

REALISING ALL COUNTABLE GROUPS AS QUASI-ISOMETRY GROUPS

PAULA HEIM, JOSEPH MACMANUS, AND LAWK MINEH

ABSTRACT. Given any countable group G , we construct uncountably many quasi-isometry classes of proper geodesic metric spaces with quasi-isometry group isomorphic to G . Moreover, if the group G is a hyperbolic group, the spaces we construct are hyperbolic metric spaces.

We make use of a rigidity phenomenon for quasi-isometries exhibited by many symmetric spaces, called strong quasi-isometric rigidity. Our method involves the construction of new examples of strongly quasi-isometrically rigid spaces, arising as graphs of strongly quasi-isometrically rigid rank-one symmetric spaces.

INTRODUCTION

Quasi-isometries are among the principal morphisms in the study of large-scale structure of metric spaces, and their use is fundamental in many areas of geometry and geometric group theory. Just as one considers the group of isometries of a metric space as a natural group of automorphisms, one can consider the *quasi-isometry group* associated to a metric space. This is the group of equivalence classes of self-quasi-isometries of a metric space, taken up to finite distance in the supremum metric. The primary goal of this article is to address the question of which groups may be realised as quasi-isometry groups.

The quasi-isometry groups of many familiar metric spaces are quite wild, and their general theory is poorly understood. For instance, we have an understanding of $\text{QI}(\mathbb{R})$: it is infinite-dimensional [20], left-orderable and virtually a direct product [36], centreless [9], and contains – for example – continuum-rank free groups and Thompson’s group F [30]. Beyond this, little is known; see [37] for a survey on $\text{QI}(\mathbb{R})$ and [27, 6] for some results on $\text{QI}(\mathbb{R}^n)$ for $n > 1$. The quasi-isometry groups of various solvable groups have also been computed [15, 14, 34]. Of course, quasi-isometries of spaces induce isomorphisms on quasi-isometry groups.

Most unbounded geodesic spaces seem to have tremendously large and complicated quasi-isometry groups. On the other hand, there are somewhat pathological metric spaces with trivial quasi-isometry groups, such as the set of factorials $\{n! : n \in \mathbb{N}\}$ equipped with the induced metric as a subset of \mathbb{R} , though this space is not even coarsely connected. To the best of our knowledge there were previously no known examples of any metric spaces with non-trivial finite quasi-isometry groups. We show that the maximum diversity is in fact achieved among countable groups.

Theorem A. *Let G be a countable group. Then there exists uncountably many quasi-isometry classes of proper geodesic metric spaces X with $G \cong \text{QI}(X)$. Moreover, if G is a hyperbolic group, then we may take every such X to be hyperbolic.*

The above result fits into a long history of so-called *representation problems*: given a class of objects \mathcal{C} , can one determine which groups G can be realised as $\text{Aut}(X)$ for an object $X \in \mathcal{C}$? An early instance of this question was posed by König [26], who asked whether any group can be realised as the automorphism group of a

Date: January 9, 2026.

graph. König's question was settled for finite groups by Frucht in 1939 [17], with an elegant and simple construction, now appearing in many standard texts on graph theory. Frucht's result was later extended to all groups independently by de Groot [10] and Sabidussi [29], and will play a large role in the proof of Theorem A.

We give a brief and non-exhaustive account of similar representation problems from geometry. Every finite group can be realised as the isometry group of a closed hyperbolic manifold in every dimension: see Greenberg for dimension 2 [19], Kojima for dimension 3 [25], and Belolipetsky–Lubotzky for higher dimensions [4]. In the dimension 2 case, this was extended to all countable groups by independently by Winkelmann [33] and Allcock [1]. Remarkably, there also exists a single surface (of infinite type), such that varying the Riemannian metric can produce any countable group as the corresponding isometry group [2]. Frigerio and Martelli have also shown that any countable group arises as the mapping class group of a hyperbolic 3-manifold [16].

Our strategy will be to exploit a rigidity phenomenon for quasi-isometries in order to reduce to a setting where we can apply a variant of Frucht's theorem. We say that a metric space X is *strongly quasi-isometrically rigid* if the natural map $\text{Isom}(X) \rightarrow \text{QI}(X)$ is a surjection. That is to say, every quasi-isometry of X is close to an isometry X . Strongly quasi-isometrically rigid spaces form one of the few classes of metric spaces for which the quasi-isometry group admits a clear description, provided one understands the isometry group. Note that essentially all spaces that are known to be strongly quasi-isometrically rigid satisfy a slightly stronger, uniform property: see Definition 1.3.

Strong quasi-isometric rigidity seems to be an intrinsic property of most symmetric spaces. Notably, Pansu showed that the octonionic plane $\mathbb{O}\mathbf{H}^2$ and the quaternionic hyperbolic spaces $\mathbb{H}\mathbf{H}^n$ are strongly quasi-isometrically rigid, as long as $n \geq 2$ [28]. Kleiner and Leeb further showed that irreducible higher-rank symmetric spaces of noncompact type and thick irreducible Euclidean buildings with cocompact affine Weyl group are also strongly quasi-isometrically rigid [24], and Bourdon–Pajot and Xie extended this to some hyperbolic buildings [7, 35]. When $n \geq 3$, the Cayley graphs of maximal non-arithmetic non-uniform lattices in n -dimensional rank one symmetric spaces and universal covers of hyperbolic n -manifolds with non-empty totally geodesic boundary also enjoy this property [31]. We also mention a beautiful, but somewhat exceptional, example of a strongly quasi-isometrically rigid space due to Kleiner and Kapovich, that is the Cayley graph of a carefully constructed hyperbolic group [22].

As an essential step in the proof of the main theorem above, we construct new examples of strongly quasi-isometrically rigid spaces. These arise from the following theorem, which states that well-behaved graphs of uniformly strongly quasi-isometrically rigid spaces are themselves quasi-isometrically rigid. We can view this as a sort of combination theorem for strong quasi-isometric rigidity.

Theorem B. *Let \mathbf{X} be a graph of spaces. Suppose that \mathbf{X} is uniformly hyperbolic, link bottlenecked, link rigid, locally congruent, and has unbounded edge spaces. If the vertex spaces of \mathbf{X} are uniformly strongly quasi-isometrically rigid, then so is the realisation of \mathbf{X} .*

We direct the reader to Section 4 for a precise definition of graph of spaces in this context, and to Sections 5 and 6 for definitions of the terms appearing in the theorem statement. We give a brief overview of this list of adjectives here. A graph of spaces is said to be *uniformly hyperbolic* if the vertex spaces are all hyperbolic with the same hyperbolicity constant, and the edge spaces in some sense diverge uniformly quickly from one another in each vertex space. *Link bottlenecked* essentially means that the edge spaces do not coarsely separate the vertex spaces, and *link rigid* means that any isometry of a vertex space that coarsely fixes its incident

edge spaces is trivial. Finally, *local congruence* means that, for any two vertices in the same orbit of an automorphism of the underlying graph, the corresponding vertex spaces and their configuration of incident edge spaces are identical.

This paper is structured as follows. Section 1 contains a collection of preliminary material. In Section 2, we discuss a strengthening of the aforementioned theorem of Frucht. In Section 3, we construct a sets of points in the boundary of a visibility manifold with a certain rigidity property. The next three sections are devoted to graphs of metric spaces: Section 4 contains definitions and basic properties, Section 5 introduces uniformly hyperbolic graphs of spaces and explores their properties, and in Section 6 we obtain control over the quasi-isometry groups of graphs of spaces to obtain Theorem B. Section 7 includes an explicit construction of a graph of spaces to which Theorem B applies, and the proof of Theorem A. Finally, Section 8 is dedicated to understanding geodesics in uniformly hyperbolic graphs of spaces, which is essential for the hyperbolicity claim of Theorem A.

Acknowledgements. The second author was supported by the Additional Funding Programme for Mathematical Sciences, delivered by EPSRC (EP/V521917/1) and the Heilbronn Institute for Mathematical Research. The third author was supported by the European Research Council (SATURN, 101076148) and the Deutsche Forschungsgemeinschaft (EXC-2047/1, 390685813).

The authors would like to thank the Isaac Newton Institute for Mathematical Sciences, Cambridge, for support and hospitality during the programme *Non-positive curvature and applications*, where much of the work on this paper was undertaken. This programme was supported by EPSRC grant EP/Z000580/1.

1. PRELIMINARIES

We begin by standardising our notation and terminology.

1.1. Metric notions. Let (X, d_X) be a metric space. Write $\text{Isom}(X)$ for the group of isometries $f: X \rightarrow X$. Given $r \geq 0$, $x \in X$, we write $B_X(x; r)$ for the closed ball of radius r centred at x . Given $A \subset X$, we define $B_X(A; r)$ as the union of all balls $B_X(x; r)$ with $x \in A$. If $A, B \subset X$, we write $d_X(A, B)$ to mean the infimal distance between A and B , i.e.

$$d_X(A, B) = \inf_{a \in A, b \in B} d_X(a, b).$$

If $A = \{x\}$ is a singleton, we abuse notation and write $d_X(x, B) := d_X(\{x\}, B)$. We denote by $\text{Haus}_X(A, B)$ the *Hausdorff distance* in X between A and B . That is,

$$\text{Haus}_X(A, B) = \inf\{r > 0 : A \subset B_X(B; r) \text{ and } B \subset B_X(A; r)\}.$$

For any set Y , denote the supremum metric on functions $f, g: Y \rightarrow X$ as

$$d_\infty(f, g) = \sup_{y \in Y} d_X(f(y), g(y)),$$

noting that this metric can take infinite values in general.

Definition 1.1 (Quasi-isometry). Let $\lambda \geq 1$ and $c \geq 0$. A map $\varphi: X \rightarrow Y$ between metric spaces is called a (λ, c) -quasi-isometric embedding if

$$\frac{1}{\lambda} d_X(x, y) - c \leq d_Y(\varphi(x), \varphi(y)) \leq \lambda d_X(x, y) + c$$

for all $x, y \in X$. If there is $K \geq 0$ such that for all $y \in Y$ there exists $x \in X$ such that $d_Y(y, \varphi(x)) \leq K$, we say that φ is a *K-coarsely surjective*.

A (λ, c) -quasi-isometry is a (λ, c) -quasi-isometry embedding that is c -coarsely surjective. A map $\varphi: X \rightarrow Y$ is simply called a *quasi-isometry* if it is a (λ, c) -quasi-isometry for some $\lambda \geq 1$ and $c \geq 0$. If $\varphi: X \rightarrow Y$, $\psi: Y \rightarrow Z$ are quasi-isometries,

we call ψ a *quasi-inverse* to φ if there is some $Q \geq 0$ with $d_\infty(\psi \circ \varphi, \text{id}_X) \leq Q$ and $d_\infty(\varphi \circ \psi, \text{id}_X) \leq Q$.

It is a standard fact that for every (λ, c) -quasi-isometry $\varphi: X \rightarrow Y$, there exists another (λ, c') -quasi-isometry $\psi: Y \rightarrow X$ with $d_\infty(\psi \circ \varphi, \text{id}_X) \leq Q$, where Q and c' depend only on λ and c . We call any such function a *quasi-inverse* to φ .

Definition 1.2 (Quasi-geodesics). Let $I \subseteq \mathbb{R}$ be a closed interval, bounded or unbounded, and X an arbitrary metric space. A (λ, c) -quasi-geodesic in X is a (λ, c) -quasi-isometric embedding $\varphi: I \rightarrow X$. Similarly to quasi-isometries, we sometimes omit mention of the particular constants. We may often abuse notation and identify a quasi-geodesic with its image in X . Further, we will call such φ a (λ, c) -quasi-geodesic segment, ray or line when I is a bounded interval, a half-bounded interval, and the entire real line respectively.

As a special case, an *geodesic* in X will for us be an isometric embedding $I \rightarrow X$ (i.e. a $(1, 0)$ -quasi-geodesic). Given two points $x, y \in X$, we will often denote by $[x, y]$ a choice of geodesic in X with endpoints x and y .

We define the equivalence relation on quasi-isometries φ, ψ by setting $\varphi \sim \psi$ if and only if $d_\infty(\varphi, \psi)$ is finite. Given metric space X and Y , let

$$\text{QI}(X, Y) = \{\text{quasi-isometries } \varphi: X \rightarrow Y\}/\sim.$$

It is straightforward to check that composition of maps gives rise to a well-defined group operation on this set. We write $\text{QI}(X)$ as shorthand for $\text{QI}(X, X)$, and we call this the *quasi-isometry group* of X . Of course, every isometry of X is a quasi-isometry, so there is a map from the isometry group of X to its quasi-isometry group. In general, this map need neither be injective nor surjective. The property of surjectivity is of particular interest, as it means that the quasi-isometries of X behave in a very rigid way.

Definition 1.3 (Strongly QI-rigid space). We say that a metric space X is *strongly QI-rigid* if the natural map $\text{Isom}(X) \rightarrow \text{QI}(X)$ is a surjection.

We say that X is *uniformly strongly QI-rigid* if for all $\lambda \geq 1, c \geq 0$, there is $\mu \geq 0$ such that for any (λ, c) -quasi-isometry $\varphi: X \rightarrow X$, there is an isometry $f \in \text{Isom}(X)$ with $d_\infty(\varphi, f) \leq \mu$.

We recall the notion of coarse connectedness for a metric space.

Definition 1.4 (Coarse connectedness). Let (X, d) be a metric space, $\kappa > 0$ a constant, and $x, y \in X$. A κ -path between x and y is a sequence of points (x_0, \dots, x_n) in X for which $x_0 = x, x_n = y$, and $d(x_{i-1}, x_i) \leq \kappa$ for every $i = 1, \dots, n$. We say that X is κ -coarsely connected if there is a κ -path between any two points of X .

The following lemma is a well-known observation about coarse separation of subsets (cf. [5, Lemma 2.3]). Important for us here is the uniformity of the constants involved: we include a proof to account for these details.

Lemma 1.5. *For every $\lambda \geq 1$ and $c \geq 0$ there is a constant $\kappa_0 = \kappa_0(\lambda, c) \geq 1$ such that for any $\kappa \geq \kappa_0$ there are $\rho = \rho(\lambda, c, \kappa)$ with $\lim_{\kappa \rightarrow \infty} \rho(\lambda, c, \kappa) = \infty$ and $L = L(\kappa) \geq \rho$ such that the following is true.*

Let $\varphi: X \rightarrow X'$ be a (λ, c) -quasi-isometry of metric spaces. Suppose that $Y \subseteq X$ is λ -coarsely connected, $a, b \in Y$, and $Z \subseteq X$ are such that a and b lie in different κ -coarsely connected components of $Y - B_X(Z; \kappa)$ and

$$d_X(a, Z) > \lambda L + \lambda c \quad \text{and} \quad d_X(b, Z) > \lambda L + \lambda c.$$

Then $\varphi(a)$ and $\varphi(b)$ lie in different ρ -coarsely connected components of the set $\varphi(Y) - B_{X'}(\varphi(Z); L)$.

Proof. Let $\psi: X' \rightarrow X$ be a quasi-inverse of φ , so that there is $Q \geq 0$ depending only on the quasi-isometry constants of φ with $d_\infty(\psi \circ \varphi, \text{id}_X) \leq Q$. We may suppose that ψ is also a (λ, c) -quasi-isometry. Write $\rho_0 = \lambda^2 + c$, and note that $\varphi(Y)$ is ρ_0 -coarsely connected. We further write $\rho'_0 = \lambda\rho_0 + c$. Define $\kappa_0 = \rho'_0 + Q$, let $\kappa \geq \kappa_0$ and set $\rho = \frac{1}{\lambda}(\kappa - c - Q)$, $\rho' = \lambda\rho + c$ and $L = \lambda\kappa + c + Q$. It is clear that ρ diverges as κ tends to infinity.

Let $p = (x_0, \dots, x_n)$ be a ρ -path in $\varphi(Y)$ whose endpoints are $\varphi(a)$ and $\varphi(b)$, which exists by ρ_0 -connectedness of Y . Then $\psi(p) = (\psi(x_0), \dots, \psi(x_n))$ is a ρ' -path whose endpoints are $\psi(\varphi(a))$ and $\psi(\varphi(b))$. Since ψ is quasi-inverse to φ , we obtain a κ -path $(a, \psi(x_0), \dots, \psi(x_n), b)$ from a and b in Y . By assumption, a and b lie in different κ -coarse components of $Y - B_X(Z; \kappa)$, so this path contains a point $z \in Y$ with $d_X(z, Z) \leq \kappa$. The assumptions imply that a and b are at a distance of greater than κ from Z , so we must have $z = \psi(x_i)$ for some $i = 0, \dots, n$. But then by the fact that φ is a (λ, c) -quasi-isometry and the choice of L ,

$$d_{X'}(x_i, \varphi(Z)) \leq \lambda\kappa + c + Q = L.$$

This means that any ρ -path in Y from $\varphi(a)$ to $\varphi(b)$ passes L -close to $\varphi(Z)$. Since a and b are a distance greater than $\lambda L + \lambda c$ from $\varphi(Z)$, the points $\varphi(a)$ and $\varphi(b)$ are a distance greater than L from $\varphi(Z)$. Therefore $\varphi(a)$ and $\varphi(b)$ lie in distinct ρ -coarsely connected components of $Y - B_{X'}(\varphi(Z); L)$ as required. \square

1.2. Hyperbolic metric spaces. We here recall the definition of a hyperbolic metric space, along with some useful facts.

Definition 1.6 (Gromov product). Let (X, d) be a metric space, and $x, y, z \in X$ be points. The *Gromov product* of x and y with respect to z is

$$\langle x, y \rangle_z = \frac{1}{2} \left(d(x, z) + d(y, z) - d(x, y) \right).$$

If $A, B \subseteq X$ are subspaces, we write

$$\langle A, B \rangle_z = \sup \{ \langle a, b \rangle_z \mid a \in A, b \in B \} \in [0, \infty]$$

Definition 1.7 (Thin triangles). Let Δ be a geodesic triangle with vertices x, y , and z in a metric space X , and let $\delta \geq 0$. Call T_Δ the tripod with leg lengths $\langle x, y \rangle_z, \langle x, z \rangle_y$, and $\langle y, z \rangle_x$. There is a unique map $\varphi: \Delta \rightarrow T_\Delta$ such that x, y , and z map to the extremal vertices of T_Δ and φ restricts to an isometry on each side of Δ . We say Δ is δ -thin if $\text{diam } \varphi^{-1}(\{t\}) \leq \delta$ for all $t \in T_\Delta$.

Definition 1.8 (Hyperbolic metric space). Let X be a geodesic metric space. If there is $\delta \geq 0$ such that every geodesic triangle in X is δ -thin, we say that X is a δ -hyperbolic metric space.

Remark 1.9. It is well-known that δ -hyperbolicity of a space X implies the following four point condition: for all $x, y, z, w \in X$

$$\min \{ \langle x, y \rangle_w, \langle y, z \rangle_w \} \leq \langle x, z \rangle_w + 2\delta.$$

Remark 1.10. It is well known that the real hyperbolic plane \mathbb{RH}^2 is δ -hyperbolic with $\delta = \log(3)$. Any CAT(-1) space is therefore also $\log(3)$ -hyperbolic.

The following lemma is essentially a restating of the thin triangles property in a form that is convenient for our later use.

Lemma 1.11. Let (X, d) be a δ -hyperbolic metric space, $x, y, z \in X$. Any geodesic $[y, z]$ contains a point t with $d_X(t, x) \leq \langle y, z \rangle_x + \delta$.

Proof. Let $\varphi: \Delta \rightarrow T_\Delta$ be the canonical map from the triangle Δ with vertices x, y, z to the associated tripod T_Δ . Call $* \in T_\Delta$ the middle point of the tripod. Take $s \in [x, y]$ and $t \in [y, z]$ to be the unique points of $[x, y]$ and $[y, z]$ in the preimage $\varphi^{-1}(\{*\})$ respectively. By definition, $d(x, s) = \langle y, z \rangle_x$ and, since Δ is δ -thin, $d(s, t) \leq \delta$. The lemma follows now by applying the triangle inequality. \square

We will use the following statement about closest point projections.

Lemma 1.12. *Let (X, d) be a δ -hyperbolic metric space, $Y \subseteq X$ a subspace, $\pi: X \rightarrow Y$ a closest point projection. For any $x \in X$ and $y \in Y$ we have*

$$\langle x, y \rangle_{\pi(x)} \leq \delta$$

and

$$\text{Haus}_X([x, y], [x, \pi(x)] \cup [\pi(x), y]) \leq 2\delta.$$

Proof. By hyperbolicity of X there exist $a \in [x, \pi(x)]$ and $b \in [\pi(x), y]$ with $d(a, b) \leq \delta$ and $d(a, \pi(x)) = d(b, \pi(x)) = \langle x, y \rangle_{\pi(x)}$. By the fact that π is a closest point projection and a lies on a geodesic between x and $\pi(x)$, we have

$$\langle x, y \rangle_{\pi(x)} = d(b, \pi(x)) = d(a, \pi(x)) \leq d(a, b) \leq \delta.$$

By the thinness of triangles, we also have that $\text{Haus}_X([x, y], [x, a] \cup [b, y]) \leq \delta$, from which it follows that $\text{Haus}_X([x, y], [x, \pi(x)] \cup [\pi(x), y]) \leq 2\delta$. \square

A hyperbolic metric space admits a natural bordification ∂X called the *Gromov boundary*. Among proper spaces, this coincides with the *geodesic boundary*, defined as equivalence classes of geodesic rays considered up to bounded Hausdorff distance; see [12, Chapter 11] for details. The Gromov product on a proper hyperbolic metric space extends naturally to its Gromov boundary. Given a (bi-)infinite geodesic $\gamma: I \rightarrow X$ in a proper hyperbolic space X , its *endpoints* are the points each of its subrays represents in ∂X .

A core feature of hyperbolic metric spaces is the stability of quasi-geodesics. This is encapsulated in the following, commonly known as the Morse Lemma.

Lemma 1.13 ([12, Theorem 11.72 and Lemma 11.105]). *Let (X, d) be a proper δ -hyperbolic metric space, $\lambda \geq 1$ and $c \geq 0$. There is $M = M(\lambda, c, \delta) \geq 0$ such that for any (λ, c) -quasi-geodesic $\varphi: I \rightarrow X$, any geodesic $\gamma: J \rightarrow X$ with the same endpoints as φ satisfies*

$$\text{Haus}_X(\varphi(I), \gamma(J)) \leq M.$$

Here, endpoints can be points in X or ∂X .

For $Y \subseteq X$, where X is a δ -hyperbolic metric space, we denote by ΛY the *limit set* of Y . We will need the following lemma about the limit sets of coarsely separating sets in hyperbolic spaces.

Lemma 1.14. *Let X be a δ -hyperbolic metric space, let $Z \subseteq X$ be a subset and let U and V be different coarse components of $X - Z$ for some $\rho \geq 0$. Then $\Lambda U \cap \Lambda V \subseteq \Lambda Z$.*

Proof. Let $\rho \geq 0$ be a constant such that U and V are in different ρ -coarse components of $X - Z$. Let $z \in \partial U \cap \partial V$, so there are sequences $(u_n) \subseteq U$ and $(v_n) \subseteq V$ with $u_n, v_n \rightarrow z$ as $n \rightarrow \infty$. Since U and V are different ρ -coarse components of $X - Z$, for every $n \in \mathbb{N}$ there is $z_n \in [u_n, v_n] \cap B(Z; \rho)$. But then the sequence of (z_n) also converges to z , so that $z \in \partial Z$. \square

2. FRUCHT'S THEOREM

In this section we will discuss Frucht's theorem. We will need a particular variant for countable groups which produces uncountably many regular graphs of arbitrary degree, and the main goal of this section is to prove this. This can essentially be pieced together from proofs already present in the literature, but we include our own for the sake of self-containment. We first introduce some terminology.

Definition 2.1 (Vertex-free and edge-free graphs). A graph Γ is said to be *vertex-free* if the stabiliser of every vertex in $\text{Aut}(\Gamma)$ is trivial. Similarly, Γ is said to be *edge-free* if the setwise stabiliser of every pair $\{u, v\}$ of adjacent vertices is trivial.

We will mostly be interested in vertex-free graphs, but it is possible for us to ensure our constructed graphs are also edge-free. A vertex-free graph Γ is edge-free if and only if $\text{Aut}(\Gamma)$ acts *without edge inversions*. That is, no isometry of Γ transposes the two endpoints of any edge.

Definition 2.2 (Labelled graphs). Let Γ be a graph. Given $e \in E(\Gamma)$ with endpoints $x, y \in V(\Gamma)$, an *orientation* of e is a choice of endpoint. A *partial orientation* of Γ is a choice of orientation for a subset of the edges. Edges without a choice of orientation are called *undirected*, and those with an orientation as *directed*.

Let S be a set. An *S -labelling* of Γ is a partial orientation of Γ together with an assignment of an element of S to every edge. We call S the labelling set. A graph with an S -labelling is called a *labelled graph*.

We define the *label-preserving automorphism group*, denoted $\text{Aut}_\ell(\Gamma)$, as the subgroup of $\text{Aut}(\Gamma)$ which preserves the labelling of Γ .

Cayley graphs are the prototypical examples of labelled graphs.

Definition 2.3 (Cayley graph). Let G be a group, and $S \subset G - \{1\}$ a generating set. Assume further that if $s, s^{-1} \in S$, then s has order 2. The *Cayley graph* of G with respect to S , denoted $\Gamma_{G,S}$, is the simplicial, labelled graph with:

- (1) $V(\Gamma_{G,S}) = G$;
- (2) for every $g \in G, s \in S$ not of order 2, there is an edge between g and gs , labelled by s , directed from g to gs ;
- (3) for every $g \in G, s \in S$ of order 2, there is an undirected edge between g and gs labelled by s , undirected.

Remark 2.4. The group of labelled automorphisms of a Cayley graph is exactly the original group. This is straightforward to see by the action of the group on its Cayley graph by (left) translation.

The following is the central construction of the section. It associates to each 3-regular labelled graph another 3-regular graph, whose automorphism group is the same as the group of label-preserving automorphisms of the original graph.

Proposition 2.5. *Let Γ_0 be a 3-regular vertex-free labelled graph. Then there exist uncountably many connected, simplicial, 3-regular, vertex-free, edge-free graphs Γ with*

$$\text{Aut}(\Gamma) \cong \text{Aut}_\ell(\Gamma_0).$$

Moreover, if Γ_0 is hyperbolic, then Γ is also hyperbolic.

Proof. Consider the set $\{0, 1\}^{\mathbb{N}}$ of one-way infinite binary sequences. for each sequence $\sigma \in \{0, 1\}^{\mathbb{N}}$, we construct a *tag graph*, denoted T_σ , as depicted in Figure 1. The construction is a sort of ladder graph with two types of rungs, whose decoration is determined by the binary sequence σ . The rungs corresponding to 1s will be single edges, while those corresponding to 0s will be a copy of the *Frucht graph*, which is a 3-regular graph with trivial automorphism group [18, Theorem 2.3].

All tags are 3-regular except for a single leaf. As such, any automorphism of a tag graph T_σ must preserve the single leaf and so also the sequence of rungs in the tag. Since the Frucht graph has no non-trivial automorphisms, this implies that T_σ has trivial automorphism group unless σ is the constant sequence of 1s. Finally, note that by the same reasoning, $T_\sigma \cong T_{\sigma'}$ implies that $\sigma = \sigma'$.




FIGURE 1. Constructing the tag T_σ .

Consider the set L of labels appearing in Γ_0 , which is countable. Fix some injection $h: L \hookrightarrow \{0, 1\}^{\mathbb{N}}$ such that no $\ell \in L$ maps to the constant sequence of 1s.

We now blow up Γ_0 to another graph Γ as follows. Given an edge $e \in E(\Gamma_0)$ with directed label $\ell \in L$, we replace this edge with the gadget depicted in Figure 2. Given an edge labelled by an undirected label replace this edge with the gadget depicted in Figure 3.

Now any automorphism φ of Γ must send tags of the form $T_{h(a)}$ to tags of the same form. As such, φ preserves the gadgets of Γ . Moreover, the only automorphism each gadget allows is an involutive automorphism, exactly when the label is undirected. Therefore φ gives rise to a labelled graph automorphism of Γ_0 . Conversely, each label preserving automorphism of Γ_0 trivially defines an automorphism of Γ .

As noted, each gadget has no non-trivial automorphisms beside the obvious involution of gadgets corresponding to undirected labels, and these automorphisms fix no vertex or edges. Therefore any vertex of Γ fixed by an automorphism must correspond to a vertex of Γ_0 fixed by a label-preserving automorphism. Since Γ_0 was vertex free, though, there are no such vertices, so Γ is vertex-free. Further, Γ is edge-free, since the only label-preserving automorphisms of Γ_0 invert an edge with an undirected label.

Recall now that a *block* is a maximal connected subgraph with no disconnecting vertices. Every graph admits a canonical decomposition into a tree of blocks (see, for example, [11, §3.1]). It is straightforward consequence that a graph is δ -hyperbolic if and only if each of its blocks are δ -hyperbolic. Each block of the gadgets is at most 5-hyperbolic, so Γ is hyperbolic whenever Γ_0 is.

Finally, note that if we chose a different labelling $h': L \hookrightarrow \{0, 1\}^{\mathbb{N}}$ with image disjoint to that of h , then we would necessarily get a different graph. Since there are uncountably many such injections with pairwise disjoint images, we are done. \square

An additional blow-up construction allows us to go from 3-regular graphs to d -regular graphs for any $d > 3$. We are able to do this blow up in such a way to preserve the large-scale geometric structure of the graph.

Lemma 2.6. *Let Γ be a 3-regular connected, simplicial, vertex-free, edge-free graph, $d > 3$. Then there exists a d -regular, connected, simplicial, edge-free graph $\Gamma^{(d)}$ such that $\text{Aut}(\Gamma) \cong \text{Aut}(\Gamma^{(d)})$, and Γ is quasi-isometric to $\Gamma^{(d)}$.*

Moreover, if Δ is another 3-regular connected, simplicial, edge-free graph such that $\Gamma^{(d)} \cong \Delta^{(d)}$, then $\Gamma \cong \Delta$.

Proof. Let Γ be a 3-regular graph, and $d > 3$. We will construct a new d -regular augmented graph $\Gamma^{(d)}$ as follows.




FIGURE 2




FIGURE 3

Let A_1, A_2, \dots, A_{d-3} be pairwise non-isomorphic $(d-1)$ -regular graphs of order a power of $d-1$, and with trivial automorphism group. Such graphs are known to be generic among regular graphs by probabilistic arguments [23, Corollary 1.2]. Any d -clique of a $(d-1)$ -regular graph is a complete connected component, which has non-trivial automorphism group. Hence no A_i contains a d -clique. Consider now a copy of the complete graph K_d on d vertices, and choose three distinguished vertices u_1, u_2, u_3 . We label the other vertices v_1, \dots, v_{d-3} .

For each $i = 1, \dots, d-3$, let $n_i > 0$ be such that A_i has order $(d-1)^{n_i}$. We define the graph A'_i as the full rooted $(d-1)$ -ary tree of depth n_i , adding edges to the $(d-1)^{n_i}$ leaves so that the subgraph of A'_i induced by the leaves is precisely A_i . The root vertex $w_i \in V(A'_i)$ is the only vertex of A'_i with degree $d-1$, and every other vertex has degree d . Note that since A_1, \dots, A_n were chosen pairwise non-isomorphic, A'_1, \dots, A'_n are pairwise non-isomorphic. We construct a new graph P as the union of K_d and A'_1, \dots, A'_{d-3} , with an extra edge between v_i and the root w_i of A'_i for each $i = 1, \dots, d-3$. Now every vertex of P has degree d , except for u_1, u_2 , and u_3 , which have degree $d-1$.

We claim that $\text{Aut}(P)$ acts 3-transitively on the vertices u_1, u_2 , and u_3 , and that every other vertex of P is fixed by every automorphism. We first note that $\text{Aut}(P)$ setwise fixes $\{u_1, u_2, u_3\}$, as these are the only vertices of degree other than d . Moreover, since $\text{Aut}(K_d) = \text{Sym}(V(K_d))$ acts d -transitively on its vertex set, $\text{Aut}(P)$ acts 3-transitively on $\{u_1, u_2, u_3\}$, by permuting these vertices and their incident edges, and acting trivially elsewhere.

We must show that every other vertex in P is fixed by every automorphism. Let $\varphi \in \text{Aut}(P)$. Since none of the A_i contain d -cliques, the base copy of K_d is the only d -clique in P . This copy of K_d must therefore be fixed by φ . Since φ permutes u_1, u_2 , and u_3 , it must also permute the remaining vertices v_1, \dots, v_{d-3} of K_d . Therefore, φ must permute the copies of A'_i in P . However, the graphs A'_1, \dots, A'_n are pairwise non-isomorphic, so this induced permutation must be trivial. Thus φ restricts to a root-fixing automorphism on A'_i for each $i = 1, \dots, n$. Any non-trivial root-fixing automorphism of A'_i restricts to a non-trivial automorphism of the copy of A_i in A'_i . Since we chose A_i to be a graph with no non-trivial automorphisms, φ must

restrict to a trivial automorphism on A'_i . Therefore φ fixes every vertex of P except u_1, u_2 , and u_3 .

For each vertex $v \in V(\Gamma)$, choose a bijection φ_v between $\{u_1, u_2, u_3\}$ and the 3 edges incident to v . We now form $\Gamma^{(d)}$ by taking a copy of a copy P_v of P for every vertex $v \in V(\Gamma)$, and joining the copy of u_i in P_v to the copy of u_j in P_w whenever v and w are adjacent and $\varphi_v(u_i) = \varphi_w(u_j)$. The resulting graph is d -regular, and the map collapsing each copy of P to a single vertex is a quasi-isometry to Γ , as P is finite.

To conclude the proof, consider the *graph of d -cliques* of $\Gamma^{(d)}$, denoted $C_d(\Gamma^{(d)})$. The vertices of $C_d(\Gamma^{(d)})$ are the d -cliques of $\Gamma^{(d)}$, and two vertices are joined by an edge if their corresponding d -cliques contain adjacent vertices. Since automorphisms of graphs permute its d -cliques, there is always a homomorphism

$$(1) \quad F: \text{Aut}(\Gamma^{(d)}) \rightarrow \text{Aut}(C_d(\Gamma^{(d)})).$$

If $\varphi \in \ker F$, then φ fixes every clique of $\Gamma^{(d)}$ and the adjacency relations between them. This means that it restricts to an automorphism on each copy of P in $\Gamma^{(d)}$. Moreover, φ must fix each copy of the vertices u_1, u_2 , and u_3 , as these determine the adjacency of cliques. We saw that every automorphism of P also fixes every vertex besides u_1, u_2 , and u_3 , so φ must be trivial. Hence F is injective.

Let us show F is surjective. Since $\text{Aut}(P)$ acts 3-transitively on the $\{u_1, u_2, u_3\}$, any automorphism of $C_d(\Gamma^{(d)})$ lifts to an automorphism of $\Gamma^{(d)}$. It follows that F is an isomorphism.

Now observe that any d -clique of $\Gamma^{(d)}$ is contained in a copy of P . Each copy of P in turn has a unique d -clique, which contains copies of the vertices u_1, u_2 , and u_3 . By construction, any two such cliques are adjacent in $C_d(\Gamma^{(d)})$ if and only if the copies of P they lie in correspond to adjacent vertices of Γ . Therefore, we have an isomorphism $\Gamma \cong C_d(\Gamma^{(d)})$, and thus an isomorphism $\text{Aut}(\Gamma^{(d)}) \cong \text{Aut}(\Gamma)$ by (1). This also shows that $\Gamma^{(d)} \cong \Delta^{(d)}$ implies $\Gamma \cong C_d(\Gamma^{(d)}) \cong C_d(\Delta^{(d)}) \cong \Delta$.

To conclude, we show that $\Gamma^{(d)}$ is edge-free. Suppose that φ is an automorphism of $\Gamma^{(d)}$ fixing an edge. If φ fixes an edge in a copy of P , then it must fix that copy of P and, therefore, its unique d -clique. But then $F(\varphi)$ fixes a vertex of $C_d(\Gamma^{(d)}) \cong \Gamma$, which is vertex-free. Similarly, if φ fixes an edge adjoining two copies of P , then $F(\varphi)$ fixes an edge of $C_d(\Gamma^{(d)}) \cong \Gamma$, which is edge-free. Hence $\Gamma^{(d)}$ is edge-free. \square

We now prove our variant of Frucht's theorem, the main result of this section.

Theorem 2.7. *Let G be a countable group, and $d \geq 3$. Then there exist uncountably many d -regular, simplicial, connected, edge-free graphs Γ such that $\text{Aut}(\Gamma) \cong G$, up to isomorphism. Moreover, if G is a hyperbolic group, then we may additionally take every such Γ to be hyperbolic.*

Proof. We first deal with some exceptional small cases. Let $\sigma \in \{0, 1\}^{\mathbb{N}}$ be an infinite binary sequence that is not the constant sequence of 1s. If G is trivial, we take Γ_0 to be a copy of the tag T_{σ} from the proof of Proposition 2.5, with a copy of the Frucht graph attached to the leaf. If G is the group of order two (respectively, three), take Γ_0 to be two (respectively, three) copies of T_{σ} with the leaf vertices joined to a vertex-free and edge-free 3-regular graph whose automorphism group has order 2 (respectively, 3) as in [18, Theorem 2.4] (respectively, [18, Theorem 3.1]).

In each case, Γ_0 is vertex-free, 3-regular, satisfies $\text{Aut}(\Gamma_0) \cong G$, and is quasi-isometric to a tree, which is hyperbolic. Moreover, there are uncountably many pairwise non-isomorphic tags, so there are uncountably many such Γ_0 . Applying Lemma 2.6 then yields the theorem in these cases.

Suppose now that $|G| \geq 4$. We will construct a labelled graph Γ_0 such that:

- (1) Γ_0 is 3-regular and vertex-free;
- (2) $\text{Aut}_\ell(\Gamma_0)$ is isomorphic to G ;
- (3) if G is finitely generated, then Γ_0 is quasi-isometric to G .

Let $S \subset G - \{1\}$ be a generating set containing at least 3 elements, with the property that if $s, s^{-1} \in S$ then s has order 2. Such an S exists, since one may take any generating set and delete one of every inverse pair that appears in the set. Let $I \subset S$ denote the subset containing all order 2 elements, and let $T = S - I$.

Let $\Gamma_{G,S}$ be the labelled Cayley graph of G with respect to S , as described in Definition 2.3. As in Remark 2.4, $\text{Aut}_\ell(\Gamma_{G,S})$ is isomorphic to G . We now modify $\Gamma_{G,S}$ as follows to make it 3-regular. To do so, we must consider two cases separately.

Firstly, let us consider the case where S is countably infinite. Since G is countable, fix a choice of bijection $f: \mathbb{Z} \rightarrow G$. We then blow up each vertex in $\Gamma_{G,S}$ by replacing it with a copy of the real line, with incoming edge attached at each integer point, where the order is determined by the bijection f . We also label these new edges with directed labels, where the edge from n to $n+1$ is equipped with a directed label of the integer n ; see Figure 4. We call the resulting labelled graph Γ_0 .

Now any label-preserving automorphism of Γ_0 must permute the copies of P . Moreover, the choice of labelling on P ensures that the group of label-preserving automorphisms of Γ_0 is isomorphic to G .




FIGURE 4. Blowing up $\Gamma_{G,S}$ to a 3-regular graph Γ_0 , in the case where S is countably infinite.

Now suppose instead that S is finite. Let k denote the degree of $\Gamma_{G,S}$. If $k = 3$, then we are done, so suppose $k \geq 4$. We now form Γ_0 by equivariantly blowing up each vertex v into a path P of length $k-3$, adjoining exactly one of the edges incident upon v to each internal vertex in this copy of P , and two to each of the edges. We equivariantly label the edges of each copy of P with the numbers 1 through $k-3$. Again call the resulting labelled graph Γ_0 for this case. As before, the group of label-preserving automorphisms of Γ_0 is isomorphic to G . Moreover, in this case, collapsing each copy of P yields a quasi-isometry from Γ_0 to $\Gamma_{G,S}$.

To conclude the proof, we apply Proposition 2.5 to obtain uncountably many 3-regular vertex- and edge-free graphs whose automorphism groups are isomorphic to $\text{Aut}_\ell(\Gamma_0) \cong G$. The proposition tells us that the graphs are hyperbolic whenever Γ_0 is, and therefore, as Γ_0 is quasi-isometric to G , whenever G is hyperbolic. Finally, applying Lemma 2.6 to each member of this uncountable family of graphs gives us the required family of graphs. \square

3. BOUNDARY SETS WITH TRIVIAL STABILISER

In this section, we construct a finite set in the visual boundary ∂M of certain non-positively curved manifolds M with trivial pointwise stabiliser in $\text{Isom}(M)$. These sets will determine the configuration of edge spaces in our graph of groups construction in Section 7. A general setting this construction makes sense in is the class of visibility manifolds.

Definition 3.1. A complete, simply connected and non-positively curved Riemannian manifold M of dimension at least two is called a *visibility manifold* if for every two distinct points $x, y \in \partial M$ there exists a bi-infinite geodesic γ with $\gamma(\infty) = x$ and $\gamma(-\infty) = y$.

Note that every simply connected symmetric space of non-compact type is a visibility manifold [13].

Roughly speaking, the idea in constructing our set of boundary points is to pick the vertices of an ideal polytope that is large enough to not be contained in any proper geodesically complete and totally geodesic submanifold. As such, there will be no non-trivial isometries stabilising its vertices. For this purpose, let us introduce the following terminology.

Definition 3.2. Let M be a visibility manifold, and let $Z \subseteq \partial M$ be a set of cardinality at least two. We define *hull* of Z , denoted $\text{hull}(Z)$, to be the minimal geodesically complete and totally geodesic submanifold of M whose limit set contains Z .

Note that $\text{hull}(Z)$ does not generally coincide with the convex hull of Z , but is rather the minimal totally geodesic submanifold containing its convex hull. We see that an isometry stabilising a set Ω with the property that its hull spans the entire visibility manifold must be the identity.

Lemma 3.3. Let M be a visibility manifold. If $\Omega \subseteq \partial M$ is a subset with $\text{hull}(\Omega) = M$, then Ω has trivial pointwise stabiliser in $\text{Isom}(M)$.

Proof. First we observe that the condition $\text{hull}(\Omega) = M$ implies that Ω has at least three points, for otherwise $\text{hull}(\Omega)$ is a geodesic with endpoints the two points of Ω .

Suppose f is an isometry of M whose induced map ∂f on ∂M fixes Ω pointwise. By well-known facts about isometries of visibility manifolds (see e.g. [3]), ∂f fixes 0, 1, 2, or infinitely many points of ∂M . As ∂f fixes Ω and $|\Omega| \geq 3$, it must be that ∂f fixes infinitely many points of ∂M . Thus f is an elliptic isometry.

By [3, Lemma 6.3], the fixed point set of f in M is a complete totally geodesic submanifold containing all geodesics between pairs of points in Ω . By the lemma hypothesis, this must be M . It follows that f is the identity map. \square

There are finite subsets in the boundary of a visibility manifold satisfying the condition of the above lemma, whose cardinality depends only on the dimension of the space.

Lemma 3.4. Let M be a visibility manifold or $M = \mathbb{R}$. Then there is a set of $\Omega \subset \partial M$ with $|\Omega| \leq \dim M + 1$ such that $\text{hull}(\Omega) = M$.

Proof. We proceed by induction on the dimension. By definition, a visibility manifold has dimension at least 2, so the base case is that $M = \mathbb{R}$. In this case, ∂M only has two points, and the geodesic spanning both of them is the entire line.

Now let $n > 1$, and suppose the statement holds for visibility manifolds of dimension at most $n - 1$. Let $N \subset M$ be a geodesically complete and totally geodesic proper submanifold of maximal dimension. Such a submanifold always exists since every bi-infinite geodesic is a complete totally geodesic submanifold. By geodesic completeness we have $\dim N < \dim M$.

We claim that N is also a visibility manifold, so that the induction hypothesis applies. Let $x, y \in \partial N$ and choose $o \in N$ and geodesic rays γ_x, γ_y based at o with $\gamma_x(\infty) = x, \gamma_y(\infty) = y$. Consider the sequence of geodesic segments $\gamma_t = [\gamma_x(t), \gamma_y(t)]$, and note that $\gamma_t \subset N$ for every t as N is totally geodesic. By completeness, the geodesic γ that $(\gamma_t)_t$ subconverges to is contained in N . Necessarily, we have that $\gamma(\infty) = x$ and $\gamma(-\infty) = y$. It follows that N is a visibility manifold.

Now let $\Omega' \subset \partial N$ be a set of boundary points with $\text{hull}(\Omega') = N$ with $|\Omega'| \leq \dim N + 1$. Take any point $\omega \in \partial M - \partial N$, which exists since $\dim N < \dim M$, and set $\Omega = \Omega' \cup \{\omega\}$. It follows that $|\Omega| \leq \dim N + 2 \leq \dim M + 1$. By definition, $\text{hull}(\Omega)$ is a totally geodesic submanifold and contains $N \subseteq \text{hull}(\Omega')$. Since $\omega \notin \partial N$ and $\text{hull}(\Omega)$ is geodesically complete, $\dim \text{hull}(\Omega) > \dim N$. As N was a proper submanifold with these properties of maximal dimension, we must have that $\text{hull}(\Omega) = M$. \square

Proposition 3.5. *Let M be a visibility manifold. Then there exists a finite set $\Omega \subset \partial M$ with $|\Omega| = n + 1$ that has trivial pointwise stabiliser in $\text{Isom}(M)$.*

Proof. Lemmas 3.4 and 3.3 imply the existence of a set of boundary points Ω with trivial pointwise stabiliser in $\text{Isom}(M)$. If Ω has cardinality less than $\dim M + 1$, we may add rays in our favourite directions, which does not change the properties of Ω . \square

Remark 3.6. Given a particular visibility manifold, it may be possible to choose such a set Ω with strictly less than $\dim M + 1$ elements. This can be done by analysing its totally geodesic submanifolds in more detail. However, for our purposes, this non-sharp statement suffices.

4. GRAPHS OF METRIC SPACES

In this section, we will define our notion of a graph of metric spaces, which is an adaptation of a construction already appearing in Scott–Wall [32, p. 155].

Let Γ be a graph. In this paper, every graph will be undirected and simplicial. That is, Γ will have no loops and no double-edges. Denote by $V(\Gamma)$ its vertex set, and $E(\Gamma)$ its set of edges. Given any graph Γ , we will always fix an orientation on its edge set, so that each edge $e \in E(\Gamma)$ has a fixed initial vertex $\iota(e)$ and terminal vertex $\tau(e)$. As an abuse of notation, we will write \bar{e} to denote the edge $e \in E(\Gamma)$ with reverse orientation, so that $\iota(\bar{e}) = \tau(e)$ and $\tau(\bar{e}) = \iota(e)$. For a vertex $v \in V(\Gamma)$ we denote by $\text{In}(v)$ the set of edges incident to the vertex v , that is, either $\iota(e) = v$ or $\tau(e) = v$.

Definition 4.1 (Graph of spaces). A *graph of metric spaces* $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ consists of the following data:

- (1) a connected, simplicial graph Γ ;
- (2) for every $v \in V(\Gamma)$, a geodesic metric space X_v ;
- (3) for every $e \in E(\Gamma)$, a length metric space Y_e ;
- (4) for every $e \in E(\Gamma)$, a pair of isometric embeddings $\alpha_e: Y_e \rightarrow X_{\tau(e)}$ and $\alpha_{\bar{e}}: Y_e \rightarrow X_{\iota(e)}$.

For brevity, we will often refer to a graph of metric spaces as simply a ‘graph of spaces’ throughout; the presence of additional metric structure is always implicit. Following Scott–Wall [32], we construct a realisation of a graph of metric spaces.

Definition 4.2 (Realisation of a graph of metric spaces). Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a graph of metric spaces. Let A denote the disjoint union

$$A = \bigsqcup_{v \in V(\Gamma)} X_v \sqcup \bigsqcup_{e \in E(\Gamma)} Y_e \times [0, 1].$$

Let \sim be the equivalence relation on A generated by relations of the form

$$(x, 0) \sim \alpha_e(x), \quad (x, 1) \sim \alpha_{\bar{e}}(x),$$

where $e \in E(\Gamma)$, $x \in Y_e$. The *realisation* of \mathbf{X} is then the quotient space $|\mathbf{X}| = A / \sim$.

We will refer to the image of each X_v in $|\mathbf{X}|$ as a *vertex space*, and images of the cylinders $Y_e \times [0, 1]$ as *edge cylinders*. Edge cylinders will always be equipped with the ℓ_2 -metric corresponding to the product. We will often write $Z_e = Y_e \times [0, 1]$ to

denote an edge cylinder, and we call the subspace $Y_e \times (0, 1)$ the *interior* of Z_e . A subset of \mathbf{X} is called a *piece* of \mathbf{X} if it is equal to either X_v for some $v \in V(\Gamma)$ or Z_e for some $e \in E(\Gamma)$.

Note that, as Γ is simplicial, the projection map $A \rightarrow |\mathbf{X}|$ in the above restricts to a topological embedding on pieces of \mathbf{X} . Hence we may view the these as genuine subspaces of $|\mathbf{X}|$.

We now describe how to metrise the realisation of a graph of metric spaces. This construction is roughly based on [8, § I.7].

Definition 4.3 (String). Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a graph of metric spaces, and let $x, y \in |\mathbf{X}|$. A *string* in \mathbf{X} between x and y is a finite sequence $S = (x_0, \dots, x_n)$, $x_0 = x$, $x_n = y$, together with a choice of piece $P_i \subset |\mathbf{X}|$ containing both x_{i-1} and x_i for each $i = 1, \dots, n$. The *length* of S , denoted $\ell(S)$, is defined as

$$\ell(S) = \sum_{i=1}^m d_{P_i}(x_{i-1}, x_i).$$

Definition 4.4 (Intrinsic metric). Let \mathbf{X} be a graph of metric spaces. We call the function $d_{\mathbf{X}} : |\mathbf{X}| \times |\mathbf{X}| \rightarrow \mathbb{R}$ defined by

$$d_{\mathbf{X}}(x, y) = \inf\{\ell(S) : S \text{ is a string with endpoints } x \text{ and } y\}.$$

the *intrinsic metric* on $|\mathbf{X}|$.

Remark 4.5. Suppose that $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ is a graph of spaces and $\Gamma' \subseteq \Gamma$ is a subgraph. Writing $\mathbf{X}' = (\Gamma', X_-, Y_-, \alpha_-)$, there is a canonical map of realisations $|\mathbf{X}'| \rightarrow |\mathbf{X}|$, which is a topological embedding. Moreover, any string in \mathbf{X}' maps to a string of the same length in \mathbf{X} under this embedding, so that $d_{\mathbf{X}} \leq d_{\mathbf{X}'}$ on the image of $|\mathbf{X}'|$ in $|\mathbf{X}|$.

The function $d_{\mathbf{X}}$ defined above is easily seen to be a pseudo-metric, at least: lengths of strings are always non-negative and the triangle inequality follows by concatenating strings. We will see shortly that it is in fact a metric.

Definition 4.6 (Reduced string). Let \mathbf{X} be a graph of metric spaces. A string S in \mathbf{X} is called *reduced* if the pieces of S alternate between vertex spaces and edge cylinders, and any consecutive triple of pieces are pairwise distinct.

Note that it is not necessary for consecutive terms in a reduced string to be distinct. It is straightforward to see that the distance between two points can be redefined as the infimum of lengths of reduced strings.

Lemma 4.7. Let \mathbf{X} be a graph of metric spaces. For any $x, y \in |\mathbf{X}|$, we have

$$d_{\mathbf{X}}(x, y) = \inf\{\ell(S) : S \text{ is a reduced string between } x \text{ and } y\}.$$

Proof. For each string S , there is a reduced string S' with the same endpoints, obtained by deleting entries of S and then, if needed, adding redundant entries. Moreover, each edge space is isometrically embedded in its adjacent vertex spaces, so $\ell(S') \leq \ell(S)$ by applying the triangle inequality in each piece. \square

On small scales, the intrinsic metric on the realisation of a graph of spaces resembles the metric on individual pieces.

Lemma 4.8. Let \mathbf{X} be a graph of metric spaces and let P be a piece of \mathbf{X} . If $x, y \in P$ are such that $d_{\mathbf{X}}(x, y) < 1$, then $d_{\mathbf{X}}(x, y) = d_P(x, y)$.

Proof. Let $\varepsilon > 0$ and let S be a reduced string from x to y with $\ell(S) < d_{\mathbf{X}}(x, y) + \varepsilon$. Since $d_{\mathbf{X}}(x, y) < 1$, we may assume without loss of generality that $\ell(S) < 1$. Recall that if $Z = Y \times [0, 1]$ is an edge cylinder, $d_Z((a, 0), (b, 1)) \geq 1$ for any $a, b \in Y$. Hence

the string S does not contain two points on opposite sides of an edge cylinder. In particular, each of the points of S belong to a single vertex space X and its adjacent edge cylinders.

Since the pieces of S alternate between edge cylinders and vertex spaces, every other piece of S must be X . As S is reduced and the first and last pieces of S are both P , this implies that S has exactly one piece. This means that $S = (x, y)$, and so $\ell(S) = d_P(x, y)$. Hence $d_P(x, y) < d_X(x, y) = \epsilon$ as required. \square

Lemma 4.9. *Let \mathbf{X} be a graph of metric spaces. The intrinsic metric $d_{\mathbf{X}}$ is a metric.*

Proof. To verify that $d_{\mathbf{X}}$ is a metric, we need only check that it is positive definite. Let x and y be points of $|\mathbf{X}|$ and suppose that $d_{\mathbf{X}}(x, y) = 0$. By Lemma 4.7, $d_{\mathbf{X}}(x, y)$ is the infimum of lengths of reduced strings between x and y .

Suppose x is contained in the interior of an edge cylinder $Z \cong Y \times (0, 1)$, so that $x = (a, r)$ for some $a \in Y$ and $r \in (0, 1)$. Let $\varepsilon = \max\{r, 1 - r\}$ and let $S = (x_0, \dots, x_n)$ be a reduced string between $x_0 = x$ and $x_n = y$ with $\ell(S) < \varepsilon$. By the choice of ε , it must be that $x_i \in Z$ for all $i = 1, \dots, n$. Lemma 4.8 then implies $d_Z(x, y) = d_{\mathbf{X}}(x, y) = 0$. As d_Z is a metric, we must have $x = y$. Similarly, if y is in the interior of an edge cylinder, we also have $x = y$.

Consider now the remaining case, that x and y are both contained in vertex spaces. Now let $S = (x_0, \dots, x_n)$ be a reduced string between $x_0 = x$ and $x_n = y$ with $\ell(S) < 1$. It must be that x and y belong to some vertex space X , since the edge cylinders have width 1. Again Lemma 4.8 implies $d_X(x, y) = d_{\mathbf{X}}(x, y) = 0$, so $x = y$ as required. \square

We describe a general construction that will be useful in describing paths in a graph of spaces.

Definition 4.10 (Path and string maps). Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a graph of spaces. Denote by $\mathcal{S}_{\mathbf{X}}$ the set of strings in X and $\mathcal{P}_{\mathbf{X}}$ the set of continuous paths in $|\mathbf{X}|$ whose image meets the interior of finitely many pieces of \mathbf{X} . We define the map $\text{Str}_{\mathbf{X}}: \mathcal{P}_{\mathbf{X}} \rightarrow \mathcal{S}_{\mathbf{X}}$ as follows. Given a continuous path $p: I \rightarrow |\mathbf{X}|$ in $\mathcal{P}_{\mathbf{X}}$, where I is a closed interval, let t_0, \dots, t_n be the partition of I that is minimal with respect to the property that $p|_{(t_{i-1}, t_i)}$ is contained in the interior of exactly one piece P_i of \mathbf{X} for each $i = 1, \dots, n$. Then we define the string

$$\text{Str}_{\mathbf{X}}(p) = (p(t_0), \dots, p(t_n)).$$

with the given piece decomposition P_1, \dots, P_n .

For each piece P of \mathbf{X} and each pair of points $x, y \in P$, fix a choice of geodesic $[x, y]$. We define the map $\text{Path}_{\mathbf{X}}: \mathcal{S}_{\mathbf{X}} \rightarrow \mathcal{P}_{\mathbf{X}}$ as follows. Suppose that $S = (x_0, \dots, x_n)$ is a string in \mathbf{X} and, for each $i = 1, \dots, n$, let P_i denote the choice of piece of \mathbf{X} containing x_{i-1} and x_i . Then define $\text{Path}_{\mathbf{X}}(S): [0, \ell(S)] \rightarrow |\mathbf{X}|$ as the path that is the concatenation of the P_i -geodesics $[x_{i-1}, x_i]$ over $i = 1, \dots, n$.

Remark 4.11. If there is $\delta \geq 0$ such that all vertex spaces of \mathbf{X} are δ -hyperbolic, then any two geodesics in a piece with the same endpoints will be a Hausdorff distance of $\delta + 1$ from one another. In such a situation, $\text{Path}_{\mathbf{X}}(S)$ is independent of the choice of geodesics made in the definition of $\text{Path}_{\mathbf{X}}$ given a string S , up to a Hausdorff distance of $\delta + 1$. Adding $\delta + 1$ to constants in the relevant statements, therefore, we may make any choice of geodesics that is convenient to us in each given situation. We will use this fact without reference in the remainder of the paper.

Under some basic assumptions on the underlying graph and vertex and edge spaces, this intrinsic metric makes the realisation into a geodesic space in an effective manner. We will invoke these assumptions fairly often in the remainder of the paper, so we collect them here.

Convention 4.12. Given a graph of spaces $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$, we will assume that the graph Γ is locally finite, and X_v and Y_e are complete and locally compact for each $v \in V(\Gamma)$ and $e \in E(\Gamma)$.

Proposition 4.13. *Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a graph of metric spaces. Under Convention 4.12, $|\mathbf{X}|$ is a proper geodesic space.*

Moreover, for any geodesic p in $|\mathbf{X}|$, there is a reduced string S_p whose points lie on p , and p is a concatenation of geodesics between the terms of S_p , each of which lie in the interior of a single piece of \mathbf{X} .

Proof. Recall that the Hopf–Rinow theorem tells us that being complete and locally compact implies being proper and geodesic, among the class of length metric spaces [8, Proposition I.3.7]. Hence we need only verify the former properties. Let $x \in |\mathbf{X}|$ be a point. As Γ is locally finite, $B = B_{\mathbf{X}}(x; 1)$ intersects at most finitely many pieces of \mathbf{X} . Thus the closure \bar{B} is the union of finitely many closed and bounded subsets in pieces of \mathbf{X} . Again by the Hopf–Rinow theorem each piece of \mathbf{X} is proper, so each of these finitely many subsets are compact. Hence \bar{B} is compact, and so x has a compact neighbourhood.

To see that $|\mathbf{X}|$ is complete, let $(a_n)_{n \in \mathbb{N}}$ be a Cauchy sequence in $|\mathbf{X}|$. For each $n \in \mathbb{N}$, let P_n be a piece of \mathbf{X} containing a_n . There is $N \in \mathbb{N}$ such that for $n, m \geq N$, we have $d_{\mathbf{X}}(a_n, a_m) < 1$. This implies that for any $n \geq N$, the terms a_n and a_N lie in adjacent pieces. As Γ is locally finite, there are only finitely many pieces adjacent to P_N , so we may pass to a subsequence $(a_{n_i})_{i \in \mathbb{N}}$ whose terms eventually lie in the same piece P . By Lemma 4.8, this subsequence is also Cauchy in P and, hence converges in P , as P is complete. It follows that $|\mathbf{X}|$ is complete, since any Cauchy sequence with a convergent subsequence converges.

For the latter claim, suppose that $p: I \rightarrow |\mathbf{X}|$ is a geodesic, where $I = [a, b]$ is a closed interval. Since p is continuous, its image $p(I)$ is a compact subset of $|\mathbf{X}|$. As $|\mathbf{X}|$ is a complete metric space, $p(I)$ is totally bounded. In particular, there is a finite cover of $p(I)$ by subsets whose diameter is at most $1/2$, say. Now since Γ is locally finite, each of these subsets meets at most finitely many pieces of \mathbf{X} . Therefore $p(I)$ meets the interior of at most finitely many pieces of \mathbf{X} . Thus $S_p = \text{Str}_{\mathbf{X}}(p)$ is defined, and we may take it to be reduced as p is geodesic and the edge spaces of \mathbf{X} are isometrically embedded in their adjacent vertex spaces. Moreover, each subpath of p is also geodesic, so that its subpaths between the finitely many terms of S are geodesics lying in pieces of \mathbf{X} , as required. \square

5. UNIFORMLY HYPERBOLIC GRAPHS OF SPACES

5.1. Definition and basic properties. We are particularly interested in graphs of hyperbolic metric spaces that are assembled in a uniformly regular way. We will call these *uniformly hyperbolic* graphs of spaces. In this section we will see that in such graphs of spaces, geodesics are well-behaved. As a consequence, these spaces satisfy auspicious metric properties such as quasiconvexity of the vertex spaces.

Definition 5.1. (Uniformly hyperbolic graph of spaces) Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a graph of metric spaces, $\delta \geq 0$, and $C \geq 0$. We say that \mathbf{X} is (δ, C) -uniformly hyperbolic if:

- (1) X_v is δ -hyperbolic for every $v \in V(\Gamma)$;
- (2) for any edge $e \in E(\Gamma)$, there is $y_e \in Y_e$ such that for any distinct edge $e' \in E(\Gamma)$ we have
 - (a) $\langle \alpha_e(Y_e), \alpha_{e'}(Y_{e'}) \rangle_{\alpha_e(y_e)} \leq C$ if $\tau(e') = \tau(e)$;
 - (b) $\langle \alpha_{\bar{e}}(Y_e), \alpha_{e'}(Y_{e'}) \rangle_{\alpha_{\bar{e}}(y_e)} \leq C$ if $\tau(e') = \iota(e)$.

Moreover, for any $e, e' \in E(\Gamma)$ with $\tau(e) = \tau(e')$, we have $\alpha_e(y_e) = \alpha_{e'}(y_{e'})$.

The last condition implies that for any given vertex space X_v there is a unique point $y_v = \alpha_e(y_e)$, which does not depend on the choice of incident edge e , satisfying condition (2). We will call this point the *basepoint* of X_v .

The second condition should be seen as a sort of acylindricity condition. It essentially states that the images of distinct edge spaces diverge from one another uniformly quickly in the vertex spaces.

Remark 5.2. Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a (δ, C) -uniformly hyperbolic graph of spaces. Since each edge space isometrically embeds into its adjacent vertex spaces, each edge space is also δ -hyperbolic, so long as it is a geodesic space.

Moreover, by hyperbolicity, any two geodesics joining the same two points of X_v are δ -Hausdorff close in X_v . Since each α_e is an isometric embedding, this implies that $\alpha_e(Z)$ is δ -quasiconvex in $X_{\tau(e)}$ for each connected component $Z \subseteq Y_e$.

The realisation of a uniformly hyperbolic graph of spaces contains a naturally embedded copy of its underlying graph Γ , as the union of subsets $\{y_e\} \times [0, 1]$ ranging over $e \in E(\Gamma)$, where the points y_e are given by uniform hyperbolicity. It is immediate from Lemma 4.8 that this embedding is a local isometry. We see that this copy of Γ is in fact globally isometrically embedded.

Lemma 5.3. Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a (δ, C) -uniformly hyperbolic graph of spaces, and consider the subspace

$$|\mathbf{Y}| = \bigcup_{e \in E(\Gamma)} \{y_e\} \times [0, 1] \subset |\mathbf{X}|,$$

where y_e are the points given by uniform hyperbolicity. Then $|\mathbf{Y}|$ equipped with the subspace metric $d_{\mathbf{Y}}$ induced from $|\mathbf{X}|$ is a geodesic and convex subspace of $|\mathbf{X}|$ and isometric to the geometric realisation $|\Gamma|$ of Γ .

Proof. Define a map

$$\begin{aligned} \Phi: V(\Gamma) &\rightarrow \{(y_e, 1) : e \in E(\Gamma)\}, \\ v &\mapsto (y_e, 1), \text{ where } e \text{ is such that } v = \tau(e). \end{aligned}$$

Note that $\Phi(v)$ exists for every vertex $v \in V(\Gamma)$ because Γ is connected. The definition of uniform hyperbolicity implies that $\Phi(v)$ is well defined, for whenever we have two edges e and e' with $\tau(e) = \tau(e')$ then $\alpha_e(y_e) = \alpha_{e'}(y_{e'})$ in $X_{\tau(e)}$. In other words, $(y_e, 1) = (y_{e'}, 1)$ in $|\mathbf{X}|$.

This can be extended to a map $\Psi: |\Gamma| \rightarrow |\mathbf{Y}|$. Let $e \in E(\Gamma)$ be an edge with $\iota(e) = u$ and $\tau(e) = v$. We map a point $t \in [0, 1]$ in the interior of the realisation of edge e to $(y_e, t) \in Z_e$. The map Ψ is a continuous bijection. We will now show that Ψ is an isometry.

Let $\varepsilon > 0$, let $x, y \in |\mathbf{Y}|$ and let $p \subset |\mathbf{X}|$ be a path from x to y such that $\ell(p) \leq d_{\mathbf{X}}(x, y) + \varepsilon$. As auxiliary terminology, we will say that p *traverses an edge cylinder* Z_e whenever $p \cap Z_e$ is connected and $p \cap X_{\iota(e)} \neq \emptyset, p \cap X_{\tau(e)} \neq \emptyset$. Of course,

$$\ell(p) \geq \#\{e \in E(\Gamma) : p \text{ traverses } Z_e\},$$

so that when $x, y \in \Psi(V(\Gamma))$ we can already conclude that

$$d_{\mathbf{Y}}(x, y) = d_{\mathbf{X}}(x, y) \geq d_{\Gamma}(\Psi^{-1}(x), \Psi^{-1}(y)) - \varepsilon.$$

In the general case one needs to add the distance to the first and last vertex, which by construction of Ψ is at least as large as the corresponding distance in the graph. The same inequality of distances follows. As $\varepsilon > 0$ was arbitrary and Ψ is a bijection, we have $d_{\mathbf{Y}}(\Psi(u), \Psi(v)) \geq d_{\Gamma}(u, v)$ for all $u, v \in |\Gamma|$.

On the other hand, for every pair of points $\Psi(u), \Psi(v)$ in $|\mathbf{Y}|$, let p be the geodesic in $|\Gamma|$ from u to v . Then $\Psi(p)$ is a path from $\Psi(u)$ to $\Psi(v)$ whose length is equal to

$\ell_\Gamma(p)$, so that $d_{\mathbf{Y}}(\Psi(u), \Psi(v)) \leq d_\Gamma(u, v)$. Thus Ψ is an isometry between $|\Gamma|$ and \mathbf{Y} . As $|\Gamma|$ is geodesic, so $|\mathbf{Y}|$ is geodesic.

It remains to show that $|\mathbf{Y}|$ is convex. We define a projection

$$\begin{aligned} \Pi: |\mathbf{X}| &\rightarrow |\mathbf{Y}| \\ x &\mapsto \begin{cases} \alpha_e(y_e) & \text{if } x \in X_{\tau(e)}, \\ (y_e, t) & \text{if } x = (z, t) \in Z_e. \end{cases} \end{aligned}$$

Now Π is continuous and length non-increasing on pieces, so that $\ell(p) \geq \ell(\Pi(p))$ for any rectifiable path p in \mathbf{X} . Further, whenever p intersects a piece P in a segment q such that $q \cap |\mathbf{Y}| \neq \emptyset$ and $q \cap (|\mathbf{X}| - |\mathbf{Y}|) \neq \emptyset$, then $\ell_P(q) > \ell_P(\Pi(q))$, so the above length inequality is strict. We already know that any two points $x, y \in |\mathbf{Y}|$ can be joined by a geodesic contained entirely in $|\mathbf{Y}|$, so every other geodesic also has to be contained entirely in $|\mathbf{Y}|$. \square

5.2. Geodesics in uniformly hyperbolic graphs of spaces. This subsection is devoted to understanding geodesics in uniformly hyperbolic graphs of spaces. We will show that geodesics in a uniformly hyperbolic graph of spaces trace out quasi-geodesics in the underlying graph. We begin by analysing strings whose lengths are close to realising the distances between points. It turns out that such strings cannot travel too deep into any vertex piece, besides the start and end terms.

Lemma 5.4. *Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a (δ, C) -uniformly hyperbolic graph of spaces, and let $x, y \in |\mathbf{X}|$. There is $D = D(\delta, C) \geq 0$ such that for any $\varepsilon > 0$, the following holds.*

Let $S = (x_0, \dots, x_m)$ be a reduced string in \mathbf{X} with $\ell(S) \leq d_{\mathbf{X}}(x_0, x_m) + \varepsilon$. For each $i = 1, \dots, m$, write P_i for the piece of \mathbf{X} containing x_{i-1} and x_i and suppose that P_1 and P_m are vertex spaces of \mathbf{X} . If $P_i = X_v$ is a vertex space then the following is true:

- if $2 \leq i \leq m-3$, then $x_i \in B_{X_v}(y_v; D + \varepsilon)$; and
- if $4 \leq i \leq m-1$, then $x_{i-1} \in B_{X_v}(y_v; D + \varepsilon)$.

where y_v is the basepoint of X_v given by the definition of uniform hyperbolicity.

Proof. Define $D = 4C + 4\delta + 1$. We will show the first inclusion in the statement, for the latter will follow from an exactly symmetrical argument.

Let $2 \leq i \leq m-3$ and suppose P_i corresponds to a vertex space X_v as in the statement. Since S is a reduced string, P_{i-1} , P_i , and P_{i+1} are distinct pieces of \mathbf{X} . Thus, P_{i-1} and P_{i+1} correspond to distinct edge cylinders Z_e and $Z_{e'}$ incident on X_v . As $i \leq m-3$ and S is reduced, P_{i+2} is a vertex space $X_{v'}$ distinct from X_v . We will first consider the former case.

Let $y_v \in X_v$ be the basepoint of X_v given by uniform hyperbolicity. Lemma 1.11 gives us a point s on the X_v -geodesic $[x_i, x_{i+1}]$ such that

$$(2) \quad d_{\mathbf{X}}(s, y_v) \leq C + \delta.$$

Moreover, observe that $i+2 \leq m-1$, and so P_{i+3} must be an edge cylinder adjacent to P_{i+2} . Applying the same argument as above to P_{i+2} , there is a point t on the $X_{v'}$ -geodesic $[x_{i+2}, x_{i+3}]$ such that

$$(3) \quad d_{\mathbf{X}}(t, y_{v'}) \leq C + \delta,$$

similarly writing $y_{v'}$ for the basepoint of $X_{v'}$. Define the string

$$S' = (x_0, \dots, x_i, s, y_v, y_{v'}, t, x_{i+3}, \dots, x_m).$$

Now S' by definition has length

$$(4) \quad \begin{aligned} \ell(S') &= \ell(S) - d_{P_i}(s, x_{i+1}) - d_{P_{i+1}}(x_{i+1}, x_{i+2}) - d_{P_{i+2}}(x_{i+2}, t) \\ &\quad + d_{P_i}(s, y_v) + d_{P_{i+1}}(y_v, y_{v'}) + d_{P_{i+2}}(y_{v'}, t). \end{aligned}$$

The triangle inequality combined with (2) implies that

$$d_{P_i}(s, y_v) - d_{P_i}(s, x_{i+1}) \leq 2C + 2\delta - d_{P_i}(x_{i+1}, y_v)$$

and similarly (3) gives that

$$d_{P_{i+2}}(t, y_{v'}) - d_{P_{i+2}}(t, x_{i+2}) \leq 2C + 2\delta - d_{P_{i+2}}(x_{i+2}, y_{v'}).$$

As $d_{P_{i+1}}(y_v, y_{v'}) = 1$ by definition, combining the above two equations with (4) and the lemma hypothesis yields

$$\ell(S') \leq d_{\mathbf{X}}(x_0, x_m) - d_{P_i}(x_{i+1}, y_v) - d_{P_{i+2}}(x_{i+2}, y_{v'}) + \varepsilon + 4C + 4\delta + 1.$$

Finally, we have $\ell(S') \geq d_{\mathbf{X}}(x_0, x_m)$ by definition of the metric, so we obtain the inequality $d_{P_i}(x_{i+1}, y_v) \leq D + \varepsilon$. \square

Reduced strings close to realising the distance between two points track a path through the underlying graph that is close to being a geodesic.

Proposition 5.5. *Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a (δ, C) -uniformly hyperbolic graph of spaces. There is a constant $K = K(\delta, C) \geq 0$ such that the following is true.*

Let $u, v \in V(\Gamma)$ and $x \in X_u, y \in X_v$, and let $\varepsilon > 0$. Suppose that S is a reduced string between x and y with $\ell(S) \leq d_{\mathbf{X}}(x, y) + \varepsilon$. Write $v_1, \dots, v_n \in V(\Gamma)$ for the sequence of vertices for which a term in S is contained in X_{v_i} , in the order they appear. Then

$$n \leq d_{\Gamma}(u, v) + K + 5\varepsilon.$$

In particular, the path in Γ with vertices v_1, \dots, v_n is a $(1, K + 5\varepsilon)$ -quasi-geodesic between u and v in Γ .

Proof. Let $S = (x_0, \dots, x_m)$ be a reduced string between $x = x_0$ and $y = x_m$ as in the statement, with $d_{\mathbf{X}}(x, y) \leq \ell(S) + \varepsilon$. For each $i = 1, \dots, m$, denote by P_i the piece of \mathbf{X} containing x_{i-1} and x_i . As S is reduced, the pieces P_i alternate between vertex spaces and edge cylinders. Since P_1 and P_m are both vertex spaces, this means that m must be odd. We thus have that $m = 2n - 1$ and P_{2i-1} is the vertex space X_{v_i} for each $i = 1, \dots, n$. Without loss of generality we may take $n \geq 3$, for otherwise setting $K = 2$ gives us the statement trivially.

Let y_{v_1}, \dots, y_{v_n} be the basepoints of vertex spaces X_{v_1}, \dots, X_{v_n} , and D the constant given by Lemma 5.4 applied to S . We will use these points to obtain our estimate for n . Take $s \in P_3$ to be a point on a P_3 -geodesic $[x_2, x_3]$ with

$$(5) \quad d_{P_3}(s, y_{v_2}) \leq C + \delta,$$

which exists by Lemma 1.11. Likewise, there is a point $t \in P_{m-2}$ on P_{m-2} -geodesic $[x_{m-3}, x_{m-2}]$ with

$$(6) \quad d_{P_{m-2}}(t, y_{v_{n-1}}) \leq C + \delta.$$

By the definition of uniform hyperbolicity, we have $d_{P_{2k}}(y_{v_k}, y_{v_{k+1}}) = 1$ for each $k = 1, \dots, n - 1$.

Recall that every edge space of \mathbf{X} is isometrically embedded in its adjacent vertex spaces, so that

$$d_{P_{2k}}(a, b) = d_{P_{2k-1}}(a, b) = d_{P_{2k+1}}(a, b) \quad \text{for all } k = 1, \dots, n - 1$$

whenever $a, b \in P_{2k} = Y \times [0, 1]$ are elements with the same second component, where Y is the edge space corresponding to the cylinder P_{2k} . We will use this fact freely below without further reference.

Observe that $d_{P_1}(x_1, y_{v_1}) \leq d_{P_2}(x_1, x_2) + d_{P_2}(x_2, y_{v_2}) + d_{P_2}(y_{v_1}, y_{v_2})$ by the triangle inequality. Moreover, $d_{P_2}(x_2, y_{v_2}) \leq d_{P_3}(x_2, s) + d_{P_3}(s, y_{v_2})$. Similar estimates

hold for $d_{P_m}(x_{m-1}, y_{v_n})$ and $d_{P_{m-2}}(x_{m-2}, y_{v_{n-1}})$. Combining these with (5) and (6), we obtain

$$(7) \quad \begin{aligned} d_{P_1}(x_1, y_{v_1}) &\leq d_{P_2}(x_1, x_2) + d_{P_3}(x_2, s) + C + \delta + 1 \quad \text{and} \\ d_{P_m}(x_{m-1}, y_{v_n}) &\leq d_{P_{m-1}}(x_{m-1}, x_{m-2}) + d_{P_{m-2}}(x_{m-2}, t) + C + \delta + 1. \end{aligned}$$

By the definition of uniform hyperbolicity, we have that $d_{\mathbf{X}}(y_{v_1}, y_{v_n}) \leq d_{\Gamma}(u, v)$. Now by the triangle inequality,

$$\begin{aligned} d_{\mathbf{X}}(x_0, x_m) &\leq d_{\mathbf{X}}(x_0, x_1) + d_{\mathbf{X}}(x_0, y_{v_1}) \\ &\quad + d_{\mathbf{X}}(y_{v_1}, y_{v_n}) + d_{\mathbf{X}}(y_{v_n}, x_{m-1}) + d_{\mathbf{X}}(x_{m-1}, x_m). \end{aligned}$$

Together with the pair of inequalities in (7), this gives us that

$$\begin{aligned} d_{\mathbf{X}}(x_0, x_m) &\leq d_{P_1}(x_0, x_1) + d_{P_2}(x_1, x_2) + d_{P_3}(x_2, s) + d_{\Gamma}(u, v) \\ &\quad + d_{P_{m-1}}(x_{m-1}, x_{m-2}) + d_{P_{m-2}}(x_{m-2}, t) + 2C + 2\delta + 1. \end{aligned}$$

By (5) and Lemma 5.4, we can also see that $d_{P_3}(x_2, s) \leq d_{P_3}(x_2, x_3) + C + \delta + D + \varepsilon$. Again, a symmetrical estimate holds for $d_{P_{m-2}}(x_{m-2}, t)$. Therefore, we can deduce from the previous inequality that

$$(8) \quad d_{\mathbf{X}}(x_0, x_m) \leq \ell(S) - \sum_{i=4}^{m-3} d_{P_i}(x_{i-1}, x_i) + d_{\Gamma}(u, v) + 4C + 4\delta + 2D + 2\varepsilon + 2.$$

On the other hand, $\ell(S) \leq d_{\mathbf{X}}(x_0, x_m) + \varepsilon$ by assumption, so that (8) yields

$$(9) \quad \sum_{i=4}^{m-3} d_{P_i}(x_{i-1}, x_i) \leq d_{\Gamma}(u, v) + 4C + 4\delta + 2D + 3\varepsilon + 2.$$

Repeated applications of the reverse triangle inequality to the left hand side give us a lower bound on the left hand side sum:

$$\sum_{i=2}^{n-2} d_{P_{2i}}(y_{v_i}, y_{v_{i+1}}) - d_{P_3}(x_3, y_{v_2}) - d_{P_{m-2}}(x_{m-3}, y_{v_{n-1}}) \leq \sum_{i=4}^{m-3} d_{P_i}(x_{i-1}, x_i).$$

Combined with (9) and Lemma 5.4, we have

$$\sum_{i=2}^{n-2} d_{P_{2i}}(y_{v_i}, y_{v_{i+1}}) \leq d_{\Gamma}(u, v) + 4C + 4\delta + 4D + 5\varepsilon + 2.$$

By definition, the sum on the left is simply equal to $n - 3$. Setting $K = 4C + 4\delta + 4D + 5$ hence gives us the required inequality. \square

5.3. Local metric properties. We will now explore the consequences of the machinery developed above. Firstly, we see that the vertex spaces in a uniformly hyperbolic graph of spaces are quasi-isometrically embedded.

Proposition 5.6. *Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a (δ, C) -uniformly hyperbolic graph of spaces satisfying Convention 4.12. There is a constant $k = k(\delta, C) \geq 0$ such that the inclusion $X_v \hookrightarrow |\mathbf{X}|$ is a $(1, k)$ -quasi-isometric embedding for all $v \in V(\Gamma)$.*

Proof. Let $v \in V(\Gamma)$ and $x, x' \in X_v$. By definition $d_{\mathbf{X}}(x, x') \leq d_{X_v}(x, x')$. Suppose then that $d_{\mathbf{X}}(x, x') < d_{X_v}(x, x')$. Since $|\mathbf{X}|$ is a geodesic space by Proposition 4.13, this means that there is a geodesic p in $|\mathbf{X}|$ between x and x' which is not contained in X_v . It follows as in (7) from the proof of Proposition 5.5 that

$$d_{X_v}(x, y_v) + d_{X_v}(x', y_v) \leq \ell(p) + 2\delta + 2C + 2,$$

where y_v denotes the basepoint of X_v . By the triangle inequality, this implies

$$d_{X_v}(x, x') - 2\delta - 2C - 2 \leq \ell(p) = d_{\mathbf{X}}(x, x')$$

so that setting $k = 2\delta + 2C + 2$ completes the proof. \square

Another consequence is that the vertex spaces in a uniformly hyperbolic graph of spaces are uniformly quasiconvex. This is, of course, an application of Proposition 5.5 to the case that the quasi-geodesic in the underlying graph is a loop.

Proposition 5.7. *Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a (δ, C) -uniformly hyperbolic graph of spaces satisfying Convention 4.12. Then there is a constant $\sigma = \sigma(\delta, C) \geq 0$ such that each vertex space of \mathbf{X} is σ -quasiconvex in $|\mathbf{X}|$.*

Proof. Let $D = D(\delta, C) \geq 0$ be the constant of Lemma 5.4 and $K = K(\delta, C) \geq 0$ the constant of Proposition 5.5. We will show that each vertex space of \mathbf{X} is σ -quasiconvex with $\sigma = K + D + 1$.

By Proposition 4.13, $|\mathbf{X}|$ is a geodesic space. Let $v \in V(\Gamma)$ and let $p: I \rightarrow |\mathbf{X}|$ be a geodesic with endpoints $x, y \in X_v$. Proposition 4.13 tells us that these is a reduced string $S_p = (x_0, \dots, x_m)$ whose terms lie on p , and p is a concatenation of geodesics $[x_{i-1}, x_i]$ in the interiors of pieces joining points of S_p . Let $v_1, \dots, v_n \in V(\Gamma)$ be the sequence of vertices for which the points of S_p are contained in X_{v_i} . As S_p is reduced, we have $m = 2n - 1$. Of course, $v_1 = v_n = v$, so by Proposition 5.5, we have $n \leq K$.

Write P_1, \dots, P_m for the pieces of S_p , so $X_{v_i} = P_{2i-1}$ for $i = 1, \dots, n$. Let y_3, \dots, y_{m-2} be the points obtained from Lemma 5.4. Write $Z_{e_i} = Y_{e_i} \times [0, 1]$ for the edge cylinder such that $\iota(e_i) = v_i$ and $\tau(e_i) = v_{i+1}$ for $i = 1, \dots, n-1$. We write $y_1 = \alpha_{e_1}(y_{e_1})$, where $y_{e_1} \in Y_{e_1}$ is provided by uniform hyperbolicity.

Now let $i = 1, \dots, m$ and let z be a point on the geodesic subsegment $[x_{i-1}, x_i]$ of p . If $i = 1$ or $i = m$ then $z \in X_v$ trivially, so suppose otherwise. If i is odd, then z is contained in a δ -neighbourhood of $[x_{i-1}, y_i] \cup [x_i, y_i]$ by uniform hyperbolicity of \mathbf{X} . Moreover, by Lemma 5.4, x_i is at distance of at most D from y_i if $i \leq m-3$, and similarly x_{i-1} from y_i if $i \geq 4$. Hence, in these cases, z is contained in a σ -neighbourhood of X_v . As the edge cylinders are equipped with the ℓ^2 -metric and have width 1, this inclusion also holds when i is even and in the same range. Finally, the cases that $i = 3$ and $i = m-2$ are dealt with by noting that $[x_2, y_3]$ and $[x_{m-2}, y_{m-2}]$ are contained in a 1-neighbourhood of X_v , since they are contained in edge pieces P_2 and P_{m-1} which are adjacent to $P_1 = P_m = X_v$. This proves that X_v is σ -quasiconvex, as required. \square

Remark 5.8. With only superficial modifications to the proof, the above also yields that the subspace of $|\mathbf{X}|$ consisting of a vertex space X_v and its adjacent edge cylinder – that is, $B_{\mathbf{X}}(X_v; 1)$ – is also σ -quasiconvex. Increasing the value of σ by 1, the same proof shows that $X_u \cup Z_e \cup X_v$ is σ -quasiconvex when X_u and X_v are adjacent vertex spaces joined by edge cylinder Z_e .

We are also able to analyse the intersections of neighbourhoods of pairs of pieces. We consider the cases that both the pieces in question are components of an edge space, or are both vertex spaces.

Lemma 5.9. *Let \mathbf{X} be a (δ, C) -uniformly hyperbolic graph of spaces. For any $r \geq 0$, there is $R = R(r, \delta, C) \geq 0$ such that for any distinct edge pieces Z_e and $Z_{e'}$ of \mathbf{X} , the set*

$$B_{\mathbf{X}}(Z_e; r) \cap B_{\mathbf{X}}(Z_{e'}; r)$$

has diameter at most R .

Proof. Let us first observe that the lemma reduces to the case that Z and Z' are both adjacent to the same vertex space X . Indeed, any path joining a point in $|X|$ to a point in an edge cylinder must pass through a vertex space adjacent to this edge cylinder. It follows that $B_{\mathbf{X}}(Z; r) \cap B_{\mathbf{X}}(Z'; r) \subseteq B_{\mathbf{X}}(Z; r) \cap B_{\mathbf{X}}(Z''; r)$, where Z'' shares an adjacent vertex space with Z , since this observation applies to a path

of length r joining a point in the former intersection to Z . Hence without loss of generality suppose that Z and Z' are such.

Let x and y be points in the given intersection. Then there are points $x_Z, y_Z \in Z \cap X$ and $x_{Z'}, y_{Z'} \in Z' \cap X$ at most r from x and y respectively. We can deduce

$$(10) \quad d_X(x, y) \leq d_X(x_Z, y_Z) + 2r$$

by the triangle inequality.

Now uniform hyperbolicity gives us that there is $z \in Z \cap Z'$ with $\langle y_Z, y_{Z'} \rangle_z \leq C$. By assumption, we have $d_X(x_Z, x_{Z'}) \leq 2r$ and $d_X(y_Z, y_{Z'}) \leq 2r$. Applying Proposition 5.6, there is $k = k(\delta, C) \geq 0$ such that the above implies $d_X(x_Z, x_{Z'}) \leq 2r + k$ and $d_X(y_Z, y_{Z'}) \leq 2\lambda r$. Combining these two inequalities implies

$$d_X(x_Z, z) + d_X(x_{Z'}, z) \leq 2C + 2r + k.$$

An identical equation holds for y_Z and $y_{Z'}$ in place of x_Z and $x_{Z'}$ respectively. Hence, all four points $x_Z, x_{Z'}, y_Z$, and $y_{Z'}$ lie in $B_X(z; 2C + 2r + k)$. In particular,

$$(11) \quad d_X(x_Z, y_Z) \leq 4C + 4\lambda r.$$

Observing that $d_X(x_Z, y_Z) \leq d_X(x_Z, y_{Z'})$, equations (10) and (11) give us

$$d_X(x, y) \leq 4C + 4\lambda r + 2r,$$

so that setting $R = 4C + 4\lambda r + 2r$ completes the lemma. \square

Lemma 5.10. *Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a (δ, C) -uniformly hyperbolic graph of spaces. For any $r \geq 0$, there is $R = R(r, \delta, C) \geq 0$ such that for any distinct vertex pieces X_u and X_v of \mathbf{X} , then either*

- (1) $d_\Gamma(u, v) = 1$ and any connected component of $B_{\mathbf{X}}(X_u; r) \cap B_{\mathbf{X}}(X_v; r)$ has diameter at most R or is contained in $B_{\mathbf{X}}(Z_e; R)$, where $e \in E(\Gamma)$ has endpoints u and v and $Z_e \subseteq Y_e$ is an edge cylinder; or
- (2) $d_\Gamma(u, v) > 1$ and $B_{\mathbf{X}}(X_u; r) \cap B_{\mathbf{X}}(X_v; r)$ has diameter at most R .

Proof. Let $z \in B_{\mathbf{X}}(X_u; r) \cap B_{\mathbf{X}}(X_v; r)$. Observe that there are paths p_u and p_v of length at most r joining z to a point of X_u and a point of X_v respectively. For the first statement, if either p_u or p_v intersect Z_e , then $z \in B_{\mathbf{X}}(Z_e; r)$ and we are done. Suppose otherwise, that then either p_u or p_v (or both) must pass through two distinct edge cylinders Z and Z' . But then z belongs to a path component of $B_{\mathbf{X}}(Z; r) \cap B_{\mathbf{X}}(Z'; r)$, which has diameter at most $R = R(r, \delta, C)$ by Lemma 5.9.

For the latter case, as $d_\Gamma(u, v) > 1$, the paths p_u and p_v must pass through at least two distinct edge cylinders Z and Z' . It follows that $z \in B_{\mathbf{X}}(Z; r) \cap B_{\mathbf{X}}(Z'; r)$, whence again applying Lemma 5.9 bounds the diameter by R . \square

6. STRONG QI-RIGIDITY FOR GRAPH OF SPACES

In this section we describe a series of additional technical properties pertaining to graphs of spaces, which will provide strong control over their quasi-isometry groups. The purpose of these properties is to guarantee strong QI-rigidity for the realisation of a graph of spaces; the results are summarised in the statement of Theorem 6.17.

6.1. Link bottlenecked pairs. The first condition we describe allows us to control what happens to vertex spaces under a quasi-isometry of graphs of spaces \mathbf{X} and \mathbf{Y} . Roughly speaking, the condition states that the edge spaces of \mathbf{Y} do not coarsely separate vertex spaces of \mathbf{X} into arbitrarily deep components. Before we define it formally, it will be helpful introduce the following notation.

Definition 6.1 (Link of a vertex space). Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a graph of spaces. Let $v \in V(\Gamma)$. Then the \mathbf{X} -link of v , denoted $\text{lk}_{\mathbf{X}}(v)$, is defined as the space

$$\text{lk}_{\mathbf{X}}(v) := \bigcup_{e \in \text{In}(v)} \alpha_e(Y_e),$$

equipped with the subspace metric induced by the metric d_{X_v} on X_v .

Definition 6.2 (Link bottlenecked pairs). Let

$$\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-), \quad \mathbf{Y} = (\Delta, X'_-, Y'_-, \beta_-)$$

be graphs of spaces. We say that the ordered pair (\mathbf{X}, \mathbf{Y}) is *link bottlenecked* if:

(LB) for all $\lambda \geq 1$ and $c \geq 0$, there are $K \geq 0$ and $\rho \geq 0$ such that for all $v \in V(\Gamma)$ and all subsets $Z \subset X_v$, if there exists $u \in V(\Delta)$ such that Z admits a (λ, c) -quasi-isometric embedding into $\text{lk}_{\mathbf{Y}}(u)$, then there exists a unique ρ -coarse component $U \subset X_v - Z$ with $U \not\subset B_{X_v}(Z; K)$.

If the pair (\mathbf{X}, \mathbf{X}) is link bottlenecked, then we simply say \mathbf{X} is *link bottlenecked*. If both \mathbf{X} and \mathbf{Y} are link bottlenecked, and both pairs (\mathbf{X}, \mathbf{Y}) and (\mathbf{Y}, \mathbf{X}) are link bottlenecked, we call (\mathbf{X}, \mathbf{Y}) a *totally link bottlenecked pair*.

Example 6.3. We describe some basic examples.

- (1) Consider a graph of spaces \mathbf{X} where every vertex space is isometric to the Euclidean plane \mathbb{R}^2 , and every \mathbf{X} -link is uniformly bounded in diameter. It is easy to verify that \mathbf{X} satisfies (LB).
- (2) Given any graph Γ , one can form a trivial graph of spaces \mathbf{X} where each vertex and edge space is just a single point, and so trivially $|\mathbf{X}|$ is isometric to Γ . This is generally **not** link bottlenecked.

Remark 6.4. The condition (LB) depends only on the geometry of the vertex spaces and their links. Thus, if \mathbf{X}, \mathbf{Y} are graphs of spaces such that all vertex spaces of \mathbf{X} and \mathbf{Y} are pairwise isometric, and all links across both \mathbf{X} and \mathbf{Y} are pairwise isometric, then the pair (\mathbf{X}, \mathbf{Y}) is totally link bottlenecked if and only if \mathbf{X} and \mathbf{Y} are link bottlenecked.

We record some basic and useful consequences of being link bottlenecked.

Lemma 6.5. Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a link bottlenecked graph of spaces, where Γ is not a single point. Then each vertex space X_v is one-ended.

Proof. If $X_v \subset |\mathbf{X}|$ were bounded then it is trivially quasi-isometric to a subset of some link, which contradicts (LB). Thus, no vertex space is bounded. Let ρ be the constant given by the definition (LB) for the coarseness of components. Suppose then that X_v is unbounded and has more than one end. Then there exists a bounded subset $Z \subset X_v$ such that $X_v - Z$ contains two ρ -coarse components containing points arbitrarily far from Z . This also contradicts (LB). \square

Lemma 6.6. Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a link bottlenecked graph of spaces. Let $v \in V(\Gamma)$. Then $\text{Haus}_{\mathbf{X}}(X_v, \text{lk}_{\mathbf{X}}(v)) = \infty$.

Moreover, if $u, v \in V(\Gamma)$ are distinct vertices then $\text{Haus}_{\mathbf{X}}(X_v, X_u) = \infty$.

Proof. If these were not the case then we would immediately contradict (LB). \square

We now give a technical, more quantitative version of Lemma 6.6, which states that (LB) implies that given any point in a vertex space, there is always a uniformly nearby point which is far away from the link.

Lemma 6.7. Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a link bottlenecked graph of spaces. Then for all $D > 0$, there exists $r = r(D, \mathbf{X}) > 0$ such that for all $v \in V(\Gamma)$ and all $x \in X_v$ there exists $y \in X_v$ such that

- (1) $d_{X_v}(x, y) \leq r$,
- (2) $d_{X_v}(y, \text{lk}_{\mathbf{X}}(v)) > D$.

Proof. Suppose this were not the case. Then this means that there exists $D > 0$ such that for all $r > 0$, there exists $v \in V(\Gamma)$ and $x \in X_v$ such that for all $y \in B_{X_v}(x; r)$, we have that $d_{X_v}(y, \text{lk}_{\mathbf{X}}(v)) \leq D$. Let $\rho = \rho(1, 2D)$ be the constant given by the definition (LB). For $r > \rho$, consider the set

$$S_r = \{z \in X_v : d_{X_v}(x, z) \in [r - \rho, r]\}.$$

By our assumption, any closest point projection of S_r to $\text{lk}_{\mathbf{X}}(v)$ is a $(1, 2D)$ -quasi-isometric embedding into $\text{lk}_{\mathbf{X}}(v)$. However, x lies in a bounded ρ -coarse component of $X_v - S_r$. Since X_v is unbounded by Lemma 6.5, this contradicts (LB) if r is chosen to be sufficiently large with respect to $K(1, 2D)$. \square

The key consequence of property (LB) is that the images of vertex spaces under quasi-isometries of link bottlenecked graphs of spaces are uniformly close to vertex spaces in the codomain.

Lemma 6.8. *Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$, $\mathbf{Y} = (\Delta, X'_-, Y'_-, \beta_-)$ be uniformly hyperbolic graphs of spaces. Suppose further that the pair (\mathbf{X}, \mathbf{Y}) is totally link bottlenecked. Then for all $\lambda \geq 1$, $c \geq 0$, there exists $C = C(\lambda, c, \mathbf{X}, \mathbf{Y}) \geq 0$ such that for all (λ, c) -quasi-isometries $\varphi: |\mathbf{X}| \rightarrow |\mathbf{Y}|$ and all $v \in V(\Gamma)$, there exists a unique $u \in V(\Delta)$ such that*

$$\text{Haus}_{\mathbf{Y}}(X'_u, \varphi(X_v)) \leq C.$$

In particular, every quasi-isometry $\varphi: |\mathbf{X}| \rightarrow |\mathbf{Y}|$ induces a map $\Phi(\varphi): V(\Gamma) \rightarrow V(\Delta)$ on the vertices of the underlying graphs of \mathbf{X} and \mathbf{Y} .

Proof. Let $\psi: \mathbf{Y} \rightarrow \mathbf{X}$ be a quasi-inverse to φ , so $d_\infty(\text{id}_{\mathbf{X}}, \psi\varphi), d_\infty(\text{id}_{\mathbf{Y}}, \varphi\psi) \leq Q$ where Q depends only on λ and c . After possibly increasing c , we may assume ψ is also a (λ, c) -quasi-isometry. Since \mathbf{X} and \mathbf{Y} are uniformly hyperbolic, Proposition 5.7 gives us a constant $k \geq 0$ such that for all $v \in V(\Gamma)$, $v' \in V(\Delta)$, the inclusions $X_v \hookrightarrow |\mathbf{X}|$ and $X'_{v'} \hookrightarrow |\mathbf{Y}|$ are $(1, k)$ -quasi-isometric embeddings.

We first show that $\varphi(X_v)$ lands arbitrarily deep inside vertex spaces of \mathbf{Y} .

Claim. *For each $a > 0$, there is $u \in V(\Delta)$ such that*

$$\varphi(X_v) \cap (X'_u - B_{X'_u}(\text{lk}_{\mathbf{Y}}(u); a)) \neq \emptyset.$$

Proof of claim. For any $x \in X_v$, there is $x' \in |\mathbf{Y}|$ in a vertex space, say $x' \in X'_u$, such that $d_{\mathbf{Y}}(x', \varphi(x)) \leq 1$ (projecting to one side of an edge cylinder if necessary). Let $r = r(a + c + k, \mathbf{Y})$ be as in Lemma 6.7 so that, given $x' \in X'_u$, there is $y' \in X'_u$ with

$$d_{X'_u}(x', y') \leq r \quad \text{and} \quad d_{X'_u}(y', \text{lk}_{\mathbf{Y}}(u)) > a + c + k.$$

Moreover, any $y' \in |\mathbf{Y}|$ is a distance at most c from the image of φ , as φ is c -coarsely surjective. Thus by the triangle inequality and the fact that the inclusion $X'_u \hookrightarrow |\mathbf{Y}|$ is a $(1, k)$ -quasi-isometric embedding, for any $x \in X_v$ there is $y \in |\mathbf{Y}|$ with

$$(12) \quad d_{X'_u}(y, \text{lk}_{\mathbf{Y}}(u)) > a \quad \text{and} \quad d_{\mathbf{Y}}(y, \varphi(x)) \leq 1 + r + c.$$

We show that for $x \in X_v$ far enough from $\text{lk}_{\mathbf{X}}(v)$, the y obtained as above lies in $\varphi(X_v)$. To this end, take $R = \lambda(1 + r + 2c) + k$. By Lemma 6.6, there is $x \in X_v$ such that $d_{X_v}(x, \text{lk}_{\mathbf{X}}(v)) > R$.

Let $z \in \varphi^{-1}(y)$. We verify that $d_{\mathbf{X}}(x, z) < d_{\mathbf{X}}(x, \text{lk}_{\mathbf{X}}(v))$, which implies that $z \in X_v$. Using that φ is a quasi-isometry and (12),

$$d_{\mathbf{X}}(x, z) \leq \lambda(d_{\mathbf{Y}}(\varphi(x), y) + c) \leq \lambda(1 + r + 2c) = R - k.$$

To conclude, again note that since $X_v \hookrightarrow |\mathbf{X}|$ is a $(1, k)$ -quasi-isometric embedding, the above equation implies

$$d_{\mathbf{X}}(x, z) \leq R - k < d_{X_v}(x, \text{lk}_{\mathbf{X}}(v)) - k \leq d_{\mathbf{X}}(x, \text{lk}_{\mathbf{X}}(v))$$

as required. \square

Next, we use (LB) to show that $\varphi(X_v)$ cannot contain points too far away from X'_u . This is one of the two needed inclusions to prove the lemma.

Claim. *There exists $A = A(\lambda, c, \mathbf{X}, \mathbf{Y}) \geq 0$ and $u \in V(\Delta)$ such that*

$$\varphi(X_v) \subset B_{\mathbf{Y}}(X'_u; A).$$

Moreover, the vertex u is unique.

Proof of claim. Suppose otherwise, so that for any $A \geq 0$ and $u \in V(\Delta)$, there is $y \in X_v$ with $d_{\mathbf{Y}}(\varphi(y), X'_u) > A$. By the first claim, there is $u \in V(\Delta)$ such that $\varphi(X_v) \cap (X'_u - B_{X'_u}(\text{lk}_{\mathbf{Y}}(u); A)) \neq \emptyset$. That is, there is $x \in X_v$ such that $\varphi(x) \in X'_u$ and $d_{\mathbf{Y}}(\varphi(x), \text{lk}_{\mathbf{Y}}(u)) > A$. By definition, we have that $\text{lk}_{\mathbf{Y}}(u)$ separates $\varphi(x)$ from $\varphi(y)$. In particular, for any $\kappa \geq 0$ we know that $\varphi(x), \varphi(y)$ are in different κ -coarse components of $|\mathbf{Y}| - B_{\mathbf{X}}(\text{lk}_{\mathbf{Y}}(u), \kappa)$. Now applying Lemma 1.5, we obtain constants $\rho = \rho(\lambda, c, \kappa)$ and $L > 0$, depending only on κ, λ and c , such that the set

$$B = B_{\mathbf{X}}(\psi(\text{lk}_{\mathbf{Y}}(u)); L)$$

separates $\psi\varphi(x)$ from $\psi\varphi(y)$ in $|\mathbf{X}|$, so long as $A \geq \lambda L + \lambda c$. Setting $Z = B \cap X_v$, then, we have that x and y lie in distinct $(\rho + 2Q)$ -coarse components of $X_v - Z$, when equipped with the subspace metric induced from $|\mathbf{X}|$, as $d_{\infty}(\psi\varphi, \text{id}_{\mathbf{X}}) \leq Q$.

Recall that all vertex spaces in \mathbf{X} and \mathbf{Y} are $(1, k)$ -quasi-isometrically embedded in $|\mathbf{X}|$ and $|\mathbf{Y}|$. By Lemma 1.5, we know that there exists $\rho' = \rho'(1, \rho + 2Q + k)$ such that x, y lie in different ρ' -coarse components of $X_v - Z$, when equipped with the metric of X_v . Note that ρ' diverges as κ tends to ∞ .

Moreover, Z as a subspace of (X_v, d_{X_v}) , admits a $(\lambda, c + \lambda k)$ -quasi-isometric embedding into $\text{lk}_{\mathbf{Y}}(u)$. Let $K = K(\lambda, c + \lambda k)$ be as in the definition of property (LB). Now if $A \geq \lambda(K + L) + c$, we have

$$\begin{aligned} d_{X_v}(x, Z) &\geq d_{\mathbf{X}}(x, B) \\ &\geq d_{\mathbf{X}}(x, \psi(\text{lk}_{\mathbf{Y}}(u))) - L \\ &\geq \frac{1}{\lambda}(d_{\mathbf{Y}}(\varphi(x), \text{lk}_{\mathbf{Y}}(u)) - c) - L \\ &> \frac{1}{\lambda}(A - c) - L \geq K \end{aligned}$$

Identical reasoning also shows that $d_{X_v}(y, Z) > K$.

Observe that we can do this for any $\kappa \geq 0$, so that there is no ρ' such that at most one of the ρ' -coarse components of $X_v - Z$ is at bounded distance from Z . This contradicts property (LB), and thus the first statement in the claim is shown. That the vertex is unique follows immediately from Lemma 6.6. \square

Finally, we conclude the proof of the lemma by symmetry. Let $A' = A(\lambda, c, \mathbf{Y}, \mathbf{X})$ and $v' \in V(\Gamma)$ be as in second claim applied to $\psi: |\mathbf{Y}| \rightarrow |\mathbf{X}|$, so that

$$\psi(X'_u) \subset B_{\mathbf{X}}(X_{v'}; A').$$

We first argue that $v' = v$. Indeed, it follows that $\psi\varphi(X_v)$ is contained in a bounded neighbourhood of $X_{v'}$. Since ψ is quasi-inverse to φ , this implies X_v is contained in a bounded neighbourhood of $X_{v'}$. Hence $v = v'$ by Lemma 6.6.

Now, as φ is a (λ, c) -quasi-isometry, $\varphi\psi(X'_u) \subseteq B_{\mathbf{Y}}(\varphi(X_v); \lambda A' + c)$. Since $d_{\mathbf{Y}}(\varphi\psi, \text{id}_{\mathbf{Y}}) \leq Q$, we have $\text{Haus}_{\mathbf{Y}}(\varphi\psi(X'_u), X'_u) \leq Q$. Combining these shows that $\varphi|_{X_v}$ is A'' -coarsely surjective onto X'_u , where $A'' = \lambda A' + c + Q$. Setting $C = \max\{A, A', A''\}$ concludes the proof. \square

Remark 6.9. Uniform hyperbolicity of \mathbf{X} is essential in the above, as it provides a bound on the distortion of the vertex spaces in $|\mathbf{X}|$. In principle, one could forgo this assumption, formulating property (LB) in terms of the global metric $d_{\mathbf{X}}$ as opposed to the local metrics d_{X_v} . However, such a condition would be unwieldy – potentially very difficult – to verify in practice.

As a consequence of the above, quasi-isometries of a link bottlenecked graph of spaces also coarsely preserve the edge cylinders.

Lemma 6.10. *Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a uniformly hyperbolic, link bottlenecked graph of spaces. Let $\varphi: |\mathbf{X}| \rightarrow |\mathbf{X}|$ be a (λ, c) -quasi-isometry such that $\Phi(\varphi) = \text{id}_{\Gamma}$. There is $T = T(\lambda, c, \mathbf{X}) \geq 0$ such that for any $e \in E(\Gamma)$,*

$$\text{Haus}_{\mathbf{X}}(\varphi(Z_e), Z_e) \leq T,$$

where $Z_e \subseteq |\mathbf{X}|$ is the edge cylinder corresponding to e .

Proof. Let $\varphi: |\mathbf{X}| \rightarrow |\mathbf{X}|$ be a (λ, c) -quasi-isometry such that $\Phi(\varphi) = \text{id}_{\Gamma}$. Let Z_e be an edge cylinder connecting vertex spaces X_u and X_v . Note that, by assumption, $u = \Phi(\varphi)(u)$ and $v = \Phi(\varphi)(v)$. We have that

$$Z_e = B_{\mathbf{X}}(X_u; 1) \cap B_{\mathbf{X}}(X_v; 1).$$

And so, since φ is a (λ, c) -quasi-isometry,

$$\varphi(Z_e) \subseteq B_{\mathbf{X}}(\varphi(X_u); \lambda + c) \cap B_{\mathbf{X}}(\varphi(X_v); \lambda + c).$$

Let $C = C(\lambda, c, \mathbf{X}, \mathbf{X})$ be as in Lemma 6.8. The lemma tells us that $\varphi(X_w)$ is a Hausdorff distance of at most C from X_w for all $w \in V(\Gamma)$, as $w = \Phi(\varphi)(w)$. Therefore,

$$\varphi(Z_e) \subseteq B_{\mathbf{X}}(X_u; \lambda + c + C) \cap B_{\mathbf{X}}(X_v; \lambda + c + C).$$

Now, by Lemma 5.10, the set on the right-hand side is contained in $B_{\mathbf{X}}(Z_e; K)$, for some $K > 0$ depending only on λ, c , and the uniform hyperbolicity constants of \mathbf{X} . Thus the above inclusion yields

$$\varphi(Z_e) \subseteq B_{\mathbf{X}}(Z_e; K).$$

Applying the same argument to a quasi-inverse ψ of φ shows that $\text{Haus}_{\mathbf{X}}(\varphi(Z_e), Z_e)$ is bounded by a constant $T \geq 0$ depending only on λ, c , and \mathbf{X} . \square

We can say more about the function Φ from Lemma 6.8.

Lemma 6.11. *Let \mathbf{X}, \mathbf{Y} be a totally link bottlenecked pair of uniformly hyperbolic graphs of spaces. For any quasi-isometry $\varphi: |\mathbf{X}| \rightarrow |\mathbf{Y}|$, the map $\Phi(\varphi): V(\Gamma) \rightarrow V(\Delta)$ is a bijection. Moreover, if \mathbf{X} and \mathbf{Y} have unbounded edge spaces, then $\Phi(\varphi)$ is a graph isomorphism and Φ is a group homomorphism $\text{QI}(\mathbf{X}, \mathbf{Y}) \rightarrow \text{Isom}(\Gamma, \Delta)$.*

Proof. We first show that $\Phi(\varphi)$ is surjective. If not, φ cannot be coarsely surjective by Lemma 6.6. Next, note that $\Phi(\varphi)$ must also be injective, as otherwise two vertex spaces would lie at finite Hausdorff distance, which again contradicts Lemma 6.6.

Now, suppose \mathbf{X} has unbounded edge spaces. Then by Lemma 5.10 we have that the intersection of any tubular neighbourhoods of any two distinct vertex spaces X_u and X_v is infinite if and only if u and v are adjacent in Γ . This implies that $\Phi(\varphi)(v)$ and $\Phi(\varphi)(u)$ are adjacent if and only if u and v are adjacent.

The above shows that $\Phi: \text{QI}(\mathbf{X}, \mathbf{Y}) \rightarrow \text{Isom}(\Gamma, \Delta)$ is a well-defined function. It only remains to check that it is a homomorphism. Let $\varphi, \varphi': |\mathbf{X}| \rightarrow |\mathbf{X}|$ be quasi-isometries. Applying Lemma 6.8, we see that $\varphi\varphi'(X_v)$ is Hausdorff close to the vertex space $X_{\Phi(\varphi\varphi')(v)}$ on one hand. However, the same lemma implies that $\varphi'(X_v)$ is close to $X_{\Phi(\varphi\varphi')(v)}$. On the other hand, then, we have $\varphi\varphi'(X_v)$ is close to $X_{\Phi(\varphi)\Phi(\varphi')(v)}$. Therefore Lemma 6.6 gives that $\Phi(\varphi\varphi')(v) = \Phi(\varphi)\Phi(\varphi')(v)$ for all $v \in V(\Gamma)$, as required. \square

Corollary 6.12. Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$, $\mathbf{Y} = (\Delta, X'_-, Y'_-, \beta_-)$ be uniformly hyperbolic graphs of spaces with unbounded edge spaces, such that the pair (\mathbf{X}, \mathbf{Y}) is totally link bottlenecked. If \mathbf{X} and \mathbf{Y} are quasi-isometric then Γ and Δ are isometric.

Let $\iota: \text{Isom}(\mathbf{X}) \rightarrow \text{QI}(\mathbf{X})$ denote the natural map, and let us write $\Psi = \Phi \circ \iota$, where Φ is the map of Lemma 6.8. We have the following commutative diagram:

$$\begin{array}{ccc}
 \text{Isom}(\mathbf{X}) & \xrightarrow{\iota} & \text{QI}(\mathbf{X}) \\
 \Psi \searrow & & \swarrow \Phi \\
 & \text{Isom}(\Gamma) &
 \end{array}$$

Throughout the rest of this section, we will continue to refer to Ψ , Φ , and ι as above.

6.2. Local congruence. Our second condition provides a criterion for isometries of the underlying graph to lift to isometries of the whole graph of spaces.

Definition 6.13 (Locally congruent). Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a graph of spaces. We say that \mathbf{X} is *locally congruent* if the following holds:

(LC) for all $f \in \text{Isom}(\Gamma)$, $v \in V(\Gamma)$, there exists an isometry $\tilde{f}_v : X_v \rightarrow X_{f(v)}$ such that for all $e \in \text{In}(v)$ we have that $\tilde{f}_v(\alpha_e(Y_e)) = \alpha_{f(e)}(Y_{f(e)})$.

Intuitively, a graph of spaces is locally congruent if, given two vertices of the underlying graph in the same automorphic orbit, the corresponding vertex spaces and their links look identical.

Proposition 6.14. *Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a locally congruent graph of spaces. Then there exists an injective function $\chi: \text{Isom}(\Gamma) \hookrightarrow \text{Isom}(\mathbf{X})$.*

Further, if \mathbf{X} is uniformly hyperbolic, link bottlenecked, and has unbounded edge spaces, then $\Psi \circ \chi = \text{id}_{\text{Isom}(\Gamma)}$. In particular, Ψ and Φ are surjective.

Proof. Let $f \in \text{Isom}(\Gamma)$. We construct an isometry $\chi(f) \in \text{Isom}(\mathbf{X})$ as follows. Given $v \in V(\Gamma)$, we define \tilde{f} on X_v via $\chi(f)|_{X_v} = i_v \circ \tilde{f}_v$, where \tilde{f}_v is the isometry instantiated by (LC), and $i_v: X_v \hookrightarrow \mathbf{X}$ is the natural inclusion. Now (LC) implies that there is an isometry $Y_e \rightarrow Y_{f(e)}$ for any edge $e \in E(\Gamma)$. We use this to extend our definition of $\chi(f)$ to the edge cylinders $Z_e = Y_e \times [0, 1]$, so that

$$\chi(f)(y, t) = (\alpha_{f(e)}^{-1} \tilde{f}_v \alpha_e(y), t)$$

for any $e \in E(\Gamma)$, $y \in Y_e$, and $t \in [0, 1]$. The condition that $\tilde{f}_v(\alpha_e(Y_e)) = \alpha_{f(e)}(Y_{f(e)})$ ensures that this extension is well-defined as a map $Z_e \rightarrow Z_{f(e)}$. The constructed map $\chi(f)$ is an isometry of \mathbf{X} , as it maps preserves the length of strings in \mathbf{X} . We have thus defined a function $\chi: \text{Isom}(\Gamma) \rightarrow \text{Isom}(\mathbf{X})$.

If f and g are distinct isometries of Γ , then they disagree on a vertex or edge of Γ . Hence $\chi(f)$ and $\chi(g)$ disagree on a vertex space or edge space of \mathbf{X} . That is, χ is injective. Now suppose that \mathbf{X} is uniformly hyperbolic, link bottlenecked, and has unbounded edge spaces. By construction $\chi(f)(X_v) = X_{f(v)}$, so Corollary 6.12 tells us that $\Psi(\chi(f))(v) = f(v)$, for all $v \in V(\Gamma)$. Therefore $\Psi \circ \chi = \text{id}_{\text{Isom}(\Gamma)}$ as required. \square

6.3. Local rigidity. The third and final property we discuss provides control over the kernel of the homomorphism $\Phi: \text{QI}(\mathbf{X}) \rightarrow \text{Isom}(\Gamma)$.

Definition 6.15 (Link rigidity). Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$. We say that \mathbf{X} is *link rigid* if the following assertion holds:

for all $v \in V(\Gamma)$, if $\psi: X_v \rightarrow X_v$ is an isometry such that

(LR) $\text{Haus}_{X_v}(\alpha_e(Y_e), \psi(\alpha_e(Y_e))) < \infty$,

for all $e \in \text{In}(v)$, then $\psi = \text{id}_{X_v}$.

That is to say, in a link rigid graph of spaces, any isometry that coarsely preserves the link of a vertex space is trivial. We will appeal to Proposition 3.5 to ensure this condition in practice.

Proposition 6.16. *Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a uniformly hyperbolic, link bottlenecked, link rigid graph of spaces with unbounded edge spaces. If the vertex spaces of \mathbf{X} are uniformly strongly QI-rigid, then the map $\Phi: \text{QI}(\mathbf{X}) \rightarrow \text{Isom}(\Gamma)$ of Lemma 6.8 is injective.*

Proof. We will show that the kernel of Φ is trivial. Let $\varphi: |\mathbf{X}| \rightarrow |\mathbf{X}|$ be a quasi-isometry such that $\Phi(\varphi) = \text{id}_\Gamma$. By Lemma 6.8, there is $C \geq 0$ such that

$$\text{Haus}(\varphi(X_v), X_v) \leq C \text{ for all } v \in V(\Gamma).$$

Therefore, there is a $(\lambda, c + C)$ -quasi-isometry $\varphi': |\mathbf{X}| \rightarrow |\mathbf{X}|$ with $d_\infty(\varphi, \varphi') \leq C$ with the property that $\varphi'(X_v) \subseteq X_v$ for all $v \in V(\Gamma)$.

By Proposition 5.6, there is $k \geq 0$ such that the inclusion $X_v \hookrightarrow |\mathbf{X}|$ is a $(1, k)$ -quasi-isometric embedding for all $v \in V(\Gamma)$, where k depends only on the uniform hyperbolicity constants of \mathbf{X} . Hence φ' restricts to a (λ', c') -quasi-isometry $\varphi'_v: X_v \rightarrow X_v$, where $\lambda' \geq 1$ and $c' \geq 0$ depend only on λ, c , and the uniform hyperbolicity constants of \mathbf{X} . By uniform strong QI-rigidity of X_v , there is an isometry $f_v \in \text{Isom}(X_v)$ and a constant $M = M(\lambda', c', \delta)$ with $d_{\infty, X_v}(f, \varphi'_v) \leq M$.

Furthermore, Lemma 6.10 tells us that there is $T = T(\lambda', c', \mathbf{X}) \geq 0$ such that $\text{Haus}_{\mathbf{X}}(\varphi'(Z_e), Z_e) \leq T$ for all $e \in E(\Gamma)$. It follows that

$$\text{Haus}_{X_v}(\alpha_e(Y_e), f_v(\alpha_e(Y_e))) \leq T + M + k \text{ for all } e \in \text{In}(v)$$

as well. Applying property (LR), we have $f_v = \text{id}_{X_v}$ and thus $d_{\infty, X_v}(\varphi'_v, \text{id}_{X_v}) \leq M$ for all $v \in V(\Gamma)$. Again using the fact that each vertex space is $(1, k)$ -quasi-isometrically embedded, we conclude that $d_\infty(\varphi', \text{id}_{\mathbf{X}}) \leq M + k$, so that φ' is a representative of the trivial class in $\text{QI}(\mathbf{X})$. Since $d_\infty(\varphi, \varphi') < \infty$, the same is true for φ , and so Φ is injective. \square

We conclude this section by summarising the results with the following theorem statement, which we note includes the statement of Theorem B.

Theorem 6.17. *Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a uniformly hyperbolic graph of spaces. Suppose that \mathbf{X} is locally congruent, link bottlenecked, and link rigid. If the vertex spaces of \mathbf{X} are uniformly strongly QI-rigid, then:*

- (1) $|\mathbf{X}|$ is strongly QI-rigid;
- (2) there is an isomorphism $\Phi: \text{QI}(\mathbf{X}) \rightarrow \text{Isom}(\Gamma)$;
- (3) if $\mathbf{Y} = (\Delta, X'_-, Y'_-, \beta_-)$ is another uniformly hyperbolic graph of spaces such that \mathbf{X} is quasi-isometric to \mathbf{Y} and (\mathbf{X}, \mathbf{Y}) is a totally link bottlenecked pair, then Γ is isometric to Δ .

Proof. Let ι, Φ , and Ψ be as in diagram (†). By Proposition 6.14, we know that Ψ and therefore also Φ are surjective. By Proposition 6.16, we have that Φ is injective, and thus Φ is an isomorphism. It follows then that ι must be surjective, by the commutativity of (†). The third item follows directly from Corollary 6.12. \square

Remark 6.18. Any visual proper hyperbolic metric space that is strongly QI-rigid is uniformly strongly QI-rigid (cf. [28, Lemme 9.11]). This is not a heavy restriction, since any hyperbolic space admitting a cocompact isometric group action is visual.

7. PROOF OF THE MAIN THEOREM

We now apply the framework of the previous sections to construct explicit examples of spaces and prove our main theorem. In particular, we construct graphs of spaces whose vertex spaces are copies of the quaternionic hyperbolic plane $\mathbb{H}\mathbb{H}^2$,

and whose edge spaces will be geodesic rays. The properties of Section 6 will be relatively straightforward to establish for these graphs of spaces.

The quaternionic hyperbolic plane is a uniformly strongly QI-rigid rank-one symmetric space. It is a contractible eight-dimensional manifold with pinched negative sectional curvature $-4 \leq \kappa \leq -1$. It may be thought of as a less symmetric version of the hyperbolic plane. By Remark 1.10, it is $\log(3)$ -hyperbolic.

We choose the edge space embeddings of our graph of spaces in such a way to ensure (LR). To do so, we appeal to Proposition 3.5. Let Ω be the set obtained by applying Proposition 3.5 to \mathbb{HH}^2 , so $\Omega \subset \partial\mathbb{HH}^2$ has cardinality $|\Omega| = 9$ and has trivial pointwise stabiliser in $\text{Isom}(\mathbb{HH}^2)$.

Fix some basepoint $x_0 \in \mathbb{HH}^2$. For each $p \in \Omega$, let γ_p be a geodesic ray based at x_0 and tending to p . Let Γ be any connected, simplicial, vertex-free, 9-regular graph. We define the graph of spaces $\mathbf{X}_\Gamma = (\Gamma, X_-, Y_-, \alpha_-)$ as follows:

- (1) every vertex space X_v is a copy of the quaternionic hyperbolic plane \mathbb{HH}^2 ;
- (2) every edge space Y_e is a copy of the half-line $[0, \infty)$;
- (3) let $\{v_i : i \in I \dots\}$ be some choice of orbit representatives for the action of $\text{Isom}(\Gamma)$ on $V(\Gamma)$. For each $i \in I$, fix a bijection $\sigma_i : \text{In}(v_i) \rightarrow \Omega$.

Now, given any $v = g \cdot v_i \in V(\Gamma)$, where $g \in \text{Isom}(\Gamma)$, and any $e \in \text{In}(v)$, let $\alpha_e : [0, \infty) \rightarrow \mathbb{HH}^2$ denote the (unique) isometry between $[0, \infty)$ and the ray $\gamma_p \subset \mathbb{HH}^2$, where $p = \sigma_i(g^{-1} \cdot e)$. This is well-defined as the action of $\text{Isom}(\Gamma)$ is free, and so the isometry g is unique.

This concludes the construction of \mathbf{X}_Γ . See Figure 5 for an illustration of part of this space. We now verify that \mathbf{X}_Γ satisfies the hypotheses of Theorem 6.17.




FIGURE 5. Illustration of a section of our graph of spaces \mathbf{X}_Γ .

Lemma 7.1. *The graph of spaces \mathbf{X}_Γ is uniformly hyperbolic, locally congruent, and has unbounded edge spaces.*

Proof. Following Remark 1.10, every vertex space of \mathbf{X}_Γ is hyperbolic with hyperbolicity constant $\delta = \log(3)$. Taking $C = \max\{\langle p, q \rangle_{x_0} : p, q \in \Omega\} < \infty$, we see that \mathbf{X}_Γ is (δ, C) -uniformly hyperbolic, using the point $0 \in [0, \infty)$ as the basepoint in each edge space. By construction, every vertex space and link is an identical isometric copy of the same space, so that \mathbf{X}_Γ is locally congruent. Of course, each edge space is $[0, \infty)$ is unbounded. \square

Lemma 7.2. *The graph of spaces \mathbf{X}_Γ is link rigid.*

Proof. Consider an isometry $f \in \text{Isom}(\mathbb{HH}^2)$. If f coarsely fixes all the incoming edge spaces, then the action of f on the boundary $\partial\mathbb{HH}^2$ fixes the set Ω pointwise. By Proposition 3.5, we have that $f = \text{id}_{\mathbb{HH}^2}$. Thus \mathbf{X}_Γ is link rigid. \square

Lemma 7.3. *The graph of spaces \mathbf{X}_Γ is link bottlenecked.*

Proof. Let S denote the metric tree consisting of nine rays with a shared endpoint. It is clear that every \mathbf{X}_Γ -link is uniformly quasi-isometric to S . We write $X = \mathbb{H}\mathbb{H}^2$.

Suppose Z is a set that admits a (λ, c) -quasi-isometric embedding into S . The image of Z separates S . We may assume that Z is coarsely connected with a uniform constant. Indeed, if it is not, pass to a subset of Z whose image is a connected: this will be coarsely connected with constants depending only on λ and c . Moreover, since any connected subset of a tree is convex, we may take Z to be quasiconvex, with constant $\sigma \geq 0$ depending on λ, c , and $\delta = \log(3)$. By Lemma 1.5, Z coarsely separates X . Of course, $|\Lambda Z| \leq 9$.

As Z is quasiconvex, we know that $\Lambda Z \subset \partial X$ is closed, so that ΛZ separates the limit sets of the ρ -coarse components $\{V_i : i \in I\}$ of $X - Z$. Now $\partial \mathbb{H}\mathbb{H}^2$ is homeomorphic to the 7-dimensional sphere S^7 , which cannot be separated by the finite set ΛZ . Thus Lemma 1.14 implies that at most one of the sets $\Lambda V_i - \Lambda Z$ can be non-empty, say ΛV_{i_0} . Then for all $i \in I - \{i_0\}$, there is $C_i \geq 0$ with $V_i \subseteq B_X(Z; C_i)$.

We show that $C_i \leq 2\sigma$ for all $i \neq i_0$. Suppose otherwise, that $C_i > 2\sigma$. Then there is $\varepsilon > 0$ and $a \in V_i$ with $d_X(a, Z) = C_i - \varepsilon > 2\sigma$. Pick a point $x \in Z$. Then the geodesic ray γ from x through a has endpoint not in $\Lambda V_i \subset \Lambda Z$, since Z is σ -quasiconvex. Thus every point on any geodesic ray with endpoint in ∂Z is at distance at most σ from Z .

We claim that $\gamma \cap Z \subseteq [x, a]$. Suppose otherwise, so there is a point $p \in \gamma \cap Z$ not on $[x, a]$. By σ -quasiconvexity, there is a point $z \in Z$ with $d_X(a, z) \leq \sigma$. This contradicts the fact that $d_X(a, Z) > 2\sigma$. It follows that the cofinal segment $\gamma - [x, a]$ of γ is contained in a single coarse component of $X - Z$, and so γ is asymptotic to a point in $\partial X - \Lambda Z$. Necessarily then, this coarse component must be V_{i_0} . But a is in the same ρ -coarse component as $\gamma - [x, a]$, a contradiction. It follows that $V_i \subseteq B_X(Z; 2\sigma)$, for all $i \neq i_0$, as required. \square

We are now ready to conclude the proof of our main theorem.

Theorem A. *Let G be a countable group. Then there exists uncountably many quasi-isometry classes of proper geodesic metric spaces X with $G \cong \text{QI}(X)$. Moreover, if G is a hyperbolic group, then we may take every such X to be hyperbolic.*

Proof. Let G be a countable group. By Theorem 2.7, there exists an uncountable collection \mathcal{F} of pairwise non-isomorphic, connected, simplicial, 9-regular, vertex-free graphs Γ such that $G \cong \text{Isom}(\Gamma)$.

For each $\Gamma \in \mathcal{F}$, form the graph of spaces \mathbf{X}_Γ described above. By Remark 6.4, we have that for any two $\Gamma, \Delta \in \mathcal{F}$, the pair $(\mathbf{X}_\Gamma, \mathbf{X}_\Delta)$ is totally link bottlenecked if \mathbf{X}_Γ and \mathbf{X}_Δ are link bottlenecked. Thus Lemmas 7.1, 7.2, 7.3 and Remark 6.18 show that \mathbf{X}_Γ satisfies the hypotheses of Theorem 6.17.

The vertex spaces of \mathbf{X}_Γ are strongly QI-rigid by [28, Théorème 1]. Therefore, by Theorem 6.17, the following hold:

- (1) if $\mathbf{X}_\Gamma, \mathbf{X}_\Delta$ are quasi-isometric then $\Gamma \cong \Delta$;
- (2) each \mathbf{X}_Γ is strongly QI-rigid;
- (3) there is an isomorphism $\text{QI}(\mathbf{X}_\Gamma) \cong \text{Isom}(\Gamma) \cong G$.

Finally, Theorem 2.7 tells us that if G is a hyperbolic group, then we may take every $\Gamma \in \mathcal{F}$ as above to be hyperbolic. By Theorem 8.11, \mathbf{X}_Γ is hyperbolic for all $\Gamma \in \mathcal{F}$. This concludes the proof of our theorem. \square

Remark 7.4. The properties we established for \mathbf{X}_Γ are not unique to graphs of spaces whose vertex spaces are copies of $\mathbb{H}\mathbb{H}^2$. Indeed, we could have chosen any strongly QI-rigid rank-one symmetric space M as the vertex spaces of \mathbf{X}_Γ , at the cost of increasing the degree of the the underlying graph Γ based on the dimension of M , as in Proposition 3.5.

Remark 7.5. In principle, one could try to apply the above method to realise uncountable groups as quasi-isometry groups. Indeed, an analogue of Theorem 2.7 holds for any group, though the degree of the resulting graphs may have arbitrary cardinality. A graph of proper hyperbolic spaces, however, can only be uniformly hyperbolic if the underlying graph is locally finite. As uniform hyperbolicity is essential in controlling the geometry of our graphs of spaces, it is not obvious how to extend the method to treat the general case.

8. HYPERBOLICITY OF UNIFORMLY HYPERBOLIC GRAPHS OF SPACES

In this section, we will prove that the metric realisation of a uniformly hyperbolic graph of spaces is hyperbolic if and only if the underlying graph is. We will prove hyperbolicity by verifying the thin triangles condition explicitly. To do this, we need a good understanding of geodesics in uniformly hyperbolic graphs of spaces.

Recall that Proposition 5.5 tells us that geodesics in the graph of spaces project onto quasigeodesics of the graph. Building on this, we implement a series of local moves that yield a set of standardised paths which are Hausdorff close to geodesics. This reduces thinness of geodesic triangles to thinness of triangles whose sides are given by such standardised paths, a much simpler condition to verify.

Convention 8.1. Given a graph of spaces $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$, we will fix a choice of nearest point projection $\pi_{v,e} : X_v \rightarrow \alpha_e(Y_e)$ of X_v on $\alpha_e(Y_e)$ for each $e \in E(\Gamma)$ and vertex $v \in V(\Gamma)$ with $v = \tau(e)$.

Definition 8.2 (Opposite point). Let $Z_e = Y_e \times [0, 1]$ be an edge cylinder in \mathbf{X} . Given a point $x = (z_0) \in Z_e$, the *opposite* of x in Z_e is the point $(z, 1) \in Z_e$, which we denote by \bar{x} . Similarly, the opposite of $x = (z, 1)$ is $\bar{x} = (z, 0)$.

The following lemmas show that geodesics, as they pass through an edge cylinder, fellow-travel paths that are in a sense orthogonal to the edge cylinder. We first show that every reduced string is at bounded Hausdorff distance from one that travels through edge cylinders parallel to the embedded copy of Γ in \mathbf{X} .

Lemma 8.3. Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a (δ, C) -uniformly hyperbolic graph of spaces, and let $u, v \in V(\Gamma)$ be adjacent vertices, with adjoining edge $e \in E(\Gamma)$. Suppose that $S = (x_0, x_1, x_2, x_3)$ is a reduced string with $x_0 \in X_u$ and $x_3 \in X_v$, and suppose that $\ell(S) = d_{\mathbf{X}}(x_0, x_3)$. If $S' = (x_0, x_1, \bar{x}_2, x_3)$, then

$$\text{Haus}_{\mathbf{X}}(\text{Path}_{\mathbf{X}}(S), \text{Path}_{\mathbf{X}}(S')) \leq 4\delta + 1.$$

Proof. As the edge cylinder Z_e is equipped with the ℓ_2 -metric and has width 1, we have $\text{Haus}_{Z_e}([x_1, x_2], [x_1, \bar{x}_2] \cup [\bar{x}_2, x_2]) \leq 1$. Moreover, the first two terms of S and S' agree, so the Hausdorff distance bound holds on X_u and on Z_e . It remains to show the bound on X_v . By definition, $d_{Z_e}(x_2, \bar{x}_2) = 1$, so that

$$\ell(S') = d_{X_u}(x_0, x_1) + 1 + d_{X_v}(\bar{x}_2, x_3).$$

By the triangle inequality, one has $d_{Z_e}(\bar{x}_2, x_2) \leq d_{Z_e}(x_1, x_2) + 1$. Thus,

$$\begin{aligned} 2\langle \bar{x}_2, x_3 \rangle_{x_2} &\leq d_{X_v}(x_2, \bar{x}_2) + d_{X_v}(x_2, x_3) - d_{X_v}(\bar{x}_2, x_3) \\ &\leq d_{Z_e}(x_1, x_2) + d_{X_v}(x_2, x_3) - d_{X_v}(\bar{x}_2, x_3) + 1 \\ &= \ell(S) - \ell(S') + 2 \\ &\leq 2, \end{aligned}$$

because S was assumed length-minimising. By δ -hyperbolicity of X_v we then have $\text{Haus}_{X_v}([\bar{x}_2, x_2] \cup [x_2, x_3], [\bar{x}_2, x_3]) \leq 4\delta + 1$. This completes the proof. \square

We see that any reduced string that passes transversely through an edge cylinder is a bounded distance from one obtained by taking a nearest point projection.

Lemma 8.4. *Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a (δ, C) -uniformly hyperbolic graph of spaces, $v \in V(\Gamma)$ and $e \in E(\Gamma)$ an edge incident to v . Let $x_0 \in X_v$ and $y \in Y_e$. Write $x_1 = \alpha_e(y)$, $x_2 = \alpha_{\bar{e}}(y)$, $x'_1 = \pi_{v,e}(x_0)$. If $S = (x_0, x_1, x_2)$, and $S' = (x_0, x'_1, x_2)$, then*

$$\text{Haus}_{\mathbf{X}}(\text{Path}_{\mathbf{X}}(S), \text{Path}_{\mathbf{X}}(S')) \leq 2\delta + 1.$$

Proof. By Lemma 1.12 we have that $\text{Haus}_{X_v}([x_0, x_1], [x_0, x'_1] \cup [x'_1, x_1]) \leq 2\delta$. As in the proof of Lemma 8.3, we have

$$\text{Haus}_{Z_e}([x'_1, x_2], [x_1, x_2] \cup [x_1, x'_1]) \leq 1.$$

It follows that $\text{Haus}_{\mathbf{X}}(\text{Path}_{\mathbf{X}}(S), \text{Path}_{\mathbf{X}}(S')) \leq 2\delta + 1$. \square

Lemma 8.5. *Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a (δ, C) -uniformly hyperbolic graph of spaces satisfying the assumptions of Convention 4.12. Suppose that p is a geodesic and let $S_p = (x_0, \dots, x_m)$ be its associated reduced string, with pieces P_1, \dots, P_m .*

Suppose that $P_i = Z_e$ is an edge cylinder with adjacent vertex spaces $P_{i-1} = X_u$ and $P_{i+1} = X_v$ for some $1 < i < m$. Let $x'_{i-1} = \pi_{u,e}(x_{i-2})$, $x'_i = \pi_{v,e}(x_{i+1})$, and write $S' = (x_0, \dots, x_{i-2}, x'_{i-1}, x'_i, x_{i+1}, \dots, x_m)$. Then we have

$$\text{Haus}_{\mathbf{X}}(p, \text{Path}_{\mathbf{X}}(S')) \leq 12\delta + 4.$$

Proof. For simplicity, we suppose $m = 3$, so that $S_p = (x_0, x_1, x_2, x_3)$, $x'_1 = \pi_{u,e}(x_0)$, and $x'_2 = \pi_{v,e}(x_3)$. Since the difference between S_p and S' is at most two consecutive terms, this assumption is without loss of generality.

Lemma 8.3 shows that $\text{Haus}_{\mathbf{X}}(p, \text{Path}_{\mathbf{X}}(x_0, x_1, \overline{x_2}, x_3)) \leq 4\delta + 1$. By Lemma 8.4,

$$\text{Haus}_{\mathbf{X}}(\text{Path}_{\mathbf{X}}(x_0, x_1, \overline{x_2}, x_3), \text{Path}_{\mathbf{X}}(x_0, x'_1, \overline{x_2}, x_3)) \leq 2\delta + 1.$$

Again by Lemma 8.3, we have

$$\text{Haus}_{\mathbf{X}}(\text{Path}_{\mathbf{X}}(x_0, x'_1, \overline{x_2}, x_3), \text{Path}_{\mathbf{X}}(x_0, x'_1, \overline{x'_1}, x_3)) \leq 4\delta + 1.$$

Finally, from Lemma 8.4 we obtain

$$\text{Haus}_{\mathbf{X}}(\text{Path}_{\mathbf{X}}(x_0, x'_1, \overline{x'_1}, x_3), \text{Path}_{\mathbf{X}}(x_0, x'_1, x'_2, x_3)) \leq 2\delta + 1.$$

Combining the four inequalities above completes the proof. \square

We can show that whenever two different geodesic reduced strings start at the same point in a vertex space but travel through different edge cylinders initially, then at least one of them passes close to the base point in that vertex space.

Lemma 8.6. *Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a (δ, C) -uniformly hyperbolic graph of spaces satisfying the assumptions of Convention 4.12.*

Let p be a geodesic in $|\mathbf{X}|$, and write $S_p = (x_0, \dots, x_m)$ for the reduced string associated to p , with pieces P_i . If $m \geq 2$ and $P_1 = X_v$ is a vertex space, then

$$\text{Haus}_{X_v}([x_0, x_1] \cup [x_1, y_v], [x_0, y_v]) \leq 14\delta + 4,$$

where y_v is the basepoint of X_v .

Proof. Let $x'_1 = \pi_{v,e}(x_0)$ and write $S' = (x_0, x'_1, x_2, \dots, x_m)$. Lemma 8.5 tells us that $\text{Path}_{\mathbf{X}}(S')$ a Hausdorff distance at most $12\delta + 4$ from p . Moreover, by Lemma 1.12, we have $\text{Haus}_{X_v}([x_0, y_v], [x_0, x'_1] \cup [x'_1, y_v]) \leq 2\delta$. Thus,

$$\text{Haus}_{X_v}([x_0, y_v], [x_0, x_1] \cup [x_1, y_v]) \leq 14\delta + 4,$$

as required. \square

We can use Lemma 8.6 to show that geodesics that pass through at least two edge cylinders fellow-travel a path through basepoints.

Proposition 8.7. *Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a (δ, C) -uniformly hyperbolic graph of spaces with the assumptions of Convention 4.12. There is a constant $F = F(\delta, C) \geq 0$ such that the following is true.*

Let $p: I \rightarrow |\mathbf{X}|$ be a geodesic with endpoints x and z . Write $v_1, \dots, v_n \in V(\Gamma)$ for the sequence of vertices for which the image of p meets vertex spaces X_{v_1}, \dots, X_{v_n} . For each $i = 1, \dots, n$, let y_{v_i} be the basepoint of X_{v_i} given by uniform hyperbolicity. If p intersects at least two edge cylinders, then

$$\text{Haus}_{\mathbf{X}}(p, \text{Path}_{\mathbf{X}}(S)) \leq F,$$

where $S = (x, y_{v_1}, \dots, y_{v_n}, z)$.

Proof. Let $S_p = (x_0, \dots, x_m)$ be the reduced string associated to p , where $x_0 = x$ and $x_m = z$. If the first piece of this string is a vertex space, define $x'_0 = x_0$. Otherwise $x_0 = (z, t)$ lies in an edge cylinder Z_e , and x_1 lies in a vertex space X_v with e incident on v . Let $u \neq v$ denote the other vertex upon which e is incident; without loss of generality suppose that $(z, 0) \in X_u$ (otherwise, $(z, 1) \in X_u$). We write $x'_0 = (z, 0)$ and, similarly define x'_m if the final piece of S_p is an edge cylinder. Of course, $d_{\mathbf{X}}(x_0, x'_0) \leq 1$ and $d_{\mathbf{X}}(x_m, x'_m) \leq 1$. Now the string $S = (x'_0, x'_0, x_1, \dots, x_{m-1}, x'_m, x'_m)$ is reduced and its first and last pieces are vertex spaces. We write $p' = \text{Path}_{\mathbf{X}}(S)$ so that

$$(13) \quad \ell(S) \leq d_{\mathbf{X}}(x'_0, x'_m) + 4,$$

and moreover, $\text{Haus}_{\mathbf{X}}(p, p') \leq 1$.

For convenience, we relabel $S = (x_0, \dots, x_m)$. Since the first and last pieces of S are vertex spaces, $m = 2n-1$ is odd. For each $i = 1, \dots, n$, let $v_i \in V(\Gamma)$ be such that x_{2i-2} and x_{2i-1} are contained in X_{v_i} . By uniform hyperbolicity and Lemma 1.11, there are points t and t' on geodesics $[x_2, x_3] \subseteq X_{v_2}$ and $[x_{m-3}, x_{m-2}] \subseteq X_{v_{n-1}}$ such that

$$d_{X_{v_2}}(t, y_{v_2}) \leq C + \delta \quad \text{and} \quad d_{X_{v_{n-1}}}(t', y_{v_{n-1}}) \leq C + \delta$$

Further, applying Lemma 5.4 with (13), we have that $d_{X_{v_i}}(x_{2i-1}, y_{v_i}) \leq D + 4$ for $2 \leq i \leq n-3$ and $d_{X_{v_i}}(x_{2i-2}, y_{v_i}) \leq D + 4$ for $4 \leq u \leq n-1$. Thinness of triangles implies that p' remains $(D + \delta)$ -close to $\text{Path}_{\mathbf{X}}(y_{v_2}, \dots, y_{v_{n-1}})$ on the subpath between t and t' (since $D \geq C + \delta$, in any case). By hyperbolicity of X_{v_2} we have

$$\text{Haus}_{X_{v_2}}([\bar{x}_2, t], [\bar{x}_2, y_{v_2}]) \leq D + \delta + 4$$

whereas the geometry of the edge cylinder guarantees

$$\text{Haus}_{Z_{e_1}}([x_1, y_{v_1}] \cup [y_{v_1}, y_{v_2}], [x_1, \bar{x}_2] \cup [\bar{x}_2, y_{v_2}]) = 1.$$

A symmetrical argument shows that analogous inequalities hold involving the final terms of the string. Thus applying Lemma 8.6 allows us to conclude that

$$\text{Haus}_{\mathbf{X}}(\text{Path}_{\mathbf{X}}(S), p') \leq D + 15\delta + 9.$$

As $\text{Haus}_{\mathbf{X}}(p, p') \leq 1$, setting $F = D + 15\delta + 10$ thus completes the proof. \square

The next lemma shows that geodesics whose endpoints lie in edge cylinders are close to geodesics with endpoints in vertex spaces. Thus, we will only have to check thinness of triangles with vertices in vertex spaces. We introduce a notation for projections which makes the statement of the following lemma much easier.

Definition 8.8 (Projection to vertex space). Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a graph of spaces. If $v \in V(\Gamma)$ is a vertex and $e \in E(\Gamma)$ is such that $\tau(e) = v$, we define the *projection from the edge cylinder* $Z_e = Y_e \times [0, 1]$ to X_v as

$$\pi_{e,v}: Z_e \rightarrow X_v, (z, t) \mapsto \alpha_e(z).$$

For x in the interior of an edge cylinder Z_e , we write $v(z) = \{\pi_{e,u}(x), \pi_{e,v}(x)\}$ and for a point x in a vertex space, we write $v(x) = \{x\}$.

Lemma 8.9. *Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a (δ, C) -uniformly hyperbolic graph of spaces, and let us assume that the assumptions of Convention 4.12 are satisfied. Suppose that Γ is δ_Γ -hyperbolic for some $\delta_\Gamma \geq 0$. There is a constant $\Theta = \Theta(\delta, C, \delta_\Gamma) \geq 0$ such that the following is true.*

Let $x, x' \in |\mathbf{X}|$ and let p be a geodesic between x and x' . If at least one of x or x' lie in an edge cylinder, then for any geodesic q between a point of $v(x)$ and $v(x')$, we have $\text{Haus}_{\mathbf{X}}(p, q) \leq \Theta$.

Proof. Suppose that there is $v \in V(\Gamma)$ such that x and x' both lie in X_v or its adjacent edge cylinders. Let q be a geodesic as in the lemma statement. Then q has endpoints in $B_{\mathbf{X}}(X_v; 1)$, which is σ -quasiconvex following Remark 5.8, where σ depends only on δ and C . Therefore q is contained in $B = B_{\mathbf{X}}(X_v; \sigma + 1)$.

As X_u is δ -hyperbolic, the set B is δ' -hyperbolic where $\delta' = \delta + 2\sigma + 2$. Now applying the Morse Lemma to the $(1, 2)$ -quasi-geodesic in B that follows p and shares endpoints with q yields the required statement.

It remains to consider the cases that x and x' are neither in adjacent pieces of \mathbf{X} , nor share any adjacent pieces. It must be that both p and q intersect at least two edge cylinders in this case. We write $z \in v(x)$ and $z' \in v(x')$ for the endpoints of q . Let X_{u_1}, \dots, X_{u_n} be the sequence of vertex spaces intersected by p , and X_{v_1}, \dots, X_{v_m} those by q . We also write $p' = \text{Path}_{\mathbf{X}}(x, y_{u_1}, \dots, y_{u_n}, x')$ and $q' = \text{Path}_{\mathbf{X}}(z, y_{v_1}, \dots, y_{v_m}, z')$.

By Proposition 8.7, there is $F = F(\delta, C)$ such that

$$(14) \quad \text{Haus}_{\mathbf{X}}(p, p') \leq F \quad \text{and} \quad \text{Haus}_{\mathbf{X}}(q, q') \leq F.$$

Moreover, by Proposition 5.5, $u_1 \dots u_n$ and $v_1 \dots v_m$ are $(1, K)$ -quasi-geodesics in Γ , where $K = K(\delta, C) \geq 0$. By assumption, Γ is δ_Γ -hyperbolic. Hence we have that $\text{Haus}_\Gamma(u_1 \dots u_n, v_1 \dots v_m) \leq 2M$, where $M = M(1, K, \delta_\Gamma) \geq 0$ is the constant of the Morse lemma. As Γ is isometric to $|\mathbf{Y}|$ by Lemma 5.3 and the edge cylinders have unit width, $\text{Haus}_{\mathbf{X}}(p', q') \leq 2M + 1$. Combining this with (14) and setting $\Theta = 2F + 2M + 1$ gives the lemma statement. \square

A final useful observation is that whenever two geodesics starting at a point in a vertex space X_v and enter different edge cylinders, then at least one of them must pass close to the basepoint of X_v .

Lemma 8.10. *Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a (δ, C) -uniformly hyperbolic graph of spaces satisfying the assumptions of Convention 4.12. Let p, q be geodesics in $|\mathbf{X}|$ with the same initial point $x_0 \in X_v$. Let $S_p = (x_0, x_1, \dots, x_m)$ and $S_q = (x_0, x'_1, \dots, x'_n)$ be the reduced strings associated to p and q . Suppose that x_1 and x'_1 are contained in distinct edge cylinders Z_e and $Z_{e'}$. Then there is a point $a \in X_v$ lying on p or q with*

$$\min\{\text{d}_{X_v}(a, y_v), \text{d}_{X_v}(b, y_v)\} \leq C + 5\delta,$$

where $y_v \in X_v$ is the basepoint of X_v .

Proof. We write $z = \pi_{v,e}(x_0)$ and $z' = \pi_{v,e'}(x_0)$. Lemma 1.12 give us that

$$(15) \quad \langle x_0, y_v \rangle_z \leq \delta \quad \text{and} \quad \langle x_0, y_v \rangle_{z'} \leq \delta.$$

By uniform hyperbolicity we know that $\langle z, z' \rangle_{y_v} \leq C$. The four-point condition of Remark 1.9 thus gives

$$\min\{\langle x_0, z \rangle_{y_v}, \langle x_0, z' \rangle_{y_v}\} \leq C + 2\delta.$$

Suppose the minimum is attained with the first inner product. Then (15) gives that

$$\text{d}_{X_v}(y_v, z) = \langle x_0, y_v \rangle_z + \langle x_0, z \rangle_{y_v} \leq C + 3\delta.$$

Moreover, Lemma 1.12 tells us that $[x_0, x_1]$ is a Hausdorff distance of at most 2δ from $[x_0, z] \cup [z, x_1]$. Hence there is a point $a \in [x_0, x_1]$ such that $d_{X_v}(z, a) \leq 2\delta$. Combining with the above inequality, we have $d_{X_v}(y_v, a) \leq C + 5\delta$ as required. Similarly, if the minimum were attained with the second inner product, we would have $a \in [x_0, x'_1]$ with the same property. \square

We are now ready to prove the main result of this section.

Theorem 8.11. *Let $\mathbf{X} = (\Gamma, X_-, Y_-, \alpha_-)$ be a (δ, C) -uniformly hyperbolic graph of spaces, and suppose that \mathbf{X} satisfies the assumptions of Convention 4.12. Then $|\mathbf{X}|$ is a hyperbolic metric space if and only if its underlying graph Γ is. More precisely:*

- (1) *If Γ is δ_Γ -hyperbolic, then there exists $\delta' = \delta'(\delta, C, \delta_\Gamma) \geq 0$ such that $|\mathbf{X}|$ is δ' -hyperbolic.*
- (2) *If $|\mathbf{X}|$ is δ' -hyperbolic, then so is Γ .*

Proof. The second statement is a direct consequence of the fact that there exists a convex subspace of $|\mathbf{X}|$ isometric to Γ , as in Lemma 5.3. The remainder of the proof is dedicated to a proof of the first statement.

Throughout this proof, let $D = D(\delta, C)$ be the constant from Lemma 5.4, $\sigma = \sigma(\delta, C)$ the quasiconvexity constant of Proposition 5.7, and $K = K(\delta, C)$ the constant of Proposition 5.5. Let $F = F(\delta, C) \geq 0$ be the constant of Proposition 8.7. Recall that $|\mathbf{Y}|$ is an isometrically embedded copy of Γ in $|\mathbf{X}|$ containing the base-points of \mathbf{X} , as in Lemma 5.3.

Assume that Γ is δ_Γ -hyperbolic and let Δ be a geodesic triangle in $|\mathbf{X}|$ with vertices x, y, z . Let γ_a be the side of Δ with endpoints x and z , γ_b with x and y , and γ_c with y and z . We distinguish cases depending on the number of edge cylinders traversed by sides of Δ : write a, b , and c for the numbers of edge cylinders intersected by the three sides $\gamma_a, \gamma_b, \gamma_c$ of Δ , respectively. By Lemma 8.9, there is $\Theta = \Theta(\delta, C, \delta_\Gamma) \geq 0$ such that the sides of Δ are a Hausdorff distance of at most Θ from geodesics with endpoints in vertex spaces. It thus suffices to consider the case when the vertices of Δ are in vertex spaces, say $x \in X_u, y \in X_v$, and $z \in X_w$. In each case, we will replace Δ with a triangle Δ' , whose sides are given by the well-behaved paths constructed earlier in the section.

Finally, we will write Δ_Γ to denote the triangle in $|\mathbf{Y}|$ whose sides are the vertices corresponding to the vertex spaces γ_a, γ_b , and γ_c intersect. By Proposition 5.5, this is a $(1, K)$ -quasi-geodesic triangle in $|\mathbf{Y}|$. As $|\mathbf{Y}|$ is δ_Γ -hyperbolic, Δ_Γ is $(\delta_\Gamma + 2M)$ -thin, where $M = M(1, K, \delta_\Gamma) \geq 0$ is the constant of the Morse lemma.

The case $a, b, c = 0$: It must be that $\Delta \subseteq X_v$ for some $v \in V(\Gamma)$. Since X_v is δ -hyperbolic, Δ is δ -thin.

The case $a, b = 1, c = 0$: Assume that $x \in X_u$ and $y, z \in X_v$ for some adjacent vertices $u, v \in V(\Gamma)$, with connecting edge $e \in E(\Gamma)$.

Observe that γ_a and γ_b must both intersect Z_e , since Γ is simplicial. By Lemma 8.5 we know that the Hausdorff distance between γ_a and $\text{Path}_{\mathbf{X}}(S_a)$ is at most $12\delta + 4$, where $S_a = (x, \pi_{u,e}(x), \pi_{v,e}(z), z)$. Likewise, the Hausdorff distance between γ_b and $\text{Path}_{\mathbf{X}}(S_b)$ is at most $12\delta + 4$, where $S_b = (x, \pi_{u,e}(x), \pi_{v,e}(y), y)$.

Now Lemma 8.3 shows that $\text{Path}_{\mathbf{X}}(S_a)$ and $\text{Path}_{\mathbf{X}}(S_b)$ are at Hausdorff distance at most $4\delta + 1$ from γ'_a and γ'_b respectively, where $\gamma'_a = \text{Path}_{\mathbf{X}}(x, \pi_{u,e}(x), \overline{\pi_{u,e}(x)}, z)$ and $\gamma'_b = \text{Path}_{\mathbf{X}}(x, \pi_{u,e}(x), \overline{\pi_{u,e}(x)}, z)$. Let Δ' be the triangle with sides γ'_a, γ'_b , and γ_c . The sides of Δ' coincide outside of X_v , so it is δ -thin as X_v is δ -hyperbolic. Moreover, the sides of Δ' are a Hausdorff distance of $16\delta + 5$ from the sides of Δ . Hence Δ is $(33\delta + 10)$ -thin.

The case $a \geq 2, b = 1, c = 0$: Let $\Gamma' \subseteq \Gamma$ be the subgraph consisting of u and v and the edge that joins them, and let \mathbf{X}' be the subgraph of spaces with underlying

graph Γ' . Following Remark 4.5, we identify $|\mathbf{X}'|$ with the corresponding subspace of $|\mathbf{X}|$. By the $a, b, c = 0$ and $a, b = 1, c = 0$ cases treated above, $|\mathbf{X}'|$ is $(33\delta + 10)$ -hyperbolic. As in Remark 5.8, $|\mathbf{X}'|$ is $(\sigma + 1)$ -quasiconvex, so that γ_a is contained in its $(\sigma + 1)$ -neighbourhood. Hence Δ , being a geodesic triangle in its $(\sigma + 1)$ -neighbourhood, is $(33\delta + 2\sigma + 12)$ -thin.

The case $a, b \geq 2, c = 0$: By Proposition 8.7 we know that γ_a and γ_b are at a Hausdorff distance of F from γ'_a and γ'_b , where $\gamma'_a = \text{Path}_{\mathbf{X}}(y, y_1, \dots, y_m, x)$ and $\gamma'_b = (z, y'_1, \dots, y'_n, x)$ with $y_1 = y'_1 = y_v$ and $y_m = y'_n = y_u$. Let Δ' be the triangle with sides γ'_a, γ'_b , and γ_c . Outside of X_u , the sides of Δ' are concatenations of $(1, K)$ -quasigeodesics in $|\mathbf{Y}|$ between y_u and y_v with $[y_u, x] \subseteq X_v$. Since X_u is δ -hyperbolic and $|\mathbf{Y}|$ is δ_Γ -hyperbolic, then, Δ' is $(\delta + M)$ -thin. Hence Δ is $(\delta + M + 2F)$ -thin.

The case $a, b, c = 1$: Necessarily, the vertices u, v , and w are distinct and adjacent. Applying Lemma 8.10 we find points $p_x \in X_u, p_y \in X_v, p_z \in X_w$ on Δ with

$$(16) \quad \max\{d_{X_u}(p_x, y_u), d_{X_v}(p_y, y_v), d_{X_w}(p_z, y_w)\} \leq C + 5\delta.$$

There are two cases, after possibly relabelling: either $p_x, p_z \in \gamma_a$ and $p_y \in \gamma_b$, or otherwise $p_x \in \gamma_a, p_y \in \gamma_b, p_z \in \gamma_c$. In any case, write $\mathbf{X}_1, \mathbf{X}_2$, and \mathbf{X}_3 for the subgraphs of spaces with underlying graphs on $\{u, w\}, \{u, v\}$, and $\{v, w\}$ respectively. As in Remark 4.5, we identify each $|\mathbf{X}_i|$ with its natural image in \mathbf{X} . By the previous cases, each $|\mathbf{X}_i|$ is δ_0 -hyperbolic, where $\delta_0 = 33\delta + 10$.

Let us treat the first case. We consider the geodesic polygons $P_1 \subseteq |\mathbf{X}_1|$ with vertices $\{p_x, y_u, y_w, p_z\}$, $P_2 \subseteq |\mathbf{X}_2|$ with vertices $\{x, p_x, y_u, y_v, p_y\}$, and $P_3 \subseteq |\mathbf{X}_3|$ with vertices $\{y, z, p_z, y_w, y_v, p_y\}$. For any $i = 1, 2, 3$, the space $|\mathbf{X}_i|$ is δ_0 -hyperbolic, so each side of P_i is contained in a $4\delta_0$ neighbourhood of the others, as the polygons have at most 6 sides. The sides of P_i are either segments of distinct sides of Δ , or are otherwise at most $2C + 10\delta + 1$ from two sides of Δ by (16). Moreover, Δ is covered by the sides of P_1, P_2 , and P_3 . Hence Δ is δ' -thin, where δ' is the maximum of $3C + 15\delta + 2$ and $4\delta_0 + 2C + 10\delta + 1$.

In the other case, we similarly divide Δ using the pentagons $P_1 \subseteq |\mathbf{X}_1|$ with vertices $\{p_x, y_u, y_w, p_z, z\}$, $P_2 \subseteq |\mathbf{X}_2|$ with vertices $\{x, p_x, y_u, y_v, p_y\}$, and $P_3 \subseteq |\mathbf{X}_3|$ with vertices $\{y, p_y, y_v, y_w, p_z\}$. Each of these is contained in $|\mathbf{X}_i|$ for some $i = 1, 2, 3$ and is therefore $3\delta_0$ -thin. By the same considerations as above, Δ is δ' -thin.

The case $a \geq 2, b, c = 1$: First, suppose that γ_b, γ_c meet the same edge cylinder, so that γ_a is a geodesic with both endpoints in a single vertex space, say X_v . In this case, $\gamma_a \subseteq B_{\mathbf{X}}(X_v; \sigma)$ by Proposition 5.7. This effectively reduces this case to that of $a = 0, b, c = 1$ as above. Hence we may suppose that γ_b and γ_c do not intersect the same edge cylinder. Necessarily v is adjacent to both u and w .

Let S_a be the string (x, y_1, \dots, y_n, z) , where y_1, \dots, y_n are the basepoints of the vertex spaces γ_a meets. By Proposition 5.5, the path $\gamma = \text{Path}_{\mathbf{X}}(y_1, \dots, y_n)$ is $(1, K)$ -quasi-geodesic and Proposition 8.7 gives $\text{Haus}_{\mathbf{X}}(\gamma_a, \text{Path}_{\mathbf{X}}(S_a)) \leq F$. Since $d_{\Gamma}(u, w) \leq 2$, we have $\ell(\gamma) \leq K + 2$. Writing $\gamma'_a = \text{Path}_{\mathbf{X}}(x, y_u, y_v, y)$, then, we have $\text{Haus}_{\mathbf{X}}(\gamma_a, \gamma'_a) \leq F + K + 2$. Let Δ' be the triangle with sides γ'_a, γ_b , and γ_c .

By Lemma 8.10, there is a point $p \in X_v$ on either γ_b or γ_c with $d_{X_v}(p, y_v) \leq C + 5\delta$; suppose it is the former without loss of generality. Let \mathbf{X}_1 and \mathbf{X}_2 be the subgraphs of spaces on vertices $\{u, v\}$ and $\{v, w\}$ respectively, and identify each $|\mathbf{X}_i|$ with its natural image in $|\mathbf{X}|$ as in Remark 4.5. Consider the geodesic rectangle $P_1 \subseteq |\mathbf{X}_1|$ with vertices $\{x, y_u, y_v, p\}$ and the geodesic pentagon $P_2 \subseteq |\mathbf{X}_2|$ with vertices $\{y, z, y_w, y_v, p\}$. As in the case $a, b, c = 1$, each side is contained in a $3\delta_0$ -neighbourhood of the others, and the sides either belong to distinct sides of Δ' or a distance of at most $C + 5\delta + 1$ from two sides of Δ' . It follows that Δ' is δ' -thin, where $\delta' = 3\delta_0 + C + 5\delta + 1$, and hence that Δ is $(\delta' + F + K + 2)$ -thin.

The case $a, b \geq 2, c = 1$: By Proposition 8.7 the sides γ_a and γ_b are F -close to paths of strings $S_a = (y, y_v, y_1, \dots, y_m, y_u, x)$ and $S_b = (z, y_w, y'_1, \dots, y'_n, y_u, x)$ respectively. By Proposition 5.5, the segments of these paths in $|\mathbf{Y}|$ form $(1, K)$ -quasigeodesics. Write $S'_b = (z, y_w, y_v, y_1, \dots, y_m, y_u, x)$ and observe that $\text{Path}_{\mathbf{X}}(S_b)$ and $\text{Path}_{\mathbf{X}}(S'_b)$ are a Hausdorff distance of at most $M' = M(1, K + 1, \delta_\Gamma)$ apart, since they differ only in the middle segments, and $|\mathbf{Y}|$ is δ_Γ -hyperbolic. We will write $\gamma'_a = \text{Path}_{\mathbf{X}}(S_a)$ and $\gamma'_b = \text{Path}_{\mathbf{X}}(S'_b)$.

Let $e \in E(\Gamma)$ be the edge joining v and w , and write Δ' for the triangle with sides γ'_a, γ'_b , and γ_c . Outside of $H = X_v \cup Z_e \cup X_w$, the sides of the triangle Δ' coincide. The intersection of Δ' with H is a triangle whose sides consist of two geodesics and one $(1, 1)$ -quasigeodesic. As we showed H is $(33\delta + 10)$ -hyperbolic in the case $a, b = 1, c = 0$, this triangle is $(33\delta + M'' + 10)$ -thin, where $M'' = M(1, 1, 33\delta + 10) \geq 0$ is obtained by the Morse lemma. Hence Δ is $(33\delta + M'' + M + 2F + 10)$ -thin.

The case $a, b, c \geq 2$: By Proposition 8.7 the sides of Δ are F -close to the sides of Δ_Γ , concatenated with $[x, y_u]$, $[y, y_v]$, and $[z, y_w]$ as appropriate. This latter triangle is $(\delta_\Gamma + 2M)$ -thin, as Δ_Γ is $(\delta_\Gamma + 2M)$ -thin and its sides coincide away from Δ_Γ . Hence Δ is $(\delta_\Gamma + 2M + 2F)$ -thin. \square

Remark 8.12. Some cases in the proof of the above theorem may be bypassed if one applies the ‘guessing geodesics’ criterion for hyperbolicity (see, for example, [21, Lemma 3.5]). However, checking the coherence condition holds among the chosen paths is tantamount to checking many of the same things as is done above. With this in mind, we present the potentially longer but more direct proof above.

REFERENCES

- [1] D. Allcock. Hyperbolic surfaces with prescribed infinite symmetry groups. *Proceedings of the American Mathematical Society*, pages 3057–3059, 2006.
- [2] T. Aougab, P. Patel, and N. G. Vlamis. Isometry groups of infinite-genus hyperbolic surfaces. *Mathematische Annalen*, 381(1):459–498, 2021.
- [3] W. Ballmann, M. Gromov, and V. Schroeder. *Manifolds of nonpositive curvature*. Springer, 1985.
- [4] M. Belolipetsky and A. Lubotzky. Finite groups and hyperbolic manifolds. *Inventiones mathematicae*, 162(3):459–472, 2005.
- [5] O. Bensaid, A. Genevois, and R. Tessera. Coarse separation and large-scale geometry of wreath products, 2024.
- [6] S. Bhowmik and P. Chakraborty. On certain elements and the center of the quasi-isometry groups of Euclidean spaces. Available at SSRN 4195761, 2022.
- [7] M. Bourdon and H. Pajot. Rigidity of quasi-isometries for some hyperbolic buildings. *Commentarii Mathematici Helvetici*, 75:701–736, 2000.
- [8] M. R. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*, volume 319. Springer Science & Business Media, 2013.
- [9] P. Chakraborty. On the center of the group of quasi-isometries of the real line. *Indian Journal of Pure and Applied Mathematics*, 50(4):877–881, 2019.
- [10] J. de Groot. Groups represented by homeomorphism groups I. *Mathematische Annalen*, 138(1):80–102, 1959.
- [11] R. Diestel. *Graph theory*, volume 173. Springer Nature, sixth edition, 2025.
- [12] C. Druțu and M. Kapovich. *Geometric group theory*, volume 63. American Mathematical Society, Providence, RI, 2018. With an appendix by Bogdan Nica.
- [13] P. Eberlein. *Geometry of nonpositively curved manifolds*. University of Chicago Press, 1996.

- [14] A. Eskin, D. Fisher, and K. Whyte. Quasi-isometries and rigidity of solvable groups. *Pure Appl. Math. Q.*, 3(4):927–947, 2007.
- [15] B. Farb and L. Mosher. A rigidity theorem for the solvable Baumslag-Solitar groups. *Invent. Math.*, 131(2):419–451, 1998. With an appendix by Daryl Cooper.
- [16] R. Frigerio and B. Martelli. Countable groups are mapping class groups of hyperbolic 3-manifolds. *Math. Res. Lett.*, 13(5-6):897–910, 2006.
- [17] R. Frucht. Herstellung von Graphen mit vorgegebener abstrakter Gruppe. *Compositio Mathematica*, 6:239–250, 1939.
- [18] R. Frucht. Graphs of degree three with a given abstract group. *Canadian Journal of Mathematics*, 1(4):365–378, 1949.
- [19] L. Greenberg. Maximal groups and signatures. *Ann. of Math. Studies*, 79:207–226, 1974.
- [20] M. Gromov and P. Pansu. Rigidity of lattices: an introduction. In *Geometric Topology: Recent Developments*, pages 39–137. Springer, 2006.
- [21] U. Hamenstädt. Geometry of the complex of curves and of Teichmüller space. In *Handbook of Teichmüller theory. Vol. I*, volume 11 of *IRMA Lect. Math. Theor. Phys.*, pages 447–467. Eur. Math. Soc., Zürich, 2007.
- [22] M. Kapovich and B. Kleiner. Hyperbolic groups with low-dimensional boundary. *Ann. Sci. École Norm. Sup. (4)*, 33(5):647–669, 2000.
- [23] J. H. Kim, B. Sudakov, and V. H. Vu. On the asymmetry of random regular graphs and random graphs. *Random Structures & Algorithms*, 21(3-4):216–224, 2002.
- [24] B. Kleiner and B. Leeb. Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings. *Publications Mathématiques de l'IHÉS*, 86:115–197, 1997.
- [25] S. Kojima. Isometry transformations of hyperbolic 3-manifolds. *Topology and its Applications*, 29(3):297–307, 1988.
- [26] D. König. *Theorie der endlichen und unendlichen Graphen*, volume 72. American Mathematical Soc., 2001.
- [27] O. Mitra and P. Sankaran. Embedding certain diffeomorphism groups in the quasi-isometry groups of Euclidean spaces. *Topology and its Applications*, 265:106833, 2019.
- [28] P. Pansu. Métriques de Carnot-Carathéodory et quasiisométries des espaces symétriques de rang un. *Annals of Mathematics*, pages 1–60, 1989.
- [29] G. Sabidussi. Graphs with given infinite group. *Monatshefte für Mathematik*, 64(1):64–67, 1960.
- [30] P. Sankaran. On homeomorphisms and quasi-isometries of the real line. *Proceedings of the American Mathematical Society*, 134(7):1875–1880, 2006.
- [31] R. E. Schwartz. The quasi-isometry classification of rank one lattices. *Inst. Hautes Études Sci. Publ. Math.*, 82:133–168, 1995.
- [32] P. Scott and T. Wall. Topological methods in group theory. In *Homological group theory (Proc. Sympos., Durham, 1977)*, volume 36, pages 137–203, 1979.
- [33] J. Winkelmann. Realizing countable groups as automorphism groups of riemann surfaces. *Documenta Mathematica*, 6:413–417, 2001.
- [34] K. Wortman. A finitely presented solvable group with a small quasi-isometry group. *Michigan Math. J.*, 55(1):3–24, 2007.
- [35] X. Xie. Quasi-isometric rigidity of fuchsian buildings. *Topology*, 45:101–169, 2006.
- [36] S. Ye and Y. Zhao. The group of quasi-isometries of the real line cannot act effectively on the line. *Algebraic & Geometric Topology*, 23(8):3835–3847, 2023.
- [37] Y. Zhao. *The group of quasi-isometries of the real line*. PhD thesis, University of Liverpool, 2024.