function of the set $\bigcup_{\ell\in\mathcal{S}_0} \mathcal{V}_\ell\backslash\mathcal{V}_\ell^0$. Then we know

$$\frac{1}{|\mathcal{V}|}\sum_{x\in H}\sum_{p\in V_{x,\ell}} \chi(p) \leq \delta_6.$$

There is a subset $\widetilde{H} \subset H$ of (relative large) positive measure such that for every $x \in \widetilde{H}$, those in $\bigcup_\ell V_{x,\ell}$ that do not belong to $\bigcup_\ell \mathcal{V}_\ell^0$ have small relative measure. Fix such a $x \in \widetilde{H}$, and now look at all the $\mathcal{V}_{\ell,x}$ for ℓ's ranging over $\mathcal{S}_0$. Without loss of generality we can assume $\phi(x)$ to be an S–vertex. Let $f_x \times g_x$ be the standard map supported around x, and let $\mathfrak{b}' = f_*(\mathfrak{b})$ be the image chamber. Each of $\mathcal{V}_{\ell,x}^0$ are good for some length in the $\mathfrak{b}'$ directions. Note that the preimages of $V_{x,\ell}$ as ℓ ranges over $\mathcal{S}_0$ all lie on a common flat. This might not be true of the vertices $\bigcup_{\ell\in\mathcal{S}_0} V_{x,\ell}$ themselves. However if this is not the case, then we can push forward two parallel paths in the domain and if the images are located far too far away from where the flats coarsely intersect each other, then we will have a contradiction. So we conclude that out of $\bigcup_\ell V_{x,\ell}^0$ is a subset $\mathcal{U}_x^0$ of positive measure that all lie on a common flat as the S vertex $\phi(x)$. Now, to each element of $\mathcal{U}_x^0$ is a large subset of good directions in $\mathfrak{b}'$. Let $\widehat{\chi}$ denotes the characteristic function of pairs $(z, \vec{u})$ where z is an element of $\mathcal{V}_x^g = \bigcup_\ell \mathcal{U}_{x,\ell}^0$ and $\vec{u}$ is a direction such that most of the edges in that direction does not belong to $\mathcal{E}_4$. Then we know

$$\sum_{z\in\mathcal{V}_x^g}\sum_{\vec{u}\in\mathfrak{b}'}\widehat{\chi}\le\delta_7$$

and

$$\sum_{\vec{u}\in\mathfrak{b}'}\sum_{z\in\mathcal{V}_x^g}\widehat{\chi}=\sum_{z\in\mathcal{V}_x^g}\sum_{\vec{u}\in\mathfrak{b}'}\widehat{\chi}\le\delta_7.$$

So there is a subset of $\mathfrak{b}'$ of large proportion that are good for a large proportion of vertices in $\mathcal{V}_x^g$. □

4.3 Some geometric lemmas

An idea behind the proof of Theorem 4.0.7 is the following observation. In H_{n+1} (as in Lemma 2.1.1) suppose two travelers leave a common starting point via segments of diverging vertical geodesics. If they are to meet up again without having to travel for long, then they can only do so in some neighborhood of the starting point. In short, this section is the same as Section 5.3 in Eskin, Fisher and Whyte [5].

Recall that for three points x, y, z in a metric space, the Gromov product is defined as

$$(y|z)_x=\frac{1}{2}(d(y,x)+d(z,x)-d(y,z)).$$

In a δ hyperbolic space X, the geodesic γ_{yz} joining y to z satisfies

$$d(\gamma_{yz},x)-\delta\le(y|z)_x\le d(\gamma_{yz},x).$$

The projection Π_ℓ Let Π_ℓ: $G\to W_\ell^-\rtimes\mathbb{R}\vec{v}_\ell$ be the projection defined by $(\mathbf{x},\mathbf{t})\mapsto([\mathbf{x}]_{W_\ell^-},\ell(\mathbf{t}))$ (where $[\mathbf{x}]_{W_\ell^-}$ denotes for the W_ℓ^- coordinate of $\mathbf{x}$). Let H be a ℓ block, and write $\rho_\ell(p,q)=(\Pi_\ell(p),\Pi_\ell(q))_{\Pi_\ell(H)}$ for the Gromov product of $\Pi_\ell(p)$ and $\Pi_\ell(q)$ with respect to $\Pi_\ell(H)$ in the negatively curved space $W_\ell^-\rtimes\mathbb{R}\vec{v}_\ell$.

For the remaining of this section, H denotes for a ℓ block. Also, until the end of Section 4, ϵ is a positive number less than 1 such that $O(\operatorname{diam}(\mathbf{B}))\ll e^{\epsilon\operatorname{diam}(\mathbf{B})}$, and $\epsilon\operatorname{diam}(\omega)\ll\epsilon\operatorname{diam}(\mathbf{B})$. The lengths of all paths in this section are less than $e^{\epsilon\operatorname{diam}(\mathbf{B})}$.

We now list some properties of ρ_ℓ. In the lemma below, $\approx$ is used to denote two quantities whose ratio depends only on κ, C, ℓ and the space G.

4.3.1 Lemma (i) *Suppose that $d(p',p)\ll d(p,H)$, $d(q',q)\ll d(q,H)$ and $\rho_\ell(p,q)\ll\min\{d(p,H),d(q,H)\}$. Then $\rho_\ell(p,q)\approx\rho_\ell(p',q')$.*

(ii) *Suppose that $\ell(p')<\ell(p)$, $\ell(q')<\ell(q)$, and each of the pairs (p,p'), (q,q') lie on common flats such that $d(\Pi_\ell(p'),\Pi_\ell(p))\ge\eta d(p,p')$, $d(\Pi_\ell(q'),\Pi_\ell(q))\ge\eta d(q,q')$. If $d(p,H),d(q,H)\gg\rho_\ell(p,q)$, then $\rho_\ell(p,q)\approx\rho_\ell(p',q')$.*

(iii) *If $\rho_\ell(p,q),\rho_\ell(q,q')\gg s$, then $\rho_\ell(p,q')\gg s$.*

4.3.2 Lemma *Suppose $p, q \in G$ can be connected by a path $\widehat{\gamma}$ of length less than $e^{\epsilon R}$ such that*

(i) $\ell(\pi_A(p)), \ell(\pi_A(q)) \le \ell(H) - \rho_4$;

(ii) *the ℓ values of $\widehat{\gamma}$ decreases for at least ϵR units at both ends, and the ℓ values of points on the remaining subsegment are no more than $\ell(H) - \epsilon R$.*

Then, $\rho_\ell(p,q) \ge \Omega(\rho_4)$.

Proof Let p' and q' be the closest points to p and q on $\widehat{\gamma}$ whose ℓ values first dip below $\ell(H) - \epsilon R$. Since the length of $\widehat{\gamma}$ is less than $e^{\epsilon R}$, so is the length of $\Pi_\ell(\widehat{\gamma})$, which connects $\Pi_\ell(p)$ and $\Pi_\ell(q)$, as well as the subsegment of $\Pi_\ell(\widehat{\gamma})$ between $\Pi_\ell(p')$ and $\Pi_\ell(q')$. This means that if $\Pi_\ell(p) \neq \Pi_\ell(q)$, then any path connecting $\Pi_\ell(p')$ to $\Pi_\ell(q')$ whose ℓ value stays ϵR units below $\ell(H)$ would have length at least $e^{\epsilon R}$, contradicting the assumption about the length of $\widehat{\gamma}$. □

4.3.3 Lemma *Let p_0, q_0 be good ℓ vertices, and γ be an $\mathcal{E}_0$ edge connecting them. Let $p, q \in \gamma$ be points that are within the same good box as p_0 and q_0 respectively. Then the following holds:*

(i) *Except near the end points, γ never pass through any irregular ℓ vertices.*

(ii) *We have $\rho_\ell(\widehat{\phi}^{-1}(p), \widehat{\phi}^{-1}(q)) \ge \Omega(\rho_4)$.*

Proof Starting from p_0, let p_1 be the first place where γ hits $\widehat{W}(H_\ell)$. As γ is a $\mathcal{E}_0$ edge, it does not hit a bad ℓ shadow vertex, and so there is a good ℓ shadow vertex $p_1' \in U_* \cap W(H_\ell)$ such that $\widehat{\phi}(p_1') = p_1$ and $d(\phi^{-1}(p_1), p_1') = O(\epsilon\, \mathrm{diam}(\omega))$. Note that p_1' and $\phi^{-1}(p_1)$ are both $\Omega(\epsilon\, \mathrm{diam}(\mathbf{B}))$ away from $\partial(\phi^{-1}(\mathbf{B}))$. Let $p_2' = \phi^{-1}(p_2)$ be the next point after p_1' when $\phi^{-1}(\gamma)$ intersects $U_* \cap W(H_\ell)$. We know there must be such a point because the length of γ is less than $e^{\epsilon\, \mathrm{diam}(\mathbf{B})}$, so whenever it moves transverse to H_ℓ, it must do so in $\ell^{-1}[h_0^i, \infty]$. Since γ does not hit a bad shadow vertex, $p_2' \in \ell^{-1}(h_0^i)$ at $i \in \mathbf{I}_g(\ell)$, so p_2 is a good shadow vertex. Therefore continuing γ after p_2 hits a vertex in $\phi(U_*) \cap H_\ell$, which is a good ℓ vertex, and must be q_0. By Lemma 4.3.1,

$$\rho_\ell(\widehat{\phi}^{-1}(p),\widehat{\phi}^{-1}(q)) = \rho_\ell(\widehat{\phi}^{-1}(p_1),\widehat{\phi}^{-1}(p_2)) \approx \rho_\ell(\phi^{-1}(p_1),\phi^{-1}(p_2)) \ge \Omega(\rho_4). \quad \square$$

4.3.4 Lemma *Suppose p_0, q_0 are good ℓ vertices and $\overline{p_0 q_0}$ is a $\mathcal{E}_0$ edge connecting them. If for some $\rho_1 \ll s \ll \rho_3 \ll \rho_4$, we have $p, q \in \overline{p_0 q_0}$ in the same good box as p_0 and q_0 satisfying*

$$(f_i)_*\ell(p) - (f_i)_*\ell(p_0) = (f_j)_*\ell(q) - (f_j)_*\ell(q_0) = s, \quad \text{where } i, j \in \mathbf{I}_g(\ell).$$

Then, there is an ℓ block H_ℓ' such that p, q are within $O(\rho_1)$ of $\widehat{\phi}(H_\ell')$.

Proof Let p', q' be points on $\overline{p_0 q_0}$ close to where it enters respective good boxes. By Lemma 4.3.3, $\rho_\ell(\hat{\phi}^{-1}(p'), \hat{\phi}^{-1}(q')) \geq \Omega(\rho_4)$. Since $s \ll \rho_3 < \rho_4$, we conclude that the ℓ values of $\hat{\phi}_i^{-1}(p)$ and $\hat{\phi}_j^{-1}(q)$ are the same: both are s lower than that of H_ℓ, and so $\Pi_\ell(\hat{\phi}_i^{-1}(p))$ and $\Pi_\ell(\hat{\phi}_j^{-1}(q))$ are within 2δ of a common point, where δ is the thin triangle constant in the negatively curved space $W_\ell^- \ltimes \mathbb{R}\vec{v}_\ell$. But this is the same as saying that $\hat{\phi}_i^{-1}(p)$ and $\hat{\phi}_j^{-1}(q)$ are within $O(\rho_1)$ of a ℓ block H'_ℓ as claimed. □

4.3.5 Lemma *Let n be 4 or 6. Suppose for $0 \leq i \leq n-1$, p_i are $\hat{S}$ vertices whose ϕ preimages support standard maps. Let $\overline{p_{i-1} p_i}$ be subsegments of $\mathcal{E}_0$ edges in H graph, where the indices are counted* mod n.

Let $r(p_i) = \min\{|(f_i)_*\ell(v) - (f_i)_*\ell(p_i)|, v \text{ is a good } \ell \text{ vertex.}\}$, where $\phi_i = g_i \times f_i$ is the standard map supported in a neighborhood of $\phi^{-1}(p_i)$.

Suppose there is an index k such that $r(p_k) \ll \rho_4$, and for all $i \neq k$, $r(p_i) > r(p_k) + 2\rho_1$. Then $\overline{p_k p_{k+1}}$ and $\overline{p_k p_{k-1}}$ cannot have only p_k in common.

Proof We can assume $k = 0$. Let H'_ℓ be the ℓ block passing through $\hat{\phi}^{-1}(p_0)$. By Lemma 4.3.4, we can consider H'_ℓ in place of H_ℓ. Namely, we can replace the appearance of H_ℓ in the definition of ℓ vertices by H'_ℓ. Let p_i^+ and p_i^- be the first and last time that $\overline{p_{i-1} p_i}$ leaves $\widehat{W}(H_\ell)$. Suppose the claim is not true, then by Lemma 4.3.3 $\rho_\ell(\hat{\phi}^{-1}(p_{i-1}^+), \hat{\phi}^{-1}(p_i^-)) \geq \Omega(\rho_4)$ for all $i \neq 0$. But we also know that $\rho_\ell(\hat{\phi}^{-1}(p_0^-), \hat{\phi}^{-1}(p_0^+)) \leq \rho_1$. This is a contradiction to Lemma 4.3.1 □

4.4 Proof of Theorem 4.0.7

Let $\mathbb{R}^k \rtimes_\tau \mathbb{R}$ be an unimodular rank 1 space with roots α_i, $-\beta_j$'s, where $\alpha_i, \beta_j > 0$, and $\tau(t)$ is a matrix consisting of blocks of the form $e^{\alpha_i t} N_i(\alpha_i t)$, $e^{-\beta_j t} N_j(\beta_j t)$ where N_i, N_j are unipotent matrices with polynomial entries. For the following two lemmas, let $B[T]$ denote for the subset $(\prod_i \Omega_i \times \prod_j U_j) \rtimes [-T, T]$, where $\Omega_i \subset V^{\alpha_i}$, and $U_j \subset V^{-\beta_j}$, and $\alpha_i, \beta_j > 0$, $|\Omega_i| e^{-\alpha_i T}, |U_j| e^{-\beta_j T} \geq 1$. We call the set

$$\left(\left(\prod_i \Omega_i \times \prod_j U_j\right), -T\right) \cup \left(\left(\prod_i \Omega_i \times \prod_j U_j\right), T\right)$$

the *top and bottom of B[T]*, denoted by $\overline{\partial} B[T]$.

4.4.1 Lemma *The total number of geodesics in $B[T]$ is*

$$e^{T(\sum_i \alpha_i + \sum_j \beta_j)} \prod_{i,j} |\Omega_i||U_j|,$$

and the number of geodesics in $B[T]$ through each vertex is $e^{T(\sum_i \alpha_i + \sum_j \beta_j)}$.

Proof To specify a geodesic in $B[T]$, we need to specify its coordinates in α_i, β_j root spaces, and for every choice of α_i and β_j coordinate, there is a unique geodesic segment in $B[T]$ going from the top to the bottom. The number of different coordinates in α_i root spaces is $\prod_i |\Omega_i| e^{\alpha_i T}$, and those in β_i root spaces is $\prod_j |U_j| e^{\beta_j T}$, so the number of geodesics is

$$\Big(\prod_i |\Omega_i| e^{\alpha_i T}\Big)\Big(\prod_j |U_j| e^{\beta_j T}\Big).$$

We know that in $B[T]$,

geodesics $\times$ # vertices on a geodesic

$=$ # vertices $\times$ # geodesics through each vertex.

The number of vertices on each geodesic is $2T$, and the number of vertices is $2T(\prod_i \Omega_i)(\prod_j U_j)$, and we now see the number of geodesics through each point is indeed as claimed. □

Since $\mathbb{R}^k \rtimes_\tau \mathbb{R}$ is unimodular, let m denotes the common values of $\sum_i \alpha_i$ and $\sum_j \beta_j$.

4.4.2 Lemma (Hypothesis as in Lemma 4.4.1) *Let $X \subset B[T]$ be a subset of vertices. If $\mathcal{F}$ is a family of geodesics in $B[T]$ with size $\sigma e^{2mT} \prod_{i,j} |\Omega_i||U_j|$, where $\sigma \gg 2T/e^{m\rho_2}$, then there is a vertex $v \in X$, and two geodesics from $\mathcal{F}$ through v that stay together for shorter than ρ_2 units.*

Proof Suppose the claim is not true. Then for each X vertex v, every pairs of geodesics through v stay together at least ρ_2 units. If v is within ρ_2 neighborhood of top and bottom of $B[T]$, the number of geodesics through v with properties is $e^{(2T-h(v))m}/e^{m\rho_2}$, where h is the height function on $\mathbb{R}^m \rtimes \mathbb{R}$. On the other hand, if v is outside of ρ_2 neighborhood of top and bottom of $B[T]$, the number of geodesics through v with this property is

$$\frac{e^{h(v)\sum_i \alpha_i}}{e^{m\rho_2}} \frac{e^{(2T-h(v))\sum_j \beta_j}}{e^{m\rho_2}} = \frac{e^{2mT}}{e^{2m\rho_2}}.$$

The number of X vertices outside of ρ_2 neighborhood of the top and bottom of $B[T]$ is at most $2(T-\rho_2)\prod_{i,j} |\Omega_i||U_j|$, and the number of X vertices within ρ_2 neighborhood of top and bottom of $B[T]$ is at most $2\rho_2 \prod_{i,j} |\Omega_i||U_j|$. So the number of geodesics satisfying this scenario is at most

$$\sum_{v \in N_{\rho_2}\bar{\partial} B[T]} \frac{e^{(2T-h(v))m}}{e^{m\rho_2}} + \frac{e^{2mT}}{e^{2m\rho_2}} 2(T-\rho_2) \prod_{i,j} |\Omega_i||U_j| \leq \frac{e^{2Tm}}{e^{m\rho_2}}(2T) \prod_{i,j} |\Omega_i||U_j|.$$

Since the size of $\mathcal{F}$ is larger than this number, it is not possible that every pairs of geodesics in $\mathcal{F}$ satisfy the scenario described above. So there is a X vertex v, and two geodesics in $\mathcal{F}$ that stay together for less than ρ_2 units after passing through v. $\square$

In the remainder of this section, given a regular vector $\vec{u}$, we write $I_{\lambda,\vec{u}}(p)$, $p \in G'$, for the subset of $pW^{+}_{\vec{u}}$ that can be reached by two geodesics of length λ in the rank 1 space $G'_{\vec{u}} = \mathbf{H}' \rtimes \mathbb{R}\vec{u}$ containing p, that first moves in the direction $\vec{u}$ followed by another one in direction $-\vec{u}$. We also denote by $I'_{\lambda,\vec{u}(p)}$ for the subset of the left $W^{-}_{\vec{u}}$ coset that can be reached from p by a geodesic in direction $\vec{u}$ (or in direction $-\vec{u}$ as viewed from p) of length λ.

For the next two propositions, we make the following assumptions.

(i) Let v be a $\mathfrak{b}$ vertex such that $\phi^{-1}(v)$ locally supports a standard map $\phi_v = f_v \rtimes g_v$. Since f_v is affine and permutes root classes of G to root classes of G', its linear part induces a permutation from the chambers of G to chambers of G' and we write $(f_v)_*$ for this permutation.

(ii) Suppose for $\ell \in (\mathbf{A})^*$ a regular linear functional, and $\vec{u} \in \mathbf{A}'$, the vectors $\vec{v}_\ell$ and $(f_v)_*^{-1}\vec{u}$ lie in a common chamber $\mathfrak{b}$ of G.

The existence of such vertex is given by Lemma 4.2.8.

4.4.3 Proposition *If λ is a number such that at least σ fraction of geodesics leaving v in direction $\vec{u}$ are unobstructed by H_ℓ vertices for length at least $\lambda + \rho_2$, where $\sigma \gg 2\lambda/e^{\mathbf{b}_{\vec{u}}\rho_2}$. Suppose $\sigma > O(\eta)$ for some $\eta < 1$, then at least $1 - O(\eta)$ fraction of the $\widehat{S}$ vertices in $I_{\lambda,\vec{u}}(v)$ are H_ℓ vertices.*

Proof Let E denote the set of geodesics leaving v in direction $\vec{u}$ that are unobstructed by ℓ vertices of length at least $\lambda + \rho_2$. Let E_λ be the subset of $I'_{\lambda,\vec{u}}(v)$ passing through an element of E. By assumption, we have

$$|E_\lambda| \geq \sigma e^{\mathbf{b}_{\vec{u}}\lambda}.$$

Let $\mathcal{F}'_0$ be the union of geodesics leaving E_λ in direction $-\vec{u}$. (as viewed from E_λ). Applying Lemma 4.4.2 to $\mathcal{F}'_0$, where X means ℓ vertices, we see that either there is an ℓ vertex whose $\vec{u}^* = \langle \vec{u}, \pi_A(.) \rangle$ value is at most ρ_2 from that of v, or that there is an ℓ vertex w whose $\vec{u}^*$ value differ from that of v by more than ρ_2, and two geodesics in $G_{\vec{u}}$ through w_1 that stay together for less than ρ_2 units after passing through w_1. Suppose the latter happens. Let $x, y \in I'_{\lambda,\vec{u}}(v)$ be two upper end points of those two geodesics, and w_1 be the first time that $\overline{wx}$ diverge from $\overline{wy}$. See Figure 1 below.

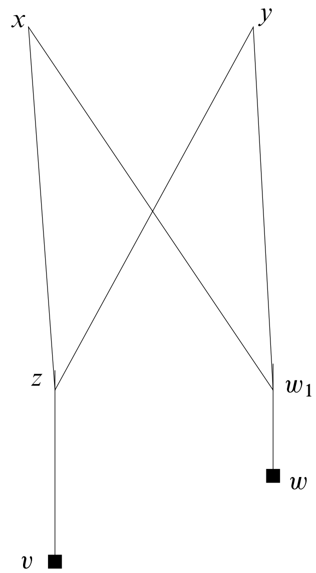

Figure 1: The loop in the proof of Proposition 4.4.3. Filled boxes are H_ℓ vertices.

Let z be the first time that $\overline{vx}$ diverge from $\overline{vy}$. Applying Lemma 4.3.5 to the loop $z-y-w_1-x-z$ creates a contradiction. So there is an ℓ vertex whose $\vec{u}^* = \langle \vec{u}, \pi_A(\cdot)\rangle$ value is at most ρ_2 from that of v. That is, there is an ℓ vertex in the ρ_2 neighborhood of $I_{\lambda,\vec{u}}(v)$.

Let $U' \subset I_{\lambda,\vec{u}}(v)$ be those vertices that can be reached by two elements of $\mathcal{F}'_0$. Since every vertex in $I_{\lambda,\vec{u}}(v)$ can be reached by at most $|E_\lambda|$ geodesics in $I'_{\lambda,\vec{u}}(v)$, it follows that the relative measure of U' in $I_{\lambda,\vec{u}}(v)$ is at least $1-O(\eta)$.

Now suppose $w \in U'$, and let $x, y \in I'_{\lambda,\vec{u}}(v)$ be such that $\overline{xw}, \overline{yw}$ are element of $\mathcal{F}'_0$ that are not obstructed by ℓ vertices. Applying Lemma 4.3.5 to the loop $v-x-w-y-v$, and noting that $r(v)=0$, $r(x), r(y) \geq \rho_2$, it follows that $r(w)=0$ (otherwise $r(v)$ would be the smallest, a contradiction), which means that w is a ℓ vertex. □

4.4.4 Proposition (Hypothesis as in Proposition 4.4.3) *Let $\mathcal{F}$ be the union of geodesics in direction $\vec{u}$ leaving a point of $I_{\lambda,\vec{u}}(v)$. Then at least $1-O(\eta)$ fraction of $\mathcal{F}$ are unobstructed by H_ℓ vertices for length $\lambda+\rho_2$.*

Proof Let $E_\lambda \subset I'_{\lambda,\vec{u}}(v)$, $U' \subset I_{\lambda,\vec{u}}(v)$ and $\mathcal{F}'_0$ be as in Proposition 4.4.3. Note that the measure of U' is at least $1-O(\eta)$ that of $I_{\lambda,\vec{u}}(v)$.

Let $\mathcal{F}''$ be the set of geodesics leaving U' in direction $\vec{u}$. Let $\mathcal{F}''_{\text{long}}$ be the set of geodesics coming from extending elements of $\mathcal{F}''$ by extending ρ_4 units on the $I_{\lambda,\vec{u}}(v)$

(ideally, we would like to apply Lemma 4.4.2 to the family $\mathcal{F}''$ in a rank 1 box of size λ, but in order to use illegal circuit we need the stub-vertical segment from H_ℓ instead of being on H_ℓ. Hence the choice of H'_ℓ lower down). Let H'_ℓ be the ℓ block whose ℓ value is ρ_4 less than $\ell(H_\ell)$. We call the resulting vertices ℓ' vertices if we replace occurrence of H_ℓ in the definition of ℓ vertices by H'_ℓ.

If all elements of $\mathcal{F}''_{\text{long}}$ are unobstructed by images of ℓ' vertices, then at least $1-O(\eta)$ proportion of all elements of $\mathcal{F}$ are unobstructed by ℓ vertices. Let U'_{long} be the set of ℓ' vertices that are within ρ_4 of U'.

We have that $|\mathcal{F}''_{\text{long}}| \geq (1-O(\eta))e^{\mathbf{b}_{\vec{u}}(2\lambda+\rho_4)}$. Lemma 4.4.2 allows us to conclude that either there is an ℓ vertex whose $\vec{u}^*$ value is within ρ_2 to $\partial\mathbf{B}\cap G'_{\vec{u}}$, or that there is an ℓ vertex q whose $\vec{u}^*$ value differ from that of $\partial\mathbf{B}\cap G'_{\vec{u}}$ by more than ρ_2 units, and two elements of $\mathcal{F}''_{\text{long}}$ that stay together for less than ρ_2 units after passing through q.

Suppose the latter scenario happens. Then there are $w_1, w_2 \in I_{\lambda,\vec{u}}(v)$ and a ℓ' vertex q such that $\overline{w_1q}$, $\overline{w_2q} \in \mathcal{F}''_{\text{long}}$. Let q_* be the first point where $\overline{w_1q}$ and $\overline{w_2q}$ come together. Then by assumption, the $d(q,q_*) < \rho_2$. Let $x_1 \in I'_{\lambda,\vec{u}}(v)$ be the first point where geodesics in direction $\vec{u}$ leaving w_1 and v first meet, and let $x_2 \in I'_{\lambda,\vec{u}}(v)$ be similarly defined for w_2 and v. Let $r(.)$ now denotes for the distance to the closest ℓ' vertex. Then in the loop $v-x_1-w_1-q_*-w_2-x_2-v$, (see Figure 2) the r value of all points but q_* are at least ρ_4, while $r(q_*) \leq \rho_2$, which is a contradiction by Lemma 4.3.5. Therefore if elements of $\mathcal{F}''_{\text{long}}$ is to contain a ℓ' vertex, this vertex is

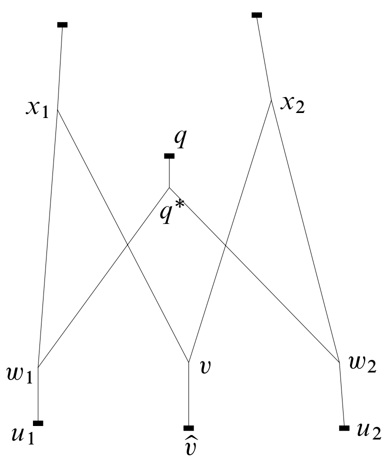

Figure 2: The loop in the proof of Proposition 4.4.4. Filled boxes are $H_{\ell'}$ vertices.

within ρ_2 neighborhood of $\partial\mathbf{B}\cap G'_{\vec{u}}$, which is just saying that no elements of $\mathcal{F}''_{\text{long}}$ are obstructed by ℓ' vertices, therefore at least $1-O(\eta)$ proportion of elements in $\mathcal{F}$ are unobstructed by ℓ vertices by construction. □

4.4.5 Theorem *Let v be a $\mathfrak{b}$ vertex and $\widehat{\phi}_v = f_v \times g_v$ be a standard map supported in a neighborhood of $\phi^{-1}(v)$. Let $\lambda_0 = d(v, \partial\mathbf{B})$. Then at least $1-O(\eta)$ proportion of vertices in $vW^+_{(f_v)_*\mathfrak{b}}$ are $\mathfrak{b}$ vertices, where $O(\eta)$ is the proportion of geodesics leaving $\phi^{-1}(v)$ that admits a geodesic approximation of length $\Omega(\mathrm{diam}(\omega))$.*

Proof Here we give a proof when rank is at least 2. The rank 1 case in general follows that of the arguments in Eskin, Fisher and Whyte [5].

Let $s_{\vec{u}}$ be the difference in $v^*_{\vec{u}}$ values between $\partial\mathbf{B} \cap G'_{\vec{u}}$ and v.

Fir a $\rho_2 = \epsilon s_{\vec{u}}$. For every $\ell \in \mathcal{S}_0$, and $w \in I_{\lambda,\vec{u}}(v)$, we let $f_\ell(w,\lambda)$ denotes the proportion of geodesics leaving w that are unobstructed by H_ℓ vertices for length at least $\lambda + \rho_2$. Let

$$f^*_\ell(v,\lambda) = \sup_{w \in I_{\lambda,\vec{u}}} f_\ell(w,\lambda).$$

In view of Propositions 4.4.3 and 4.4.4, if $f^*_\ell(v,\lambda) \geq O(\eta)$, then $f^*_\ell(v,\lambda) \geq 1-O(\eta)$.

Then, either for all $\lambda \leq s_{\vec{u}}$, we have $f^*_\ell \geq 1-O(\eta)$; or that there is a maximal number $\lambda_{\vec{u},\ell}-1$ such that $f^*_\ell(v,\lambda_{\vec{u},\ell}-1) \geq 1-O(\eta)$, but $f^*_\ell(v,\lambda_{\vec{u},\ell}) < O(\eta)$. We are done if the latter does not happen. We now proceed to show that this is indeed the case.

In the second scenario, we know that $\lambda_{\vec{u},\ell} \geq \epsilon\,\mathrm{diam}(\omega)$, and at least $1-O(\eta)$ proportion of vertices in $I_{\lambda_{\vec{u},\ell},\vec{u}}(v)$ and $I'_{\lambda_{\vec{u},\ell},\vec{u}}(v)$ are H_ℓ vertices. That is, they are ϕ images of $U_* \cap (x_0 W^+_{\mathfrak{b}} \rtimes \ker(\ell))$.

Claim *If $\lambda_{\vec{u},\xi} > \lambda_{\vec{u},\ell}$ for $\xi, \ell \in \mathcal{S}_0$, then $\lambda_{\vec{u},\xi} - \lambda_{\vec{u},\ell} > O(\epsilon\,\mathrm{diam}(\omega))$.*

Suppose not. Then we will have subsets, one in $\ker(\xi)$ and one in $\ker(\ell)$ that are within $O(\epsilon\,\mathrm{diam}(\omega))$ Hausdorff distance from each other. This can only happen if the subsets are within $O(\epsilon\,\mathrm{diam}(\omega))$ of $\ker(\xi) \cap \ker(\ell)$. But this would mean that most of $I'_{\lambda_{\vec{u},\ell},\vec{u}}(v)$ come from ϕ images of $x_0 W^+_{\mathfrak{b}} \rtimes (\ker(\xi) \cap \ker(\ell))$, contradicting the assumption that $\lambda_{\vec{u},\xi}$ is the minimal height t where most of the $I'_{t,\vec{u}}(v)$ are obstructed by H_ξ vertices.

In this way, the image of the map $\mathcal{S}_0 \to [0, s_{\vec{u}}]$ defined by sending $\xi \to \lambda_{\vec{u},\xi}$ is a $O(\epsilon\,\mathrm{diam}(\omega))$ discrete set. Let $\widehat{\lambda}_{\vec{u}}$ be the minimal image value whose preimages has positive measure. This means that most elements of $I'_{\widehat{\lambda}_{\vec{u}},\vec{u}}(v)$ and $I_{\widehat{\lambda}_{\vec{u},\vec{u}}}(v)$ are ϕ images of $U_* \cap x_0 W^+_{\mathfrak{b}}$, thus the subset of $\mathcal{S}_0$ consisting of elements ξ such that $\lambda_{\vec{u},\xi} > \widehat{\lambda}_{\vec{u}}$ is empty. Since for all $t < \widehat{\lambda}_{\vec{u}}$, the preimages of t in $\mathcal{S}_0$ has zero measure, this means not only does the preimages of $\widehat{\lambda}_{\vec{u}}$ has positive measure, it has full measure.

Now pick another direction $\vec{u}'$ in the same chamber as $\vec{u}$, but not in the $\epsilon^{1/2}$ neighborhood of the $\vec{u}$ orbit under the finite group of affine maps permuting R_g to $R_{G'}$, and repeat the same argument as above to obtain a number $\widehat{\lambda}_{\vec{u}'}$ such that most of $I'_{\widehat{\lambda}_{\vec{u}'},\vec{u}'}(v)$ come from ϕ images of $x_0 W_{\mathfrak{b}}^+$.

Pick $y_0 \in I'_{\widehat{\lambda}_{\vec{u}}}(v)$ and $y \in I'_{\widehat{\lambda}_{\vec{u}'}}(v)$ so that each locally supports a standard map. Take two geodesics leaving y_0 in direction $\vec{u}$ (as viewed from v) that stay together for $t_{\vec{u}}$ units (where $\epsilon^{1/2}\operatorname{diam}(\omega) \ll t_{\vec{u}} \ll \operatorname{diam}(\omega)$) after they leave y_0, followed by a short segment $\operatorname{diam}(\omega)$ away from v, before joining the geodesics connecting v to y. Let's say they stay together for $t_{\vec{u}'}$ units before coming to a stop at y. See Figure 3.

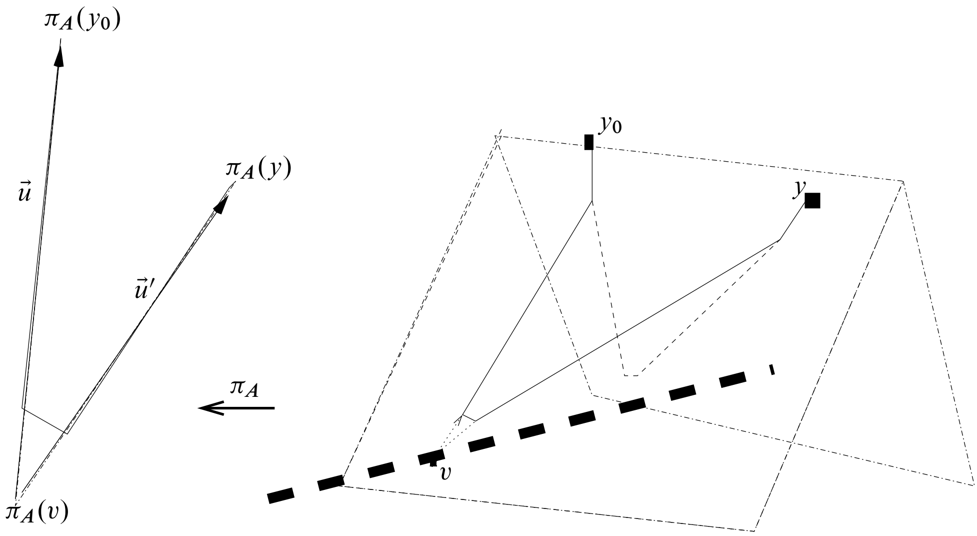

Figure 3: The loop that prevents blocking. Filled boxes represent $\mathfrak{b}$ vertices. The left hand is the π_A projection image of the loop on the right.

As most of $I'_{\widehat{\lambda}_{\vec{u}}}(v)$ and $I'_{\widehat{\lambda}_{\vec{u}'}}(v)$ come from $U_* \cap W_{\mathfrak{b}}^+$, for a full measure of $\ell \in \mathfrak{b}$, Lemma 4.3.5 requires

$$\Pi_\ell \phi_{y_0}^{-1}(t_{\vec{u}}\vec{u}) \geq \Pi_\ell \phi_y^{-1}(t_{\vec{u}'}\vec{u}'), \quad \text{as well as} \quad \Pi_\ell \phi_{y_0}^{-1}(t_{\vec{u}}\vec{u}) \leq \Pi_\ell \phi_y^{-1}(t_{\vec{u}'}\vec{u}').$$

This means that $f_{y_0}^{-1}(t_{\vec{u}}\vec{u}) = f_y^{-1}(t_{\vec{u}'}\vec{u}')$, where f_{y_0} and f_y are linear part of standard maps ϕ_{y_0} and ϕ_y. That is, $t_{\vec{u}}/t_{\vec{u}'} \in [1/(1-\epsilon), 1+\epsilon]$, and $t_{\vec{u}}\vec{u}$ lies in the $\epsilon\operatorname{diam}(\omega)$ neighborhood of the orbit of $t_{\vec{u}'}\vec{u}'$ under the finite group of linear maps that permutes R_G to $R_{G'}$. But this contradicts our choice of $\vec{u}'$ and $t_{\vec{u}}$. □

Proof of Theorem 4.0.7 As any root class horocycle is the intersection of finitely many left translates of $W_{\mathfrak{b}}^+$, where $\mathfrak{b}$ is a chamber, it suffices to show that the claim holds for left translates of $W_{\mathfrak{b}'}^+$ of G' in place of left translates of $V_{[\alpha]'}$ of G'.

We start with $\mathbf{B}(\Sigma)$ sufficiently large so that we can apply Theorem 3.0.5 to $\phi^{-1}(\mathbf{B}(\Sigma))$ and obtaining a tiling as in Equation (14). Let $U_* = \bigcup_{i \in \mathbf{I}_0} U_i$. Since U_* has large measure relative to $\phi^{-1}(\mathbf{B}(\Sigma))$, for a fixed chamber $\mathfrak{b}$, we can find a large subset $U_\mathfrak{b} \subset U_*$ with the property that for every point $p \in U_\mathfrak{b}$, there is a subset $\mathcal{S}_0 \subset \mathfrak{b}$ of relative proportion at least $1 - \vartheta$ such that $E_\ell(p)$ is very favourable for every $\ell \in \mathcal{S}_0$. By constructing the corresponding $\widehat{S}$ and H graph, application of Theorem 4.4.5 to a point $v \in \phi(U_\mathfrak{b})$ shows that the ϕ^{-1} image of $vW_\mathfrak{c}^+ \cap \mathbf{B}(\Sigma)$ is $O(\epsilon \operatorname{diam}(\mathbf{B}(\Sigma))$ away from a left translate of $W_\mathfrak{b}^+$, where $\mathfrak{c}$ is the image of $\mathfrak{b}$ under the linear part of the standard map supported in a neighborhood of $\phi^{-1}(v)$. □

5 Patching

In the previous section, we aligned the linear part of standard maps appeared in Theorem 3.0.5 by showing that they are all the same. In this last section, we remove the condition that standard maps are only defined for a subset of relative large measure and align the translational part of the standard maps by adopting the procedure used to achieve this in Eskin, Fisher and Whyte [4].

5.1 A weak version of an affine map

We have by now seen that given a box, there is a subset of large measure supporting a standard map. In this section, by controlling the sizes of increasingly larger and larger boxes, we remove the constraint of "subset of large relative measure"s and extend the result to all pairs of points $p, q \in G$ on the same flat. The precise statement is the following.

5.1.1 Theorem *Let G, G' be nondegenerate, unimodular, split abelian-by-abelian Lie groups and $\phi\colon G \to G'$ be a (κ, C) quasi-isometry between them. Given $0 < \delta$, $\eta \ll \tilde{\eta} < 1$, then there exists $\tau < 1$, M depending on δ, η, $\tilde{\eta}$, and (κ, C) such that whenever x, y belong to the same left coset of $\mathbf{H}$,*

$$|\pi_A(\phi(x)) - \pi_A(\phi(y))| \leq \tau d(x, y) + M, \tag{15}$$

where $\tau \to 0$, $M \to \infty$ as the input parameters approach zeros.

The setup to the proof of Theorem 5.1.1 follows the same sequence of steps as the analogue result in Section 6.1 of Eskin, Fisher and Whyte [4].

Fix $0 < \delta, \eta \ll \tilde{\eta} < 1$. Let $\Omega \subset \mathbf{A}$ be a product of intervals of size L_0 with barycenter located at the origin of $\mathbf{A}$. By Corollary 4.0.9, there is a subset $\mathcal{P}^0 \subset \mathcal{P}(\Omega)$ of relative large measure which is the support of a standard map $g \times f$ where f is affine. Let $\vartheta \ll 1$ satisfies $|\mathcal{P}^0| \geq (1-\vartheta)|\mathcal{P}(\Omega)|$, and set $\varrho = \sqrt{\vartheta}$. Let $\mathbf{P}$ be a left translate of $\mathbf{H}$, and we can assume $\mathbf{P}$ contains the identity element. Let $R(\Omega) = \mathbf{B}(\Omega) \cap \mathbf{P}$, and $\mathbf{R}(\Omega) = \bigcup_{g \in R(\Omega)} g(\varrho\Omega)$. The following is a rehash of Corollary 4.0.9.

5.1.2 Corollary *There is a standard map $f \times g$ where f is affine with linear part A_f, defined on $\mathcal{P}^0 \subset \mathcal{P}(\mathbf{R}(\Omega))$, with $|\mathcal{P}^0| \geq (1-\varrho)|\mathcal{P}(\mathbf{R}(\Omega))|$, such that $d(\phi|_{\mathcal{P}^0}, f \times g) \leq \epsilon \operatorname{diam}(\mathbf{B}(\Omega))$. Furthermore, if $p \in \mathcal{P}^0$, there is a subset $L^0(p) \subset L(p)$ of relative large measure such that the ϕ image of every element $\zeta \in L^0(p)$ is within ϵ–linear neighborhood of a geodesic segment.*

The tiling Choose $\epsilon \ll \varsigma \ll 1$. For each $j \in \mathbb{N}$, set $\Omega_j = (1+\varsigma)^j \Omega$. We tile $\mathbf{P}$ by $R(\Omega_j)$, where each tile is denoted by $R_{j,\iota}$, $\iota \in \mathbb{N}$. For $x \in G$, we write $R_j[x]$ for the tile in the j–th tiling containing the point in $\mathbf{P}$ that lies on the same flat as x. Note that the number $R_{j,k}$'s needed to cover a rectangle $R_{j+1}[p]$ is on the order of $e^{\sum_{i=1}^{|\triangle|} \varsigma \max\{\alpha_i(\Omega)\}}$.

The sets U_j For each tile $R_{j,k}$ in the j–th tiling of $\mathbf{P}$, Corollary 5.1.2 produces a subset $\mathcal{P}^0_{j,k} \subset \mathcal{P}(\mathbf{R}_{j,k})$ of relative large measure. We set

$$U_j = \bigcup_{k \in \mathbb{N}} \mathcal{P}^0_{j,k}.$$

In view of Corollary 5.1.2, for any $x \in U_j$,

$$\sup_{y \in \mathbf{R}[x] \cap U_j} |\pi_A(\phi(x)) - \pi_A(\phi(y))| \leq \epsilon \operatorname{diam}(\mathbf{B}(\Omega_j)). \tag{16}$$

Recall that $\triangle'$ is the set of roots in G', and we write n for the rank of G (which is also the rank of G'). We also have the following generalization:

5.1.3 Lemma *For any $x \in U_j$, and $y \in \mathbf{R}_{j+1}[x] \cap U_j$,*

$$|\pi_A(\phi(x)) - \pi_A(\phi(y))| \leq 4n|\triangle'|\epsilon \operatorname{diam}(\mathbf{B}(\Omega_j)).$$

Proof Without loss of generality, we can assume that there is a horocycle $H_{[\alpha]}$ that intersect both $\mathbf{R}[x] \cap U_j$ and $\mathbf{R}[y] \cap U_j$. (If not, then we can find a sequence of points $x = p_0, p_1, \ldots, p_l = y$ where $l \leq |\triangle'|$ such that for each pair of consecutive points, there is a horocycle intersecting $\mathbf{R}[p_\iota] \cap U_j$ and $\mathbf{R}[p_{\iota+1}] \cap U_j$). Let $x_1 \in H_{[\alpha]} \cap \mathbf{R}[x] \cap U_j$, $y_1 \in H_{[\alpha]} \cap \mathbf{R}[y] \cap U_j$ be points of intersection. Therefore by Equation (16),

$$|\pi_A(\phi(x)) - \pi_A(\phi(x_1))| \leq \varrho \operatorname{diam}(\mathbf{B}(\Omega_j)),$$
$$|\pi_A(\phi(y)) - \pi_A(\phi(y_1))| \leq \varrho \operatorname{diam}(\mathbf{B}(\Omega_j)).$$

Since $d(x_1, y_1) \le \operatorname{diam}(\mathbf{B}(\Omega_{j+1}))$, for $\iota = 1, 2$ we can find geodesic segments $\gamma_{x_1,\iota}$, $\gamma_{y_1,\iota}$ leaving x_1, y_1 respectively such that $Q = \{\gamma_{x_1,\iota}, \gamma_{y_1,\iota}\}_{\iota=1,2}$ is a 0–quadrilateral. Additionally, because $x_1, y_1 \in U_j \cap H_{[\alpha]}$, we can assume for $\iota = 1, 2$, $* = x_1, y_1$, the subsegment $\hat{\gamma}_{*,\iota} \subset \gamma_{*,\iota}$ containing $*$ and satisfies $|\hat{\gamma}_{*,\iota}| = (1/(1+\varsigma))|\gamma_{*,\iota}|$ admit geodesic approximation to its ϕ image. That is, $\phi(\hat{\gamma}_{*,\iota})$ is within $\eta|\hat{\gamma}_{*,\iota}|$ Hausdorff neighborhood of another geodesic segment.

Let $l_{*,\iota}$ be a geodesic approximation to $\phi(\hat{\gamma}_{*,\iota})$. Then angle between the direction of $l_{*,\iota}$ and that of $\hat{\gamma}_{*,\iota}$ is at most $\sin^{-1}(\epsilon)$. Therefore by modifying each $l_{*,\iota}$ by an amount at most $\epsilon 2\kappa \operatorname{diam}(Q) \le \epsilon \operatorname{diam}(\mathbf{B}(\Omega_{j+1}))$, we can assume $l_{*,\iota}$ all have parallel directions. Since $\varsigma \ll 1$, the four geodesic segment $l_{*,\iota}$ do constitute a quadrilateral $\tilde{Q}$ and [12, Lemma 4.1.10] applied to $\tilde{Q}$ yields

$$\left|\pi_{\vec{v}} \circ \Pi_{\vec{v}}(\phi(x_1)) - \pi_{\vec{v}} \circ \Pi_{\vec{v}}(\phi(y_1))\right| \le \epsilon \operatorname{diam}(\mathbf{B}(\Omega_j),$$

where $\vec{v}$ is parallel to edge directions of $\tilde{Q}$. Therefore

$$\begin{aligned}\left|\pi_{\vec{v}} \circ \Pi_{\vec{v}}(\phi(x)) - \pi_{\vec{v}} \circ \Pi_{\vec{v}}(\phi(y))\right| &\le \epsilon \operatorname{diam}(\mathbf{B}(\Omega_j) + 2\varrho \operatorname{diam}(\mathbf{B}(\Omega_j)) \\ &\le 4\epsilon \operatorname{diam}(\mathbf{B}(\Omega_j))\end{aligned}$$

since $\varrho \le \epsilon$. The claim now follows by constructing quadrilaterals whose image is approximated by quadrilaterals whose edge direction ranges over at least n many linearly independent directions. □

5.1.4 Lemma *Suppose $p \in \mathbf{R}_j[x] \cap U_j$, $q \in \mathbf{R}_{j+1}[x] \cap U_{j+1}$. Then,*

$$|\pi_A(\phi(x)) - \pi_A(\phi(y))| \le (4|\triangle|n + 2)\epsilon \operatorname{diam}(\mathbf{B}(\Omega_{j+1})).$$

Proof Both the relative measures of the projections of the sets $\mathbf{R}_{j+1}[x] \cap U_j$ and $\mathbf{R}_{j+1}[x] \cap U_{j+1}$ to $R_{j+1}[x]$ are at least $1-\varrho$. Therefore we can find $p' \in \mathbf{R}_{j+1}[x] \cap U_j$, $q' \in \mathbf{R}_{j+1}[x] \cap U_{j+1}$, and $|\pi_A(\phi(p')) - \pi_A(\phi(q'))| \le \varrho \operatorname{diam}(\mathbf{B}(\Omega_{j+1}))$.

By Lemma 5.1.3, $|\pi_A(\phi(p)) - \pi_A(\phi(p'))| \le 4n|\triangle|\epsilon \operatorname{diam}(\mathbf{B}(\Omega_j))$,

and by (16), $|\pi_A(\phi(q')) - \pi_A(\phi(q))| \le \epsilon \operatorname{diam}(\mathbf{B}(\Omega_{j+1}))$.

The claim now follows by triangle inequality. □

Proof of Theorem 5.1.1 We have

$$R_0[x] \subset R_1[x] \subset R_2[x] \subset \cdots$$

and

$$R_0[y] \subset R_1[y] \subset R_2[x] \subset \cdots.$$

There exists N with $\operatorname{diam}(\mathbf{B}(\Omega_N))$ is comparable to $d(x,y)$ (after possibly shifting the N's grid by a bit) such that $R_N[x]=R_N[y]$. Now for $0\le j\le N$, pick $x_j\in\mathbf{R}_j[x]\cap U_j$, $y_j\in\mathbf{R}_j[y]\cap U_j$. We may assume that $x_N=y_N$. By Lemma 5.1.4,

$$\begin{aligned}
&|\pi_A(\phi(x_0))-\pi_A(\phi(y_0))| \\
&\qquad\le \sum_{j=0}^{N-1}|\pi_A(\phi(x_{j+1}))-\pi_A(\phi(x_j))| + \sum_{j=0}^{N-1}|\pi_A(\phi(y_{j+1}))-\pi_A(\phi(y_j))| \\
&\qquad\le 2(4|\triangle|n+2)\sum_{j=0}^{N-1}\epsilon\operatorname{diam}(\mathbf{B}(\Omega_{j+1})) \\
&\qquad\le 4(2|\triangle|n+1)\frac{\epsilon}{\varsigma}\operatorname{diam}(\mathbf{B}(\Omega_N)),
\end{aligned}$$

where last inequality comes from $\operatorname{diam}(\mathbf{B}(\Omega_{j+1}))=(1+\varsigma)\operatorname{diam}(\mathbf{B}(\Omega_j))$. Now since $x_0\in R_0[x]$ $|\pi_A(\phi(x))-\pi_A(\phi(x_0))|\le 2\kappa d(x,x_0)\le 2\kappa\operatorname{diam}(\mathbf{B}(\Omega))=M$, and similarly $|\pi_A(\phi(y))-\pi_A(\phi(y_0))|\le M$. The claim now follows by noting that we chose $\operatorname{diam}(\mathbf{B}(\Omega_N))/d(x,y)\in[1/2,2]$ and $\varsigma\gg\epsilon$. □

5.2 Consequence of weak height preservation – flats go to flats

Theorem 5.1.1 is the first statement we have that places no additional constraints on the points other than their natural relation in G. In this subsection we show that as a first consequence, the image of a flat is within $O(1)$ of another flat, which eventually culminating in the proof of Theorem 5.3.6 in the next section.

5.2.1 Proposition *The quasi-isometry ϕ sends a flat to within $O(1)$ of a flat.*

We now proceed to establish some necessary observations.

5.2.2 Lemma *There is a linear map $A_0\colon \mathbf{A}\to\mathbf{A}'$ and numbers $\hat\tau<1$, $\widehat{M}>0$ such that for any $x,y\in G$,*

$$\left|(\pi_A(\phi(x))-\pi_A(\phi(y)))-(A_0(\pi_A(x))-A_0(\pi_A(y)))\right|\le\hat\tau d(x,y)+\widehat{M}. \tag{17}$$

Proof Let $\mathbf{B}$ be a box such that $\operatorname{diam}(\mathbf{B})/d(x,y)\in[1/2,2]$. By Corollary 4.0.9, there is a subset $\mathcal{P}^0\subset\mathcal{P}(\mathbf{B})$ of relative measure at least $(1-\sqrt{\epsilon})$ that supports a standard map which is $\epsilon\operatorname{diam}(\mathbf{B})$ away from $\phi|_{\mathcal{P}^0}$. Write A_0 for the linear part of the $\mathbf{A}'$ part of the standard map. Without loss of generality we can assume that $x\mathbf{H}\cap\mathcal{P}^0\ne\varnothing$, and $y\mathbf{H}\cap\mathcal{P}^0\ne\varnothing$. Let $\hat x\in x\mathbf{H}\cap\mathcal{P}^0$, and $\hat y\in y\mathbf{H}\cap\mathcal{P}^0$. Then by Theorem 5.1.1,

$$\left|(\pi_A(\phi(\hat y))-\pi_A(\phi(y)))-(A_0(\pi_A(\hat y))-A_0(\pi_A(y)))\right|\le\tau d(\hat y,y)+M.$$

Since $\pi_A(\hat{y}) = \pi_A(y)$. By Corollary 4.0.9 we also have

$$\left|(\pi_A(\phi(\hat{x})) - \pi_A(\phi(\hat{y}))) - (A_0(\pi_A(\hat{x})) - A_0(\pi_A(\hat{y})))\right| \le \epsilon \operatorname{diam}(\mathbf{B}),$$

and by Theorem 5.1.1 again, we have

$$\left|(\pi_A(\phi(x)) - \pi_A(\phi(\hat{x}))) - (A_0(\pi_A(x)) - A_0(\pi_A(\hat{x})))\right| \le \tau d(x, \hat{x}) + M.$$

Summing all three equations and apply triangle inequality to the left hand side we have

$$\begin{aligned}\left|(\pi_A(\phi(x)) - \pi_A(\phi(y))) - (A_0(\pi_A(x)) - A_0(\pi_A(y)))\right| &\le (2\tau + \eta) \operatorname{diam}(\mathbf{B}) + 2M \\ &\le 2(2\tau + \epsilon) d(x, y) + 2M,\end{aligned}$$

since ϵ and τ depends on our initial ϵ, δ, and approach zero as the initial inputs approach zero, we can assume that $2(2\tau + \epsilon) < 1$, and we set $\hat{\tau} = 2(2\tau + \epsilon)$, $\widehat{M} = 2M$. □

5.2.3 Corollary *There is a number M_0 such that if $p, q \in G$ are two points on the same flat and $\pi_A(\phi(p)) = \pi_A(\phi(q))$, then $d(x, y) \le M_0$.*

Proof In Equation (17) substitute x, y by p, q and note that since p, q lies on the same flat, $|\pi_A(p) - \pi_A(q)| = d(p, q)$. So we have $(1 - \hat{\tau}) d(p, q) \le \widehat{M}$, to which the result follows. Alternatively, this can be obtained from Theorem 5.1.1 applied to the inverse map ϕ^{-1} and $\phi(p)$, $\phi(q)$, for then Equation (15) becomes

$$\begin{aligned}|\pi_A \circ \phi^{-1}(\phi(p)) - \pi_A \circ \phi^{-1}(\phi(q))| &\le \tau d(\phi(p), \phi(q)) + M \le \tau\kappa d(p, q) + M + C \\ |\pi_A(p) - \pi_A(q)| = d(p, q) &\le \tau\kappa d(p, q) + M + C \\ d(p, q) &\le \frac{M + C}{1 - \tau\kappa}.\end{aligned}$$

□

5.2.4 Definition A subset L of $\mathbf{A}' \simeq \mathbb{R}^n$ is called a "grid" if it is the image of an injective homomorphism $\psi\colon \mathbb{Z}^n \to \mathbf{A}'$. A line in L refers to a subset of the form $\{\psi(\mathbf{c} + t\mathbf{u}) : t \in \mathbb{Z}\}$ for some $\mathbf{c}, \mathbf{u} \in \mathbb{Z}^n$, and each coordinate of $\mathbf{u}$ is either $+1, -1$ or 0. A grid is said to be good if none of its lines are parallel to root kernels.

5.2.5 Lemma *Let $\{x_i\} \subset G'$ be a sequence of points with the following properties:*

(i) $\pi_A(x_j) \neq \pi_A(x_i)$ *if* $i \neq j$.

(ii) $\{\pi_A(x_i)\} \subset \mathbf{A}'$ *is a good grid.*

(iii) *For any subsequence $\{x_{i_j}\}$ such that $\{\pi_A(x_{i_j})\}$ is a line, $\{x_{i_j}\}$ is within $O(1)$ of a (bi-infinite) geodesic.*

Then $\{x_i\}$ is within $O(1)$ of a flat.

Proof Write $\{\pi_A(x_i)\} = L$. Let $\{x_{1_j}\}$, and $\{x_{2_{j'}}\}$ be two subsequences such that their π_A images are two parallel lines, and let l_1, l_2 be two geodesics within $O(1)$ of $\{x_{1_j}\}$ and $\{x_{2_j}\}$ respectively. We note that to every x_{1_j}, there is an x_{2_j} such that $\pi_A(x_{2_j})$ is closest to $\pi_A(x_{1_j})$ amongst $\{\pi_A(x_{2_j})\}$. Furthermore, there is a line in L containing $\pi_A(x_{1_j})$ and $\pi_A(x_{2_j})$, and by assumption, this means that there is a geodesic $\hat{l}_j$ (whose direction is the same for all j) within $O(1)$ of x_{1_j} and x_{2_j}. Now, if l_1 and l_2 are not in the same flat, then there must be a root Ξ such that $\pi_\Xi(l_1)$ and $\pi_\Xi(l_2)$ fork out with respect to the orientation we previously fixed on l_1 and l_2. When x_{1_j}, x_{2_j} are far away from the fork point, this causes a contradiction because the existence of $\hat{l}$ mean that $\pi_\Xi(x_{1_j})$ and $\pi_\Xi(x_{2_j})$ can be connected by a vertical geodesic in V_Ξ, but they lie far away from the forking point of two geodesics. Therefore if $\{x_{i_j}\}$ is a subsequence for which their π_A image is a affine 2–subspace, then $\{x_{i_j}\}$ lie within $O(1)$ of an affine 2–subspace in a flat.

Now suppose whenever $\{x_{i_j}\}$ is a subset whose π_A image is an affine I–subspace, $\{x_{i_j}\}$ is within $O(1)$ of a flat. Let $\{x_{1_j}\}$, $\{x_{2_j}\}$ be two subsets such that $\{\pi_A(x_{1_j})\}$, $\{\pi_A(x_{2_j})\}$ are two parallel I–hyperplane, and h_1, h_2 be two affine I subspace within $O(1)$ of $\{x_{1_j}\}$ and $\{x_{2_j}\}$ respectively.

Then we know for every x_{1_j}, there is a x_{2_j} such that there is a line in L containing $\pi_A(x_{1_j})$ and $\pi_A(x_{2_j})$, so by assumption, x_{1_j}, x_{2_j} are within $O(1)$ of a (straight) geodesic $\hat{l}_j$. Furthermore we can assume without loss of generality that the direction of $\hat{l}_j$ are the same for all j. Therefore if h_1 and h_2 lie on two different flats, then for some root Ξ, $\pi_\Xi(h_1)$ and $\pi_\Xi(h_2)$ are two vertical geodesic that fork apart. Therefore for x_{1_j}, x_{2_j} such that $\pi_\Xi(x_{1_j})$, $\pi_\Xi(x_{2_j})$ lying very far from the fork point where $\pi_\Xi(h_1)$ and $\pi_\Xi(h_2)$ diverge from each other, this is a contradiction to the existence of $\hat{l}_j$ within $O(1)$ of x_{1_j} and x_{2_j}, for the latter would imply that $\pi_\Xi(x_{1_j})$ and $\pi_\Xi(x_{2_j})$ are within $O(1)$ of a vertical geodesic in V_Ξ. □

Proof of Proposition 5.2.1 Let L be a good grid in $\mathbf{A}'$ under a group isomorphism ψ. We can further make sure that for each basis $e_i \in \mathbb{Z}^n$, $|\psi(e_i) - \psi(\vec{0})| = 4M_0$, where M_0 is as in Corollary 5.2.3. The same corollary also implies that for every point $\mathbf{b} \in L$, the subset $s_\mathbf{b} = \{x \in \phi(\mathcal{F}) : \pi_A(x) = \mathbf{b}\}$ of G' is contained in a ball of radius at most M_0. Therefore for distinct points $\mathbf{b}, \mathbf{d} \in L$, $s_\mathbf{b} \cap s_\mathbf{d} = \varnothing$. Furthermore, modifying ϕ a bounded amount if necessary, we can make sure that $s_\mathbf{b} \neq \varnothing$ for all $\mathbf{b} \in L$.

Let $\{x_j\} \subset \phi(\mathcal{F})$ be a subset such that $x_j \in s_\mathbf{b}$ for some $\mathbf{b} \in L$. We choose x_j's so that if $j \neq i$, $x_j \in s_\mathbf{b}$, $x_i \in s_\mathbf{c}$ for $\mathbf{b} \neq \mathbf{c}$. Such choice of $\{x_j\}$ is said to be adapted to the grid L. Now let $\{x_{i_j}\}$ be a subsequence whose π_A images is a line in L, and $\{y_{i_j}\} \subset \mathcal{F}$ be their ϕ–preimages in $\mathcal{F}$. Then by Equation (17), $d(y_{i_j}, y_{i_{j+1}}) \leq \hat{C}(L)$, for all j, for some constant $\widehat{C}$ depending on L. This means that $d(x_{i_j}, x_{i_{j+1}}) \le 2\kappa\widehat{C}$ for all j. Applying (i) of [12, Lemma 3.3.3] to $\{x_{i_j}\}$ shows that the sequence is within $O(\widehat{C})$ of a (bi-infinite) geodesic. The desired claim now follows by applying Lemma 5.2.5 to all possible choices of $\{x_j\}$ adapted to the grid L. □

5.3 Consequences of flats go to flats

In this section we will see that a quasi-isometry of G induces quasi-similarities among the $(\partial_{[\alpha]}, D_{[\alpha]})$'s.

5.3.1 Definition We call a map $F\colon X \to Y$ between metric spaces (N, K)*–quasi-similarities* if

$$N/Kd(x,y) \le d(F(x),F(y)) \le NKd(x,y).$$

When $K = 1$, the map is called a *similarity*. We write the group of quasi-similarities of a space by $\mathrm{QSim}(X)$.

5.3.2 Definition Let $V = V_1 \supset V_2 \supset \cdots \supset V_s$ be a finite filtration of subspaces in a vector space V. A map $f\colon V \to V$ is called an *almost translation* with respect to the filtration $\{V_i\}$ if it sends left cosets of V_i to left cosets of V_i (ie respects the foliation structures induced by the filtration) and induces a translation map on each V_i/V_{i+1}.

If one wants to express in terms of coordinates, then by choosing a basis in each $V_i \backslash V_{i+1}$, an almost translation is a map of the form

$$(\vec{x}_1, \vec{x}_2, \dots, \vec{x}_s) \mapsto (\vec{x}_1 + B_1, \vec{x}_2 + B_2(\vec{x}_1), \vec{x}_3 + B_3(\vec{x}_1, \vec{x}_2), \dots, \vec{x}_s + B_s(\vec{x}_1, \vec{x}_2, \dots, \vec{x}_{s-1})),$$

where the B_i's are functions.

5.3.3 Remark We will use Definition 5.3.2 to the subspaces $V_{[\Xi]} = \bigoplus_{\alpha\in[\Xi]} V_\alpha$ corresponding to each root class $[\Xi]$. Since roots of the same root class are all positive multiples of each other, there is a well-defined linear order on them which in turn induces a filtration in $V_{[\Xi]}$. More explicitly, let $\alpha_1 > \alpha_2 > \cdots$ be the linear order on the roots in $[\Xi]$, then we have the filtration

$$\bigoplus_i V_{\alpha_i} \supset \bigoplus_{i\ge 2} V_{\alpha_i} \supset \bigoplus_{i\ge 3} V_{\alpha_i} \supset \cdots.$$

5.3.4 Definition An *almost similarity* is a composition of a similarity followed by an almost translation. Later we will use the notation $\mathrm{ASim}(\partial^-_{[}\Xi])$ to denote the group of almost similarities whose almost translations are with respect to the filtration on $V_{[\Xi]}$ as in Remark 5.3.3.

5.3.5 Proposition *Let $\psi\colon G = \mathbf{H} \rtimes_{\varphi} \mathbf{A} \to G' = \mathbf{H}' \rtimes_{\varphi'} \mathbf{A}'$ be a quasi-isometry such that the image of left translate of $\mathbf{A}$ is within $O(1)$ neighborhood of a left translate of $\mathbf{A}'$.*

(i) *There is a bijection σ between the root classes of G and the root classes of G'.*

(ii) *ψ sends foliations by $[\Xi]$ root kernel of G to the foliations by $\sigma([\Xi])$ root kernel of G', for all root classes $[\Xi]$.*

(iii) *ψ sends foliations by $[\Xi]$ root class horocycles of G to foliations by $\sigma([\Xi])$ root class horocycles of G', for all root classes $[\Xi]$.*

(iv) *ψ is within $O(1)$ of a product map $f \times g$, where $f\colon \mathbf{A} \to \mathbf{A}'$ is affine whose linear part is some scalar multiple of a finite order element of $O(n)$; while $g = (g_1, g_2, \dots, g_{s_0})$, each g_i is a bilipschitz map from $V_{[\alpha]}$ to $V_{\sigma([\alpha])}$.*

(v) *For each root class $[\Xi]$, if we list the roots $\xi_1 < \xi_2 < \xi_3 < \dots < \xi_l$, then f induces a map on the roots such that $f_*\xi_1 < f_*\xi_2 i < \dots < f_*\xi_l$, where $\{f_*\xi\} = \sigma([\Xi])$. Furthermore, with respect to this order of roots, $g|_{V_{[\Xi]}}\colon V_{[\Xi]} \to V_{f_*[\Xi]}$ respects the filtration induced by the ordering of the roots as well as the degree of nilpotency induced by the Jordan blocks within each eigenspace. In coordinates,*

$$(\vec{x}_1, \vec{x}_2, \dots, \vec{x}_l) \mapsto (g_1(\vec{x}_1, \vec{x}_2, \dots, \vec{x}_l), \dots, g_{l-1}(\vec{x}_{l-1}, \vec{x}_l), g_l(\vec{x}_l)),$$

where $\vec{x}_j \in V_{\xi_j}$ such that for each i, any $x_j \in V_{\xi_j}$ for $j > i$, the map

$$\bullet \mapsto g_i(\bullet, x_{i+1}, \dots, x_l)$$

on V_{ξ_i}, is a $(e^{f_*\xi_i(\mathbf{t}_0)}, 2\kappa e^c)$ quasi-similarity, where $\mathbf{t}_0$ is the constant part of f, and c is a constant depending on ψ.

(vi) *ψ induces an quasi-similarity between root boundaries of G and G'.*

Proof Since two flats are within a finite Hausdorff distance of each other if and only if they are the same flat, our assumption means that a flat $\hat{\mathcal{F}}$ within $O(1)$ Hausdorff neighborhood of $\psi(\mathcal{F})$ for $\mathcal{F}$ a flat, is unique. Moreover, as

$$d_H(\phi(\partial(\mathcal{F}_1 \cap \mathcal{F}_2)), \partial(\hat{\mathcal{F}}_1 \cap \hat{\mathcal{F}}_2)) \le M_0,$$

and two flats intersect at a set bounded by hyperplanes parallel to root kernels, claim (i) and (ii) follow since any root kernels have codimension 1 in a flat and hyperplanes parallel to two distinct root kernels must intersect.

To see (iii), we consider two cases depending on the number root kernels in G. In the case that G has $\mathrm{rank}(G)$ distinct root kernels, we take two points p, q in the same root class horocycle, say of root class $[\alpha]$, and build a quadrilateral Q with p, q as two if its vertices (see [12, Definition 4.1.5]) in the direction $\bigcap_{[\Xi]:[\Xi]\neq[\alpha]} \ker([\Xi])$.

As each one of the hyperplanes parallel to a root kernel $[\Xi]$ arise as the intersection between two flats, the image of Q is close to another quadrilateral $\widehat{Q}$ whose direction is parallel to $\bigcap_{[\Xi]:[\Xi]\neq\sigma([\alpha])}\ker([\Xi])$, and this shows that $\psi(p),\psi(q)$ lies within $O(1)$ neighborhoods of a $\sigma([\alpha])$ root horocycle.

In the case that G has more than $\operatorname{rank}(G)$ many distinct root kernels, we already know that the restriction of ψ on each flat has to be affine since it has to preserve at least $\operatorname{rank}(G)+1$ many parallel families of hyperplanes, and in particular, geodesics go to bounded neighborhoods of geodesics. So for two points p,q in $[\alpha]$ root horocycle, we build quadrilaterals using finitely many (and actually a positive measure many of them) nonsingular vectors $\vec{u}_i\in\mathbf{A}$ such that $\bigcap W^{*_i}_{\vec{u}_i}=V_{[\alpha]}$, where $*_i\in\{+,-\}$, depending on whether $V_{[\alpha]}$ expands or contracts with respect to $\vec{u}_i$. Since geodesics are roughly preserved, the images of those quadrilaterals with p,q as part of the vertices are also roughly preserved, and are bounded distances from quadrilaterals in the direction $\vec{v}_i$. By choosing sufficiently many $\vec{u}_i$ we can ensure that $\bigcap W^{*_i}_{\vec{v}_i}=V_{[\beta]}$ for some root class $[\beta]$, which will have to be $\sigma([\alpha])$ by (i) and (ii).

This also shows that left translate of $\mathbf{H}$ are taken to bounded neighborhoods of left translates of $\mathbf{H}$, and it follows that the restrictions of ψ to flats based at different points in $\mathbf{H}$ are the same. Denote this common map by f. We need to show that f is affine in the case that G has $\operatorname{rank}(G)$ many root kernels. To this end we observe that left translate of $\ker(\alpha_i)$ in G (resp. left translates of $\ker(\beta_i)$ in G') can be identified with a rank 1 space $G_{\alpha_i}=\mathbf{H}\rtimes\mathbb{R}(\vec{v}_{\alpha_i})$ (resp. $G'_{\beta_i}=\mathbf{H}'\rtimes\mathbb{R}(\vec{v}_{\beta_i})$), and so (ii) implies that ψ induces a quasi-isometry from G_{α_i} to G'_{β_i} that sends left translates of $\mathbf{H}$ to $O(1)$ neighborhood of left translates of G'_{β_i}. By Proposition 5.8 of Farb and Mosher [6], we conclude that the f–induced map from $\mathbf{A}/\ker(\alpha_i)$ to $\mathbf{A}'/\ker(\beta_i)$ is bounded distance from an affine map. Since this is true for all root classes, it follows that f itself is bounded distance from an affine map. So now we know that regardless of the number of root kernels, ψ splits into $f\times g$, where $f\colon\mathbf{A}\to\mathbf{A}'$ affine and respects root kernels, while $g\colon\mathbf{H}\to\mathbf{H}'$ takes root class horocycles to root class horocycles. Furthermore, the permutation on root classes induced by f and g agree.

We now proceed to show that the f actually induces a bijection between roots of G and G' (not just root classes). Since we know now the ψ restricted to any flat is the map $\mathbf{t}\mapsto M(\mathbf{t})+\mathbf{t}_0$. This means that (straight) geodesics are taken to straight geodesics, and we can compare the rate of divergence between two geodesics in the same direction but at based at different points of $\mathbf{H}$.

Specifically, take Ξ a root, and let p,q be two points on a common Ξ horocycle. Pick some $\vec{v}\in\mathbf{A}$, and let $l_x=x(t\vec{v})$, $l_y=y(t\vec{v})$ be two geodesic rays in direction $\vec{v}$

leaving x, y respectively. Then $\psi(l_x)$, $\psi(l_y)$ are within $O(1)$ of $l_{x'} = x'(tM(\vec{v}))$ and $l_{y'} = y'(tM(\vec{v}))$ respectively, where x', y' are within $O(1)$ of a left translate of $V_{[\beta]}$.

The rate of divergence between l_x and l_y is $P(t\,\Xi(\vec{v}))e^{t\,\Xi(\vec{v})}$, for some polynomial P; the divergence rate between $l_{x'}$ and $l_{y'}$ is $Q(t\xi(M(\vec{v})))e^{t\xi(M(\vec{v}))}$, for some $\xi \in [\beta]$, and some polynomial Q. Since the two rates are coarsely equivalent to each other, there must be some $\widetilde{\beta} \in [\beta]$ such that $\Xi(\vec{v}) = \widetilde{\beta}(M(\vec{v}))$, and that the degree of $P(t\,\Xi(\vec{v}))$ (as a function of t) and $Q(t\xi(M(\vec{v})))$ are the same. But $\vec{v}$ is arbitrary, so $\Xi = \widetilde{\beta} \circ M$, and degree of P is the same as degree of Q, and that x', y' are in the same left translate of $\bigoplus_{\xi \in [\beta] : \xi \leq \widetilde{\beta}} V_\xi$ (but not in any proper subspace). What this shows is that the quasi-isometry $f \times g$ not only sends root horocycles to root horocycles, but that it also respects the of filtrations induced by the Jordan blocks (ordered by the degree of nilpotency) in each root horocycle.

For $x, y \in V_\Xi$, write $g_\Xi(x), g_\Xi(y) \in V_{\sigma(\Xi)}$ for the $V_{\sigma(\Xi)}$ coordinate of $g(x), g(y)$, and for each $\xi < \sigma(\Xi)$, write $\widetilde{g}_\xi(x), \widetilde{g}_\xi(y) \in V_\xi$ for their V_ξ components, so that $g(x) = g_\Xi(x) + \sum_\xi \widetilde{g}_\xi(x)$, $g(y) = g_\Xi(y) + \sum_\xi \widetilde{g}_\xi(y)$. Pick a $\mathbf{t} \in \mathbf{A}$, then the distance between $(\mathbf{t}, x)$ and $(\mathbf{t}, y)$ with respect to path metric in $\mathbf{tH}$ is $e^{-\Xi(\mathbf{t})}|x-y|$. The ψ images of $(\mathbf{t}, x)$, $(\mathbf{t}, y)$ is c away from $(M\mathbf{t}+\mathbf{t}_0, g(x))$ and $(M\mathbf{t}+\mathbf{t}_0, g(y))$, so we have the inequality

$$\begin{aligned}\frac{1}{2\kappa}e^{-c}P(\Xi(\mathbf{t}))e^{-\Xi(\mathbf{t})}|x-y| &\leq \left(e^{-\sigma(\Xi)(M\mathbf{t}+\mathbf{t}_0)}Q(\sigma(\Xi)(M\mathbf{t}+\mathbf{t}_0))|g_\Xi(x)-g_\Xi(y)|\right) \\ &\quad + \left(\sum_{\xi<\sigma(\Xi)} e^{-\xi(M\mathbf{t}+\mathbf{t}_0)}Q_\xi(\xi(M\mathbf{t}+\mathbf{t}_0))|\widetilde{g}_\xi(x)-\widetilde{g}_\xi(y)|\right) \\ &\leq 2\kappa e^c P(\Xi(\mathbf{t}))e^{-\Xi(\mathbf{t})}|x-y|.\end{aligned}$$

Since $\sigma(\Xi) \circ M = \Xi$ and the degrees of P and Q are the same, dividing both sides by $e^{-\Xi(\mathbf{t})}P(\Xi(\mathbf{t}))$ and let $\Xi(\mathbf{t}) \to \infty$ produces

$$\frac{1}{2\kappa}e^{-c}e^{\sigma(\Xi)(\mathbf{t}_0)}|x-y| \leq |g_\Xi(x)-g_\Xi(y)| \leq 2\kappa e^c e^{\sigma(\Xi)(\mathbf{t}_0)}|x-y|.$$

So the restriction of $g|_{V_\Xi}\colon V_\Xi \to V_{\sigma(\Xi)}$ is a $(e^{\sigma(\Xi)(\mathbf{t}_0)}, 2\kappa e^c)$–quasi-similarity.

To summarize, ψ is $O(1)$ from a map of the form $(\mathbf{x}, \mathbf{t}) \mapsto (g(\mathbf{x}), mA_f(\mathbf{t})+\mathbf{t}_0)$, where $m > 0$, $A_f\colon \mathbf{A} \to \mathbf{A}'$ is a finite order element in $O(n)$ that preserves foliations by root kernels, while $g\colon \mathbf{H} \to \mathbf{H}'$ sends root horocycles to root horocycles and furthermore respects the graded foliations in each root horocycle.

Let σ be the permutation on root classes induced by M. Then as ψ sends negative $[\alpha]$ half planes to bounded neighborhood of negative $\sigma([\alpha])$ half planes, ψ induces a map from $\partial^-_{[\alpha]}$ to $\partial^-_{\sigma([\alpha])}$ for every root class $[\alpha]$ of G. Furthermore, the map

$\widehat{q}\colon \mathbb{R} \sim \mathbf{A}/\ker(\alpha_0) \to \mathbf{A}/\ker(\sigma(\alpha)_0) \sim \mathbb{R}$ is bounded distance from an affine map with linear term as m and constant term as $\sigma(\alpha)_0(\mathbf{t}_0)$. We now show the induced map on lower root boundaries is a quasi-similarity. Take $p, q \in \partial_{[\alpha]} \sim V_{[\alpha]}$. Then, up to quasi-symmetry we can take $D_{[\alpha]}(p,q)$ as

$$D_{[\alpha]}(p,q) = e^{m t_{p,q}}$$

and under ϕ, we have

$$D_{\sigma([\alpha])}(g(p), g(q)) = e^{\widehat{q}(t_{p,q})},$$

therefore

$$\begin{aligned} m t_{p,q} + \sigma(\alpha)_0(\mathbf{t}_0) - c &\le \widehat{q}(t_{p,q}) \le m t_{p,q} + \sigma(\alpha)_0(\mathbf{t}_0) + c \\ e^{m t_{p,q}} \frac{e^{\sigma(\alpha)_0(\mathbf{t}_0)}}{e^c} &\le e^{\widehat{q}(t_{p,q})} \le e^{m t_{p,q}} e^{\sigma(\alpha)_0(\mathbf{t}_0)} e^c \\ \frac{1}{e^c}(e^{\sigma(\alpha)_0(\mathbf{t}_0)} D_{[\alpha]}(p,q)) &\le D_{\sigma([\alpha])}(g(p), g(q)) \le e^c (e^{\sigma(\alpha)_0(\mathbf{t}_0)} D_{[\alpha]}(p,q)). \end{aligned}$$ □

5.3.6 Theorem *Let G, G' be a nondegenerate, unimodular, split abelian-by-abelian solvable Lie group, and $\phi\colon G \to G'$ a (κ, C) quasi-isometry. Then ϕ is bounded distance from a composition of a left translation followed by a map of the form $(\mathbf{x}, \mathbf{t}) \to (g(\mathbf{x}), f(\mathbf{t}))$, where f is affine whose linear part is a positive of a finite order element $A_f \in O(n)$ (n is the rank of G) that preserves foliations by root kernels, while $g = (g_1, g_2, \dots, g_\sharp)$, g_i is a bilipschitz map from $V_{[\alpha_i]}$ to $V_{[\alpha_i]}$ with bilipschitz constants depending only on κ, C.*

Proof This follows from Proposition 5.2.1 and Proposition 5.3.5. □

5.3.7 Corollary *If two nondegenerate, unimodular, split abelian-by-abelian solvable Lie groups $G = \mathbf{H} \rtimes_\varphi \mathbf{A}$, $G' = \mathbf{H}' \rtimes_{\varphi'} \mathbf{A}'$ are quasi-isometric, then there is an isomorphism $f\colon \mathbf{A} \to \mathbf{A}'$ such that φ and $\varphi' \circ f$ have the same absolute Jordan form.*

5.3.8 Corollary *Let $G = \mathbf{H} \rtimes_\varphi \mathbf{A}$ be a nondegenerate, unimodular, split abelian-by-abelian group such that $\varphi(\mathbf{A})$ is diagonalizable, while $G' = \mathbf{H}' \rtimes_{\varphi'} \mathbf{A}'$ is another nondegenerate, unimodular, split abelian-by-abelian group such that $\varphi'(\mathbf{A}')$ is not diagonalizable. Then G and G' are not quasi-isometric.*

Proof If there were, then Theorem 5.3.6 implies that geodesics are taken to geodesics and the induced height function on the geodesics are affine, which means that we can compare rates of divergence between two geodesics in the same direction. In G', we would detect exponential polynomial growth while in G, only exponential growth can be detected, and those two growth types cannot be related via a quasi-isometry. □

Uniform subgroups of quasi-similarities of the root boundaries are analyzed in Dymarz [2] and Dymarz and Peng [3], and we can now say something about an arbitrary finitely generated group in its quasi-isometric class.

5.3.9 Corollary *Let $G = \mathbf{H} \rtimes_{\psi} \mathbf{A}$ be a nondegenerate, unimodular, split abelian-by-abelian solvable Lie group. If Γ is a finitely generated group quasi-isometric to G, then Γ is virtually polycyclic.*

Proof Let $\varphi \colon \Gamma \to G$ be a (κ, C) quasi-isometry. For each $\gamma \in \Gamma$, write L_γ for the left translation of γ, and $\widetilde{L}_\gamma = \varphi \circ L_\gamma \circ \varphi^{-1}$. Then $\widetilde{\Gamma} = \{\widetilde{L}_\gamma\}_{\gamma \in \Gamma}$ constitute an uniform subgroup of $\mathrm{QI}(G)$, all with the same quasi-isometry constants (they are all (κ, C) quasi-isometries).

By Theorem 5.3.6, each $\widetilde{L}_\gamma$ induces a permutation on root classes. Therefore the map from $\widetilde{\Gamma}$ into the permutations on root classes of G is a well-defined homomorphism, whose kernel, $\widetilde{\Gamma}_0$ is finite index in $\widetilde{\Gamma}$.

Since $\widetilde{\Gamma}$ is a uniform subgroup of $\mathrm{QI}(G)$, by Proposition 5.3.5 $\widetilde{\Gamma}_0$ is a uniform subgroup of $\prod_{[\alpha]} \mathrm{QSim}(\partial^-_{[\alpha]})$. Applying Theorem 2 in Dymarz [2] and Dymarz and Peng [3] to the image of $\widetilde{\Gamma}_0$ in each $\mathrm{QSim}(\partial^-_{[\alpha]})$ factor, we can conjugate $\widetilde{\Gamma}_0$ into $\prod_{[\alpha]} \mathrm{ASim}(\partial^-_{[\alpha]})$, the group of almost similarities. Denote the image of $\widetilde{\Gamma}_0$ in $\prod_{[\alpha]} \mathrm{ASim}(\partial^-_{[\alpha]})$ by $\widehat{\Gamma}_0$.

For each $\widetilde{L}_\gamma \in \widetilde{\Gamma}_0$ write $g_{[\Xi],\gamma}$ for the almost similarity on $(\partial^-_{[\Xi]}, D_{[\Xi]})$ and $t_{[\Xi],\gamma}$ the corresponding similarity constant, as induced by the image of $\widehat{L}_\gamma$ in $\widehat{\Gamma}_0$.

Claim *For each $\widetilde{L}_\gamma \in \widetilde{\Gamma}_0$, there is a $\mathbf{s}_\gamma \in \mathbf{A}$ such that $t_{[\Xi],\gamma} = e^{\Xi_0(\mathbf{s}_\gamma)}$ for each root class $[\Xi]$.*

We know that for each $\widetilde{L}_\gamma$, there is a $\mathbf{t}_{0,\gamma} \in \mathbf{A}$ such that $\widetilde{L}_\gamma$ induces $(e^{\Xi_0(\mathbf{t}_{0,\gamma})}, e^c)$ quasi-similarity $\widetilde{g}_{[\Xi],\gamma}$ on $(\partial_{[\Xi]}, D_{[\Xi]})$. Theorem 2 of Dymarz [2] and Dymarz and Peng [3] says that we can find $F_{[\Xi]} \in \mathrm{Bilip}(\partial_{[\Xi]}, D_{[\Xi]})$ with bilipschitz constant K' such that for every root class $[\Xi]$ and every $\widetilde{L}_\gamma \in \widetilde{\Gamma}_0$,

$$F_{[\Xi]} \widetilde{g}_{[\Xi],\gamma} F^{-1}_{[\Xi]} = g_{[\Xi],\gamma}.$$

Therefore $t_{[\Xi],\gamma}$, the similarity constant of $g_{[\Xi],\gamma}$ satisfies

$$t_{[\Xi],\gamma} \in \Big[e^{\Xi_0(\mathbf{t}_{0,\gamma})} \frac{1}{e^c K'^2}, e^{\Xi_0(\mathbf{t}_{0,\gamma})} e^c K'^2\Big] \quad \text{for all root classes } [\Xi]. \tag{18}$$

Since the sum of roots is zero and

$$0 = \sum_{\alpha \text{ roots}} \alpha = \sum_{[\Xi]} \Xi_0 \frac{l_{[\Xi]}}{\Xi_0},$$

it follows that

$$\prod_{[\Xi]}(e^{\Xi_0(\mathbf{t}_{0,\gamma})})^{l_{[\Xi]}/\Xi_0}=1. \tag{19}$$

Therefore

$$\prod_{[\Xi]}(t_{[\Xi],\gamma})^{l_{[\Xi]}/\Xi_0}=1 \tag{20}$$

because the ratio of the left hand sides of (19) and (20) lies in an interval those end points are constants independent of γ, so if the left hand side of (20) was not 1, then the ratio of left hand sides of (19) and (20) for sufficiently high powers of γ would escape the interval.

On the other hand, for a generic linear functional ℓ, we know that

$$0=\sum_{[\Xi]}l_{[\Xi]}(\vec{v}_\ell)=\sum_{[\Xi]}\Xi_0(\vec{v}_\ell)\frac{l_{[\Xi]}}{\Xi_0}=\sum_{[\Xi]}\Xi_0(\mathbf{t}_{0,\gamma})\frac{\Xi_0(\vec{v}_\ell)}{\Xi_0(\mathbf{t}_{0,\gamma})}\frac{l_{[\Xi]}}{\Xi_0},$$

and so

$$\begin{aligned}\prod_{[\Xi]}(1/e^cK'^2)^{\Xi_0(\vec{v}_\ell)/\Xi_0(\mathbf{t}_{0,\gamma})l_{[\Xi]}/\Xi_0}&\le\prod_{[\Xi]}(t_{[\Xi],\gamma})^{\Xi_0(\vec{v}_\ell)/\Xi_0(\mathbf{t}_{0,\gamma})l_{[\Xi]}/\Xi_0}\\&\le\prod_{[\Xi]}(e^cK'^2)^{\Xi_0(\vec{v}_\ell)/\Xi_0(\mathbf{t}_{0,\gamma})l_{[\Xi]}/\Xi_0}.\end{aligned} \tag{21}$$

Because $g_{[\Xi],\gamma}$ is an almost similarity, $t_{[\Xi],\gamma^n}=t^n_{[\Xi],\gamma}$. By (18) we must then have $\mathbf{t}_{0,\gamma^n}=n\mathbf{t}_{0,\gamma}+\vec{u}_n$, with $\|\vec{u}_n\|=o(n)$. So now (21) becomes

$$\begin{aligned}\prod_{[\Xi]}\left(\frac{1}{e^cK'^2}\right)^{O(1)(\Xi_0(\vec{v}_\ell)/(n\Xi_0(\mathbf{t}_{0,\gamma})))(l_{[\Xi]}/\Xi_0)}&\\&\le\prod_{[\Xi]}(t_{[\Xi],\gamma})^{\left(\Xi_0(\vec{v}_\ell)/\Xi_0(\mathbf{t}_{0,\gamma})\right)(l_{[\Xi]}/\Xi_0)}\\&\le\prod_{[\Xi]}(e^cK'^2)^{O(1)\left(\Xi_0(\vec{v}_\ell)/(n\Xi_0(\mathbf{t}_{0,\gamma}))\right)(l_{[\Xi]}/\Xi_0)}.\end{aligned}$$

As the middle term doesn't depend on n, we must have

$$\prod_{[\Xi]}(t_{[\Xi],\gamma})^{(\Xi_0(\vec{v}_\ell)/\Xi_0(\mathbf{t}_{0,\gamma}))(l_{[\Xi]}/\Xi_0)}=1. \tag{22}$$

By letting ℓ ranging over a subset of positive measure, Equation (22) and (20) means that the $t_{[\Xi],\gamma}$ must be of the form $e^{\Xi_0(\mathbf{s}_\gamma)}$ for some $\mathbf{s}_\gamma\in\mathbf{A}$.

This means that for each $\widetilde{L}_\gamma \in \widetilde{\Gamma}_0$, $\{g_{[\Xi],\gamma}\}_{[\Xi]}$ determines an element $\psi_\gamma \in \mathrm{QI}(G)$ of the form

$$\psi_\gamma((\mathbf{x}_{[\Xi]})_{[\Xi]}, \mathbf{t}) = \big((g_{[\Xi],\gamma}(\mathbf{x}_{[\Xi]}))_{[\Xi]}, \mathbf{t} + \mathbf{s}_\gamma\big)$$

and we can define a homomorphism $h\colon \widehat{\Gamma}_0 \to \mathbf{A}$ as $\gamma \mapsto \mathbf{s}_\gamma$.

The kernel of h consist of elements with no translations, so they leave the subgroup $\mathbf{H}$ invariant.

Since Γ is quasi-isometric to G, the quasi-action of $\widetilde{\Gamma}_0$ on G is proper, which means $\widehat{\Gamma}_0$ and $\ker(h)$ quasi acts properly on $\prod_{[\Xi]}(\partial^-_{[\Xi]}, D_{[\Xi]})$. By combining Theorem 18 of Dymarz [2] and Theorem 11 in Dymarz and Peng [3], we conclude that Γ is virtually polycyclic. □

In a group G, an element $x \in G$ is called *exponentially distorted* if there are numbers c, ϵ such that for all $n \in \mathbb{Z}$,

$$\frac{1}{c}\log(|n|+1) - \epsilon \le \|x^n\|_G \le c\log(|n|+1) + \epsilon,$$

where $\|x^n\|_G$ is the distance between the identity and x^n in G. In the case of a connected, simply connected solvable Lie group G, Osin showed in [11] that the set of exponentially distorted elements forms a normal subgroup $R_{\exp}(G)$.

5.3.10 Lemma *Let G be a connected, simply connected solvable Lie group such that*

$$1 \to R_{\exp}(G) \to G \to \mathbb{R}^s \to 1,$$

where $R_{\exp}(G)$ is abelian. Then the above sequence splits and G is a semidirect product of $R_{\exp}(G)$ and $\mathbb{R}^s$.

Proof Let $\mathfrak{h}$ be a Cartan subalgebra of $\mathfrak{g}$, the Lie algebra of G. Then $\mathfrak{v}$, the Lie algebra of $R_{\exp}(G)$, is generated by root spaces in the decomposition of $\mathfrak{g}$ with respect to $\mathfrak{h}$. Since this is abelian, it means that $\mathfrak{g}$ is a semidirect product of $\mathfrak{h}$ and $\mathfrak{v}$. Since $\mathfrak{g}/\mathfrak{v}$ is abelian, $\mathfrak{h}$ is abelian. □

5.3.11 Corollary *Let $G = \mathbf{H} \rtimes_\psi \mathbf{A}$ be a nondegenerate, unimodular, split abelian-by-abelian solvable Lie group where ψ is diagonalizable. If Γ is a finitely generated group quasi-isometric to G, then Γ is virtually a lattice in a unimodular semidirect product of $\mathbf{H}$ and $\mathbf{A}$.*

Proof By Corollary 5.3.9, Γ contains a finite index subgroup that is polycyclic. By a theorem of Mostow (see Theorem 4.28 of Raghunathan [13]) which says that a polycyclic group contains a finite index subgroup that embeds as a lattice in a connected, simply connected Lie group, we have $\mathcal{L}$, a connected, simply connected solvable Lie group The crux of the proof consists of showing that $\mathcal{L}$ satisfies the short exact sequence in Lemma 5.3.10, and the argument is practically that of Section 4.3 in [2] with minor modifications. We reproduce the skeleton of the proof below, and refer the readers to the relevant sections in [2] for details.

We are now going to construct a continuous homomorphism $\widetilde{h}\colon \mathcal{L} \to \mathbf{A} = \mathbb{R}^n$ that is onto, whose Lie group kernel is not only quasi-isometric to $\mathbf{H}$, but also coincides with the exponential radical of $\mathcal{L}$. As any Lie group in the same quasi-isometric class as $\mathbb{R}^n$ must be virtually $\mathbb{R}^n$, by Lemma 5.3.10, $\mathcal{L}$ is virtually a semidirect product of $\mathbf{A}$ and $\mathbf{H}$. If this semidirect product were not unimodular, then $\mathcal{L}$ would be nonamenable, which is a contradiction because amenability is preserved under quasi-isometry.

The $\widetilde{h}$ is going to be the composition of the following three homomorphisms:

- $A\colon \mathcal{L} \to \prod_{[\alpha]} \mathrm{QSim}(\partial^-_{[\alpha]})$
- B, conjugation of a uniform subgroup of $\prod_{[\alpha]} \mathrm{QSim}(\partial^-_{[\alpha]})$ into $\mathrm{AIsom}(G)$, where $\mathrm{AIsom}(G)$ consists of maps of the form
$$\psi_\gamma((\mathbf{x}_{[\Xi]})_{[\Xi]}, \mathbf{t}) = \big((g_{[\Xi],\gamma}(\mathbf{x}_{[\Xi]}))_{[\Xi]}, \mathbf{t} + \mathbf{s}_\gamma\big)$$
- $C\colon h\colon \prod_{[\alpha]} \mathrm{AIsom}(\partial^-_{[\alpha]}) \to \mathbf{A} = \mathbb{R}^n$

The homomorphism A We can assume without loss of generality, that Γ itself is a lattice in $\mathcal{L}$. We start with the following construction which can be found in Section 3.2 of Furman [7]. Choose some open subset $E \subset \mathcal{L}$ with compact closure, such that $\mathcal{L}$ is the union of left translates of E by Γ. Also fix a function $p\colon \mathcal{L} \to \Gamma$ such that $x \in p(x)E$ for every $x \in \mathcal{L}$. Then by defining $q_h\colon \Gamma \to \Gamma$ as $q_h(\gamma) = p(h\gamma)$ for every $h \in \mathcal{L}$, we obtain a homomorphism from $\mathcal{L}$ into $\mathrm{QI}(\Gamma)$. Since Γ is quasi-isometric to G, conjugating by the quasi-isometry between Γ and G, we obtain a homomorphism from $\mathcal{L}$ into $\mathrm{QI}(G)$, where the images have uniform quasi-isometric constants. By Proposition 5.3.5(v), we can realize $\mathrm{QI}(G)$ as a subgroup of $\prod_{[\alpha]} \mathrm{QSim}(\partial^-_{[\alpha]})$. By passing to a finite index subgroup of $\mathcal{L}$ if necessary, we now have the homomorphism A from $\mathcal{L}$ to $\prod_{[\alpha]} \mathrm{QSim}(\partial^-_{[\alpha]})$, whose image is a uniform subgroup of quasi-similarities. Continuity of the homomorphism A follows from Proposition 26 of Dymarz [2] where continuity in each factor was obtained.

The homomorphism B Theorem 2 of Dymarz [2] and Dymarz and Peng [3] says that we can conjugate the image of the homomorphism A into $\prod_{[\alpha]} \mathrm{ASim}(\partial^-_{[\alpha]})$. That elements of $\prod_{[\alpha]} \mathrm{ASim}(\partial^-_{[\alpha]})$ can be realized as elements of $\mathrm{AIsom}(G)$ follows from the Claim in the proof of Corollary 5.3.9. The homomorphism B is continuous because conjugation is continuous.

The homomorphism C The definition of $\mathrm{AIsom}(G)$ means that we have a well-defined homomorphism into $\mathbf{A} = \mathbb{R}^n$, which is our homomorphism C. Now if q_i is a sequence in $\mathrm{AIsom}(G)$ approaching to identity, then the map each one of them induces on the $\mathbf{A}$ factor also has to approach that of what the identity does. Since the identity map produces no change in the $\mathbf{A}$ factor, it follows that the image of q_i's under the homomorphism C approaches $\vec{0} \in \mathbf{A}$.

Since Γ is quasi-isometric to G, the quasi-action (conjugating each left translation of Γ by the quasi-isometry between Γ and G gives a quasi-action on G) of Γ, and therefore $\mathcal{L}$, on G is cobounded, it follows that $\widetilde{h}$ must be onto because it is continuous.

We now claim that $R_{\exp}(\mathcal{L}) = \ker(\widetilde{h})$. To this end, we need the following from [11].

5.3.12 Lemma [11, Lemma 2.1] *Suppose G, H are locally compact groups generated by some symmetric compact neighborhoods of the identities, $\|\cdot\|_G$, $\|\cdot\|_H$ are canonical norms on G and H, and dist_G, dist_H are the induced metrics. Assume $\phi\colon G \to H$ is a continuous surjective homomorphism, then there is a constant K such that*

$$\mathrm{dist}_H(\phi(g_1), \phi(g_2)) \le K\, \mathrm{dist}_G(g_1, g_2).$$

Applying the lemma above to $\widetilde{h}$ gives us that

$$\|\widetilde{h}(\gamma)\| \le K|\gamma|_{\mathcal{L}} \quad \text{for all } \gamma \in \mathcal{L}.$$

Now let $\gamma \in R_{\exp}(\mathcal{L})$ such that $|\widetilde{h}| = c$. Then for any n,

$$cn = |\widetilde{h}(\gamma^n)| \le K|\gamma^n|_{\mathcal{L}} = K \log(n+1).$$

So we must have $\widetilde{h} = \vec{0}$, hence $R_{\exp}(\mathcal{L}) \subset \ker(\widetilde{h})$.

de Cornulier showed in [1] that for a connected, simply connected solvable Lie group X, the *asymptotic dimension*, defined as the dimension of $X/R_{\exp}(X)$ is a quasi-isometric invariant. This means that

$$\dim \mathcal{L}/R_{\exp}(\mathcal{L}) = \dim G/R_{\exp}(G) = \dim \mathbf{A} = n.$$

However as $\widetilde{h}$ is onto, the dimension of $\mathcal{L}/\ker(\widetilde{h})$ also equals n. So $\ker(\widetilde{h})$ cannot be strictly bigger than $R_{\exp}(\mathcal{L})$.

By construction, $\mathcal{L}$ quasi-acts properly on G as a uniform group of quasi-isometries, which means $\ker(\widetilde{h})$ quasi-acts properly on $\mathbf{H}$ as a uniform group of quasi-similarities, so $\ker(\widetilde{h})$ is finitely generated by Proposition 20 in [2]. Fix a $p \in G$. Then $\gamma \mapsto B \circ A(\gamma)(p)$ is a quasi-isometry from $\mathcal{L}$ to G. Here B, A refers to the homomorphisms mentioned above. The restriction of this map to $\ker(\widetilde{h}) = R_{\exp}(\mathcal{L})$ produces a quasi-isometric embedding of $R_{\exp}(\mathcal{L})$ into $\mathbf{H}$. However since the cohomological dimension is a quasi-isometry invariant by Gersten [8], the dimension of $R_{\exp}(\mathcal{L})$ must equal that of $\mathbf{H}$. Now by Theorem 7.6 of Farb and Mosher [6], this embedding must be coarsely onto, which means $R_{\exp}(\mathcal{L})$ is quasi-isometric to $\mathbf{H}$, so $R_{\exp}(\mathcal{L})$ must be virtually $\mathbf{H}$ since the latter is abelian. □

An unimodular solvable Lie group not quasi-isometric to any finitely generated groups Here we reproduce from Hasegawa [9], an example of a simply connected, unimodular solvable Lie group that has no lattices. (The example is not present in the journal version [10] of [9].)

Let G be a semidirect product between $\mathbb{R}$ and the (3–dimensional) Heisenberg group, where the $\mathbb{R}$ acts on the Heisenberg group by the diagonal matrix with entries e^{1t}, e^{1t}, e^{-2t}. It has no lattice because if an elements of $SL_3(\mathbb{Z})$ has two distinct eigenvalues and one of them with multiplicity 2 then they have to be $-1, -1, 1$. (See Lemma 2.2 of [9]. Alternatively one can see this by examining a characteristic polynomial of an element of $SL(\mathbb{Z})$.) But the diagonal matrix with those eigenvalues as entries is not conjugate to the action of $\mathbb{R}$ the Heisenberg group. (ie the diagonal matrix with entries e^1, e^1, e^{-2}.)

The following result is stated in Eskin, Fisher and Whyte [4]. Its proof is completed by Dymarz in [2].

5.3.13 Theorem [2, Theorem 1] *If the rank of a nondegenerate, unimodular, split abelian-by-abelian solvable group is* 1*, then a finitely generated group* Γ *quasi-isometric to it is virtually a lattice in it.*

So now Theorem 5.3.13 says that the group G in the example above cannot be quasi-isometric to any finitely generated groups.

References

[1] **Y de Cornulier**, *Dimension of asymptotic cones of Lie groups*, J. Topol. 1 (2008) 342–361 MR2399134

[2] **T Dymarz**, *Large scale geometry of certain solvable groups*, Geom. Funct. Anal. 19 (2010) 1650–1687 MR2594617

[3] **T Dymarz**, **I Peng**, *Bilipschitz maps of boundaries of certain negatively curved homogeneous spaces*, Geom. Dedicata 152 (2011) 129–145

[4] **A Eskin**, **D Fisher**, **K Whyte**, *Coarse differentiation of quasi-isometries I: Spaces not quasi-isometric to Cayley graphs* `arXiv:math/0607207`

[5] **A Eskin**, **D Fisher**, **K Whyte**, *Coarse differentiation of quasi-isometries II: Rigidity for Sol and Lamplighter groups* `arXiv:0706.0940`

[6] **B Farb**, **L Mosher**, *On the asymptotic geometry of abelian-by-cyclic groups*, Acta Math. 184 (2000) 145–202 MR1768110

[7] **A Furman**, *Mostow–Margulis rigidity with locally compact targets*, Geom. Funct. Anal. 11 (2001) 30–59 MR1829641

[8] **S M Gersten**, *Quasi-isometry invariance of cohomological dimension*, C. R. Acad. Sci. Paris Sér. I Math. 316 (1993) 411–416 MR1209258

[9] **K Hasegawa**, *Four dimensional compact solvmanifold with and without complex analytic structures* `arXiv:math.CV/0401413`

[10] **K Hasegawa**, *Complex and Kähler structures on compact solvmanifolds*, J. Symplectic Geom. 3 (2005) 749–767 MR2235860 Conference on Symplectic Topology

[11] **D V Osin**, *Exponential radicals of solvable Lie groups*, J. Algebra 248 (2002) 790–805 MR1882124

[12] **I Peng**, *Coarse differentiation and quasi-isometries of a class of solvable Lie groups I*, Geom. Topol. 15 (2011) 1883–1925

[13] **M S Raghunathan**, *Discrete subgroups of Lie groups*, Ergebnisse der Math. und ihrer Grenzgebiete 68, Springer, New York (1972) MR0507234

Department of Mathematics, Indiana University
831 E 3rd St, Bloomington IN 47401, USA

`ipeng@indiana.edu`

Proposed: Benson Farb — Received: 13 April 2009
Seconded: Danny Calegari, Martin R Bridson — Revised: 3 August 2011