

MAPS BETWEEN RELATIVELY HYPERBOLIC SPACES AND BETWEEN THEIR BOUNDARIES

JOHN M. MACKAY AND ALESSANDRO SISTO

ABSTRACT. We study relations between maps between relatively hyperbolic groups/spaces and quasisymmetric embeddings between their boundaries. More specifically, we establish a correspondence between (not necessarily coarsely surjective) quasi-isometric embeddings between relatively hyperbolic groups/spaces that coarsely respect peripherals, and quasisymmetric embeddings between their boundaries satisfying suitable conditions. Further, we establish a similar correspondence regarding maps with at most polynomial distortion.

We use this to characterise groups which are hyperbolic relative to some collection of virtually nilpotent subgroups as exactly those groups which admit an embedding into a truncated real hyperbolic space with at most polynomial distortion, generalising a result of Bonk and Schramm for hyperbolic groups.

CONTENTS

1. Introduction	1
2. Statement of results	7
3. Preliminaries on horoballs and relative hyperbolicity	12
4. Subexponential distortion and (relative) hyperbolicity	26
5. Extending quasi-isometries to the boundary	36
6. Characterisation of boundary extensions	44
References	53

1. INTRODUCTION

The boundary at infinity of a Gromov hyperbolic group is a crucial invariant for studying the properties of the group. While its topology is already important, one can say more by looking at its metric properties: a result of Paulin [Pau96] shows that two hyperbolic groups G and H are quasi-isometric if and only if their boundaries are quasisymmetric. This additional quasisymmetric structure on the boundary

Date: September 6, 2023.

2010 Mathematics Subject Classification. 20F65, 20F67, 20F69, 31L10, 51F99.

Key words and phrases. Relatively hyperbolic groups and spaces, boundary at infinity, quasisymmetric map.

allows one to distinguish quasi-isometry types of hyperbolic groups which have the same topological boundary, amongst other applications. Bonk and Schramm [BS00] further established a correspondence of quasi-isometric embeddings between (certain) Gromov hyperbolic spaces to quasisymmetric embeddings between their boundaries, with refinements for classes of maps as we shall discuss below.

Our goal is to establish a similar correspondence for relatively hyperbolic groups (Definition 3.9). The motivating example of a relatively hyperbolic group G is the fundamental group of a finite volume non-compact hyperbolic manifold; G is a non-uniform lattice in $\text{Isom}(\mathbb{H}^n)$. If $n \geq 3$ then G is not Gromov hyperbolic, as it contains \mathbb{Z}^{n-1} parabolic subgroups $\{H_i\}$ due to the cusps, but one still can say the pair $(G, \{H_i\})$ is relatively hyperbolic (with peripheral subgroups $\{H_i\}$). That is, roughly, gluing in ‘horoballs’ to the left cosets of the H_i one obtains a *cusped space* $\text{Cusp}(G, \{H_i\})$ which is Gromov hyperbolic (see Definition 3.8), and so has a boundary $\partial_\infty(G, \{H_i\})$, which we call its *Bowditch boundary*. This boundary was defined by Bowditch [Bow12] (see also [GM08]) who showed that it depends only on $(G, \{H_i\})$ up to homeomorphism, not on the other choices involved in its construction. In this example $\partial_\infty(G, \{H_i\})$ is homeomorphic to \mathbb{S}^{n-1} .

Out of the many relevant works on relatively hyperbolic groups, we point out that the topology of the Bowditch boundary has been well-studied, in particular in connection to splittings, see e.g. [Bow99a, Bow99b]. We also highlight that maps between Bowditch boundaries have been previously studied by Groff [Gro13]. Also, Gerasimov–Potyagailo [GP13] study metric properties of the related Floyd boundaries, and in previous work we established metric properties of Bowditch boundaries [MS20].

Finally, as this paper was being finalised Healy and Hruska released a preprint with an extensive study of the properties of boundary maps for relatively hyperbolic groups [HH20]. In Remark 1.7, we elaborate upon how the results in [HH20] and in this paper complement each other.

Since a given group can have different relatively hyperbolic structures (with different Bowditch boundaries), to have a sensible theory we have to restrict our attention to maps which respect the peripheral structures, meaning that images and preimages of cosets of peripheral subgroups are close to cosets of peripheral subgroups. For a precise statement see Definition 2.2 below. The following example illustrates why we need this notion.

Example 1.1. *Let $G = \langle x, y \rangle$ be the fundamental group of a punctured torus, which is hyperbolic relative to the parabolic subgroup $H = \langle [x, y] \rangle$. Then $\partial_\infty(G, \emptyset)$ is a Cantor set, while $\partial_\infty(G, \{H\})$ is a circle. So the identity map from $(G, \emptyset) \rightarrow (G, \{H\})$ cannot induce a homeomorphism*

of the boundaries, while the identity map from $(G, \{H_i\}) \rightarrow (G, \emptyset)$ cannot induce anything sensible on the boundary at all.

The peripheral structures are often automatically respected. When $(G, \{H_i\})$ and $(G', \{H'_i\})$ are relatively hyperbolic groups and no peripheral group is itself properly relatively hyperbolic, any quasi-isometry from G to G' will automatically respect the peripheral structures (see Drutu [Dru09, Theorem 5.7] and also Groff [Gro13, Theorem 6.3]). In this context, Groff shows that the two Bowditch boundaries are homeomorphic [Gro13, Corollary 6.5].

A simplified statement of our first result (Theorem 2.8) is the following. A function $f : G \rightarrow G'$ *coarsely respects peripherals* if, roughly, images and preimages of peripheral sets are in bounded neighbourhoods of peripheral sets, see Definition 2.2.

Corollary 1.2. *Suppose that $(G, \{H_i\})$ and $(G', \{H'_i\})$ are relatively hyperbolic groups, and that $f : G \rightarrow G'$ is a quasi-isometric embedding that coarsely respects peripherals. Then f extends to a quasi-isometric embedding $f_{\text{Cusp}} : \text{Cusp}(G, \{H_i\}) \rightarrow \text{Cusp}(G', \{H'_i\})$. This map induces a quasisymmetric embedding $\partial_\infty f_{\text{Cusp}} : \partial_\infty(G, \{H_i\}) \rightarrow \partial_\infty(G', \{H'_i\})$.*

Recall that $\text{Cusp}(G, \{H_i\})$ denotes the cusped space where one takes the Cayley graph of G and glues horoballs to the left cosets of groups in $\{H_i\}$ (see Definition 3.8 and [GM08]).

We note that if f is a quasi-isometry then a large part of the proof of this result can be replaced by a much shorter argument; this is done in the proof of [Gro13, Theorem 6.3], which contains a proof of the theorem in this case. In both cases one first proves that the extension f_{Cusp} is coarsely Lipschitz, but then the difference is that, when f is a quasi-isometry, using a quasi-inverse of f one can see that f_{Cusp} has a quasi-inverse. When f is not coarsely surjective, a more sophisticated argument is needed, and more specifically we will rely on detailed understanding of geodesics in relatively hyperbolic spaces, which is not needed for the quasi-isometry case.

The quasisymmetric embedding $\partial_\infty f$ is not arbitrary; for example, it has to send parabolic points in $\partial_\infty(G, \{H_i\})$ to parabolic points in $\partial_\infty(G', \{H'_i\})$. We find conditions on $\partial_\infty f$ which are sufficient to give a converse statement; i.e., any suitable quasisymmetric embedding $h : \partial_\infty(G, \{H_i\}) \rightarrow \partial_\infty(G', \{H'_i\})$ will induce a quasi-isometric embedding $\hat{h} : G \rightarrow G'$ so that $\partial_\infty \hat{h} = h$ (see Theorem 2.9), but for now we state a simpler result.

Corollary 1.3. *Suppose that $(G, \{H_i\})$ and $(G', \{H'_i\})$ are relatively hyperbolic groups with infinite proper peripheral groups, and that $h : \partial_\infty(G, \{H_i\}) \rightarrow \partial_\infty(G', \{H'_i\})$ is a shadow-respecting quasisymmetry. Then there exists a quasi-isometry $\hat{h} : G \rightarrow G'$ so that \hat{h} extends to give $\partial_\infty(\hat{h})_{\text{Cusp}} = h$.*

The ‘shadow’ of a horoball in the boundary is, roughly, the ball in the boundary which consists of the end points of geodesic rays from $e \in G$ which meet the horoball. Given this, ‘shadow-respecting’ means that parabolic points are mapped to parabolic points, with the corresponding shadows approximately preserved, and in addition that near parabolic points the map \hat{h} looks like a ‘snowflake’ map; see Definition 2.6.

In the case of hyperbolic groups G, G' , a quasisymmetric embedding $\partial_\infty G \rightarrow \partial_\infty G'$ induces a quasi-isometric embedding $G \rightarrow G'$ by work of Bonk and Schramm [BS00] (see also Paulin [Pau96]). However, when we weaken h in Corollary 1.3 to be a shadow-respecting quasisymmetric embedding, we will find an induced map $\hat{h} : G \rightarrow G'$ that is *polynomially distorted*: there exists $C > 0, \alpha \in (0, 1]$ so that for all $x, y \in G$,

$$\frac{1}{C}d_G(x, y)^\alpha - C \leq d_{G'}(\hat{h}(x), \hat{h}(y)) \leq Cd_G(x, y) + C.$$

This is reasonable to expect: Bonk and Schramm [BS00] give us an extension of h to a quasi-isometric embedding between the cusped spaces. The shadow-respecting condition gives that this maps horoballs to horoballs. However, for pairs of points in the same left coset of a peripheral group their distance in the group is essentially the exponential of their distance in the cusped space. Thus the multiplicative distortion of the quasi-isometry is exponentiated to a polynomial distortion in the group.

As we will see in Theorem 2.9, one can correct for this by requiring finer control around parabolic points. We also strengthen Corollary 1.2 to apply to certain polynomially distorting maps (Definition 2.1), giving a correspondence between maps on the group and maps on the boundary. Polynomially distorting maps have appeared in work of many authors, for example in the study of embeddings into Hilbert spaces, where they are used to define compression exponents [GK04].

Polynomially distorting maps, while not necessarily quasi-isometric embeddings, do preserve some geometric properties of a group, and in fact the results below provide examples of this.

If G has a polynomially (or indeed subexponentially) distorted embedding f into a hyperbolic group H , then f is in fact a quasi-isometric embedding, and G is hyperbolic. This follows from easily from standard facts about hyperbolic groups; see also Proposition 4.4. In this paper we show an analogous result for relatively hyperbolic groups and spaces (Theorem 4.1), generalising known results about quasi-isometric embeddings. Here we state the version for groups.

Corollary 1.4. *Let G be hyperbolic relative to $\{H_i\}$. If H is a subexponentially distorted subgroup of G then H is relatively quasiconvex in G , and so H is hyperbolic relative to subgroups $\{K_j\}$ each of which is the intersection of a conjugate (in G) of some H_i with H .*

For the case of polynomial distortion see also [GP13, Theorem C].

The reason for proving Theorem 4.1 here is that we will use it in conjunction with Theorem 2.9 to generalise a result of Bonk and Schramm which characterises hyperbolic groups (amongst finitely generated groups with a word metric) as those which admit a quasi-isometric embedding into some \mathbb{H}^n [BS00]. In an analogous way, we characterise groups that are hyperbolic relative to virtually nilpotent subgroups. A *truncated real hyperbolic space* is the complement of a collection of (disjoint) horoballs in some \mathbb{H}^n which is then given the resulting length metric [BH99, page 362]. A truncated real hyperbolic space is hyperbolic relative to the collection of boundary horospheres, in the sense of Definition 3.8.

Theorem 1.5. *Suppose G is finitely generated group with a word metric. Then the following are equivalent:*

- G is hyperbolic relative to some collection of virtually nilpotent subgroups,
- G admits a polynomially distorted embedding into some truncated real hyperbolic space, and
- G admits a subexponentially distorted embedding into some truncated real hyperbolic space.

The inclusion $X \rightarrow \mathbb{H}^n$ of a truncated real hyperbolic space X into its ambient hyperbolic space is a coarse embedding (that is, a uniformly proper coarse Lipschitz map, Definition 2.1), but with exponential distortion. So if G is hyperbolic relative to virtually nilpotent groups then the above theorem gives a coarse embedding of $G \rightarrow \mathbb{H}^n$, which, as already mentioned, cannot be subexponentially distorted unless G is hyperbolic. The general problem of which groups admit coarse embeddings into real hyperbolic spaces (with any distortion function) is also of interest, see work of Hume and the second author [HS17] and of Tessera [Tes20]; in the latter paper it is shown that the only amenable groups which coarsely embed into real hyperbolic spaces are the virtually nilpotent ones.

The embedding in Theorem 1.5 is built in roughly the same way as Bonk–Schramm’s result: one applies Assouad’s embedding theorem to embed the boundary into some Euclidean space, and then extends the map inside using Theorem 2.9 (the stronger version of Corollary 1.3 which applies to relatively hyperbolic spaces such as a truncated real hyperbolic space). Given such an embedding, to find the relatively hyperbolic structure on G we use that a truncated real hyperbolic space is hyperbolic relative to the horospheres of the removed horoballs, and pull back the relatively hyperbolic structure to the embedded group (via Corollary 4.2, a consequence of Theorem 4.1 and work of Druţu [Dru09] relating ‘metric’ relative hyperbolicity,

‘asymptotic tree-gradedness’ in her terminology, to ‘algebraic’ relative hyperbolicity).

Theorem 1.5 cannot be extended to quasi-isometric embeddings: if G is a non-uniform lattice in $\text{Isom}(\mathbb{C}\mathbb{H}^2)$ then a quasi-isometric embedding into a truncated real hyperbolic space would induce a quasi-isometric embedding of the real Heisenberg group with its Carnot metric into a Euclidean space, which is impossible. However, Theorem 1.5 and some consideration of the geometry of the horoballs motivates the following.

Question 1.6. *Is a finitely generated group hyperbolic relative to some collection of virtually abelian subgroups if and only if the group admits a quasi-isometric embedding into some truncated real hyperbolic space?*

Remark 1.7. As this paper was being finished, in independent work Healy–Hruska released a preprint studying metric properties of boundaries of relatively hyperbolic groups [HH20]. Similar to us, they show that quasi-isometries of relatively hyperbolic groups extend to quasi-isometries of the cusped spaces and quasisymmetries of the Bowditch boundary [HH20, Theorem 1.2].

While on the surface the results look similar to results in this paper, we would like to point out that the contents of the papers are actually rather different. In fact, the main point of [HH20] is to show, roughly, that using different models of horoballs does not affect the cusped space, and they dedicate a lot of work to the study of different types of horoballs in order to do so. They are then able to extend maps on peripheral sets to horoballs, which, as indicated below Corollary 1.2, is basically all one needs to extend quasi-isometries of relatively hyperbolic groups to quasi-isometries of cusped spaces. In contrast, we regard as our main contribution in Corollary 1.2 to be the case of non-coarsely-surjective quasi-isometric embeddings, while we only use horoballs very similar to one fixed model, so that our extensions come with less work.

The papers complement each other in that, roughly, thanks to [HH20] we can apply our results to different models of cusped spaces, while we strengthen certain results of [HH20], like [HH20, Theorem 1.2], in the case of our particular cusped spaces.

1.1. Outline. In Section 2 we give the full statements of our results, with definitions of the classes of maps involved on the spaces and their boundaries. In particular, we highlight the “shadow-respecting quasisymmetric embeddings” in Definition 2.6, which is the appropriate class of maps on boundaries in our context. The main results are the following. First, Theorem 2.8 gives extensions of maps between relatively hyperbolic spaces to maps between their Bowditch boundaries, which leads to Corollary 1.2. Second, Theorem 2.9 gives the converse, extending maps between the boundaries to maps between the spaces, which leads to Corollary 1.3.

Section 3 contains preliminary material on the metric geometry of cusped spaces and their boundaries, and basic properties of the classes of quasimetric maps introduced in Section 2. In Section 4 we prove Corollary 1.4 and Theorem 1.5, assuming Theorem 2.9.

The proofs of Theorems 2.8 and 2.9 are contained in Sections 5 and 6 respectively.

1.2. Notation. $A \preceq B$ means $A \leq \lambda B$ for some $\lambda \geq 1$; $A \asymp B$ means $A \preceq B$ and $B \preceq A$.

$A \lesssim B$ means $A \leq B + C$ for some C ; $A \approx B$ means $A \lesssim B$ and $B \lesssim A$.

$A \lesssim_{\mu} B$ means $A \leq \lambda B + \mu$ for some $\lambda \geq 1$ and $\mu \geq 0$; $A \cong B$ means $A \lesssim_{\mu} B$ and $B \lesssim_{\mu} A$. Moreover $A \cong_{\lambda, \mu} B$ means that $A/\lambda - \mu \leq B \leq \lambda A + \mu$.

For U a subset of a metric space X , and $C > 0$, we denote the C -neighbourhood of U by $N_C(U) = \{x \in X : d(x, U) \leq C\}$.

1.3. Acknowledgements. We thank Christopher Hruska, David Hume and referee(s) for helpful comments. The first author was partially supported by EPSRC grant EP/K032208/1.

2. STATEMENT OF RESULTS

There are various characterisations of relatively hyperbolic spaces (X, \mathcal{P}) , where X is a geodesic metric space and \mathcal{P} is a collection of (coarsely connected) subspaces of X called ‘peripheral’ subsets. (Note that we allow peripheral sets to be bounded, which is not always the case in the literature. Also, we always assume X is geodesic.) In this paper, the main characterisation we use is that, roughly, X is hyperbolic relative to \mathcal{P} (i.e. (X, \mathcal{P}) is relatively hyperbolic) if the cusped space $\text{Cusp}(X, \mathcal{P})$ formed by gluing in horoballs to each $P \in \mathcal{P}$ is Gromov hyperbolic. We allow (and later need) different models for these horoballs, entailing some preliminary discussion, so we defer the full definition to Section 3, see Definition 3.8. In a few papers, including [DS05, Dru09], relative hyperbolicity is referred to as ‘asymptotic tree-gradedness’.

Given a finitely generated group G and subgroups $\{H_i\}$ we say that the pair $(G, \{H_i\})$ is relatively hyperbolic if any Cayley graph X of G is relatively hyperbolic with peripheral sets \mathcal{P} given by all left cosets of all $\{H_i\}$. We give the precise definition in the next section (see Definition 3.9).

To state our results, we have to first define the varieties of coarse maps and of quasimetric maps that we work with.

2.1. Coarse maps. We denote the metric on a metric space X by d_X , or just d if there is no possibility of confusion.

Definition 2.1. A map $f : X \rightarrow X'$ between metric spaces is coarsely Lipschitz if there exists $C > 0$ so that for all $x, y \in X$,

$$d_{X'}(f(x), f(y)) \leq Cd_X(x, y) + C.$$

We say f is τ -uniformly proper, for some $\tau : [0, \infty) \rightarrow \mathbb{R}$ with $\lim_{t \rightarrow \infty} \tau(t) = \infty$, if it is coarsely Lipschitz and if for all $x, y \in X$,

$$\tau(d_X(x, y)) \leq d_{X'}(f(x), f(y)).$$

We say a τ -uniformly proper map f is

- (1) a quasi-isometric embedding if we can take τ linear,
- (2) polynomially distorted if we can take $\tau(t) = t^\alpha - C$ for some $\alpha \in (0, 1]$ and $C > 0$, or
- (3) subexponentially distorted if we can take τ so that we have $\lim_{t \rightarrow \infty} \tau(t)/\log t = \infty$.

These notions are progressively weaker: (1) \Rightarrow (2) \Rightarrow (3).

When mapping between relatively hyperbolic spaces, our maps should interact well with the peripheral sets.

Definition 2.2. A map $f : X \rightarrow X'$ between relatively hyperbolic spaces $(X, \mathcal{P}), (X', \mathcal{P}')$ coarsely respects peripherals if

- (1) $\exists C > 0$ so that for all $P \in \mathcal{P}$ there exists $P' \in \mathcal{P}'$ so that $f(P) \subset N_C(P')$,
- (2) $\forall C' > 0 \exists C > 0$ so that for all $P' \in \mathcal{P}'$ either $\text{diam } f^{-1}(N_{C'}(P')) \leq C$, or there exists $P \in \mathcal{P}$ so that $f^{-1}(N_{C'}(P')) \subset N_C(P)$,

We are also interested in polynomially distorting maps which have better control on peripheral sets, namely they behave like coarse snowflake maps on each.

Definition 2.3. A polynomially distorted map $f : X \rightarrow X'$ between relatively hyperbolic spaces $(X, \mathcal{P}), (X', \mathcal{P}')$ which coarsely respects peripherals is a snowflake on peripherals if $\exists C > 0$ so that for every $P \in \mathcal{P}$ there exists $\lambda \in (0, 1]$ so that for all $x, y \in P$,

$$\frac{1}{C}d_X(x, y)^\lambda - C \leq d_{X'}(f(x), f(y)) \leq Cd_X(x, y)^\lambda + C.$$

Assuming (as we often will) that each $P \in \mathcal{P}$ is unbounded, λ is uniformly bounded away from zero by a constant depending on the polynomial distortion of f . A natural example of a snowflake on peripherals map is a quasi-isometric embedding which coarsely respects peripherals; in this case each $\lambda = 1$. Fairly immediately, relatively hyperbolic spaces with unbounded peripherals and snowflake on peripheral maps form a category, see Proposition 5.11.

2.2. Quasisymmetric maps. We now consider boundaries and the maps between them. As discussed in Section 3, gluing a horoball $\mathcal{H}(P)$ to each $P \in \mathcal{P}$ gives a Gromov hyperbolic space $\text{Cusp}(X, \mathcal{P})$. This space has Gromov boundary $\partial_\infty \text{Cusp}(X, \mathcal{P}) = \partial_\infty(X, \mathcal{P})$ which comes decorated with the following data:

$$(\partial_\infty(X, \mathcal{P}), \rho, \{(a_P, r_P)\}_{P \in \mathcal{P}}),$$

where ρ is a visual metric with visual parameter $\epsilon > 0$ with respect to a basepoint $o \in X$, each $a_P = \partial_\infty P \in \partial_\infty(X, \mathcal{P})$ is the parabolic point associated to $\mathcal{H}(P)$, and each $r_P = e^{-\epsilon d(o, P)}$ is the radius of the ‘‘shadow’’ of P in $\partial_\infty(X, \mathcal{P})$. We formalise this as follows, considering \mathcal{P} as an abstract index set for points a_P and radii r_P of balls in a given metric space (Z, ρ) .

Definition 2.4. *A shadow decorated metric space is a tuple*

$$(Z, \rho, \mathcal{B} = \{(a_P, r_P)\}_{P \in \mathcal{P}}),$$

where (Z, ρ) is a metric space, and for $P \in \mathcal{P}$ we have $a_P \in Z, r_P > 0$, with all a_P distinct and $\sup r_P \preceq \text{diam } Z$.

The notion of quasisymmetric mappings was generalised to metric spaces by Tukia and Väisälä [TV80].

Definition 2.5. *A distortion function is a homeomorphism $\eta : [0, \infty) \rightarrow [0, \infty)$. A topological embedding (respectively homeomorphism) $h : (Z, \rho) \rightarrow (Z', \rho')$ between metric spaces (Z, ρ) and (Z', ρ') is called an η -quasisymmetric embedding (resp. homeomorphism), for some distortion function η , if for every triple of points $x, y, z \in Z$ and $t \geq 0$,*

$$\rho(x, y) \leq t\rho(x, z) \implies \rho'(h(x), h(y)) \leq \eta(t)\rho'(h(x), h(z)).$$

If η is not specified we just call h a quasisymmetric embedding (resp. homeomorphism).

For maps between shadow decorated metric spaces, we refine this definition as follows. The reader should have in mind boundaries of relatively hyperbolic spaces, and as we will see the definition characterises the extension to the boundary of a quasi-isometric embedding between spaces coarsely respecting peripherals, in particular the behaviour of such extensions around parabolic points. To explain the terminology we use in the last property below, in analysis on metric spaces, for a metric space (Z, ρ) and $\epsilon > 0$ such that (Z, ρ^ϵ) is also a metric space (for example, $\epsilon \in (0, 1]$), the identity map $(Z, \rho) \rightarrow (Z, \rho^\epsilon)$ is called a ‘snowflake’ transformation.

Definition 2.6. *Suppose $(Z, \rho, \mathcal{B}), (Z', \rho', \mathcal{B}')$ are shadow decorated metric spaces. A shadow-respecting η -quasisymmetric embedding h is an η -quasisymmetric embedding $h : (Z, \rho) \rightarrow (Z', \rho')$ so that, for a constant $C \geq 1$,*

- (1) for all $(a, r) \in \mathcal{B}$, there exists $(a', r') \in \mathcal{B}'$ where $a' = h(a)$ and for all $b \in Z$,

$$\rho(a, b) \leq r \implies \rho'(a', h(b)) \leq Cr';$$

- (2) for all $(a', r') \in \mathcal{B}'$ if $a' = h(a)$, $(a, r) \in \mathcal{B}$ then for all $b \in Z$,

$$r \leq \rho(a, b) \implies r' \leq C\rho'(a', h(b)),$$

and if no such $(a, r) \in \mathcal{B}$ exists, $\rho'(a', h(Z)) \geq r'/C$;

- (3) for all $(a, r) \in \mathcal{B}$, there exists $\lambda_a \in (0, \infty)$ so that, writing (a', r') as in (1), for all $b, c \in Z$ if

$$\left(\frac{\rho(a, b)}{r}\right) \rho(a, b) \leq \rho(b, c) \leq \rho(a, b) \leq r,$$

then

$$\frac{\rho'(h(b), h(c))}{r'} \asymp_C \left(\frac{\rho(b, c)}{r}\right)^{\lambda_a}.$$

If h is a homeomorphism we say that it is a shadow-respecting η -quasisymmetry.

If λ_a equals a fixed λ for all $(a, r) \in \mathcal{B}$, we say h asymptotically λ -snowflakes.

The first two properties here say that h should respect the sizes of shadows, while the third says that as we get close to the centre of a shadow the map should look more and more like a suitably scaled snowflake embedding. A basic example of a shadow-respecting quasisymmetry that asymptotically snowflakes is given by a change of visual metric on the boundary of a fixed relatively hyperbolic space: if ρ and ρ' are visual metrics with parameters ϵ and ϵ' respectively, then $\rho' \asymp \rho^{\epsilon'/\epsilon}$, and the radii $r_P = e^{-\epsilon d(o, P)}$ and $r'_P = e^{-\epsilon' d(o, P)}$ also satisfy $r'_P = (r_P)^{\epsilon'/\epsilon}$. Definition 2.6 then holds with each $\lambda_a = \epsilon'/\epsilon$; note that the hypothesis in property (3) is superfluous.

Remark 2.7. In Definition 2.6, if h is a homeomorphism, we can state (1), (2) more simply as: h induces a bijection $\mathcal{B} \rightarrow \mathcal{B}'$, $(a, r) \mapsto (h(a), r'_{h(a)}) \in \mathcal{B}'$, and for all $(a, r) \in \mathcal{B}$

$$B(a', r'_{h(a)}/C) \subset h(B(a, r)) \subset B(a', Cr'_{h(a)}).$$

It is not clear that shadow-respecting quasisymmetric maps form a category, however they do when the spaces involved are “uniformly perfect”, see Proposition 3.25.

2.3. A correspondence between maps. We can now state our two main theorems.

Theorem 2.8. *Suppose $(X, \mathcal{P}), (X', \mathcal{P}')$ are relatively hyperbolic spaces, where each $P \in \mathcal{P}$ is unbounded.*

Suppose $f : X \rightarrow X'$ is a map which is a snowflake on peripherals (and so also polynomially distorting and coarsely respecting peripherals). Then f extends to a quasi-isometric embedding $f_{\text{Cusp}} : \text{Cusp}(X, \mathcal{P}) \rightarrow \text{Cusp}(X', \mathcal{P}')$, quantitatively.

The boundary map $\partial_\infty f_{\text{Cusp}} : \partial_\infty(X, \mathcal{P}) \rightarrow \partial_\infty(X', \mathcal{P}')$ is a shadow-respecting quasisymmetric embedding, quantitatively.

Moreover, if f is a quasi-isometric embedding, then $\partial_\infty f_{\text{Cusp}}$ asymptotically $\frac{\epsilon'}{\epsilon}$ -snowflakes, where ϵ, ϵ' denote the visual parameters of the boundary metrics on $\partial_\infty(X, \mathcal{P}), \partial_\infty(X', \mathcal{P}')$, respectively.

Finally, if (X'', \mathcal{P}'') is also a relatively hyperbolic space and $g : X' \rightarrow X''$ is a snowflake on peripherals, then $(g \circ f)_{\text{Cusp}}$ is bounded distance to $g_{\text{Cusp}} \circ f_{\text{Cusp}}$, and $\partial_\infty(g \circ f)_{\text{Cusp}} = (\partial_\infty g_{\text{Cusp}}) \circ (\partial_\infty f_{\text{Cusp}})$. That is, $(X, \mathcal{P}) \mapsto \partial_\infty \text{Cusp}(X, \mathcal{P}), f \mapsto \partial_\infty f_{\text{Cusp}}$ is a functor from relatively hyperbolic spaces having unbounded peripherals with snowflake on peripheral maps to shadow decorated metric spaces with shadow-respecting quasisymmetric embeddings.

Corollary 1.2 is an immediate consequence, since a quasi-isometry that coarsely respects peripherals is automatically a snowflake on peripherals (with each $\lambda = 1$), and we may discard any finite peripheral groups from $\{H_i\}$. (As an aside, the assumption above on unbounded peripherals is necessary in order to get the uniqueness of extensions required for functoriality.)

The proof of Theorem 2.8 is in Section 5.

Theorem 2.9. *Suppose $(X, \mathcal{P}), (X', \mathcal{P}')$ are two relatively hyperbolic spaces that are visually complete (Definition 3.21), and $h : \partial_\infty(X, \mathcal{P}) \rightarrow \partial_\infty(X', \mathcal{P}')$ is a shadow-respecting, η -quasisymmetric embedding. Then there exists a polynomially distorted embedding $\hat{h} : X \rightarrow X'$, which is a snowflake on peripherals, so that $h = \partial_\infty(\hat{h})_{\text{Cusp}}$, and \hat{h} is unique up to bounded distance.*

Moreover, if h also asymptotically $\frac{\epsilon'}{\epsilon}$ -snowflakes then \hat{h} is a quasi-isometric embedding, where ϵ, ϵ' denote the visual parameters of the boundary metrics on $\partial_\infty(X, \mathcal{P}), \partial_\infty(X', \mathcal{P}')$, respectively.

Finally, if (X'', \mathcal{P}'') is also a relatively hyperbolic space that is visually complete, and $j : \partial_\infty(X', \mathcal{P}') \rightarrow \partial_\infty(X'', \mathcal{P}'')$ is a shadow-respecting, quasisymmetric embedding, then $\widehat{j \circ \hat{h}}$ and $\hat{j} \circ \hat{h}$ are at bounded distance. That is, $\partial_\infty(X, \mathcal{P}) \mapsto X, h \mapsto \hat{h}$ is a functor from shadow decorated boundaries of relatively hyperbolic spaces with shadow-respecting quasisymmetric embeddings, to relatively hyperbolic spaces with equivalence classes of snowflake on peripheral maps, where two such maps are considered equivalent if they are bounded distance apart.

The property of being ‘visually complete’ (Definition 3.21) is satisfied by relatively hyperbolic groups having proper infinite peripheral groups (Proposition 3.23), and is equivalent to the Bowditch boundary

being uniformly perfect (Lemma 3.22). The importance of this condition is two-fold: first, by work of Tukia–Väisälä it allows us to upgrade quasisymmetric maps to “power” quasisymmetric maps, which have quasi-isometric extensions by Bonk–Schramm, see Theorem 6.6 and discussion. Second, without it we cannot hope for uniqueness of extensions \hat{h} : consider the many quasi-isometries of \mathbb{Z} which induce the same boundary maps. Healy–Hruska also show that boundaries of relatively hyperbolic groups are uniformly perfect, see [HH20, Section 1.1] and references within, particularly to work of Meyer [Mey09].

Corollary 1.3 now follows from Theorem 2.9.

Proof of Corollary 1.3. If $h : \partial_\infty(G, \{H_i\}) \rightarrow \partial_\infty(G', \{H'_i\})$ is a shadow-respecting quasisymmetry, then by Proposition 3.25 h^{-1} is also a shadow-respecting quasisymmetry, and so $\text{id}_G \approx \text{id}_{\widehat{\partial_\infty(G, \{H_i\})}} = \widehat{h \circ h^{-1}} \approx \widehat{\hat{h} \circ \hat{h}^{-1}}$, where \approx means the maps agree up to bounded error. Thus as $\hat{h}, \widehat{h^{-1}}$ are both coarsely Lipschitz, they must be quasi-isometries. \square

The proof of Theorem 2.9 is in Section 6 (with the map \hat{h} being $\hat{f}|_X$ where \hat{f} is a map between cusped spaces constructed via Theorem 6.6 due to Bonk and Schramm).

3. PRELIMINARIES ON HOROBALLS AND RELATIVE HYPERBOLICITY

In this section we define models of horoballs and the notions of relatively hyperbolic groups and spaces, and their (Bowditch) boundaries. We then collect some preliminary results on quasi-centres in hyperbolic spaces, separation of horoballs, visual completeness and transient sets.

3.1. Horoballs. There are many definitions of relatively hyperbolic groups and metric spaces. We will give one in terms of cusped spaces defined in a way that allows us to consider both real hyperbolic horoballs and the following combinatorial horoballs.

Definition 3.1. *Suppose Γ is a connected graph with vertex set V and edge set E , where every edge has length one. The horoball $\mathcal{H}(\Gamma)$ is defined to be the graph with vertex set $V \times \mathbb{N}$ and edges $((v, n), (v, n+1))$ of length 1, for all $v \in V, n \in \mathbb{N}$, and edges $((v, n), (v', n))$ of length e^{-n} , for all $(v, v') \in E$.*

Note that $\mathcal{H}(\Gamma)$ is quasi-isometric to the metric space constructed from Γ by gluing to each edge in E a copy of the strip $[0, 1] \times [1, \infty)$ in the upper half-plane model of \mathbb{H}^2 , where the strips are attached to each other along $v \times [1, \infty)$.

As is well known, these horoballs are hyperbolic with boundary a single point. Moreover, it is easy to estimate distances in horoballs. (Recall that we write $A \approx B$ if $|A - B|$ is bounded by some constant.)

Lemma 3.2. *Suppose Γ and $\mathcal{H}(\Gamma)$ are defined as above. Let d_Γ and $d_{\mathcal{H}}$ denote the corresponding path metrics. Then for each $(x, m), (y, n) \in \mathcal{H}(\Gamma)$, we have*

$$d_{\mathcal{H}}((x, m), (y, n)) \approx 2 \log(d_\Gamma(x, y)e^{-\max\{m, n\}} + 1) + |m - n|.$$

Proof. We may assume that $m \leq n$. We can suppose $d_\Gamma(x, y) \geq 2e^n$, as if not $|m - n| \leq d_{\mathcal{H}}((x, m), (y, n)) \leq 2 + |m - n|$. Consequently $\log(d_\Gamma(x, y)e^{-n} + 1) \approx \log d_\Gamma(x, y) - n$. By construction of $\mathcal{H}(\Gamma)$, any geodesic γ between x and y in $\mathcal{H}(\Gamma)$ must, for some $t \geq n$, go from (x, m) to (x, t) changing only the second coordinate, then follow a geodesic $\gamma' \subset \Gamma \times \{t\} \subset \mathcal{H}(\Gamma)$ to (y, t) , then back to (y, n) . Thus

$$(3.3) \quad d_{\mathcal{H}}((x, m), (y, n)) \leq 2(t - n) + |n - m| + e^{-t}d_\Gamma(x, y),$$

with equality for the best choice of t . It is readily seen that this value is attained for the least t so that $l_t = e^{-t}d_\Gamma(x, y)$ satisfies $l_t/e + 2 \geq l_t$, that is, $l_t \leq C := 2e/(e - 1)$. So the best choice of t is $t = \lceil \log(d_\Gamma(x, y)/C) \rceil$ which is $\geq n$, and the right hand side of (3.3) is $\approx_{2+C} 2(\log d_\Gamma(x, y) - n) - 2 \log C + |n - m|$. \square

This metric distance estimate is all we really need in our results, so we abstract it as follows. A metric space P is (ϵ) -coarsely connected if there exists $\epsilon > 0$ so that for any $x, y \in P$ there exists $x = x_0, x_1, \dots, x_n = y$ in P with $d(x_{i-1}, x_i) \leq \epsilon$ for each $i = 1, \dots, n$.

Definition 3.4. *Let (P, d_P) be a coarsely connected metric space. An admissible horoball is a geodesic metric space $(\mathcal{H}(P), d_{\mathcal{H}(P)})$ which is δ -hyperbolic with a single boundary point $\{a_P\} = \partial_\infty \mathcal{H}(P)$, and for which there exists $C > 0, L \geq 1$ satisfying the following properties. There exists a subset $\hat{P} \subset P$ with $d_P(x, \hat{P}) \leq C$ for all $x \in P$, and \hat{P} is identified with a topologically embedded subset of $\mathcal{H}(P)$. Moreover we require:*

- (1) *for each $x \in \hat{P} \subset \mathcal{H}(P)$ there exists a geodesic ray $\gamma_x : [0, \infty) \rightarrow \mathcal{H}(P)$ with $\gamma_x(0) = x$ and (necessarily) $\gamma_x(\infty) = a_P$.*
- (2) *for every $p \in \mathcal{H}(P)$ there exists $x \in \hat{P}$ with $d_{\mathcal{H}(P)}(p, \gamma_x) \leq C$,*
- (3) *for any $x, y \in \hat{P}$ with $d_{\mathcal{H}(P)}(x, y) \leq 1$ we have $d_{\mathcal{H}(P)}(x, y) \geq d_P(x, y)/L$,*
- (4) *and for any $x, y \in \hat{P}$ and $m, n \geq 0$,*

$$d_{\mathcal{H}(P)}(\gamma_x(m), \gamma_y(n)) \approx_C 2 \log(d_P(x, y)e^{-\max\{m, n\}} + 1) + |m - n|.$$

A collection of admissible horoballs on a collection of spaces is uniform if the constants δ, C, L may be chosen uniformly.

In our definition of a relatively hyperbolic space below, and in all our results, we require uniformly admissible horoballs in order to control the horoball geometry up to uniformly bounded errors.

Example 3.5 ([Sis12, Definitions 3.1,3.2]). *For a k -coarsely connected metric space (P, d_P) , define a graph approximation Γ_P by choosing as vertex set a maximal ϵ -separated subset \hat{P} in (P, d) , and connecting by an edge any $x, y \in \hat{P}$ with $d_P(x, y) \leq R$ for a suitable R chosen to make Γ_P connected. Then let $\mathcal{H}(P)$ be the graph horoball (Definition 3.1) on Γ_P . Properties (1,2,3) of Definition 3.4 follow by construction, and Definition 3.4(4) by Lemma 3.2.*

Other than graph horoballs, the only other example we work with directly is the following.

Proposition 3.6. *Let H be a horoball in \mathbb{H}^N ; in the upper half plane model \mathbb{R}_+^N we may take $H = \{x = (x_i) \in \mathbb{R}_+^N : x_N \geq 1\}$. Then H is an admissible horoball for $\partial_\infty H \cong \mathbb{R}^{N-1} \subset \overline{\mathbb{R}_+^N} \setminus H$.*

Proof. Write $\|\cdot\|_2$ for the standard norm on \mathbb{R}^{N-1} . Definition 3.4(1,2) holds trivially. The hyperbolic distance between $(x, e^m), (y, e^n) \in \mathbb{R}^{N-1} \times [1, \infty) = H$ satisfies:

$$\begin{aligned} & d_H((x, e^m), (y, e^n)) \\ &= 2 \log \left(\frac{\sqrt{\|x - y\|_2^2 + (e^m - e^n)^2} + \sqrt{\|x - y\|_2^2 + (e^m + e^n)^2}}{2\sqrt{e^m e^n}} \right) \\ (3.7) \quad & \approx \log \left(\frac{\|x - y\|_2^2 + (e^m + e^n)^2}{e^m e^n} \right). \end{aligned}$$

Without loss of generality suppose $m \leq n$. If $\|x - y\|_2 \leq e^n$, (3.7) is $\approx \log e^{n-m} = |m - n|$. Otherwise, $\|x - y\|_2 \geq e^n$ and (3.7) is $\approx \log(\|x - y\|_2^2 e^{-(m+n)}) = 2 \log(d_{\mathbb{R}^{N-1}}(x, y) e^{-n}) + |n - m|$. In either case, this approximately equals the distance estimate of Definition 3.4(4).

Finally, the shortest path distance between $(x, 1), (y, 1) \in \partial_\infty H \subset \overline{\mathbb{H}^N} \setminus H$ is $\|x - y\|_2$, so

$$\begin{aligned} \frac{d_H((x, 1), (y, 1))}{\|x - y\|_2} &= \frac{2 \log \left(\|x - y\|_2 + \sqrt{\|x - y\|_2^2 + 4} \right) - 2 \log(2)}{\|x - y\|_2} \\ &\rightarrow 1 \text{ as } \|x - y\|_2 \rightarrow 0, \end{aligned}$$

and Definition 3.4(3) follows. \square

3.2. Cusped spaces.

Definition 3.8 (cf. [Sis12, Definitions 3.3–3.4]). *Suppose (X, d) is a geodesic metric space, \mathcal{P} a collection of uniformly coarsely connected subsets of X with metrics $d_P = d|_P$ for $P \in \mathcal{P}$, and $\{\mathcal{H}(P)\}_{P \in \mathcal{P}}$ a collection of uniformly admissible horoballs for \mathcal{P} .*

The cusped space $\text{Cusp}(X, \mathcal{P}) = \text{Cusp}(X, \{\mathcal{H}(P)\}_{P \in \mathcal{P}})$ is the path metric space obtained from X and $\{\mathcal{H}(P)\}_{P \in \mathcal{P}}$ by identifying each \hat{P} with its copies in $\mathcal{H}(P)$ and in $P \subset X$.

If $\text{Cusp}(X, \mathcal{P})$ is a Gromov hyperbolic geodesic space, we say that (X, \mathcal{P}) is relatively hyperbolic, or that X is hyperbolic relative to the collection of subsets \mathcal{P} . The elements of \mathcal{P} are called peripheral sets. We also assume X and $\text{Cusp}(X, \mathcal{P})$ are proper.

The original definition in [Sis12] used only the admissible horoballs of Example 3.5, and showed that (X, \mathcal{P}) being relatively hyperbolic was independent of the choice of graph approximation horoballs [Sis12, Proposition 4.9].

Except for the application to truncated hyperbolic spaces (Theorem 1.5), we are mainly interested in relatively hyperbolic group pairs $(G, \{H_i\})$ as defined below, and the reader may primarily keep this case in mind.

Definition 3.9. *Suppose G is a finitely generated group, and $\{H_i\}_{i=1}^n$ a collection of finitely generated subgroups of G . Let S be a finite generating set for G , so that $S \cap H_i$ generates H_i for each $i = 1, \dots, n$.*

We say that $(G, \{H_i\})$ is relatively hyperbolic if a Cayley graph X of G with respect to S , is hyperbolic relative to the collection \mathcal{P} of copies of Cayley graphs of H_i in X on each left coset of H_i , for each i . (We may use graph horoballs, as in Definition 3.1, on the corresponding Cayley graphs.)

Definitions 3.8 and 3.9 are equivalent to the other usual definitions of (strong) relative hyperbolicity; see [Sis12, Proposition 4.9], [GM08, Theorem 3.25], and also [Bow12]. It is also equivalent to being ‘asymptotically tree-graded’, as in [DS05] and [Dru09].

Note that, as previously remarked, we do not assume peripheral sets are unbounded. We remark that the hyperbolicity of $\text{Cusp}(X, \mathcal{P})$ implies that neighbourhoods of different peripheral sets have bounded intersection, in particular two unbounded peripheral sets have infinite Hausdorff distance.

The fact that we may take any uniformly admissible horoballs in Definition 3.8 is allowed by the following result.

Proposition 3.10. *Suppose (X, d) is a geodesic metric space, \mathcal{P} a collection of uniformly coarsely connected subsets of X , and $\{\mathcal{H}(P)\}_{P \in \mathcal{P}}$ and $\{\mathcal{H}'(P)\}_{P \in \mathcal{P}}$ are two uniformly admissible collections of horoballs, with resulting cusped spaces $\text{Cusp}(X, \{\mathcal{H}(P)\})$ and $\text{Cusp}(X, \{\mathcal{H}'(P)\})$. Then there exists C so that the following holds.*

For $x \in P \cap \mathcal{H}(P)$ let $f(x) \in P \cap \mathcal{H}'(P)$ be a point with $d_X(x, f(x)) \leq C$ for uniform C . Define $F : \text{Cusp}(X, \{\mathcal{H}(P)\}) \rightarrow \text{Cusp}(X, \{\mathcal{H}'(P)\})$ by letting $F|_X = \text{id}_X$, and for $p \in \mathcal{H}(P) \subset \text{Cusp}(X, \{\mathcal{H}(P)\})$ with $d_{\mathcal{H}(P)}(p, \gamma_x(m)) \leq C$ for some $x \in P \cap \mathcal{H}(P)$, let $F(p) = \gamma'_{f(x)}(m)$. Here $\{\gamma_x\}_{x \in P \cap \mathcal{H}(P)}$, $\{\gamma'_x\}_{x \in P \cap \mathcal{H}'(P)}$ are the two families of geodesic rays of the construction.

Then $F : \text{Cusp}(X, \{\mathcal{H}(P)\}) \rightarrow \text{Cusp}(X, \{\mathcal{H}'(P)\})$ is a quasi-isometry; in particular, $\text{Cusp}(X, \{\mathcal{H}(P)\})$ is Gromov hyperbolic if and only if $\text{Cusp}(X, \{\mathcal{H}'(P)\})$ is.

Proof. We first show F is coarsely Lipschitz. Suppose we have $p, q \in \text{Cusp}(X, \{\mathcal{H}(P)\})$ connected by a geodesic $\gamma = [p, q]$. We wish to estimate $d_{\text{Cusp}(X, \{\mathcal{H}'(P)\})}(F(p), F(q))$ from above by building a curve $\gamma' \subset \text{Cusp}(X, \{\mathcal{H}'(P)\})$ from $F(p)$ to $F(q)$.

Let $\gamma_1, \dots, \gamma_n$ be the closed essentially disjoint subgeodesics of γ which have $\text{length}(\gamma_i) \geq 1$, each lie in some $\mathcal{H}(P_i)$, and have endpoints in $\mathcal{H}(P_i) \cap P_i$. The length condition ensures $n \leq \text{length}([p, q])$. If $p \in \mathcal{H}(P_0)$ for some P_0 let γ_0 be the subsegment of γ from p to the first time it meets X , otherwise let $\gamma_0 = \{p\}$. Likewise if $q \in \mathcal{H}(P_{n+1})$ for some P_{n+1} let γ_{n+1} be the last subsegment joining X to q , else let $\gamma_{n+1} = \{q\}$.

For each $i = 1, \dots, n$, replace γ_i by a curve γ'_i in $X \cup \mathcal{H}'(P_i)$ consisting of bounded length segments in X joining the endpoints of γ_i to nearby points in $\mathcal{H}'(P_i) \cap X$, and a geodesic in $\mathcal{H}'(P_i)$ joining those points. By Definition 3.4(4) applied to $\mathcal{H}(P_i)$ and $\mathcal{H}'(P_i)$, we have $\text{length}(\gamma'_i) \leq \text{length}(\gamma_i) + C_1$ for some constant C_1 .

If $p \in \mathcal{H}(P_0)$, then let $x \in P_0 \cap \mathcal{H}(P_0)$ be given with $d_{\mathcal{H}(P_0)}(p, \gamma_x(m)) \leq C$ for some $m \geq 0$. Write $F(p) = \gamma'_{f(x)}(m)$ as hypothesised. Suppose γ_0 joins p to some $z \in P_0 \cap \mathcal{H}(P_0)$, so

$$\text{length}(\gamma_0) \approx_C 2 \log(d_{P_0}(x, z)e^{-m} + 1) + m.$$

Now

$$\begin{aligned} d_{\mathcal{H}'(P_0)}(F(p), F(z)) &= d_{\mathcal{H}'(P_0)}(\gamma'_{f(z)}(m), z) \\ &\approx_C 2 \log(d_{P_0}(f(x), z)e^{-m} + 1) + m \approx \text{length}(\gamma_0) \end{aligned}$$

since $d_{P_0}(f(x), x) \leq C$. So we can replace γ_0 by a curve γ'_0 joining $F(p)$ to z with $\text{length}(\gamma'_0) \leq \text{length}(\gamma_0) + C_1$, possibly increasing C_1 . Do likewise with γ_{n+1} .

The rest of the curve $\gamma \setminus \bigcup \gamma_i$ consists of points of X and (countably many) non-trivial short paths in different horoballs $\mathcal{H}(P)$. By Definition 3.4(3), any path $[x, y] \subset \mathcal{H}(P)$ with endpoints in $\mathcal{H}(P) \cap X$ and $d_{\mathcal{H}}(x, y) \leq 1$ can be replaced by a path in X with the same endpoints of length $\leq L \text{length}([x, y])$, where $L \geq 1$.

Thus we have found a curve γ' joining $F(p)$ to $F(q)$ with

$$\begin{aligned} \text{length}(\gamma') &\leq \sum_{i=0}^{n+1} \text{length}(\gamma'_i) + \text{length}\left(\gamma' \setminus \bigcup \gamma'_i\right) \\ &\leq \sum_{i=0}^{n+1} (\text{length}(\gamma_i) + C_1) + L \text{length}\left(\gamma \setminus \bigcup \gamma_i\right) \\ &\leq L \text{length}(\gamma) + (n+2)C_1 \leq (L+C_1) \text{length}(\gamma) + 2C_1 \end{aligned}$$

as desired.

If we define $G : \text{Cusp}(X, \{\mathcal{H}'(P)\}) \rightarrow \text{Cusp}(X, \{\mathcal{H}(P)\})$ in a similar way to F , it also will be coarsely Lipschitz, and $G \circ F$ and $F \circ G$ are a bounded distance to the identity, since by hyperbolicity and Definition 3.4(4) of $\mathcal{H}(P)$, if $x, y \in P \cap \mathcal{H}(P)$ have $d_P(x, y) \approx 0$ then $d_{\mathcal{H}(P)}(\gamma_x(m), \gamma_y(m)) \approx 0$. So F is a quasi-isometry. \square

Remark 3.11. A very similar argument, replacing this time each γ_i by a path in X , allows us to show that the inclusion $\iota : X \rightarrow \text{Cusp}(X, \{\mathcal{H}(P)\})$ is τ -uniformly proper, with $\tau(t) = 2 \log(t + 1) - C$ for some C .

3.3. Visual metrics. Let Y be a proper, geodesic, Gromov hyperbolic space, with fixed basepoint $o \in Y$. One equivalent definition of the boundary $\partial_\infty Y$ is as the set of equivalence classes of geodesic rays $\gamma : [0, \infty) \rightarrow Y$, with $\gamma(0) = o$, where two rays are equivalent if they are at finite Hausdorff distance from each other. Let $(a|b) = (a|b)_o$ denote the Gromov product on $\partial_\infty Y$ with respect to o . Up to an additive error, $(a|b)$ equals the infimal distance from o to some (any) geodesic line from a to b , which exists since we assume Y is proper.

A metric ρ on $\partial_\infty Y$ is a visual metric with visual parameter $\epsilon > 0$ if there exists $C_0 > 0$ so that $\frac{1}{C_0} e^{-\epsilon(a|b)} \leq \rho(a, b) \leq C_0 e^{-\epsilon(a|b)}$ for all $a, b \in \partial_\infty Y$. Boundaries of cusped spaces will always be endowed with a visual metric.

Any two visual metrics ρ, ρ' on $\partial_\infty Y$ are *snowflake equivalent*: there exists $\lambda \geq 1$ and $\alpha > 0$ so that $\rho' \asymp_\lambda \rho^\alpha$ (in fact α is the ratio of their visual parameters).

Suppose that X is hyperbolic relative to a collection of subsets \mathcal{P} , and let $\text{Cusp}(X, \mathcal{P})$ be a cusped space. Endow $\partial_\infty \text{Cusp}(X, \mathcal{P})$ with a visual metric ρ with parameter $\epsilon > 0$. For each $P \in \mathcal{P}$ we have a horoball with a unique limit point $a_P \in \partial_\infty \text{Cusp}(X, \mathcal{P})$. We let $r_P = e^{-\epsilon d(o, P)}$ and call it the *radius* of the ‘‘shadow’’ of P in $\partial_\infty \text{Cusp}(X, \mathcal{P})$. (It is the approximate size of the set of limit points of all geodesic rays from o that pass close to P .) This makes

$$(\partial_\infty \text{Cusp}(X, \mathcal{P}), \rho, \{(a_P, r_P)\}_{P \in \mathcal{P}})$$

a shadow decorated metric space (see Definition 2.4).

For results on the geometric properties of $\partial_\infty \text{Cusp}(X, \mathcal{P})$, see [MS20].

3.4. Quasi-centres. We make frequent use of the notion of a ‘quasi-centre’.

Definition 3.12. *Suppose Y is a δ_Y -Gromov-hyperbolic geodesic space and $x, y, z \in Y \cup \partial_\infty Y$ are distinct. Then a quasi-centre for x, y, z is a point $p \in Y$ which lies within $2\delta_Y$ of geodesics joining x to y , y to z , and z to y .*

We record three elementary properties of quasi-centres.

Lemma 3.13 (Quasi-centres exist and are coarsely unique). *If Y is δ_Y -hyperbolic and proper (to ensure existence of bi-infinite geodesics between boundary points), any distinct $x, y, z \in Y \cup \partial_\infty Y$ have a quasi-centre p . Moreover, there exists C so that if a point $q \in Y$ lies within a distance A of geodesics $[x, y]$, $[y, z]$ and $[z, x]$, then $d(q, p) \leq C + A$.*

Proof. Consider $p \in [x, y]$ as p moves from x to y ; at a point where p switches from being $2\delta_Y$ -close to $[x, z]$ to being $2\delta_Y$ -close to $[z, y]$, p is a quasi-centre. For the last claim, consider a tree approximation $f : V \rightarrow T$ where V is the union of x, y, z, p, q and the geodesics between them, T is a tree, and f is a rough isometry, that is a quasi-isometry with only additive errors. Let $b \in T$ be the unique point on geodesics $[f(x), f(y)]$, $[f(y), f(z)]$ and $[f(z), f(x)]$. Then $f(p)$ is within C of b , and $f(q)$ is within $C + A$ of b for some C independent of x, y, z, q, A . \square

Lemma 3.14 (Quasi-centres preserved by quasi-isometries). *Let Y, Y' be Gromov hyperbolic and proper, and let $f : Y \rightarrow Y'$ be a quasi-isometric embedding. Then there exists C so that if p is a quasi-centre of $x, y, z \in Y \cup \partial_\infty Y$, then $f(p)$ is at most C from a quasi-centre of $\bar{f}(x), \bar{f}(y), \bar{f}(z)$. (Here \bar{f} denotes the extension of f to $\bar{f} : Y \cup \partial_\infty Y \rightarrow Y' \cup \partial_\infty Y'$.)*

Proof. Since p is within $2\delta_Y$ from geodesics between x, y, z , we have $f(p)$ is within bounded distance of quasi-geodesics between $\bar{f}(x), \bar{f}(y), \bar{f}(z)$; by the Morse lemma these are bounded Hausdorff distance from geodesics with the same endpoints. The distance estimate follows from Lemma 3.13. \square

Lemma 3.15 (Quasi-centres and boundary distances). *Suppose we have a proper Gromov hyperbolic space Y with basepoint o and a visual metric ρ with parameter ϵ . There exists C so that for any distinct $a, b \in \partial_\infty X$, and quasi-centre s of o, a, b , we have*

$$d(o, s) \approx_C \frac{-1}{\epsilon} \log \rho(a, b).$$

Proof. Both sides of the equation are $\approx (a|b)_o$. \square

3.5. Separation of horoballs. We now characterise the shadow decorations which arise from disjoint horoballs; both directions of this characterisation are needed in the proof of Theorem 1.5.

Recall that the *Busemann function* on a hyperbolic space Y corresponding to $a \in \partial_\infty Y, o \in Y$ is the function $\beta_a(\cdot, o) : Y \rightarrow \mathbb{R}$ defined by

$$\beta_a(x, o) = \sup_{\gamma} \left(\limsup_{t \rightarrow \infty} (d(\gamma(t), x) - t) \right),$$

where the supremum is taken over all geodesic rays $\gamma : [0, \infty) \rightarrow Y$ from o to a .

Remark 3.16. Suppose $Y = \mathbb{H}^2$ is the upper half-plane model, with $a = \infty$ the point at infinity, and $o = (0, 1)$. The geodesic ray from o to a is $\gamma(t) = (0, e^t)$. Given $x = (x_1, x_2)$ and $t \geq \log x_2$, the vertical path from x to (x_1, e^t) has length $t - \log x_2$, and the horizontal path from (x_1, e^t) to $(0, e^t)$ has length $|x_1|e^{-t}$, so

$$t - \log x_2 \leq d(\gamma(t), x) \leq t - \log x_2 + |x_1|e^{-t},$$

thus $\beta_a(x, o) = -\log x_2$. Therefore for any $a \in \partial_\infty \mathbb{H}^2$, $o \in \mathbb{H}^2$ a set of the form $\{y \in Y : \beta_a(y, o) \leq C\}$ defines a (classical) horoball, and likewise for $Y = \mathbb{H}^n$.

For a space X hyperbolic relative to a collection of subsets \mathcal{P} , Definition 3.4(4) gives that for $P \in \mathcal{P}$ and $x \in P$, $\beta_{a_P}(x, o) \approx -\log d(o, P)$.

Lemma 3.17. *If a space X is hyperbolic relative to a collection of subsets \mathcal{P} and $Z = \partial_\infty \text{Cusp}(X, \mathcal{P})$ is a boundary with visual metric ρ and the shadow decoration $\{(a_P, r_P)\}_{P \in \mathcal{P}}$ described in Subsection 2.2, then there exists C so that for all $P \neq Q \in \mathcal{P}$,*

$$(3.18) \quad \rho(a_P, a_Q) \succeq_C \sqrt{r_P r_Q}.$$

Conversely, suppose Y is Gromov hyperbolic with basepoint o and boundary $Z = \partial_\infty Y$ having visual metric ρ with parameter ϵ . Suppose further that $\{(a_P, r_P)\}_{P \in \mathcal{P}}$ is any (abstract) shadow decoration on Z that satisfies (3.18) for some C . Then there exists t_0 so that for any $t \geq t_0$ the horoballs defined by

$$H_{P,t} = \left\{ y \in Y : \beta_{a_P}(y, o) \leq -t + \frac{1}{\epsilon} \log r_P \right\}, \quad P \in \mathcal{P}$$

are pairwise separated by a distance $\geq 2(t - t_0)$.

Proof. First, suppose $P \neq Q \in \mathcal{P}$, so by construction the horoballs $\mathcal{H}(P), \mathcal{H}(Q) \subset X$ are essentially disjoint (being combinatorial horoballs attached to distinct peripheral cosets). Let $p \in [o, a_P), q \in [o, a_Q)$ be points with $r_P = e^{-cd(o,p)}, r_Q = e^{-cd(o,q)}$, so they are roughly the closest points to o in $\mathcal{H}(P), \mathcal{H}(Q)$ respectively. Suppose $r_Q \leq r_P$.

If $[o, a_Q)$ passes through $\mathcal{H}(P)$ then the geodesics $[o, a_P), [o, a_Q)$ branch in $\mathcal{H}(P)$ and $d(o, q) \gtrsim d(o, p) + 2((a_P|a_Q) - d(o, p))$ (see Figure 1). Thus $(a_P|a_Q) \lesssim \frac{1}{2}(d(o, p) + d(o, q))$ and (3.18) follows.

If $[o, a_Q)$ does not pass through $\mathcal{H}(P)$ then $(a_P|a_Q) \lesssim d(o, p)$ and moreover $0 \leq d(p, q) \approx (d(o, p) - (a_P|a_Q)) + (d(o, q) - (a_P|a_Q))$ so $(a_P|a_Q) \lesssim \frac{1}{2}(d(o, p) + d(o, q))$ and (3.18) follows.

For the converse, it suffices to find t_0 so that the horoballs are pairwise disjoint, since in that case for any $t \geq t_0$, as $\beta_{a_P}(\cdot, o)$ is 1-Lipschitz $d(H_{P,t}, Y \setminus H_{P,t_0}) \geq t - t_0$, and likewise for $H_{Q,t}, H_{Q,t_0}$. Considering a geodesic from $H_{P,t}$ to $H_{Q,t}$ we have $d(H_{P,t}, H_{Q,t}) \geq 2(t - t_0)$.

Suppose for some t_0 we can find $P, Q \in \mathcal{P}$ so that $H_{P,t_0} \cap H_{Q,t_0} \neq \emptyset$. By the quasiconvexity of the horoballs, there must be a point z along the geodesic (a_P, a_Q) almost in the horoballs, i.e., with $\beta_{a_P}(z, o) \lesssim$




FIGURE 1. Disjoint horoballs

$-t_0 + \frac{1}{\epsilon} \log r_P$ and $\beta_{a_Q}(z, o) \lesssim -t_0 + \frac{1}{\epsilon} \log r_Q$. Let $y \in (a_P, a_Q)$ be a quasi-centre for o, a_P, a_Q ; without loss of generality, a_P, z, y, a_Q appear in this order. Then $\beta_{a_P}(z, o) \approx -d(o, y) - d(y, z)$ and $\beta_{a_Q}(z, o) \approx -d(o, y) + d(y, z)$ so by (3.18)

$$\begin{aligned} -2(a_P | a_Q) &\approx -2d(o, y) \approx \beta_{a_P}(z, o) + \beta_{a_Q}(z, o) \\ &\lesssim -2t_0 + \frac{1}{\epsilon}(\log r_P + \log r_Q) \lesssim -2t_0 - 2(a_P | a_Q), \end{aligned}$$

thus $t_0 \lesssim 0$. So choosing t_0 large enough we have a contradiction. \square

3.6. Uniform perfectness and visual completeness. When studying quasymmetric maps it is useful to have, near any point, points of approximately any possible distance away.

Definition 3.19. A metric space (Z, ρ) is uniformly perfect if for some $\lambda > 1$, for any $z \in Z$ and $0 < r \leq \text{diam } Z$ we have that $B(z, r) \setminus B(z, r/\lambda) \neq \emptyset$.

An easy observation is that

Lemma 3.20. (Z, ρ) is uniformly perfect if and only if there exists $\lambda \geq 1$ so that for any $B(z, r) \subset Z, r \leq \text{diam}(Z)$, we have $\text{diam } B(z, r) \asymp_\lambda r$.

For cusped spaces, there is the following related characterisation.

Definition 3.21. Suppose a space X is hyperbolic relative to \mathcal{P} . We will say that X is visually complete if for some (hence, any) $w \in \text{Cusp}(X, \mathcal{P})$ there exists $C \geq 0$ so that for any $x \in \text{Cusp}(X, \mathcal{P})$ there exist geodesic rays $[w, a), [w, b)$, for some distinct $a, b \in \partial_\infty \text{Cusp}(X, \mathcal{P})$,

so that x is within distance C of $[w, a)$, $[w, b)$ and some bi-infinite geodesic from a to b .

Recall that $\text{Cusp}(X, \mathcal{P})$ is *visual* if for some (hence any) basepoint w there exists $C \geq 0$ so that every point $x \in \text{Cusp}(X, \mathcal{P})$ is within distance C from some infinite geodesic ray from w .

Lemma 3.22. *Suppose a space X is hyperbolic relative to \mathcal{P} . Then X is visually complete if and only if $\text{Cusp}(X, \mathcal{P})$ is visual and $\partial_\infty \text{Cusp}(X, \mathcal{P})$ is uniformly perfect.*

Proof. We may assume that o is the basepoint of $\text{Cusp}(X, \mathcal{P})$, and that the visual metric on $Z = \partial_\infty \text{Cusp}(X, \mathcal{P})$ satisfies $\rho(a, b) \asymp e^{-\epsilon(a|b)}$.

If X is visually complete, $\text{Cusp}(X, \mathcal{P})$ is certainly visual. Let t be a constant to be determined. Given $a \in \partial_\infty \text{Cusp}(X, \mathcal{P})$ and $r \leq \text{diam}(Z)$, let x be a point on $[o, a)$ at distance $-\frac{1}{\epsilon} \log(r) + t$ from o ; this is always possible for $t \geq \frac{1}{\epsilon} \log(\text{diam}(Z))$. Then by visual completeness there exist $b, c \in \partial_\infty \text{Cusp}(X, \mathcal{P})$ so that $(b|a)$ and $(c|a)$ are $\gtrsim t - \frac{1}{\epsilon} \log(r)$ and $(b|c) \approx t - \frac{1}{\epsilon} \log(r)$. Therefore $\rho(a, b)$ and $\rho(a, c)$ are $\leq e^{-\epsilon t}$. We now fix t large enough to ensure that $b, c \in B(a, r)$. Since $\rho(b, c) \asymp e^{-\epsilon t}$ we have $\text{diam } B(a, r) \asymp_\lambda r$ for λ independent of a, r as desired.

Conversely, given $x \in \text{Cusp}(X, \mathcal{P})$, find a geodesic ray $[o, a)$ within C of x . By uniform perfectness, there exists $b \in \partial_\infty \text{Cusp}(X, \mathcal{P})$ with $\rho(a, b) \asymp e^{-\epsilon d_{\text{Cusp}}(o, x)}$ thus $(a|b) \approx d_{\text{Cusp}}(o, x)$ and so a, b satisfy the desired condition. \square

Relatively hyperbolic groups are always visually complete, and hence have uniformly perfect boundaries (for the latter fact and a variation on visual completeness, see also [HH20, Proposition 6.3]).

Proposition 3.23. *Let G be a group hyperbolic relative to its proper, infinite subgroups H_1, \dots, H_n . Then if G is not virtually cyclic it is visually complete.*

Proof. We will show that there exists D so that for any $x \in \text{Cusp}(G, \{H_i\})$ there exist three geodesic rays γ_i towards $a_i \in \partial_\infty(G, \{H_i\})$ emanating from x so that $(a_i|a_j)_x \leq D$ for any $i \neq j$. To check visual completeness for given o, x , observe that for (at least) two of the rays as above, say γ_1, γ_2 , we have $(a_1|o)_x, (a_2|o)_x \leq D'$, for a suitable D' . Setting $a_1 = a, a_2 = b$, visual completeness easily follows from the fact that concatenating geodesics in a hyperbolic space yields a quasi-geodesic with constants only depending on the hyperbolicity constant and the Gromov product at the concatenation point.

The rays can be chosen as follows. If $n = 0$ then, as is well known, as G is not virtually cyclic the boundary at infinity has at least three points, and we can move the quasi-centre of these three points to a bounded distance from any x using the action of G .

If $n \geq 1$ then for $x \in G$, we can take 3 rays towards the points at infinity of distinct horoballs each within uniformly bounded distance of x .

For $x = (g, m)$ in a graph horoball $\mathcal{H}(P)$, for convenience we describe 3 quasi-geodesic rays instead of 3 geodesic rays, and they are the following:

- a ray $\{g\} \times [m, \infty)$ towards the point at infinity of the horoball,
- the concatenation of a geodesic from (g, m) to g and a ray from g towards the point at infinity of some horoball $\mathcal{H}(P') \neq \mathcal{H}(P)$ within uniformly bounded distance of g , and
- the concatenation of a horizontal geodesic of length l with $1 \leq l \leq e$ connecting (g, m) to, say, (g', m) , a geodesic from (g', m) to g' and a ray from g' towards the point at infinity of some horoball $\mathcal{H}(P'') \neq \mathcal{H}(P), \mathcal{H}(P')$ within uniformly bounded distance of g' . \square

3.7. Uniform perfectness and shadow decorations. There are consequences of uniform perfectness for shadow decorated spaces. In the proof of Theorem 2.9, we use the lower bound in the following proposition to get polynomial distortion of the extension of the boundary map to the inside.

Proposition 3.24. *If $(Z, \rho, \mathcal{B}), (Z', \rho', \mathcal{B}')$ are shadow decorated metric spaces, and (Z, ρ) is bounded and uniformly perfect, then any shadow-respecting quasisymmetric embedding $h : Z \rightarrow Z'$ will have the exponents $\lambda_a, (a, r) \in \mathcal{B}$, uniformly bounded away from zero and infinity.*

Proof. The uniform perfectness and boundedness of Z implies h is bi-Hölder [TV80, Theorem 3.14]: there exists $\alpha \geq 1$ so that for all $x, y \in Z$,

$$\rho(x, y)^\alpha \preceq \rho'(h(x), h(y)) \preceq \rho(x, y)^{1/\alpha}.$$

Take $(a, r) \in \mathcal{B}$, with corresponding $(a', r') \in \mathcal{B}'$. By (uniform) perfectness there exists a sequence (b_i) in $Z \setminus \{a\}$ with $b_i \rightarrow a$. Applying Definition 2.6(3) with $b = b_i, c = a$ for all large i we have $\rho'(h(b_i), h(a)) \asymp r' r^{-\lambda_a} \rho(b_i, a)^{\lambda_a}$, so

$$\rho(b_i, a)^\alpha \preceq r' r^{-\lambda_a} \rho(b_i, a)^{\lambda_a} \preceq \rho(b_i, a)^{1/\alpha}.$$

Letting $i \rightarrow \infty$, we deduce that $\alpha \geq \lambda_a \geq 1/\alpha$. \square

The following proposition shows that shadow-respecting quasisymmetric embeddings between uniformly perfect, shadow decorated spaces form a category, motivating the definition of such maps. The only places we use this proposition are in the statement of Theorem 2.9 and its application to Corollary 1.3 (if h is a shadow-respecting quasisymmetry then so is h^{-1} and $\widehat{h \circ h^{-1}} \approx \widehat{\hat{h} \circ \hat{h}^{-1}}$). Therefore, as the proof is somewhat tedious, we suggest the reader skips it on a first reading.

Proposition 3.25. *Let $(Z_i, \rho_i, \mathcal{B}_i)$ for $i = 1, 2, 3$ be uniformly perfect, shadow decorated metric spaces.*

If $h_i : Z_i \rightarrow Z_{i+1}$ are shadow-respecting quasisymmetric embeddings for $i = 1, 2$, then $h_2 \circ h_1 : Z_1 \rightarrow Z_3$ is a shadow-respecting quasisymmetric embedding, quantitatively. The identity map is a shadow-respecting quasisymmetric embedding, so uniformly perfect shadow decorated spaces and shadow-respecting quasisymmetric embeddings form a category.

If $h : Z_1 \rightarrow Z_2$ is a shadow-respecting quasisymmetric embedding that is a quasisymmetry, then $h^{-1} : Z_2 \rightarrow Z_1$ is a shadow-respecting quasisymmetry.

Proof. We first show $h_2 \circ h_1$ is shadow-respecting. If h_i is an η_i -quasisymmetric embedding for $i = 1, 2$ then $h_2 \circ h_1$ is an $\eta_2 \circ \eta_1$ -quasisymmetric embedding.

(1,2) Given $(a_1, r_1) \in \mathcal{B}_1$, there exists $(a_2, r_2) \in \mathcal{B}_2$ and $(a_3, r_3) \in \mathcal{B}_3$ with $h_1(a_1) = a_2$ and $h_2(a_2) = a_3$. By uniform perfectness, there exists $p, q \in Z_2$ with $\rho_2(p, a_2) \leq r_2 \leq \rho_2(q, a_2) \asymp \rho_2(p, a_2)$. For (1), if $\rho_1(a_1, b) \leq r_1$, then $\rho_2(a_2, h_1(b)) \preceq r_2 \asymp \rho_2(a_2, p)$, so by quasisymmetry and (1) for h_2 ,

$$\rho_3(a_3, h_2 h_1(b)) \preceq \rho_3(a_3, h_2(p)) \preceq r_3.$$

Likewise for (2), if $\rho_1(a_1, b) \geq r_1$, then $\rho_2(a_2, h_1(b)) \succeq r_2 \asymp \rho_2(a_2, q)$, so by quasisymmetry and (2) for h_2 ,

$$\rho_3(a_3, h_2 h_1(b)) \succeq \rho_3(a_3, h_2(q)) \succeq r_3.$$

If $(a_3, r_3) \in \mathcal{B}_3$ has no $(a_1, r_1) \in \mathcal{B}_1$ with $h_2 h_1(a_1) = a_3$, then either $a_3 \notin h_2(Z_2)$ so $\rho_3(a_3, h_2 h_1(Z_1)) \geq \rho_3(a_3, h_2(Z_2)) \succeq r_3$, or $a_3 = h_2(a_2)$ for some $(a_2, r_2) \in \mathcal{B}_2$ but $a_2 \notin h_1(Z_1)$, in which case $\rho_2(a_2, h_1(Z_1)) \succeq r_2$, which again implies $\rho_3(a_3, h_2 h_1(Z_1)) \succeq r_2$.

(3) For $(a_i, r_i) \in \mathcal{B}_i$, $i = 1, 2, 3$ as in (1), suppose $b, c \in Z_1$ satisfy

$$(3.26) \quad \frac{\rho_1(a_1, b)^2}{r_1} \leq \rho_1(b, c) \leq \rho_1(a_1, b) \leq r_1.$$

If

$$(3.27) \quad \frac{\rho_2(a_2, h_1(b))^2}{r_2} \leq \rho_2(h_1(b), h_1(c)) \leq \rho_2(a_2, h_1(b)) \leq r_2$$

we are done, since

$$(3.28) \quad \frac{\rho_3(h_2 h_1(b), h_2 h_1(c))}{r_3} \asymp \left(\frac{\rho_2(h_1(b), h_1(c))}{r_2} \right)^{\lambda_{a_2}} \asymp \left(\frac{\rho_1(b, c)}{r_1} \right)^{\lambda_{a_1} \lambda_{a_2}}.$$

So it remains to consider when (3.27) fails.

We claim that if $\rho_1(a_1, b) \asymp r_1$ we are done. Indeed, if $\rho_1(a_1, b) \asymp r_1$ we can find $e \in Z_1$ with $\rho_1(a_1, e) \geq r_1$ and $\rho_1(a_1, e) \asymp r_1 \asymp \rho_1(a_1, b)$, so

by (1,2) and quasisymmetry we have

$$\rho_2(a_2, h_1(e)) \succeq r_2 \succeq \rho_2(a_2, h_1(b)) \asymp \rho_2(a_2, h_1(e)),$$

thus $\rho_2(a_2, h_1(b)) \asymp r_2$. Moreover, $\rho_1(a_1, b) \asymp r_1$ and (3.26) imply that $\rho_1(b, c) \asymp \rho_1(a_1, b)$ so by quasisymmetry,

$$\rho_2(h_1(b), h_1(c)) \asymp \rho_2(a_2, h_1(b)) \asymp r_2.$$

The same argument applied to h_2 then gives that $\rho_3(h_2h_1(b), h_2h_1(c)) \asymp \rho_3(a_3, h_2h_1(b)) \asymp r_3$. So both $\rho_3(h_2h_1(b), h_2h_1(c))/r_3$ and $\rho_1(b, c)/r_1$ are $\asymp 1$ giving (3.28), proving the claim that $\rho_1(a_1, b) \asymp r_1$ suffices.

Thus if $\rho_2(a_2, h_1(b)) > r_2$ by (1) we have $\rho_1(a_1, b) \asymp r_1$ and by the claim we are done. So suppose $\rho_2(a_2, h_1(b)) \leq r_2$. In particular, applying (3) to $h_1(b), a_2$ and b, a_1 we have

$$\frac{\rho_3(h_2h_1(b), a_3)}{r_3} \asymp \left(\frac{\rho_2(h_1(b), a_2)}{r_2} \right)^{\lambda_{a_2}} \asymp \left(\frac{\rho_1(b, a_1)}{r_1} \right)^{\lambda_{a_1}\lambda_{a_2}}.$$

Assume from now on that $\rho_2(a_2, h_1(b)) \leq r_2$.

If $\rho_2(h_1(b), h_1(c)) > \rho_2(a_2, h_1(b))$ then by quasisymmetry and (3.26) we have $\rho_2(h_1(b), h_1(c)) \asymp \rho_2(a_2, h_1(b))$. Thus by uniform perfectness there exists $e \in Z_2$ so that $\rho_2(h_1(b), h_1(c)) \asymp \rho_2(h_1(b), e) \leq \rho_2(a_2, h_1(b))$. If $\rho_2(a_2, h_1(b))^2/r_2 > \rho_2(h_1(b), e)$ then $\rho_2(a_2, h_1(b))/r_2 \succeq 1$ and again the claim above implies we are done. So by quasisymmetry and (3) applied to $h_1(b), e$ and b, c we have

$$\begin{aligned} (3.29) \quad & \frac{\rho_3(h_2h_1(b), h_2h_1(c))}{r_3} \asymp \frac{\rho_3(h_2h_1(b), h_2(e))}{r_3} \asymp \left(\frac{\rho_2(h_1(b), e)}{r_2} \right)^{\lambda_{a_2}} \\ & \asymp \left(\frac{\rho_2(h_1(b), h_1(c))}{r_2} \right)^{\lambda_{a_2}} \asymp \left(\frac{\rho_1(b, c)}{r_1} \right)^{\lambda_{a_1}\lambda_{a_2}}. \end{aligned}$$

Assume from now on that $\rho_2(h_1(b), h_1(c)) \leq \rho_2(a_2, h_1(b))$.

Finally, if $\rho_2(h_1(b), h_1(c)) < \rho_2(h_1(b), a_2)^2/r_2$ then since by (3.26)

$$\frac{\rho_2(h_1(b), h_1(c))}{r_2} \asymp \left(\frac{\rho_1(b, c)}{r_1} \right)^{\lambda_{a_1}} \geq \left(\frac{\rho_1(a_1, b)}{r_1} \right)^{2\lambda_{a_1}} \asymp \left(\frac{\rho_2(a_2, h_1(b))}{r_2} \right)^2$$

we have $\rho_2(h_1(b), h_1(c)) \asymp \rho_2(a_2, h_1(b))^2/r_2$. So again we can find $e \in Z_2$ with $\rho_2(h_1(b), e) \asymp \rho_2(h_1(b), h_1(c))$ and $\rho_2(a_2, h_1(b))^2/r_2 \leq \rho_2(h_1(b), e) \leq \rho_2(a_2, h_1(b)) \leq r_2$, and a similar argument to (3.29) shows (3.28) holds. This completes the proof of property (3) for h_2h_1 .

Now suppose $h : Z_1 \rightarrow Z_2$ is a shadow-respecting quasisymmetric embedding that is a quasisymmetry. As is standard, if h is an η -quasisymmetry, then h^{-1} is a η' -quasisymmetry, where we set $\eta'(t) = 1/\eta^{-1}(1/t)$. It remains to check (1,2,3).

(1,2) For $(a_1, r_1) \in \mathcal{B}_1, (a_2, r_2) \in \mathcal{B}_2$ with $h(a_1) = a_2$, by uniform perfectness pick $p, q \in Z_1$ with $\rho_1(p, a_1) \leq r_1 \leq \rho_1(q, a_1)$ and $\rho_1(p, a_1) \asymp \rho_1(q, a_1)$. Then (1,2) for h and quasisymmetry give

$$\rho_2(a_2, h(q)) \asymp \rho_2(a_2, h(p)) \preceq r_2 \preceq \rho_2(a_2, h(q)).$$

To see (1) for h^{-1} , given $b \in Z_2$, if $\rho_2(a_2, b) \leq r_2$ then $\rho_2(a_2, b) \asymp \rho_2(a_2, h(q))$ so by the quasisisymmetry of h^{-1} we have $\rho_1(a_1, h^{-1}(b)) \asymp \rho_1(a_1, q) \asymp r_1$ as desired. A similar proof gives (2).

(3) Suppose $\rho_2(a_2, h(b))^2/r_2 \leq \rho_2(h(b), h(c)) \leq \rho_2(a_2, h(b)) \leq r_2$. If $\rho_1(a_1, b)^2/r_1 \leq \rho_1(b, c) \leq \rho_1(a_1, b) \leq r_1$ then by (3) for h we have

$$(3.30) \quad \frac{\rho_1(b, c)}{r_1} \asymp \left(\frac{\rho_2(h(b), h(c))}{r_2} \right)^{1/\lambda_{a_1}}.$$

If $r_1 \leq \rho_1(a_1, b)$, then we have $r_2 \leq \rho_2(a_2, h(b)) \leq r_2$, so $\rho_2(h(b), h(c)) \asymp r_2 \asymp \rho_2(a_2, h(b))$, and by quasisisymmetry $\rho_1(b, c) \asymp \rho_1(a_1, b)$ and again (3.30) holds with both sides $\asymp 1$ so we are done.

Assume from now on that $\rho_1(a_1, b) \leq r_1$ so by (3) applied to b, a_1 we have $\rho_2(h(b), a_2)/r_2 \asymp (\rho_1(b, a_1)/r_1)^{\lambda_{a_1}}$.

If $\rho_1(b, c) > \rho_1(a_1, b)$ then by quasisisymmetry $\rho_1(b, c) \asymp \rho_1(a_1, b)$ and $\rho_2(h(b), h(c)) \asymp \rho_2(a_2, h(b))$, so

$$\frac{\rho_1(b, c)}{r_1} \asymp \frac{\rho_1(b, a_1)}{r_1} \asymp \left(\frac{\rho_2(h(b), a_2)}{r_2} \right)^{1/\lambda_{a_1}} \asymp \left(\frac{\rho_2(h(b), h(c))}{r_2} \right)^{1/\lambda_{a_1}}.$$

Finally, suppose $\rho_1(b, c) < \rho_1(a_1, b)^2/r_1$. By uniform perfectness we can find $e \in Z_1$ with $\rho_1(b, a_1)^2/r_1 \leq \rho_1(b, e) \leq \rho_1(b, a_1)$ and $\rho_1(b, a_1)^2/r_1 \asymp \rho_1(b, e)$, unless $\rho_1(b, a_1) \geq r_1$ which as we have seen suffices to give (3.30). Since $\rho_1(b, c) < \rho_1(b, e)$, quasisisymmetry and (3) for h applied to b, e gives:

$$(3.31) \quad \begin{aligned} \left(\frac{\rho_2(h(b), a_2)}{r_2} \right)^2 &\leq \frac{\rho_2(h(b), h(c))}{r_2} \leq \frac{\rho_2(h(b), h(e))}{r_2} \asymp \left(\frac{\rho_1(b, e)}{r_1} \right)^{\lambda_{a_1}} \\ &\asymp \left(\frac{\rho_1(b, a_1)}{r_1} \right)^{2\lambda_{a_1}} \asymp \left(\frac{\rho_2(h(b), a_2)}{r_2} \right)^2. \end{aligned}$$

Thus all \leq are \asymp , and $\rho_2(h(b), h(c)) \asymp \rho_2(h(b), h(e))$, which by quasisisymmetry gives $\rho_1(b, c) \asymp \rho_1(b, e)$. Then (3.31) implies (3.30) and we are done. \square

3.8. Transient sets. We consider again a space X hyperbolic relative to a collection of subsets \mathcal{P} .

Let μ, R be constants and α a geodesic in X . Denote by $deep_{\mu, R}(\alpha)$ the set of points p of α that belong to some subgeodesic $[x, y]$ of α with endpoints in $N_\mu(P)$ for some $P \in \mathcal{P}$ and so that $d(p, x), d(p, y) > R$. Denote by $trans_{\mu, R}(\alpha) = \alpha \setminus deep_{\mu, R}(\alpha)$ the set of *transient points* [Hru10, Definition 8.9], [Sis12, Definition 3.9].

We collect the following properties of transient and deep sets from [Hru10, Sis12, Sis13], using as well results in [DS05] which are however not phrased in terms of these notions.

Lemma 3.32. *Suppose a space X is hyperbolic relative to \mathcal{P} . For each $P \in \mathcal{P}$ denote by $\pi_P : X \rightarrow P$ a coarse closest point projection, i.e.*

a function so that $d(x, \pi_P(x)) \leq d(x, P) + 1$. There exist μ, R, D, t, C with the following properties: For all $x, y, z \in X$,

(1) [Relative Rips condition] We have

$$\text{trans}_{\mu, R}[x, y] \subseteq N_D(\text{trans}_{\mu, R}[x, z] \cup \text{trans}_{\mu, R}[z, y]).$$

(2) $\text{deep}_{\mu, R}[x, y]$ is contained in a disjoint union of subgeodesics of $[x, y]$ each contained in $N_{t\mu}(P)$ for some $P \in \mathcal{P}$ and called a deep component along P . Moreover, any geodesic connecting points of $N_L(P)$, $L \geq 1$ is contained in $N_{tL}(P)$.

(3) The endpoints of the deep component of $[x, y]$ along $P \in \mathcal{P}$ (if it exists) are C -close to $\pi_P(x), \pi_P(y)$.

(4) If for some $P \in \mathcal{P}$ we have $d(\pi_P(x), \pi_P(y)) > C$, then $[x, y]$ has a deep component along P of length at least $d(\pi_P(x), \pi_P(y)) - C$.

Proof. (1) For groups, this follows combining [Hru10, Proposition 8.13] and [Osi06, Theorem 3.26], while for general spaces this is [Sis12, Proposition 4.6(3)], see also [DS05, Corollary 4.27].

(2) This follows from uniform quasiconvexity of peripheral sets [DS05, Lemma 4.15] and the fact that neighbourhoods of distinct peripheral sets have bounded intersection [DS05, Theorem 4.1(α_1)], see [Sis12, Proposition 5.7(1)].

(3) By [Sis13, Lemma 1.13(1)] a geodesic from x to y that has points near P must travel from x , enter $N_\mu(P)$ near $\pi_P(x)$ and leave $N_\mu(P)$ near $\pi_P(y)$.

(4) By [Sis13, Lemma 1.15(1)] if $d(\pi_P(x), \pi_P(y))$ is large enough, any geodesic from x to y must pass close to $\pi_P(x)$ and $\pi_P(y)$, which implies the existence of a suitable deep component. \square

Convention 3.33. When we refer to transient points without explicit mention of the constants we always imply a choice of constants as in the lemma.

4. SUBEXPONENTIAL DISTORTION AND (RELATIVE) HYPERBOLICITY

The main goal of this section is to prove the following.

Theorem 4.1. *Let $f : X \rightarrow Y$ be a subexponentially distorted embedding, and suppose that Y is hyperbolic relative to a collection of subsets \mathcal{P} . Then there exists $\rho_0 \geq 0$ so that for all $\rho \geq \rho_0$ we have that X is hyperbolic relative to $\{f^{-1}(N_\rho(P))\}_{P \in \mathcal{P}}$.*

As we will see at the end of this section, this statement about relative hyperbolicity of metric spaces implies the group case stated in Corollary 1.4. It also gives one part of Theorem 1.5.

Corollary 4.2. *Let Y be a truncated real hyperbolic space, that is, \mathbb{H}^n with a family of open disjoint horoballs removed, endowed with its path metric. If a group G admits a subexponentially distorted embedding into Y then G is hyperbolic relative to virtually nilpotent subgroups.*

Proof. Recall that [Dru09, Theorem 1.5] says that if a group G has a Cayley graph which is, as a space, hyperbolic relative to some collection of subsets \mathcal{P} then G is also hyperbolic relative to a collection of subgroups, each of which is contained in a uniform neighbourhood of one of the subsets in \mathcal{P} . (Recall that Druţu’s terminology of being ‘asymptotically tree-graded’ is equivalent to being relatively hyperbolic.)

Thus by Theorem 4.1 we get that G is hyperbolic relative to subgroups each of which admits a uniformly proper embedding into some (polynomially growing) horospheres of the truncated space. In particular, each peripheral subgroup has (at most) polynomial growth, and is therefore virtually nilpotent [Gro81, vdDW84]. \square

4.1. Uniformly proper maps into hyperbolic spaces. Towards proving Theorem 4.1, we begin with two essentially standard facts.

Lemma 4.3. *Suppose Y' is a geodesic metric space. Then any uniformly proper, coarsely Lipschitz and coarsely surjective map $f : Y \rightarrow Y'$ is a quasi-isometry, quantitatively.*

A proof of this may be found, for example, in [DK18, Lemma 8.29]. This has the following well-known consequence (cf. [Kap01, Proposition 2.6]), which we include as a warm-up to Proposition 4.6.

Proposition 4.4. *Suppose Y and Y' are geodesic metric spaces, with Y' Gromov hyperbolic. If $f : Y \rightarrow Y'$ is a subexponentially distorted map, then f is a quasi-isometric embedding.*

Proof. It suffices to show that f restricted to any geodesic $[p, q] \subset Y$ is a quasi-isometric embedding, quantitatively. By the definition of subexponential distortion (Definition 2.1), f is coarsely Lipschitz with some constant C_1 , and there exists $\tau : [0, \infty) \rightarrow \mathbb{R}$ with $\lim_{t \rightarrow \infty} \tau(t)/\log(t) = \infty$ and for any $x, y \in X$,

$$\tau(d_Y(x, y)) \leq d_{Y'}(f(x), f(y)).$$

We now follow the proof of [BH99, Theorem III.H.1.7] to show that the Hausdorff distance between $f([p, q])$ and a geodesic $[f(p), f(q)] \subset Y'$ is at most R , for some R depending only on C_1, τ and the hyperbolicity constant $\delta_{Y'}$ of Y' . Let $c : [0, b] \rightarrow Y'$, where $b = d_Y(p, q)$, be the composition of the unit-speed parametrisation of $[p, q] \subset Y$ and f . For each interval I in $[0, 1], [1, 2], \dots, [[b], b]$, replace $c|_I$ with a geodesic segment in Y' with the same endpoints, to find a continuous rectifiable path $c' : [0, b] \rightarrow Y'$. This path c' satisfies $d_{Y'}(c(t), c'(t)) \leq C_2$ for all t , and is τ' -uniformly proper for $\tau'(t) = \tau(t) - C_2$, where C_2 is a constant which only depends on C_1 and τ . Moreover, for any $s, t \in [0, b]$, we have the length bound $l(c'|_{[s, t]}) \leq C_2|s - t|$.

Suppose $x_0 \in [f(p), f(q)]$ is a point which maximises $d_{Y'}(x, \text{im}(c'))$ for $x \in [f(p), f(q)]$, and let D be this maximal distance. Let $y \in [f(p), x_0] \subset [f(p), f(q)]$ be the point at distance $2D$ from x (if $d_{Y'}(f(p), x_0) <$

$2D$, set $y = f(p)$, and choose $z \in [x_0, f(q)]$ likewise. Now choose points $y' = c'(s), z' = c'(t)$ with $d_{Y'}(y, y'), d_{Y'}(z, z') \leq D$.

Let γ be the path obtained concatenating a geodesic $[y, y']$, the path $c'|_{[s,t]}$ and a geodesic $[z', z]$. This path has length $l(\gamma) \leq 2D + l(c'|_{[s,t]}) \leq 2D + C_2|s - t|$. Observe that

$$\tau'(|s - t|) \leq d_{Y'}(y', z') \leq d_{Y'}(y', y) + d_{Y'}(y, z) + d_{Y'}(z, z') \leq 6D,$$

while

$$C_2|s - t| \geq l(c'|_{[s,t]}) \geq d_{Y'}(y', z') \geq 2D.$$

By [BH99, Proposition III.H.1.6] we have

$$D - 1 \leq \delta_{Y'} |\log_2 l(\gamma)|.$$

So for large D (and hence large $|s - t|$) these bounds combine to get

$$\tau'(|s - t|) \leq 6D \leq 7(D - 1) \leq 7\delta_{Y'} \log_2(2C_2|s - t|)$$

which implies that D is bounded by some D_0 depending only on τ', C_2 and $\delta_{Y'}$.

To bound the Hausdorff distance between $f([p, q])$ and $[f(p), f(q)]$, it remains to show that c' lies in a D_1 -neighbourhood of $[f(p), f(q)]$ for some D_1 . Suppose that some $x = c'(u)$ satisfies $d_{Y'}(x, [f(p), f(q)]) > D_0$. Every point in $[f(p), f(q)]$ is within a distance of D_0 from c' , so by connectedness there exists a point $w \in [f(p), f(q)]$ and values $s < u$ and $t > u$ so that $d_{Y'}(w, c'(s)), d_{Y'}(w, c'(t)) \leq D_0$.

Therefore, $d_{Y'}(c'(s), c'(t)) \leq 2D_0$, and so by uniform properness $|s - t| \leq C_3$ for some C_3 depending only on D_0 and τ' . Because $s < u < t$, we have $|s - u| \leq C_3$ also, thus the distance between x and $[f(p), f(q)]$ is at most

$$d_{Y'}(x, w) \leq d_{Y'}(x, c'(s)) + d_{Y'}(c'(s), w) \leq C_2|u - s| + D_0 \leq C_2C_3 + D_0;$$

setting $D_1 = C_2C_3 + D_0$, we are done.

Since c' and $[f(p), f(q)]$ are at Hausdorff distance at most D_1 , we can adjust the values c' by at most D_1 to find $c'' : [0, d_Y(p, q)] \rightarrow [f(p), f(q)]$. This map is uniformly proper and coarsely Lipschitz. Since it maps endpoints to endpoints it is coarsely surjective, so by Lemma 4.3 it is a quasi-isometry, quantitatively. Because c'' is within finite distance of c , $f|_{[x,y]}$ is also a quasi-isometry, quantitatively. \square

4.2. Transient points are close to image paths. We fix the notation of Theorem 4.1 from now on. In particular, Y denotes a fixed relatively hyperbolic space with fixed constants for transient/deep points as in Lemma 3.32.

The basic observation that we will use to prove Theorem 4.1 is contained in the following lemma.

Lemma 4.5. *For every L there exists a constant K with the following property. Let γ be a geodesic in Y and let $\alpha : ([a, b] \subseteq \mathbb{R}) \rightarrow Y$ be*

a coarsely L -Lipschitz path with the same endpoints x, y as γ . Then every transient point $p \in \gamma$ satisfies $d(p, \alpha) \leq K \log(|b - a| + 1) + K$.

The proof is a minor variation on that of [BH99, Proposition III.H.1.6]: α splits into two subpaths with domain half the size of that of α and one can consider the triangle with endpoints the “middle” point of α and the endpoints of γ , apply the relative Rips condition, and then repeat. The details are left to the reader.

In this subsection we improve the simple bound one obtains by applying Lemma 4.5 to images in Y of quasi-geodesics in X . The idea is to reapply the lemma to subpaths.

Proposition 4.6. *Let γ be a geodesic in Y and let α be the image of a subexponentially distorted coarsely Lipschitz map $[0, T] \rightarrow Y$ with the same endpoints as γ . Then every transient point $p \in \gamma$ satisfies $d(p, \alpha) \leq E$, where E is a constant depending only on the distortion of α and on Y .*

Proof. Let $p \in \gamma$ be “the worst” transient point, meaning one maximising the distance from α . Set $\eta = d(p, \alpha)$. We can assume $\eta \geq 1$. Also, we assume that α is minimal in the sense that for any transient point p' on a geodesic with endpoints on α there is a point q' on α with $d(p', q') \leq \eta + 1$. Let x^\pm be the endpoints of γ .

Because α is subexponentially distorted, there exists $\tau : [0, \infty) \rightarrow \mathbb{R}$ so that for all $s, t \in [0, T]$ we have $\tau(|s - t|) \leq d(\alpha(s), \alpha(t)) \leq L|s - t| + L$, where $\tau(t)/\log t \rightarrow \infty$ as $t \rightarrow \infty$.

We will consider two cases (the second one of which is irrelevant to the special case of Theorem 4.1 for Y hyperbolic). The constant $M \geq 1$ is a large enough constant, depending on Y, \mathcal{P}, τ, L only, to allow us enough “space” for the arguments in Case 2 to go through.

Case 1. There exist transient points $p^\pm \in \gamma$ respectively preceding and following p so that $d(p, p^\pm) \in [\min\{d(x^\pm, p), 10\eta\}, M\eta]$.

In this case we choose q^\pm on α at distance at most $\eta + 1$ from p^\pm (choose $q^\pm = p^\pm$ if $p^\pm = x^\pm$). Consider the path β obtained concatenating in the suitable order geodesics from p^\pm to q^\pm and a subpath of α between q^- and q^+ .

Suppose $q^\pm = \alpha(t^\pm)$. Assuming γ is at least 20η -long, at least one of p^-, p^+ is at distance at least 10η from p . This combines with the fact that α is coarsely L -Lipschitz to give that

$$L|t^+ - t^-| + L \geq d(q^-, q^+) \geq 10\eta - 2(\eta + 1) \geq 6\eta,$$

so for large η we have $|t^+ - t^-| \geq (5/L)\eta$. On the other hand, certainly $d(p, \beta) \geq \eta$, and we can apply Lemma 4.5 to β to see that

$$(4.7) \quad \eta \leq d(p, \beta) \leq K \log(|t^+ - t^-| + 2\eta + 2 + 1) + K \leq K \log(|t^+ - t^-|) + K',$$

for constants K, K' . Using the distortion bound on α and the estimate $d(p^-, p^+) \leq 2M\eta$, we have

$$\tau(|t^+ - t^-|) \leq d(q^-, q^+) \leq d(p^-, p^+) + 2\eta + 2 \leq 6M\eta.$$

Our hypothesis on τ gives that $\log |t^+ - t^-| \leq o_{\eta \rightarrow \infty}(1)\tau(|t^+ - t^-|)$, where $o_{\eta \rightarrow \infty}(1)$ is a function that goes to zero as η (and hence $|t^+ - t^-|$) go to infinity. Combined with (4.7) this gives

$$\eta \leq K o_{\eta \rightarrow \infty}(1)\tau(|t^+ - t^-|) + K' \leq 6MK\eta o_{\eta \rightarrow \infty}(1) + K',$$

which in turn implies that η is bounded depending only on required data.

Case 2. γ contains a deep component along, say, $P \in \mathcal{P}$ on one side of p with one endpoint at distance $> M\eta$ from p and the other one at distance $< 10\eta$ from p .

Suppose that the deep component in the statement lies before p . We would like to reduce to Case 1 by suitably changing α and γ . As π_P is coarsely Lipschitz there exists $x' \in \alpha$ so that $d(p, \pi_P(x')) \in [M\eta/4, M\eta/2]$ (in fact, the values of π_P along α vary between $\pi_P(x^-)$ and $\pi_P(x^+)$, and the latter coarsely coincides with $\pi_P(p)$, which in turn coarsely coincides with p). Choose $q \in \alpha$ so that $d(p, q) \leq \eta + 1$. Any geodesic $[x', q]$ contains a transient point p^- close to $\pi_P(x')$, and p^- is within distance $\eta + 1$ from some $x \in \alpha$ (recall that we chose a minimal α).

Also, any geodesic γ' from p^- to the final point of γ contains some transient point p' close to p . Notice that the hypotheses of Case 1 are now satisfied by p and γ' “on one side” of p' . Up to possibly reapplying the argument above to “the other side” of γ' to obtain a new geodesic γ'' and point p'' on it, we end up in the situation of Case 1. \square

4.3. Morse lemma for transient sets. As seen in Proposition 4.4, the image of a subexponentially distorting map from an interval into a hyperbolic space lies within a bounded Hausdorff distance of a geodesic with the same endpoints. The following proposition shows an analogous property for the transient sets of geodesics. We will not need this to prove Theorem 4.1 (but we will use it later). The proof builds on Proposition 4.6. Recall Definition 2.2 for the notion of a map coarsely respecting peripherals.

Proposition 4.8. *Suppose $f : X \rightarrow X'$ is a subexponentially distorting map between relatively hyperbolic spaces $(X, \mathcal{P}), (X', \mathcal{P}')$ that coarsely respects peripherals, and fixed deep/transient set constants as in Lemma 3.32. Then for any $x, y \in X$, we have that $f(\text{trans}(x, y))$ and $\text{trans}(f(x), f(y)) \subset X'$ are at Hausdorff distance $\leq C$, where C is independent of $x, y \in X$, and depends only on the data of relative hyperbolicity, subexponential distortion, coarsely respecting peripherals and deep/transient constants.*

We require the following estimate on projections.

Lemma 4.9. *Suppose we are in the situation of Proposition 4.8. Then there exists D with the following property: for each $x \in X$ and $P \in \mathcal{P}$, we have $d(f(\pi_P(x)), \pi_{P'}(f(x))) \leq D$, where $P' \in \mathcal{P}'$ is given by Definition 2.2(1).*

Proof. Denote $y = \pi_P(x)$, and consider a geodesic $[x, y]$. Then $f \circ [x, y]$ connects $f(x)$ to $f(y)$, and the latter lies in the C -neighbourhood of P' , where C is as in the definition of f coarsely respecting peripherals. For later purposes, we increase C so that Lemma 3.32 applies with constant C . Suppose by contradiction that we have $d(f(y), \pi_{P'}(f(x))) > D$, where D is a sufficiently large constant to be determined below (we note the order of choice of constants is C, E, C', D). Whenever $D > C$, any geodesic $[f(x), f(y)]$ contains a transient point C -close to $\pi_{P'}(f(x))$ by Lemma 3.32-(3)-(4). In particular, $f \circ [x, y]$ contains a point $f(z)$ which is $(E + C)$ -close to $\pi_{P'}(f(x))$, where E is as in Proposition 4.6. Note that $f(z)$ is $(D - E - C)$ -far from $f(y)$.

Now, since f coarsely respects peripherals we have $f^{-1}(N_{E+C}(P')) \subseteq N_{C'}(P)$, for some C' . Therefore, we have that z lies in the neighbourhood $N_{C'}(P)$, and since f is uniformly proper, for D sufficiently large we have $d(z, y) \geq C' + 3$. But then we see that $d(x, \pi_P(z)) < d(x, y) - 1$, which contradicts $y = \pi_P(x)$. \square

Proof of Proposition 4.8. The argument is a variation on the proof of Proposition 4.4.

Given $x, y \in X$, fix geodesics $[x, y] \subset X$ and $[f(x), f(y)] \subset X'$. Let $\alpha = f([x, y])$; then α is a subexponentially-distorted, coarse Lipschitz ‘quasi’-geodesic. By Proposition 4.6 every transient point of $[f(x), f(y)]$ lies at distance of at most E from α .

Let $\Pi : \text{trans}(f(x), f(y)) \rightarrow \alpha$ be a map which displaces each point by at most E . Suppose $[p^-, p^+]$ is a (maximal) connected interval of $\text{trans}(f(x), f(y))$. We claim that $[p^-, p^+]$ and $\alpha[\Pi(p^-), \Pi(p^+)]$ are at bounded Hausdorff distance from each other, where $\alpha[\Pi(p^-), \Pi(p^+)]$ denotes the subpath of α from $\Pi(p^-)$ to $\Pi(p^+)$. Indeed, the same argument as in the proof of Proposition 4.4 shows that as we follow z along $[p^-, p^+]$, the point $\Pi(z)$ has to follow α along $\alpha[\Pi(p^-), \Pi(p^+)]$ only omitting subpaths of controlled size.

What happens as we cross large deep components of $[f(x), f(y)]$? Suppose $[q^-, q^+]$ is a connected component of $\text{deep}(f(x), f(y))$ of length $\geq N$, where N is sufficiently large as determined below. Let $x^-, x^+ \in [x, y]$ be points with $f(x^\pm) = \Pi(q^\pm)$. Since f is coarsely Lipschitz, we have $d(x^-, x^+) \geq cN - C_1$; here all constants c, C_1, C_2, \dots are independent of N .

By the definition of deep components, q^- and q^+ lie C_2 -close to a peripheral set $P' \in \mathcal{P}'$. Since f coarsely respects peripherals, x^-, x^+ lie C_3 -close to a peripheral set $P \in \mathcal{P}$, and as $d(x^-, x^+)$ is large P and P'

uniquely correspond to each other in Definition 2.2. By Lemma 4.9 and Lemma 3.32(3) we have that $d(f(\pi_P(x)), q^-)$ and $d(f(\pi_P(y)), q^+)$ are at most C_4 , so the uniform properness of f gives $d(\pi_P(x), x^-), d(\pi_P(y), x^+) \leq C_5$. Therefore $d(\pi_P(x), \pi_P(y)) \geq cN - C_1 - 2C_5$, so for sufficiently large N we can apply Lemma 3.32(4) to $[x^-, x^+]$ to find that $[x^-, x^+]$ is within Hausdorff distance C_6 of a (the) deep component of $[x, y]$ along P .

By Lemma 3.32(2), $[x^-, x^+]$ lies in the C_7 -neighbourhood of P . Again, since f coarsely respects peripherals, $f([x^-, x^+])$ lies in the C_8 -neighbourhood of P' .

So, follow z along $[f(x), f(y)]$, and consider $f^{-1} \circ \Pi(z) \in [x, y]$: as we go along components of $\text{trans}(f(x), f(y))$ we have a coarsely Lipschitz map, while if we jump over a large deep component of $[f(x), f(y)]$ we jump over a corresponding deep component of $[x, y]$. This shows that $f(\text{trans}(x, y))$ is contained in a bounded neighbourhood of $\text{trans}(f(x), f(y))$.

It remains to check the converse inclusion. Suppose $[w^-, w^+]$ is a deep component in $[x, y]$ along $P \in \mathcal{P}$ of length $\geq N'$, where N' is sufficiently large to be determined, and P has not already been considered above, i.e. P does not correspond to a $P' \in \mathcal{P}'$ with a long deep component in $[f(x), f(y)]$. Then since $[x, y]$ is connected, $[w^-, w^+]$ coarsely lies in the image of $f^{-1} \circ \Pi$, in particular in the image of $\text{trans}(f(x), f(y))$. So we have points $y^\pm \in [f(x), f(y)]$ with $d(y^\pm, f(w^\pm)) \leq E$. By Lemma 3.32(3) we have $d(w^-, \pi_P(x)) \leq C$ and $d(w^+, \pi_P(y)) \leq C$. Thus by Lemma 4.9 we have

$$\begin{aligned} d(\pi_{P'}(f(x)), \pi_{P'}(f(y))) &\approx_{2D} d(f(\pi_P(x)), f(\pi_P(y))) \\ &\geq \tau(d(w^-, w^+) - 2C) \geq \tau(N' - 2C), \end{aligned}$$

where τ is the uniform properness function of f . Provided N' is large enough that $\tau(N' - 2C) - 2D \geq N$, we have a contradiction.

Thus, $\text{trans}(f(x), f(y))$ is contained in a bounded neighbourhood of the image of $[x, y]$ under f , omitting large deep components of $[x, y]$, which in turn lies within finite Hausdorff distance from $f(\text{trans}(x, y))$. \square

4.4. Proof of Theorem 4.1. We are now ready to prove Theorem 4.1.

It will be convenient to use the following equivalent characterisation of relative hyperbolicity. We note that it might be possible to use the cusped space characterisation of relative hyperbolicity directly, but this runs into issues with controlling geodesics in $\text{Cusp}(X, \mathcal{P})$ in terms of paths in X *without* knowing yet that X is relatively hyperbolic. Recall that being ‘asymptotically tree-graded’ with respect to a collection of subsets \mathcal{P} is equivalent to being hyperbolic relative to \mathcal{P} (in the sense of Definition 3.8).

Theorem 4.10 ([Dru09, Theorem 1.7]). *A geodesic metric space X is hyperbolic relative to a collection of subsets \mathcal{P} if and only if*

- for each K there exists B so that $\text{diam}(N_K(P) \cap N_K(Q)) \leq B$ for each distinct $P, Q \in \mathcal{P}$,
- there exists $\epsilon \in (0, 1/2)$ and $M \geq 0$ so that for each $P \in \mathcal{P}$ and $x, y \in X$ with $x, y \in N_{\epsilon d(x,y)}(P)$ any geodesic from x to y intersects $N_M(P)$,
- for every $c \geq 0$ there exist σ, δ so that every triangle with continuous $(1, c)$ -quasi-geodesic edges $\gamma_1, \gamma_2, \gamma_3$ satisfies either
 - (1) there exists a ball of radius δ intersecting all sides of the triangle,
 - or
 - (2) there exists $P \in \mathcal{P}$ with $N_\sigma(P)$ intersecting all sides of the triangle and the entrance (resp. exit) points x_i (resp. y_i) of the sides γ_i in (from) $N_\sigma(P)$ satisfy $d(y_i, x_{i+1}) \leq \delta$.

Proof of Theorem 4.1. We have a subexponentially distorting embedding $f : X \rightarrow Y$ into a space Y hyperbolic relative to a collection of subsets \mathcal{P} . We fix constants as in Lemma 3.32 for Y .

We will use a variation of the third property in Theorem 4.10 for Y .

Lemma 4.11. *For a geodesic space Y hyperbolic relative to a collection of subsets \mathcal{P} , there exists δ_0 so that the following holds. For any geodesic triangle in Y either*

- (1') there exists a ball of radius δ_0 intersecting the transient sets of all sides, or
- (2') there exists $P \in \mathcal{P}$ so that all sides have deep components along P .

Proof. Consider a geodesic triangle with sides $\gamma_i = [p_i, p_{i+1}]$, $i = 0, 1, 2$ modulo 3. By Lemma 3.32(1), every transient point along γ_0 is D -close to a transient point on either γ_1 or γ_2 , and hence one of the following holds.

(a) There are two transient points x_1, x_2 on γ_0 within distance $10C$ and so that $d(x_i, \text{trans}(\gamma_i)) \leq D$ (here C is given by Lemma 3.32). In this case, (1') holds.

(b) There is a deep component $[x_1, x_2]$ of γ_0 , say along P , of length at least $10C$ so that $d(x_i, \text{trans}(\gamma_i)) \leq D$. By Lemma 3.32(3) x_1 is C -close to $\pi_P(p_1)$ and x_2 is C -close to $\pi_P(p_0)$. Since $\pi_P(p_2)$ cannot be C -close to both $\pi_P(p_0)$ and $\pi_P(p_1)$, without loss of generality, γ_1 has a deep component along P as well, by Lemma 3.32(4). If γ_2 does as well, we are in case (2'), so suppose that it does not. Then $\pi_P(p_2)$ is C -close to $\pi_P(p_0)$. In particular, x_2 is close to $\pi_P(p_0)$, which is close to $\pi_P(p_2)$, which is close to a transient point of γ_1 . So, x_2 is within uniformly bounded distance from the transient sets of both γ_1 and γ_2 , and hence (1') holds. \square

For X , we first prove the third item in the equivalent definition of relative hyperbolicity given by Theorem 4.10, since this is the part of the proof where ρ_0 gets chosen.

Suppose we have $p_0, p_1, p_2 \in X$ and γ_i is a $(1, c)$ -quasi-geodesic connecting p_i to p_{i+1} (modulo 3) for $i = 0, 1, 2$, and consider a geodesic triangle in Y with vertices $f(p_1), f(p_2), f(p_3)$. If such triangle is as in Lemma 4.11(1') then by the lemma and Proposition 4.6, we see that each $f(\gamma_i)$ intersects a specified ball of radius $\delta_0 + E$, which easily implies that a ball of some radius δ meets each γ_i .

If the triangle is as in Lemma 4.11(2'), we show that (2) holds for the γ_i . Since all sides have deep components along P , and the endpoints of a deep component are transient points, in view of Proposition 4.6 we see that each γ_i intersects $Q = f^{-1}(N_\rho(P))$, for any $\rho \geq \rho_0$ where ρ_0 is chosen suitably large. For later purposes, we assume $\rho_0 \geq C + E$ with C from Lemma 3.32 and E from Proposition 4.6. Let now x_i, y_i be the entrance/exit points with respect to Q as in the statement of (2) (setting $\sigma = 0$). We will prove that $d(f(x_i), \pi_P(f(p_i)))$ can be uniformly bounded by, say, K , and similarly for y_i . This easily implies that $d(y_i, x_{i+1})$ is uniformly bounded, as required. Suppose that $d(f(x_i), \pi_P(f(p_i)))$ is large and take a subpath γ'_i of γ_i with final point x_i . Increasing K , we can assume $d(\pi_P(f(x_i)), \pi_P(f(p_i))) \geq C$ (if $f(x_i)$ is far from $\pi_P(f(p_i))$ then so is its projection, since $f(x_i)$ is close to P). In particular, $\pi_P(f(p_i))$ is $(C + E)$ -close to some point in $f(\gamma'_i)$, necessarily different from $f(x_i)$ if $K > C + E$, because $\pi_P(f(p_i))$ is C -close to a transient point on $[f(p_i), f(x_i)]$ and such transient point is E -close to $f(\gamma'_i)$ by Proposition 4.6. This contradicts the fact that x_i is the entrance point in $f^{-1}(N_\rho(P))$. This concludes the proof of the third property of Theorem 4.10.

The first item of Theorem 4.10 easily follows from the fact that the analogous property holds in Y and the uniform properness of f .

Let us now show the second property in Theorem 4.10. First of all, we claim that there exists K with the following property. Let $Q = f^{-1}(N_\rho(P))$ for some $P \in \mathcal{P}$ and let $x \in X$. Then for any $\bar{x} \in Q$ that satisfies $d(x, \bar{x}) \leq d(x, Q) + 1$, we have $d(f(\bar{x}), \pi_P(f(x))) \leq K$.

Let us first argue that the claim implies the second property, for $\epsilon = 1/3$ and a suitable M . In fact, let $Q = f^{-1}(N_\rho(P))$ for some $P \in \mathcal{P}$ and let x, y satisfy $x, y \in N_{d(x,y)/3}(P)$. We can also assume that $d(x, y)$ is large enough so that

$$d(f(\bar{x}), f(\bar{y})) \geq \tau(d(\bar{x}, \bar{y})) \geq \tau\left(\frac{1}{3}d(x, y) - 2\right) > 2K + C,$$

where τ is the uniform properness function of f . In this case by the claim we have $d(\pi_P(f(x)), \pi_P(f(y))) > C$, which in view of Proposition 4.6 and Lemma 3.32 implies that $f([x, y])$ passes $(C + E)$ -close to P , whence we get the second property.

Let us now prove the claim. If the claim did not hold, then any geodesic from $f(x)$ to $f(\bar{x})$ would have a transient point C -close to $\pi_P(f(x))$, and hence $f([x, \bar{x}])$ would contain a point $f(p)$ that is $(C + E)$ -close to P . Thus $d(f(p), f(\bar{x})) \geq K - (C + E)$ so as f is coarse

Lipschitz $d(p, \bar{x}) \geq \phi(K)$, for some diverging (linear) function $\phi(K)$. Also, $d(p, Q) \leq C'$, for some C' . But then

$$d(x, Q) \leq d(x, p) + C' \leq d(x, \bar{x}) + (C' - \phi(K)),$$

which for K large enough contradicts the choice of \bar{x} . \square

4.5. Truncated hyperbolic spaces. We now show our main application of Theorem 4.1. Namely, assuming Theorem 2.9, we now prove Theorem 1.5: a finitely generated group is hyperbolic relative to virtually nilpotent groups if and only if it admits a polynomially distorted embedding into some truncated real hyperbolic space, if and only if it admits a subexponentially distorted embedding into some such space.

Proof of Theorem 1.5. If G is a finitely generated group and G admits a subexponentially distorted embedding into some truncated real hyperbolic space, then by Corollary 4.2 G is hyperbolic relative to some collection of virtually nilpotent subgroups.

Conversely, suppose G is hyperbolic relative to virtually nilpotent subgroups H_1, \dots, H_n . It remains to construct a polynomially distorted embedding into some truncated real hyperbolic space. Let X be the Cayley graph of G and $\mathcal{P} = \{gH_i\}$ the collection of left cosets of each H_i , let $\text{Cusp}(X, \mathcal{P})$ be the corresponding cusped space with basepoint o , and endow $Z = \partial_\infty \text{Cusp}(X, \mathcal{P})$ with a visual metric $\tilde{\rho}$ with parameter $2\epsilon > 0$.

By [MS20, Proposition 4.5] $(Z, \tilde{\rho})$ is a doubling metric space (see the citation for details), so by Assouad's embedding theorem [Ass83] $(Z, \tilde{\rho}^{1/2})$ admits a bi-Lipschitz embedding into some \mathbb{R}^N , and hence (as it is bounded) into \mathbb{S}^N .

So we may assume that there is a bi-Lipschitz embedding $h : (Z, \rho) \rightarrow \mathbb{S}^N$, where $\rho = \tilde{\rho}^{1/2}$ is a visual metric on $Z = \partial_\infty \text{Cusp}(X, \mathcal{P})$ with parameter ϵ . Let $\mathcal{B} = \{(a_P, r_P = e^{-\epsilon d(o, P)})\}_{P \in \mathcal{P}}$ be the usual shadow decoration on Z . By Lemma 3.17, $\rho(a_P, a_Q) \succeq \sqrt{r_P r_Q}$ uniformly for all $P \neq Q \in \mathcal{P}$.

Recall that $\mathbb{S}^N = \partial_\infty \mathbb{H}^{N+1}$, and the Euclidean metric ρ' on \mathbb{S}^N is a visual metric with parameter 1 and any fixed basepoint o' . For $P \in \mathcal{P}$, let $a'_P = h(a_P)$. Since h is bi-Lipschitz, there exists C so that $\rho'(a'_P, a'_Q) \succeq_C \sqrt{r_P r_Q}$ for all $P \neq Q \in \mathcal{P}$. Apply Lemma 3.17 to find t_0 so that the horoballs

$$H_P = \{y \in \mathbb{H}^{N+1} : \beta_{a'_P}(y, o') \leq -t_0 + \log r_P\}, P \in \mathcal{P}$$

are separated in \mathbb{H}^{N+1} by a distance greater than 1. Let $X' = \mathbb{H}^{N+1} \setminus \bigcup_{P \in \mathcal{P}} \text{int}(H_P)$ (viewing \mathcal{P} now as an abstract index set), and consider $\{H_P\}_{P \in \mathcal{P}}$ as a uniformly admissible set of horoballs for X' (see Proposition 3.6), so then $\text{Cusp}(X', \{H_P\}) = \mathbb{H}^{N+1}$. We give $\mathbb{S}^N = \partial_\infty \text{Cusp}(X', \{H_P\})$ the standard shadow structure $\mathcal{B}' = \{(a'_P, r'_P)\}_{P \in \mathcal{P}}$ by setting $r'_P = e^{-d_{\mathbb{H}^{N+1}}(o', H_P)}$.

By the definition of H_P , $r'_P \asymp r_P$ uniformly. So if $\rho(a_P, b) \leq r_P$ for some $P \in \mathcal{P}$, $b \in Z$, then

$$\rho'(a'_P, h(b)) \asymp \rho(a_P, b) \leq r_P \asymp r'_P.$$

Likewise if $\rho(a_P, b) \geq r$ then $\rho'(a'_P, h(b)) \geq r'_P$, and h maps $\{a_P\}$ bijectively to $\{a'_P\}$, so properties (1) and (2) of Definition 2.6 hold. Property (3) holds easily with each $\lambda_{a_P} = 1$ as h is bi-Lipschitz and each $r'_P \asymp r_P$. Thus h is shadow-respecting as a map from $\partial_\infty(X, \mathcal{P}) \rightarrow \partial_\infty \text{Cusp}(X', \{H_P\})$.

We have $(Z, \rho, \mathcal{B}) = \partial_\infty(X, \mathcal{P})$ is visually complete by Lemma 3.23, as is $\text{Cusp}(X', \{H_P\}) = \mathbb{H}^{N+1}$. Therefore Theorem 2.9 gives a polynomially distorted embedding from X (the Cayley graph of G) to X' (a truncated hyperbolic space). Theorem 1.5 is proved. \square

4.6. Relative quasiconvexity. Finally, we apply the metric relative hyperbolicity result of Theorem 4.1 in the case of groups.

Corollary 1.4 *Let G be hyperbolic relative to $\{H_i\}$. If H is a subexponentially distorted subgroup of G then H is relatively quasiconvex in G , and so H is hyperbolic relative to subgroups $\{K_j\}$ each of which is the intersection of a conjugate (in G) of some H_i with H .*

Proof. We will use the characterisation (QC – 3) of relative quasiconvexity from [Hru10], and more specifically we will show that there is a constant C so that, given a geodesic γ in $\text{Cusp}(G, \{H_i\})$ with endpoints in H , we have that $\gamma \cap G$ is contained in the C -neighbourhood of H .

By Theorem 1.5, H is (as a metric space) relatively hyperbolic, and it is readily seen that the inclusion map coarsely respects peripherals. Therefore, Proposition 4.8 implies that, given a geodesic in G with endpoints in H , its transient set lies in a controlled neighbourhood of H . Now, $\text{trans}(x, y)$ coarsely coincides (in $\text{Cusp}(G, \{H_i\})$) with the intersection with G of a geodesic $[x, y]$ from x to y in $\text{Cusp}(G, \{H_i\})$ by [Hru10, Propositions 7.9 and 8.13]. Therefore, we see that the aforementioned characterisation applies.

Having shown H is relatively quasiconvex, the corollary follows from [Hru10, Theorem 9.1]. \square

5. EXTENDING QUASI-ISOMETRIES TO THE BOUNDARY

In this section, we prove Theorem 2.8. We fix relatively hyperbolic spaces (X, \mathcal{P}) and (X', \mathcal{P}') , and a map $f : X \rightarrow X'$ which is a snowflake on peripherals. As hypothesised in Theorem 2.8, throughout this section we assume that each $P \in \mathcal{P}$ is unbounded.

5.1. Extending to cusped spaces. Our first goal is to extend f to a quasi-isometric embedding $f_{\text{Cusp}} : \text{Cusp}(X, \mathcal{P}) \rightarrow \text{Cusp}(X', \mathcal{P}')$ between their cusped spaces.

First we observe that a rough λ -snowflake map between graphs (i.e. a map with (5.2) holding) extends to a rough λ -similarity between the corresponding horoballs (i.e. a map with (5.3) holding).

Lemma 5.1. *Suppose P and P' are coarsely connected metric spaces and $f : P \rightarrow P'$ satisfies*

$$(5.2) \quad \frac{1}{C}d_P(x, y)^\lambda - C \leq d_{P'}(f(x), f(y)) \leq Cd_P(x, y)^\lambda + C.$$

Consider admissible horoballs $\mathcal{H}(P)$ and $\mathcal{H}(P')$ with corresponding geodesic rays $\{\gamma_x\}_{x \in P \cap \mathcal{H}(P)}$ and $\{\gamma'_x\}_{x \in P' \cap \mathcal{H}(P')}$. Define a map $\bar{f} : \mathcal{H}(P) \rightarrow \mathcal{H}(P')$ by setting $\bar{f}(p) = \gamma'_{x'}(\lambda m)$ where for each $p \in \mathcal{H}(P)$ we choose $x \in P \cap \mathcal{H}(P)$ and $m \geq 0$ with $d_{\mathcal{H}(P)}(p, \gamma_x(m)) \leq C$, and $x' \in P'$ with $d_{P'}(f(x), x') \leq C$. Then \bar{f} is a quasi-isometry, in fact it is a rough λ -similarity: there exists C' so that for all $p, q \in \mathcal{H}(P)$,

$$(5.3) \quad d_{\mathcal{H}(P')}(\bar{f}(p), \bar{f}(q)) \approx_{C'} \lambda d_{\mathcal{H}(P)}(p, q)$$

Proof. By hyperbolicity and Definition 3.4(2) the map \bar{f} is coarsely onto. For $p, q \in \mathcal{H}(P)$, suppose we have $d_{\mathcal{H}(P)}(p, \gamma_x(m)) \leq C$ and $d_{\mathcal{H}(P)}(q, \gamma_y(n)) \leq C$ with $\bar{f}(p) = \gamma'_{x'}(\lambda m)$, $\bar{f}(q) = \gamma'_{y'}(\lambda n)$, $d_{P'}(f(x), x') \leq C$ and $d_{P'}(f(y), y') \leq C$. Then by Definition 3.4(4) we see that

$$\begin{aligned} d_{\mathcal{H}(P')}(\bar{f}(p), \bar{f}(q)) &\approx 2 \log \left(d_{P'}(x', y') e^{-\max\{\lambda m, \lambda n\}} + 1 \right) + |\lambda m - \lambda n| \\ &\approx 2 \log \left(d_P(x, y)^\lambda e^{-\lambda \max\{m, n\}} + 1 \right) + \lambda |m - n| \\ &\approx \lambda \left(2 \log \left(d_P(x, y) e^{-\max\{m, n\}} + 1 \right) + |m - n| \right) \\ &\approx \lambda d_{\mathcal{H}(P)}(p, q). \quad \square \end{aligned}$$

We now build f_{Cusp} .

Construction 5.4. Since $f : X \rightarrow X'$ coarsely respects peripherals, there exists $C > 0$ so that for all $P \in \mathcal{P}$ there exists $P' \in \mathcal{P}'$ with $f(P) \subset N_C(P')$. (Since $f(P)$ is unbounded, the choice of P' is unique.) So after adjusting f by a bounded distance, we have a rough snowflake map from P to P' . We then use Lemma 5.1 to define a map f_{Cusp} from $\mathcal{H}(P) \subset \text{Cusp}(X, \mathcal{P})$ to $\mathcal{H}(P') \subset \text{Cusp}(X', \mathcal{P}')$. These extensions combine to define $f_{\text{Cusp}} : \text{Cusp}(X, \mathcal{P}) \rightarrow \text{Cusp}(X', \mathcal{P}')$.

Proposition 5.5. *Let $(X, \mathcal{P}), (X', \mathcal{P}'), f : X \rightarrow X'$ be given by Theorem 2.8. Then f_{Cusp} as given by Construction 5.4 is a quasi-isometric embedding.*

Towards proving the proposition, we start with the following.

Lemma 5.6. *Under the assumptions of Proposition 5.5, f_{Cusp} is uniformly proper.*

Proof. We know that f is uniformly proper, and f_{Cusp} restricted to the horoballs is uniformly proper, so f_{Cusp} is also uniformly proper.

In slightly more detail, suppose $d_{\text{Cusp}(X', \mathcal{P}')} (f_{\text{Cusp}}(x), f_{\text{Cusp}}(y)) \leq C$ for some $x, y \in X$ and large C .

From the definition of f_{Cusp} , if f_{Cusp} maps points into a horoball $\mathcal{H}(P')$, $P' \in \mathcal{P}'$, then these points come from a horoball $\mathcal{H}(P)$, $P \in \mathcal{P}$. The choice of P is unique: for any such P , we have that $f(P)$ is unbounded and contained in a neighbourhood of P' , so $f^{-1}(N_{C'}(P'))$ is also unbounded for suitable C' . Since f coarsely respects peripherals, $f^{-1}(N_{C'}(P'))$ is contained in a neighbourhood of P , and therefore the choice of P is unique.

If $f_{\text{Cusp}}(x)$ is in a horoball $P' \in \mathcal{P}'$, and $d(f_{\text{Cusp}}(x), \partial P') \geq 2C$, then $f_{\text{Cusp}}(x)$ and $f_{\text{Cusp}}(y)$ both lie within the image of a single horoball $\mathcal{H}(P)$, $P \in \mathcal{P}$. Since f_{Cusp} restricted to $\mathcal{H}(P)$ is uniformly proper by Lemma 5.1, $d_{\text{Cusp}(X, \mathcal{P})}(x, y)$ is bounded by a uniform constant depending on C .

So we are reduced to the case that $f_{\text{Cusp}}(x), f_{\text{Cusp}}(y)$ lie within $2C$ of $X' \subset \text{Cusp}(X', \mathcal{P}')$, and so x and y both lie within C' of $X \subset \text{Cusp}(X, \mathcal{P})$, for suitable C' . Since X is uniformly distorted inside $\text{Cusp}(X, \mathcal{P})$ (Remark 3.11) we have that f_{Cusp} restricted to X is uniformly proper, and so $d_{\text{Cusp}(X, \mathcal{P})}(x, y)$ is uniformly bounded depending on C . \square

Proof of Proposition 5.5. Since f is coarsely Lipschitz and each horoball extension is coarsely Lipschitz, f_{Cusp} is coarsely Lipschitz.

We begin by considering points $x, y \in X$, with the goal of showing that f_{Cusp} is a quasi-isometry (with uniform constants) when restricted to any geodesic $[x, y] \subset \text{Cusp}(X, \mathcal{P})$. We will use that $\text{trans}(x, y)$ coarsely coincides (in $\text{Cusp}(X, \mathcal{P})$) with the geodesic $[x, y]$ from x to y in $\text{Cusp}(X, \mathcal{P})$ intersected with $X \subset \text{Cusp}(X, \mathcal{P})$; for groups this follows from [Hru10, Propositions 7.9 and 8.13] and for general space from [Sis12, Proposition 4.9].

By Proposition 4.8 the image of $\text{trans}(x, y)$ under f coarsely coincides with $\text{trans}(f(x), f(y)) \subset X'$. So $f_{\text{Cusp}}(\text{trans}(x, y))$ is contained in a neighbourhood of a geodesic $[f(x), f(y)] \subset \text{Cusp}(X', \mathcal{P}')$, again by the relation between geodesics in X' and $\text{Cusp}(X', \mathcal{P}')$ explained above.

On the other hand, by Construction 5.4 any maximal subgeodesic α of $[x, y]$ contained in a horoball is mapped to a quasi-geodesic in the corresponding horoball in $\text{Cusp}(X', \mathcal{P}')$ with endpoints close to $[f(x), f(y)]$. In particular $f(\alpha)$ lies close to $[f(x), f(y)]$.

Combined, we get that $f([x, y])$ lies in a uniformly bounded neighbourhood of $[f(x), f(y)]$. We can then think of a perturbation of $f_{\text{Cusp}}|_{[x, y]}$ as a map $g : [x, y] \rightarrow [f(x), f(y)]$. This map g is coarsely surjective, since the endpoints of $[f(x), f(y)]$ are in the image and f is coarsely Lipschitz. Because f_{Cusp} is uniformly proper (Lemma 5.6), g is also uniformly proper. So by Lemma 4.3 g is a quasi-isometric embedding, quantitatively.

Thus f_{Cusp} is a quasi-isometry when restricted to any $[x, y] \subset \text{Cusp}(X, \mathcal{P})$ for arbitrary $x, y \in X \subset \text{Cusp}(X, \mathcal{P})$, with uniform constants.

If we have, say, $x \in X$ and $y \in \text{Cusp}(X, \mathcal{P}) \setminus X$, then $y \in \mathcal{H}(P)$ for some $P \in \mathcal{P}$. The geodesic $[x, y]$ passes within bounded distance of $\pi_P(x) \in P$ by Lemma 3.32(3); let $z \in \mathcal{H}(P)$ be the last point of $[x, y] \cap X$. Again, $d_X(\pi_P(x), z)$ is uniformly bounded.

Choose some $z' \in P \cap \mathcal{H}(P)$ so that $d(y, \gamma_{z'}) \leq C$ where $\{\gamma_x\}_{x \in P \cap \mathcal{H}(P)}$ are the geodesic rays in the horoball (Definition 3.4). Let γ be the geodesic segment $[x, y] \cap \mathcal{H}(P)$ from z to y , and let $w \in \mathcal{H}(P)$ be a quasi-centre of z, y, a_P . By the construction of the horoballs, up to bounded distances γ follows γ_z from z to w , then follows $\gamma_{z'}$ from w back to y .

Thus y lies close to a geodesic from z to z' . By the hyperbolicity of $\text{Cusp}(X, \mathcal{P})$, there is a constant C_1 so that if $d(y, z) \geq C_1$ then both z and y are within bounded distance of a geodesic $[x, z'] \subset \text{Cusp}(X, \mathcal{P})$. As both x and z' are in $X \subset \text{Cusp}(X, \mathcal{P})$, the first case considered above shows that f_{Cusp} restricted to $[x, z']$ is a quasi-isometry, and hence so is f_{Cusp} restricted to $[x, y]$.

If y is within C_1 from z , then f_{Cusp} restricted to $[x, y]$ is within bounded distance of f_{Cusp} restricted to $[x, z]$ which again is a quasi-isometry by the argument of the first case.

Similar arguments apply to show that $f_{\text{Cusp}}|_{[x, y]}$ is a quasi-isometry when x and y lie in different horoballs in $\text{Cusp}(X, \mathcal{P})$. If x and y lie in the same horoball, the conclusion follows from Lemma 5.1. \square

5.2. Extending to boundaries. Having extended f to a quasi-isometric embedding $f_{\text{Cusp}} : \text{Cusp}(X, \mathcal{P}) \rightarrow \text{Cusp}(X', \mathcal{P}')$, we now show that the boundary map is controlled.

Recall that $Z = \partial_\infty(X, \mathcal{P})$ comes with a shadow decoration: ρ is a visual metric on $\partial_\infty(X, \mathcal{P})$, with visual parameter ϵ and constant C , and $\mathcal{B} = \{(a_P, r_P)\}_{P \in \mathcal{P}}$ is the collection of balls with for each $P \in \mathcal{P}$, $a_P \in \partial_\infty(X, \mathcal{P})$ is the boundary point corresponding to some (any) geodesic ray $\{\gamma_x\}_{x \in P \cap \mathcal{H}(P)}$ in the horoball $\mathcal{H}(P) \subset \text{Cusp}(X, \mathcal{P})$, and we set $r_P = e^{-\epsilon d_{\text{Cusp}(X, \mathcal{P})}(o, P)} \in (0, \infty)$.

Likewise, $\partial_\infty(X', \mathcal{P}')$ comes with the analogous data $\rho', \epsilon', C', \mathcal{B}' = \{(a_{P'}, r_{P'})\}_{P' \in \mathcal{P}'}$. It makes sense to ask whether $\partial_\infty f_{\text{Cusp}}$ is shadow-respecting with respect to these shadow decorations, and that is what we show in the following proposition.

Proposition 5.7. *Let $(X, \mathcal{P}), (X', \mathcal{P}'), f : X \rightarrow X'$ be given by Theorem 2.8, and f_{Cusp} by Construction 5.4. Then the induced map $\partial_\infty f_{\text{Cusp}} : \partial_\infty(X, \mathcal{P}) \rightarrow \partial_\infty(X', \mathcal{P}')$ is a shadow-respecting quasisymmetric embedding, quantitatively. If f is a quasi-isometry, then $\partial_\infty f_{\text{Cusp}}$ asymptotically $\frac{\epsilon'}{\epsilon}$ -snowflakes, where as above ϵ, ϵ' are the visual parameters of the metrics on $\partial_\infty(X, \mathcal{P}), \partial_\infty(X', \mathcal{P}')$.*

Proof. Since f_{Cusp} is a quasi-isometric embedding (Proposition 5.5) by [BS00, Theorem 6.5(2)], f_{Cusp} extends to a (power) quasisymmetric embedding $\partial_\infty f_{\text{Cusp}} : \partial_\infty(X, \mathcal{P}) \rightarrow \partial_\infty(X', \mathcal{P}')$. We may assume $f_{\text{Cusp}}(o) = o'$, by allowing our constants to depend on $d'(f_{\text{Cusp}}(o), o')$; changing the basepoint in the image space introduces controlled multiplicative errors in ρ' and $r_{P'}$. We prove the properties of Definition 2.6 in turn.

Claim (1):

First note that for any horoball $\mathcal{H}(P) \subset \text{Cusp}(X, \mathcal{P})$, f_{Cusp} is defined to send geodesic rays $\{\gamma_x\}$ in $\mathcal{H}(P)$ to geodesic rays $\{\gamma'_{x'}\}$ in a horoball $\mathcal{H}(P') \subset \text{Cusp}(X', \mathcal{P}')$, therefore $\partial_\infty f_{\text{Cusp}}$ maps a_P to $a_{P'}$.

For simplicity, we write $d = d_{\text{Cusp}(X, \mathcal{P})}$, and likewise $d' = d_{\text{Cusp}(X', \mathcal{P}')}$.

Assume we have $(a_P, r_P) \in \mathcal{B}$ and $b \in \partial_\infty(X, \mathcal{P}) = \partial_\infty \text{Cusp}(X, \mathcal{P})$, with $\rho(a_P, b) \leq r_P$. Let p be a quasi-centre of o, a_P, b and let x be a closest point of P to o . Let $b' = \partial_\infty f_{\text{Cusp}}(b)$, $p' = f_{\text{Cusp}}(p)$ and $x' = f_{\text{Cusp}}(x) = f(x)$.

Let $P' \in \mathcal{P}'$ be the peripheral space corresponding to $f(P)$, with $a_{P'} \in \partial_\infty(X, \mathcal{P})$ the limit point of $\mathcal{H}(P') \subset \text{Cusp}(X', \mathcal{P}')$. Since $f(P)$ is in a bounded neighbourhood of P' , $x' = f(x)$ is close to P' and by the Morse lemma also close to $[o', a_{P'})$, thus as $\mathcal{H}(P')$ is quasiconvex, x' is bounded distance from a closest point in P' to o' .

Since $\rho(a_P, b) \leq r_P$, we have $d(o, p) \approx (a_P|b) \gtrsim d(o, x)$. Since p, x are both close to a geodesic $[o, a_P)$, which is mapped by f_{Cusp} within bounded distance of $[o', a_{P'})$, we have $d'(o', p') \gtrsim d'(o', x')$. By Lemma 3.14 p' is approximately a quasi-centre for $o', a_{P'}, b'$. Thus $(a_{P'}|b') \approx d'(o', p') \gtrsim d'(o', x') \approx d'(o', P')$ so $\rho'(a_{P'}, b') \leq r_{P'}$ as desired.

Claim (2):

Suppose $\rho(a_P, b) \geq r_P$ for some $(a_P, r_P) \in \mathcal{B}, b \in \partial_\infty(X, \mathcal{P})$. Then the argument of Claim (1) applies, except this time the quasi-centre p of o, a_P, b has $d(o, p) \approx (a_P|b) \lesssim d(o, x)$ and so $(a_{P'}|b') \approx d'(o', p') \lesssim d'(o', x') \approx d'(o', P')$ and $\rho'(a_{P'}, b') \geq r_{P'}$.

Before completing the proof of (2), we state a useful lemma.

Lemma 5.8. *Suppose X is a geodesic metric space hyperbolic relative to \mathcal{P} and $\partial_\infty(X, \mathcal{P})$ has a metric with visual parameter ϵ . Suppose we have $P \in \mathcal{P}$ and $b \in \partial_\infty \text{Cusp}(X, \mathcal{P})$ with $b \neq a_P$. Let p denote a quasi-centre of o, a_P, b , and let $x \in P \subset \text{Cusp}(X, \mathcal{P})$ denote the closest point in P to o .*

There exists $C > 0$ so that if $\rho(b, a_P) \leq r_P/C$, we have $d(o, p) \geq d(o, x)$, any geodesic $[o, b)$ must intersect P , and

$$\rho(a_P, b) \asymp r_P \cdot e^{-\epsilon d(x, q)/2},$$

where q is the last point in P on $[o, b)$.

Proof. Since $\log \rho(b, a_P) \approx -\epsilon(b|a_P)$ and $\log r_P \approx -\epsilon d(o, P) = -\epsilon d(o, x)$, for any C' any C sufficiently large will have $(b|a_P) \geq d(o, x) + C'$. Since $(b|a_P) \approx d(o, p)$, choosing C' large enough ensures that $d(o, p) \geq d(o, x)$.

Denote by $\beta_{a_P}(\cdot, o)$ the Busemann function corresponding to a_P and o . As P is approximately a level set of $\beta_{a_P}(\cdot, o)$, if q is the last point in a geodesic $[o, b]$ to meet P then we have $\beta_{a_P}(q, o) \approx \beta_{a_P}(x, o) \approx -d(o, x)$. On the other hand, using that p is the quasi-centre of o, a_P, b , we have $\beta_{a_P}(q, o) \approx d(q, p) - d(p, x) - d(x, o)$.

Combined, these show that $d(p, q) \approx d(p, x)$. Since p approximately lies on a geodesic from x to q , we have $d(p, x) \approx d(q, x)/2$.

Applying this to $\log \rho(b, a_P) \approx -\epsilon(b|a_P)$, we have $(b|a_P) \approx d(o, p) \approx d(o, x) + d(x, p) \approx d(o, x) + d(q, x)/2$, and the conclusion follows since $-\epsilon d(o, x) \approx \log r_P$. \square

Suppose now we have $b \in \partial_\infty \text{Cusp}(X, \mathcal{P})$, and $P' \in \mathcal{P}'$ with $a_{P'} \notin h(Z)$. Let p' be a quasi-centre of $o', a_{P'}, b' = \partial_\infty f_{\text{Cusp}}(b)$, and let x' be a closest point in P' to o' . Then, in view of Lemma 5.8, we have either $\rho'(a_{P'}, b') \geq r_{P'}/C$, proving (2), or $\rho'(a_{P'}, b') \asymp r_{P'} e^{-\epsilon' d'(x', q')/2}$, where q' is the last point in a geodesic $[o', b']$ to meet P' .

Suppose that the latter holds. Since $[o', b']$ and $f_{\text{Cusp}}([o, b])$ are at bounded Hausdorff distance, there exist $x, q \in \text{Cusp}(X, \mathcal{P})$ with $d'(f_{\text{Cusp}}(x), x')$ and $d'(f_{\text{Cusp}}(q), q')$ uniformly bounded. Since f_{Cusp} sends points far in horoballs in $\text{Cusp}(X, \mathcal{P})$ far into horoballs in $\text{Cusp}(X, \mathcal{P})$, and $x', q' \in X' \subset \text{Cusp}(X', \mathcal{P}')$, we may assume that x and q lie in $X \subset \text{Cusp}(X, \mathcal{P})$ by moving them a bounded distance.

Thus for a suitable C' , by Lemma 5.8,

$$\text{diam}(f(X) \cap N_{C'}(P')) \gtrsim d'(x', q') \approx \frac{2}{\epsilon'} \log \left(\frac{r_{P'}}{\rho'(a_{P'}, b')} \right).$$

If, by contradiction, $r_{P'}/\rho'(a_{P'}, b') > A$ for some large A then by the uniform properness and coarsely respecting peripherals properties of f_{Cusp} and f there exists $P \in \mathcal{P}$ with $x, q \in f^{-1}(N_{C'}(P')) \subset N_{C''}(P)$. In the definition of f_{Cusp} , we extend f by mapping $\mathcal{H}(P)$ to some $\mathcal{H}(\hat{P}')$, $\hat{P}' \in \mathcal{P}'$. But as x', q' are in a bounded neighbourhood of \hat{P}' also, by hyperbolicity we have $\hat{P}' = P'$ assuming A is large enough. This is a contradiction, so $\rho'(a_{P'}, b') \geq r_{P'}/A$.

Claim (3):

Suppose, as hypothesised, that we have $(a, r) = (a_P, r_P) \in \mathcal{B}$ and points $b, c \in \partial_\infty \text{Cusp}(X, \mathcal{P})$ so that

$$\left(\frac{\rho(a, b)}{r} \right) \rho(a, b) \leq \rho(b, c) \leq \rho(a, b) \leq r.$$

This implies that

$$(5.9) \quad 2(a|b) - d(o, P) \gtrsim (b|c) \gtrsim (a|b) \gtrsim d(o, P).$$

Let $x \in \text{Cusp}(X, \mathcal{P})$ be the closest point to o of P . Let x_b, x_c be the last points of $[o, b), [o, c)$ to meet P . Let $p_{a,b}$ be a quasi-centre of o, a, b , and likewise for other pairs of points.

By Lemma 5.8 and (5.9) we have that

$$d(x_b, o) \gtrsim d(p_{b,c}, o) \gtrsim d(p_{a,b}, o) \gtrsim d(o, P).$$

Let a', b', c' , etc., be the images under f_{Cusp} of the corresponding points in $\text{Cusp}(X', \mathcal{P}') \cup \partial_\infty \text{Cusp}(X', \mathcal{P}')$. By Lemma 3.14 we have that $p'_{a,b}$, the image of $p_{a,b}$, is within bounded distance of the quasi-centre of o', a', b' , and likewise for the other quasi-centres.

Moreover, x is characterised up to bounded distance as the point of intersection of P and $[o, a)$. Since f_{Cusp} maps $[o, a)$ within bounded Hausdorff distance of $[o', a')$, and also P close to P' , we have that $x' = f_{\text{Cusp}}(x)$ is within bounded distance the intersection of P' and $[o', a')$, and thus of the closest point to o' in P' . Similarly, since x_b is the final point of intersection of P and $[o, b)$, it is sent to roughly the final point of intersection of P' and $[o', b')$.

Putting this together, we have two cases for the configuration of a tree approximating o, a, b, c : in case (i) we have that $[o, c)$ breaks off from $[o, a)$ after $p_{a,b}$, while in case (ii) we have that $[o, c)$ breaks off from $[o, b)$ after $p_{a,b}$ but before x_b . See Figure 2.




FIGURE 2. Configurations of a, b, c

In case (i) we have $p_{a,b} \approx p_{b,c}$, and so $(a|b) \approx (b|c)$, thus

$$\frac{1}{-\epsilon} \log \frac{\rho(b, c)}{r} \approx (b|c) - d(o, P) \approx \frac{1}{2} d(x, x_b),$$

while since this configuration is preserved in the image we also have

$$\frac{1}{-\epsilon'} \log \frac{\rho'(b', c')}{r'} \approx (b'|c') - d(o', P') \approx \frac{1}{2} d'(x', x'_b).$$

Now since f is a snowflake on peripherals, there exists λ_P so that f restricted to P is a rough λ_P -snowflake map. This means that

$$d_{X'}(x', x'_b) \cong d_X(x, x_b)^{\lambda_P}.$$

(Recall that the notation $A \cong_{\lambda, \mu} B$ means $A/\lambda - \mu \leq B \leq \lambda A + \mu$.) However, distances in $\text{Cusp}(X, \mathcal{P}), \text{Cusp}(X', \mathcal{P}')$ for points in the boundary of horoballs are approximately double the logarithm of the corresponding distances in X, X' . Thus

$$d'(x', x'_b) \approx \lambda_P d(x, x_b).$$

Combining these results, we have

$$(5.10) \quad \frac{\rho'(b', c')}{r'} \asymp \left(\frac{\rho(b, c)}{r} \right)^{\lambda_P \epsilon' / \epsilon}.$$

Case (ii) is similar, but here

$$\frac{1}{-\epsilon} \log \frac{\rho(b, c)}{r} \approx (b|c) - d(o, P) \approx d(x, x_b) - d(x_b, x_c)/2,$$

and

$$\frac{1}{-\epsilon'} \log \frac{\rho'(b', c')}{r'} \approx (b'|c') - d(o', P') \approx d'(x', x'_b) - d'(x'_b, x'_c)/2.$$

Since distances between pairs of points in a horoball are again roughly scaled by λ_P , we have $d'(x', x'_b) - d'(x'_b, x'_c)/2 \approx \lambda_P (d(x, x_b) - d(x_b, x_c)/2)$, and so (5.10) again holds.

Finally, observe that if $\lambda_P = 1$, e.g. if f is a quasi-isometry, then by (5.10) $\partial_\infty f_{\text{Cusp}}$ asymptotically $\frac{\epsilon'}{\epsilon}$ -snowflakes. \square

Before completing the proof of Theorem 2.8, we note some properties of Definition 2.3.

Proposition 5.11. *If $(X, \mathcal{P}), (X', \mathcal{P}'), (X'', \mathcal{P}'')$ are relatively hyperbolic spaces, and $f : (X, \mathcal{P}) \rightarrow (X', \mathcal{P}')$ and $g : (X', \mathcal{P}') \rightarrow (X'', \mathcal{P}'')$ are snowflakes on peripherals then so is the composition $g \circ f : (X, \mathcal{P}) \rightarrow (X'', \mathcal{P}'')$. The identity map on a relatively hyperbolic space is a snowflake on peripherals. Therefore we have a category where objects are relatively hyperbolic spaces, and morphisms are snowflake on peripheral maps, where we consider $f : X \rightarrow X'$ and $g : X \rightarrow X'$ equivalent if we have $\sup_{x \in X} d_{X'}(f(x), g(x)) < \infty$.*

Moreover, if $f : (X, \mathcal{P}) \rightarrow (X', \mathcal{P}')$ is a snowflake on peripherals and is coarsely onto, then a (any) coarse inverse $g : (X', \mathcal{P}') \rightarrow (X, \mathcal{P})$ is also a snowflake on peripherals.

Proof. The properties of being polynomially distorted, coarsely respecting peripherals and being a snowflake on peripherals (which implies the former two properties), are all easily seen to be preserved under composition.

The only subtle point is the coarse inverse: if $f : (X, \mathcal{P}) \rightarrow (X', \mathcal{P}')$ is a snowflake on peripheral map that is coarsely surjective, then by Lemma 4.3 f is a quasi-isometry, and so any coarse inverse will also satisfy the snowflake on peripheral conditions (with $\lambda = 1$). \square

We now conclude our study of the properties of $\partial_\infty(\cdot)_{\text{Cusp}}$.

Proof of Theorem 2.8. The extension of f to a quasi-isometry f_{Cusp} is given by Proposition 5.5. That $\partial_\infty f_{\text{Cusp}}$ is a shadow-respecting quasisymmetric embedding, and asymptotically $\frac{\epsilon'}{\epsilon}$ -snowflakes if f is a quasi-isometry, follows from Proposition 5.7.

It remains to check the functorial properties. For $\text{id}_X : (X, \mathcal{P}) \rightarrow (X, \mathcal{P})$, the extension into horoballs satisfies $(\text{id}_X)_{\text{Cusp}} = \text{id}_{\text{Cusp}}$ and so $\partial_\infty(\text{id}_X)_{\text{Cusp}} = \text{id}_{\partial_\infty(X, \mathcal{P})}$. Suppose $f : (X, \mathcal{P}) \rightarrow (X', \mathcal{P}')$ and $g : (X', \mathcal{P}') \rightarrow (X'', \mathcal{P}'')$ are snowflake on peripheral maps. The construction of the extensions into cusped spaces implies that $g_{\text{Cusp}} \circ f_{\text{Cusp}}$ and $(g \circ f)_{\text{Cusp}}$ are equal on $X \subset \text{Cusp}(X, \mathcal{P})$. Given a peripheral set $P \in \mathcal{P}$, coarsely respecting peripherals implies that there is $P' \in \mathcal{P}'$ with $f(P)$ coarsely in P' , and $P'' \in \mathcal{P}''$ with $g(P')$ coarsely in P'' . Suppose f is a rough λ -snowflake on P , and g is a rough λ' -snowflake on P' ; consequently $g \circ f$ is a rough $\lambda\lambda'$ -snowflake on P . For $x \in P$, let $\gamma_x : [0, \infty) \rightarrow \text{Cusp}(X, \mathcal{P})$, $\gamma'_x : [0, \infty) \rightarrow \text{Cusp}(X', \mathcal{P}')$ and $\gamma''_x : [0, \infty) \rightarrow \text{Cusp}(X'', \mathcal{P}'')$ be unit-speed geodesic rays from x to a_P , from $f(x)$ to $a_{P'}$, and from $g(f(x))$ to $a_{P''}$ respectively. Construction 5.4 and hyperbolicity imply that for any $t \geq 0$, we have uniform bounds on $d_{\text{Cusp}(X', \mathcal{P}')} (f_{\text{Cusp}}(\gamma_x(t)), \gamma'_x(\lambda t))$, on $d_{\text{Cusp}(X'', \mathcal{P}'')} (g_{\text{Cusp}}(\gamma'_x(t)), \gamma''_x(\lambda' t))$, and on $d_{\text{Cusp}(X'', \mathcal{P}'')} ((g \circ f)_{\text{Cusp}}(\gamma_x(t)), \gamma''_x(\lambda\lambda' t))$. Consequently, $(g \circ f)_{\text{Cusp}}(x)$ and $g_{\text{Cusp}} \circ f_{\text{Cusp}}(x)$ are at bounded distance from each other.

Being at bounded distance, the boundary extensions $\partial_\infty(g \circ f)_{\text{Cusp}}$ and $(\partial_\infty g_{\text{Cusp}}) \circ (\partial_\infty f_{\text{Cusp}})$ are equal. \square

6. CHARACTERISATION OF BOUNDARY EXTENSIONS

In this section we prove Theorem 2.9. Namely, given a shadow-respecting η -quasisymmetric embedding $h : \partial_\infty(X, \mathcal{P}) \rightarrow \partial_\infty(X', \mathcal{P}')$ between boundaries of relatively hyperbolic spaces, we construct a polynomially distorted embedding $\hat{h} : X \rightarrow X'$ inducing h , and we show various further properties of this correspondence.

We fix the notation of Theorem 2.9. As before, we choose basepoints o, o' of $\text{Cusp}(X, \mathcal{P})$, $\text{Cusp}(X', \mathcal{P}')$, and denote the metrics on $\text{Cusp}(X, \mathcal{P})$, $\text{Cusp}(X', \mathcal{P}')$ by d, d' respectively.

6.1. Quasi-centre estimates. We now collect technical lemmas needed for the proof of Theorem 2.9.

Lemma 6.1. *Suppose a space X is hyperbolic relative to \mathcal{P} and is visually complete. Let $o \in \text{Cusp}(X, \mathcal{P})$ be a basepoint. For any $C' \geq 0$*

there exists C with the following property. Let $P \in \mathcal{P}$, $x \in \text{Cusp}(X, \mathcal{P})$, and s be a quasi-centre of a_P, x, o .

(1) If $x \in N_{C'}(\mathcal{H}(P))$, then

$$d(o, x) - d(o, P) \lesssim_C 2(d(o, s) - d(o, P));$$

if moreover $x \in N_{C'}(P)$ then we can replace \lesssim_C by \approx_C .

(2) If $x \in \text{Cusp}(X, \mathcal{P})$ satisfies for some K that

$$d(o, x) - d(o, P) \lesssim_K 2(d(o, s) - d(o, P)),$$

then $x \in N_{C+K}(\mathcal{H}(P))$. Moreover, if $d(o, x) - d(o, P) \approx_K 2(d(o, s) - d(o, P))$ then $x \in N_{C+K}(P)$.

Proof. (1) Any geodesic from o to x is the concatenation of a geodesic of length approximately $d(o, P)$ and a geodesic coarsely contained in $\mathcal{H}(P)$. The quasiconvexity of $\mathcal{H}(P)$ means that s coarsely lies in $\mathcal{H}(P)$, and considering tree approximations as in Figure 3 the conclusion follows.




FIGURE 3. Points in horoballs

(2) The Busemann function $\beta_{a_P}(\cdot, o)$ satisfies $\beta_{a_P}(x, o) \approx d(s, x) - d(o, s) \approx d(o, x) - 2d(o, s)$. We used that s is at distance coarsely equal to $(a_P|x)_o$ from o and coarsely equal to $(a_P|o)_x$ from x in the first coarse equality, and that it is close to geodesics from o to x in the second one. Using the hypothesis with \lesssim_K (the proof when we have \approx_K being analogous), we get $\beta_{a_P}(x, o) \lesssim_K -d(o, P)$, which implies $x \in N_{C+K}(\mathcal{H}(P))$, as required. \square

Lemma 6.2. *Suppose a space X is hyperbolic relative to \mathcal{P} and is visually complete. Let $o \in \text{Cusp}(X, \mathcal{P})$ be a basepoint. For any $C' \geq 0$ there exists C so that for every $x \in \text{Cusp}(X, \mathcal{P})$, there exist $a, b \in \partial_\infty \text{Cusp}(X, \mathcal{P})$ with the properties that*

$$(1) \rho(a, b) \asymp_C e^{-\epsilon d(o, x)},$$

(2) x is within distance C from a quasi-centre of a, b, o ,

and if $x \in N_{C'}(\mathcal{H}(P))$ for some $P \in \mathcal{P}$ then one can arrange that

$$(3) \rho(a, a_P)^2 / r_P \leq \rho(a, b) \leq \rho(a, a_P) \leq r_P, \text{ where } r_P = e^{-\epsilon d(o, P)} \text{ for } \epsilon \text{ the visual parameter of } \rho.$$

Proof. If x lies in no horoball, then as the first and second items follow directly from visual completeness we are done.

Now suppose $x \in N_{C'}(\mathcal{H}(P))$ for some $P \in \mathcal{P}$. Since horoballs are coarsely sub-level sets of Busemann functions, we have $\beta_{a_P}(x, o) \lesssim -d(o, P)$. We have two cases. First, let s be a quasi-centre for x, a_P, o and suppose for a large constant K_1 to be determined we have

$$(6.3) \quad d(s, o) + 2K_1 \leq d(x, o).$$

Let $x' \in [o, x]$ be chosen with $d(x', o) = d(x, o) - K_1 \geq d(s, o) + K_1$. Apply visual completeness to x' to find $a, b \in \partial_\infty \text{Cusp}(X, \mathcal{P})$ so that x' is within bounded distance from a quasi-centre for a, b, o , so (1,2) are satisfied. By (6.3) and Lemma 6.1(1)

$$(6.4) \quad d(o, s) + 2K_1 - d(o, P) \leq d(o, x) - d(o, P) \lesssim 2(d(o, s) - d(o, P)),$$

so $-d(o, P) \gtrsim 2K_1 - d(o, s)$ and $r_P = e^{-\epsilon d(o, P)} \gtrsim e^{-\epsilon d(o, s)} e^{\epsilon 2K_1}$. By the choice of x' we have $\rho(a, b) \asymp_{M_1} e^{-\epsilon d(x', o)} = e^{-\epsilon d(x, o)} e^{\epsilon K_1}$ for some M_1 . By Lemma 3.15 $\rho(a, a_P) \asymp_{M_2} e^{-\epsilon d(o, s)}$ for some M_2 . By (6.4),

$$\frac{\rho(a, a_P)^2}{r_P} \asymp \frac{e^{-\epsilon 2d(o, s)}}{e^{-\epsilon d(o, P)}} \preceq e^{-\epsilon d(o, x)}.$$

So for some M_3 we have

$$\begin{aligned} \frac{\rho(a, a_P)^2}{r_P} &\leq M_3 e^{-\epsilon d(o, x)}, \\ \frac{1}{M_1} e^{-\epsilon d(o, x)} e^{\epsilon K_1} &\leq \rho(a, b) \leq M_1 e^{-\epsilon d(o, x)} e^{\epsilon K_1} \leq M_1 e^{-\epsilon d(o, s)} e^{-\epsilon K_1}, \\ \frac{1}{M_2} e^{-\epsilon d(o, s)} &\leq \rho(a, a_P) \leq M_2 e^{-\epsilon d(o, s)} \leq M_2 e^{-\epsilon d(o, P)} e^{-2\epsilon K_1}, \text{ and} \\ e^{-\epsilon d(o, P)} &\leq r_P. \end{aligned}$$

Fixing K_1 so that $e^{\epsilon K_1} \geq \max\{M_1 M_3, M_1 M_2, \sqrt{M_2}\}$ we have shown (3) holds.

Second, if (6.3) fails, then x is within a bounded distance of $[o, a_P)$ (depending on K_1 , which we now regard as a constant). For a large K_2 to be determined, let $x' \in [o, a_P)$ be chosen with $d(x', o) = d(x, o) + K_2$. Applying visual completeness to x' , we find $a', b' \in \partial_\infty \text{Cusp}(X, \mathcal{P})$ so that x' is within bounded distance from a quasi-centre of a', b', o . Since $d(x', o) \approx (a'|b')_o \gtrsim \min\{(a'|a_P)_o, (a_P|b')_o\}$, swapping a' and b' if necessary, we may assume $(a'|a_P)_o \lesssim d(x', o)$.

Now find $x'' \in [o, a']$ with $d(x'', o) = d(x', o) + K_3 = d(x, o) + K_2 + K_3$ for some large K_3 to be determined. Let s'' be a quasi-centre for x'', a_P, o ; by construction it is coarsely x' , so $d(s'', o) \approx_{M_4} d(x', o)$ and $d(x'', o) \approx_{M_4} d(x', o) + K_3$ for some M_4 , so fixing $K_3 \geq 2K_1 + 2M_4$ we have that x'' and s'' satisfy (6.3). Moreover,

$$\begin{aligned} \beta_{a_P}(x'', o) &\approx -d(x', o) + d(x'', x') \approx -d(x, o) - K_2 + K_3 \\ &\lesssim -d(o, P) - K_2 + K_3, \end{aligned}$$

so provided we choose K_2 large enough depending on K_3 we can ensure that $x'' \in \mathcal{H}(P)$. Thus x'' , which is coarsely x , satisfies the hypotheses of the first case and so we can find the desired points a, b . \square

Lemma 6.5. *Suppose a space X is hyperbolic relative to \mathcal{P} and is visually complete, and fix a basepoint $o \in \text{Cusp}(X, \mathcal{P})$. There exists C with the following property. If the ray from o to $a \in \partial_\infty \text{Cusp}(X, \mathcal{P})$ contains a point $x \in \mathcal{H}(P)$, for some $P \in \mathcal{P}$, then*

$$\rho(a_P, a) \leq C r_P e^{-cd(x, X)}.$$

Proof. Using the definition of visual metric, we see that we have to prove $(a|_{a_P}) \gtrsim d(o, P) + d(x, X)$. Let $y, z \in X \cap \mathcal{H}(P)$ be the first and last points of the ray in $\mathcal{H}(P)$. Let s be a quasi-centre for o, a, a_P . There exists m, n so that s is bounded distance to both $\gamma_y(m)$ and $\gamma_z(n)$ where γ_y, γ_z are the geodesic rays to a_P in $\mathcal{H}(P)$ given by Definition 3.4. (By Definition 3.4(4), $m \approx n$.) Using tree approximations, we see there exists C so that either $d(x, \gamma_y(m')) \leq C$ for some $m' \leq m$ or $d(x, \gamma_z(n')) \leq C$ for some $n' \leq n$. Suppose the former; the latter case is analogous. Then by Definition 3.4(4) we have $m' + C \geq d(x, X) \gtrsim m'$ and so

$$(a|_{a_P}) \approx d(o, s) \approx d(o, y) + m \geq d(o, y) + m' \approx d(o, P) + d(x, X). \quad \square$$

6.2. Embedding of cusped spaces. We return to the proof of Theorem 2.9. In the remainder of this section, we denote by C_i, K_i constants that depend on the data only (including h). The index of the K_i 's will be reset to 1 at the end of each proof.

First of all, work of Bonk–Schramm allow us to extend h to the cusped spaces.

Theorem 6.6 (Bonk–Schramm). *Suppose $(X, \mathcal{P}), (X', \mathcal{P}')$ are two visually complete relatively hyperbolic spaces, and that $h : \partial_\infty(X, \mathcal{P}) \rightarrow \partial_\infty(X', \mathcal{P}')$ is a quasisymmetric embedding. Then there exists a (C_1, C_1) -quasi-isometric embedding $\hat{f} : \text{Cusp}(X, \mathcal{P}) \rightarrow \text{Cusp}(X', \mathcal{P}')$ so that $\partial_\infty \hat{f} = h$ and that $\hat{f}(o) = o'$.*

Proof. By Lemma 3.22 $\partial_\infty(X, \mathcal{P}) = \partial_\infty \text{Cusp}(X, \mathcal{P})$ (and $\partial_\infty(X', \mathcal{P}') = \partial_\infty \text{Cusp}(X', \mathcal{P}')$) are uniformly perfect, hence by [TV80, Corollary 3.12] h is a “power” quasisymmetric embedding in the sense of [BS00, §6]. Therefore there is a quasi-isometric embedding $\hat{f} : \text{Cusp}(X, \mathcal{P}) \rightarrow$

$\text{Cusp}(X', \mathcal{P}')$ so that $\partial_\infty \hat{f} = h$ by [BS00, Theorems 7.4, 8.2]. We change \hat{f} by defining $\hat{f}(o) = o'$; the resulting quasi-isometry constant C_1 depends only on h and the data of the spaces. \square

6.3. Images of horoballs. Recall that throughout this subsection, $h : \partial_\infty(X, \mathcal{P}) \rightarrow \partial_\infty(X', \mathcal{P}')$ is the shadow-respecting quasisymmetric embedding of Theorem 2.9, with $\text{Cusp}(X, \mathcal{P}), \text{Cusp}(X', \mathcal{P}')$ having basepoints o, o' and metrics d, d' . Additionally, now we fix the quasi-isometric embedding extension $\hat{f} : \text{Cusp}(X, \mathcal{P}) \rightarrow \text{Cusp}(X', \mathcal{P}')$ of Theorem 6.6.

For $P \in \mathcal{P}$, we denote by $h_\#(P)$ the element of \mathcal{P}' so that $h(a_P) = a_{h_\#(P)}$.

Lemma 6.7. *For each C there exists C_2 so that for any $P \in \mathcal{P}$ and $x \in N_C(\mathcal{H}(P))$, with $\lambda_P = \lambda_{a_P}$ from Definition 2.6(3), we have*

$$d'(o', \hat{f}(x)) \approx_{C_2} d'(o', h_\#(P)) + \frac{\epsilon}{e'} \lambda_P (d(o, x) - d(o, P)).$$

This lemma will only be applied with C depending only on the quasiconvexity of horoballs, hence C_2 will depend only on the data.

Proof. Let γ_P be a ray from o to a_P . By Lemma 6.2 there exist $a, b \in \partial_\infty \text{Cusp}(X, \mathcal{P})$ so that $\rho(a, b) \asymp e^{-\epsilon d(o, x)}$, x is K_1 -close to a quasi-centre of a, b, o , and $\rho(a, a_P)^2 / r_P \leq \rho(a, b) \leq \rho(a, a_P) \leq r_P$. Thus by Definition 2.6(3), we have

$$\frac{\rho'(h(a), h(b))}{r_{h_\#(P)}} \asymp_{K_2} \left(\frac{\rho(a, b)}{r_P} \right)^{\lambda_P}.$$

Hence, since quasi-centres are coarsely mapped to quasi-centres (see Lemma 3.14) and the distance of the quasi-centre of two boundary points and the basepoint from the basepoint controls the visual distance (Lemma 3.15), we get

$$\begin{aligned} d'(o', \hat{f}(x)) &\approx_{K_3} \frac{-1}{e'} \log \rho'(h(a), h(b)) \\ &\approx_{K_4} d'(o', P') + \frac{\epsilon \lambda_P}{e'} (d(o, x) - d(o, P)), \end{aligned}$$

as required. \square

Corollary 6.8. *There exists C_3 so that for each $P \in \mathcal{P}$ we have*

- (1) $\hat{f}(\mathcal{H}(P)) \subseteq N_{C_3}(\mathcal{H}(h_\#(P)))$,
- (2) $\hat{f}(P) \subseteq N_{C_3}(h_\#(P))$,
- (3) $\hat{f}(\mathcal{H}(P)^c) \subseteq N_{C_3}(\mathcal{H}(h_\#(P))^c)$.

Proof. If $x \in \mathcal{H}(P)$, then by Lemma 6.1 with $C' = 0$, we have $d(o, x) - d(o, P) \lesssim_{K_1} 2(d(o, s) - d(o, P))$, where s is a quasi-centre of a_P, x, o . By Lemma 6.7 with $C = 0$, we have

$$d'(o', \hat{f}(x)) \approx_{K_2} d'(o', h_\#(P)) + \frac{\epsilon \lambda_P}{e'} (d(o, x) - d(o, P)).$$

In view of Lemma 3.14, if s' is a quasi-centre of $\hat{f}(x), a_{h_\#(P)}, o'$ then s' coarsely coincides with $\hat{f}(s)$, so that again by Lemma 6.7 (with C chosen uniformly so that $s \in N_C(\mathcal{H}(P))$) we have

$$d'(o', s') \approx_{K_3} d'(o', h_\#(P)) + \frac{\epsilon\lambda_P}{\epsilon'}(d(o, s) - d(o, P)).$$

So,

$$\begin{aligned} d'(o', \hat{f}(x)) - d'(o', h_\#(P)) &\approx_{K_2} \frac{\epsilon\lambda_P}{\epsilon'}(d(o, x) - d(o, P)) \\ &\lesssim_{K_4} \frac{2\epsilon\lambda_P}{\epsilon'}(d(o, s) - d(o, P)) \\ &\approx_{2K_3} 2(d'(o', s') - d'(o', h_\#(P))). \end{aligned}$$

By Lemma 6.1(2) with $K = K_2 + K_4 + 2K_3$, we have that $\hat{f}(x)$ lies within a distance C_3 of $\mathcal{H}(h_\#(P))$.

The second part can be proven in the same way replacing \lesssim by \approx .

Finally, if the third part did not hold (for a much larger C_3 than in (2)), then there would be $x \in \mathcal{H}(P)^c$ whose image is “well-inside” $\mathcal{H}(h_\#(P))$. A geodesic ray from x towards a_P has image a quasi-geodesic ray with controlled constants, that starts “well-inside” $\mathcal{H}(h_\#(P))$, then enters $N_{C_3}(h_\#(P))$ (by (2)) and then goes to $a_{h_\#(P)}$. Such a quasi-geodesic ray does not exist. \square

In a similar spirit to parts (1) and (2) of the previous corollary, one also has:

Corollary 6.9. *For each $P \in \mathcal{P}$, $\hat{f}|_{\mathcal{H}(P)}$ is a C_4 -rough $(\epsilon\lambda_P/\epsilon')$ -similarity.*

Proof. For $x, y \in \mathcal{H}(P)$, the quasiconvexity of horoballs implies that a quasi-centre s of x, y, o is also contained in (a controlled neighbourhood of) $\mathcal{H}(P)$. By hyperbolicity,

$$d(x, y) \approx_{K_1} d(o, x) + d(o, y) - 2d(o, s),$$

and the same holds for pairs of points in $\text{Cusp}(X', \mathcal{P}')$.

If s' is a quasi-centre of $\hat{f}(x), \hat{f}(y), o'$, then bearing in mind Lemma 3.14, we get from Lemma 6.7 (with C chosen so that s, s' are in the C -neighbourhoods of $\mathcal{H}(P), \mathcal{H}(P')$):

$$\begin{aligned} d'(\hat{f}(x), \hat{f}(y)) &\approx_{K_1} \left(d'(o', \hat{f}(x)) - d'(o', h_\#(P)) \right) + \left(d'(o', \hat{f}(y)) \right. \\ &\quad \left. - d'(o', h_\#(P)) \right) - 2 \left(d'(o', s') - d'(o', h_\#(P)) \right) \\ &\approx_{K_2} \frac{\epsilon\lambda_P}{\epsilon'}(d(o, x) - d(o, P)) + \frac{\epsilon\lambda_P}{\epsilon'}(d(o, y) - d(o, P)) \\ &\quad - 2 \frac{\epsilon\lambda_P}{\epsilon'}(d(o, s) - d(o, P)) \\ &\approx_{\epsilon\lambda_P K_1/\epsilon'} \frac{\epsilon\lambda_P}{\epsilon'} d(x, y) \end{aligned}$$

as required. \square

Finally, Corollary 6.8(3) can be strengthened to:

Corollary 6.10. *If $P' \in \mathcal{P}'$ is so that $a_{P'} \notin h(\bigcup_{\mathcal{P}} a_P)$, then $f(X) \cap \mathcal{H}(P') \subseteq N_{C_5}(X')$. Moreover, $f(X) \subseteq N_{C_5}(X')$.*

Proof. Consider any $x \in \text{Cusp}(X, \mathcal{P})$. Then x lies within uniformly bounded distance from a geodesic ray from o to, say, $a \in \partial_\infty \text{Cusp}(X, \mathcal{P})$. Suppose that $\hat{f}(x)$ lies inside the horoball $\mathcal{H}(P')$ in $\text{Cusp}(X', \mathcal{P}')$ and at distance R from $X' \subseteq \text{Cusp}(X', \mathcal{P}')$.

By Lemma 6.5, we have $\rho(a_{P'}, h(a)) \leq K_1 r_{P'} e^{-\epsilon R}$. For R large enough, this contradicts Definition 2.6(2) which gives $\rho(a_{P'}, h(Z)) \geq r_{P'}/C$, proving the first part of the corollary.

Suppose now $x \in X$ and let us show that its image lies in a controlled neighbourhood of X' . If we had that $\hat{f}(x)$ lies well inside a horoball, say $\mathcal{H}(P')$, then the argument above would provide $P \in \mathcal{P}$ so that $h(a_P) = a_{P'}$. But then Corollary 6.8(3) implies that x cannot lie in $\mathcal{H}(P)^c$, contradicting $x \in X$. This shows the ‘‘moreover’’ part. \square

Announcement: We perturb \hat{f} up to bounded distance and assume from now on that it maps X into X' .

6.4. Projection terms of the distance formula. Our goal is now to show that $\hat{f}|_X : X \rightarrow X'$ has the required metric properties. In order to do so, we will need a version of the distance formula for relatively hyperbolic spaces, as described below. Since we will have to be careful about distinguishing between distances in a relatively hyperbolic space and in the corresponding cusped space, we start using the notation $d_X, d_{\text{Cusp}(X, \mathcal{P})}$ and similar to emphasise the difference.

We require some notation to describe the distance formula. Suppose a space Y is hyperbolic relative to \mathcal{Q} , with a fixed cusped space $\text{Cusp}(Y, \mathcal{Q})$. For $Q \in \mathcal{Q}$ and $x, y \in Y$, we set $\theta_Q(x, y)$ to be the following value. Choose any geodesic $[x, y]$ in $\text{Cusp}(Y, \mathcal{Q})$ from x to y . If $[x, y]$ does not intersect $\mathcal{H}(Q)$, then set $\theta_Q(x, y) = 0$. If it does, then consider the entrance and exit points $x', y' \in Q$, and let $\theta_Q(x, y) = d_Y(x, y)$. We emphasize that the distance is taken in Y , not in $\text{Cusp}(Y, \mathcal{Q})$.

Let $\{\{A\}\}_L$ denote A if $A > L$, and 0 otherwise.

Theorem 6.11 (Distance formula, [Sis13, Theorem 0.1]). *In the notation above, there exists L_0 so that for each $L \geq L_0$ there exist λ, μ so that the following holds. If $x, y \in Y$ then*

$$(6.12) \quad d_Y(x, y) \cong_{\lambda, \mu} \sum_{Q \in \mathcal{Q}} \{\{\theta_Q(x, y)\}\}_L + d_{\text{Cusp}(Y, \mathcal{Q})}(x, y).$$

Returning to our study of $\hat{f}|_X : X \rightarrow X'$, we will make a term-by-term comparison of the distance formulas in X and X' . The following lemma will be used to compare the θ terms.

Lemma 6.13. *Under the assumptions of Theorem 2.9 with extension $\hat{f}|_X : X \rightarrow X'$ as above, there exists C_6 with the following property. Let $P \in \mathcal{P}$. Then for each $x, y \in X$ we have*

$$\theta_{h_{\#}(P)}(\hat{f}(x), \hat{f}(y)) \approx_{C_6, C_6} \theta_P(x, y)^{\epsilon\lambda_P/\epsilon'}.$$

Moreover, if $P' \in \mathcal{P}'$ is so that $a_{P'} \notin h(\bigcup_{\mathcal{P}} a_P)$, then for each $x, y \in X$ we have $\theta_{P'}(\hat{f}(x), \hat{f}(y)) \leq C_6$.

Proof. Consider $x, y \in X$ and suppose that a geodesic $[x, y]$ goes deep into a horoball $\mathcal{H}(P)$ for $P \in \mathcal{P}$. Then, Corollary 6.8 (together with the fact that $\hat{f}([x, y])$ is a quasi-geodesic and hence stays close to a geodesic with the same endpoints) implies that the entrance and exit points of $[x, y]$ get mapped close to the entrance and exit point of $[\hat{f}(x), \hat{f}(y)]$ in $h_{\#}(P)$. Moreover, if $[x, y]$ does not go deep into $\mathcal{H}(P)$ then $[\hat{f}(x), \hat{f}(y)]$ does not go deep into $h_{\#}(P)$ either.

Hence, the conclusion follows using that \hat{f} is a rough similarity on $\mathcal{H}(P)$ (Corollary 6.9) and Definition 3.4(4).

To show the “moreover” part, notice that the geodesic from $\hat{f}(x)$ to $\hat{f}(y)$ cannot go deep into $\mathcal{H}(P')$ because of Corollary 6.10. \square

Corollary 6.14. *Under the assumptions of Lemma 6.13, for any $P \in \mathcal{P}$ we have $\epsilon\lambda_P/\epsilon' \leq 1$.*

Proof. Each $P \in \mathcal{P}$ is uniformly coarsely connected, and is unbounded since (X, \mathcal{P}) is visually complete: were $P \in \mathcal{P}$ bounded a_P would be an isolated point. Take a quasi-geodesic ray $\gamma : [0, \infty) \rightarrow P$. Then using the bounds of Lemma 6.13, taking large $t > 2C_6$ and large N , we have

$$\begin{aligned} d_X(\gamma(0), \gamma(Nt))^{\epsilon\lambda_P/\epsilon'} &\asymp d_{X'}(\hat{f}(\gamma(0)), \hat{f}(\gamma(Nt))) \\ &\leq \sum_{i=1}^N d_{X'}(\hat{f}(\gamma((i-1)t)), \hat{f}(\gamma(it))) \\ &\asymp \sum_{i=1}^N d_X(\gamma((i-1)t), \gamma(it))^{\epsilon\lambda_P/\epsilon'}. \end{aligned}$$

The left-hand side is $\asymp (Nt)^{\epsilon\lambda_P/\epsilon'}$, and the right-hand side is $\asymp Nt^{\epsilon\lambda_P/\epsilon'}$, so fixing t and letting $N \rightarrow \infty$ we have $\epsilon\lambda_P/\epsilon' \leq 1$. \square

Proof of Theorem 2.9. It is easily seen that $\hat{f}|_X$ is coarsely Lipschitz. By Proposition 3.24 we have $\lambda_P \geq \alpha > 0$ for all $P \in \mathcal{P}$, and some $\alpha \in (0, 1]$. In the following we use that if $r_i \geq 1$ and $\lambda_i \geq \alpha$ for $i = 1, 2, \dots$ then

$$(6.15) \quad \sum_i r_i^{\lambda_i} \geq \sum_i r_i^{\alpha} \geq \left(\sum_i r_i \right)^{\alpha}.$$

Using the distance formula Theorem 6.11, Lemma 6.13, and choosing T large enough we get

$$\begin{aligned}
d_{X'}(\hat{f}(x), \hat{f}(y)) &\cong_{K_1, K_1} \sum_{P' \in \mathcal{P}'} \{ \{ \theta_{P'}(\hat{f}(x), \hat{f}(y)) \} \}_T + d_{\text{Cusp}(X', \mathcal{P}')}(\hat{f}(x), \hat{f}(y)) \\
&= \sum_{P \in \mathcal{P}} \{ \{ \theta_{h_{\#}(P)}(\hat{f}(x), \hat{f}(y)) \} \}_T + d_{\text{Cusp}(X', \mathcal{P}')}(\hat{f}(x), \hat{f}(y)) \\
&\succsim_{K_2, K_2} \sum_{P \in \mathcal{P}} \{ \{ \theta_P(x, y)^{\epsilon \lambda_P / \epsilon'} \} \}_{T'} + d_{\text{Cusp}(X, \mathcal{P})}(x, y) \\
&\geq \left(\sum_{P \in \mathcal{P}} \{ \{ \theta_P(x, y) \} \}_{T'} + d_{\text{Cusp}(X, \mathcal{P})}(x, y) \right)^{(\epsilon / \epsilon') \inf_{P \in \mathcal{P}} \lambda_P} \\
&\cong_{K_2, K_2} d_X(x, y)^{(\epsilon / \epsilon') \inf_{P \in \mathcal{P}} \lambda_P},
\end{aligned}$$

where the fourth inequality follows from (6.15). Notice that we changed the threshold when passing from X to X' . Thus $\hat{h} = \hat{f}|_X$ is polynomially distorting, and coarsely respects peripherals by Corollary 6.8. That \hat{h} is a snowflake on peripherals follows from Lemma 6.13 and Corollary 6.14.

We have $h = \partial_\infty(\hat{h})_{\text{Cusp}}$: by construction $h = \partial_\infty \hat{f}$ and $\hat{f} = \hat{h}$ on X . For any geodesic ray $[o, a) \subset \text{Cusp}(X, \mathcal{P})$, either $a = a_P$ for some $P \in \mathcal{P}$, in which case $\partial_\infty(\hat{h})_{\text{Cusp}}(a) = h(a)$ by the construction of $(\hat{h})_{\text{Cusp}}$, or a is a limit of points in X , and so again $\partial_\infty(\hat{h})_{\text{Cusp}}(a) = \partial_\infty \hat{f}(a) = h(a)$.

Moreover, if h asymptotically $\frac{\epsilon'}{\epsilon}$ -snowflakes then for all $P \in \mathcal{P}$, $\lambda_P = \epsilon' / \epsilon$, so the bound above shows \hat{f} has at most linear distortion, i.e., \hat{f} is a quasi-isometric embedding.

Finally, we verify the uniqueness of \hat{h} and its functorial properties. Since (X, \mathcal{P}) is visually complete, any $x \in \text{Cusp}(X, \mathcal{P})$ is bounded distance to a quasi-centre for some $a, b \in \partial_\infty(X, \mathcal{P})$ and o . By Lemma 3.14 any quasi-isometry $\text{Cusp}(X, \mathcal{P}) \rightarrow \text{Cusp}(X', \mathcal{P}')$ which sends a, b to $h(a), h(b)$ must coarsely send a quasi-centre for a, b, o to a quasi-centre for $h(a), h(b), o'$; in other words, $\hat{f}(x)$ is uniquely defined up to bounded ($\text{Cusp}(X', \mathcal{P}')$) distance. Since $X' \subset \text{Cusp}(X', \mathcal{P}')$ is uniformly embedded, this means that $\hat{h} = \hat{f}|_X$ is uniquely defined up to bounded distance too.

Likewise, if $j : \partial_\infty(X', \mathcal{P}') \rightarrow \partial_\infty(X'', \mathcal{P}'')$ is another shadow-respecting quasisymmetric embedding, then by this same uniqueness argument $\widehat{j \circ h}$ and $\hat{j} \circ \hat{h}$ are at bounded distance from each other. Finally, for $\text{id} : \partial_\infty(X, \mathcal{P}) \rightarrow \partial_\infty(X, \mathcal{P})$, by uniqueness again, $\hat{\text{id}}$ and id_X are at bounded distance from each other. \square

REFERENCES

- [Ass83] P. Assouad. Plongements lipschitziens dans \mathbf{R}^n . *Bull. Soc. Math. France*, 111(4):429–448, 1983.

- [BH99] M. R. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.
- [Bow99a] B. H. Bowditch. Boundaries of geometrically finite groups. *Math. Z.*, 230(3):509–527, 1999.
- [Bow99b] B. H. Bowditch. Connectedness properties of limit sets. *Trans. Amer. Math. Soc.*, 351(9):3673–3686, 1999.
- [Bow12] B. H. Bowditch. Relatively hyperbolic groups. *Internat. J. Algebra Comput.*, 22(3):1250016, 66, 2012.
- [BS00] M. Bonk and O. Schramm. Embeddings of Gromov hyperbolic spaces. *Geom. Funct. Anal.*, 10(2):266–306, 2000.
- [DK18] Cornelia Druțu and Michael Kapovich. *Geometric group theory*, volume 63 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 2018. With an appendix by Bogdan Nica.
- [Dru09] Cornelia Druțu. Relatively hyperbolic groups: geometry and quasi-isometric invariance. *Comment. Math. Helv.*, 84(3):503–546, 2009.
- [DS05] C. Druțu and M. Sapir. Tree-graded spaces and asymptotic cones of groups. *Topology*, 44(5):959–1058, 2005. With an appendix by D. Osin and Sapir.
- [GK04] Erik Guentner and Jerome Kaminker. Exactness and uniform embeddability of discrete groups. *J. London Math. Soc. (2)*, 70(3):703–718, 2004.
- [GM08] D. Groves and J. F. Manning. Dehn filling in relatively hyperbolic groups. *Israel J. Math.*, 168:317–429, 2008.
- [GP13] Victor Gerasimov and Leonid Potyagailo. Quasi-isometric maps and Floyd boundaries of relatively hyperbolic groups. *J. Eur. Math. Soc. (JEMS)*, 15(6):2115–2137, 2013.
- [Gro81] Mikhael Gromov. Groups of polynomial growth and expanding maps. *Inst. Hautes Études Sci. Publ. Math.*, (53):53–73, 1981.
- [Gro13] B. W. Groff. Quasi-isometries, boundaries and JSJ-decompositions of relatively hyperbolic groups. *J. Topol. Anal.*, 5(4):451–475, 2013.
- [HH20] Brendan Burns Healy and G. Christopher Hruska. Cusped spaces and quasi-isometries of relatively hyperbolic groups, 2020.
- [Hru10] G. C. Hruska. Relative hyperbolicity and relative quasiconvexity for countable groups. *Algebr. Geom. Topol.*, 10(3):1807–1856, 2010.
- [HS17] David Hume and Alessandro Sisto. Groups with no coarse embeddings into hyperbolic groups. *New York J. Math.*, 23:1657–1670, 2017.
- [Kap01] I. Kapovich. The combination theorem and quasiconvexity. *Internat. J. Algebra Comput.*, 11(2):185–216, 2001.
- [Mey09] J. B. T. Meyer. *Uniformly perfect boundaries of Gromov hyperbolic spaces*. PhD thesis, University of Zurich, 2009.
- [MS20] J. M. Mackay and A. Sisto. Quasi-hyperbolic planes in relatively hyperbolic groups. *Ann. Acad. Sci. Fenn. Math.*, 45:139–174, 2020.
- [Osi06] D. V. Osin. Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems. *Mem. Amer. Math. Soc.*, 179(843):vi+100, 2006.
- [Pau96] F. Paulin. Un groupe hyperbolique est déterminé par son bord. *J. London Math. Soc. (2)*, 54(1):50–74, 1996.
- [Sis12] A. Sisto. On metric relative hyperbolicity. *arXiv e-prints:arXiv:1210.8081*, October 2012.

- [Sis13] Alessandro Sisto. Projections and relative hyperbolicity. *Enseign. Math. (2)*, 59(1-2):165–181, 2013.
- [Tes20] Romain Tessera. On coarse embeddings of amenable groups into hyperbolic graphs. *arXiv e-prints:arXiv:2010.07205*, October 2020.
- [TV80] P. Tukia and J. Väisälä. Quasisymmetric embeddings of metric spaces. *Ann. Acad. Sci. Fenn. Ser. A I Math.*, 5(1):97–114, 1980.
- [vdDW84] L. van den Dries and A. J. Wilkie. Gromov’s theorem on groups of polynomial growth and elementary logic. *J. Algebra*, 89(2):349–374, 1984.

SCHOOL OF MATHEMATICS, UNIVERSITY OF BRISTOL, BRISTOL, UK
Email address: `john.mackay@bristol.ac.uk`

DEPARTMENT OF MATHEMATICS, HERIOT-WATT UNIVERSITY, EDINBURGH, UK
Email address: `a.sisto@hw.ac.uk`