Cross Ratios on Cube Complexes and Length-Spectrum Rigidity



Elia Fioravanti

Merton College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2019

To my parents Giovanni and Patrizia

Acknowledgements.

Above all, I thank my supervisor Cornelia Druţu for her constant guidance over these years, for her invaluable advice on mathematics and mathematicians and for having taught me to have confidence in myself.

I also want to express my most heartfelt gratitude to Talia Fernós for passing on to me her enthusiasm for cubical and median geometry and for her encouragement during our time at MSRI.

I am grateful to my collaborator Jonas Beyrer for having introduced me to cross ratios and for our many interesting conversations. I also thank all the people with whom I have discussed topics related to this thesis: Pierre-Emmanuel Caprace, Christopher Cashen, Matt Cordes, Ruth Charney, Ilya Gekhtman, Anthony Genevois, Mark Hagen, Tobias Hartnick, Merlin Incerti-Medici, John Mackay, Gabriel Pallier, Beatrice Pozzetti, Michah Sageev, Sam Shepherd, Alessandro Sisto and Ric Wade.

I gratefully acknowledge the financial support received from the Clarendon Fund and the Moussouris Scholarship of Merton College. This work would not have been possible without them.

It is impossible for me to mention all the amazing people that I have met during the past four years. Still, I would like to offer my warmest thanks to each and every one of you for the time spent together. Also, sorry for the countless cringy moments.

In random order, special thanks go: to Nicolaus Heuer for giving my pet pheasant something to read; to Federico Vigolo for sharing Johanna's sour tangerines with me; to Tom Zeman and Luciana Basualdo Bonatto for being simply the best housemates that one could wish for (and for ruthlessly correcting my English); to Ronja Kuhne for having been such a thoughtful and selfless friend when I needed one; to Esmee Te-Winkel and Alex Wendland for never saying no to a hiking opportunity; to Alice Kerr for Fawlty Towers, Cabin Pressure and Buffy, but also for having slowly and patiently shown me how to be a better human; to Carla Groenland, Jan Steinebrunner, Chris Weis and Jordan Wheeler for a bit of dancing and a lot of climbing; to Ric Wade for not being able to resist a comfy couch; to Joe Scull for the most successful hell cake; to Ido Grayevsky for the unexpected momos; to Robert Kropholler for always letting me find his belly button and for never showing me the goat; to Ilaria Castellano, Carlos De la Cruz Mengual and the squirrels for a scarring afternoon; to David and Laura Hume for sharing my affection for overly buttered turkeys; to Claudio Llosa Isenrich for his love of asymptotic cones; to Giles Gardam for having spurred me to

work on my German with his fluency; to Antonio De Capua for never having had a jar of honey at hand; and also to Benjamin Brück, Thibaut Dumont, Renee Hoekzema, XiXi Liu, Alex Margolis, Danny Scarponi, Sam Shepherd, Davide Spriano, Gareth Wilkes, Thomas Wasserman, Matthias Wink and all those not explicitly mentioned here. The Whale of Darkness also thanks Aura Raulo for the bucket of oranges.

I am deeply grateful to my partner Maria Martini (and also to Pebbles, the Green Monster, the Literary Koala, Lord Feathers, the Topology Group, the Armed Mice, the Hungry Squirrel, the Hennaed Axolotl, Marta the Rabbittess, Mr. Nuts and to all the other creatures that accompany her, including the one in Figure 8.1). Thank you for having shared with me the last four years, and as many before. Thank you for your constant unwavering support, never showing the slightest trace of doubt in me. Thank you for having always lent me words of comfort and advice, even when my problems seemed to disappear in the face of your own. I could not imagine my life without you.

Finally, I want to thank my parents for having taught me to always be curious and love knowledge. I really owe you everything I have today.

Abstract.

We prove two versions of the *marked length-spectrum rigidity conjecture* for a large class of non-positively curved spaces: CAT(0) cube complexes. Under weak assumptions, we show that proper cocompact actions of Gromov hyperbolic groups on CAT(0) cube complexes are fully determined by their combinatorial length functions. If the cube complexes are irreducible and have the geodesic extension property, the same result holds for non-proper, non-cocompact actions of arbitrary groups.

While some of our arguments are well-known in *negative* curvature, until now no length-spectrum rigidity result had been obtained in a setting of *non-positive* curvature (except for some particular cases in dimension 2). Our spaces are allowed to have arbitrarily high dimension and can contain flats.

Fixing a group Γ, our work sets the basis for the study of the *space of all actions* of Γ on CAT(0) cube complexes. As a first step in this direction, we use length functions to compactify large spaces of cubulations of Γ. This can be viewed as an analogue of Thurston's compactification of Teichmüller space.

As a special case of our results, we construct a compactification of the Charney-Stambaugh-Vogtmann Outer Space for the group of untwisted outer automorphisms of an (irreducible) right-angled Artin group. This generalises the length function compactification of the classical Culler-Vogtmann Outer Space.

Throughout the thesis, our main tool is a new notion of cross ratio on Roller boundaries of CAT(0) cube complexes. We develop a general framework reducing length-spectrum rigidity questions to the problem of extending cross-ratio preserving maps between (subsets of) Roller boundaries. We then study under what conditions such boundary maps extend to cubical isomorphisms. In the hyperbolic case, our results extend part of the work of Bourdon and Xie on Fuchsian buildings.

Along the way, we also analyse the relationship between Roller boundaries, contracting boundaries and visual boundaries of CAT(0) cube complexes. For a Gromov hyperbolic group Γ, this allows us to embed the space of cubulations of Γ into the space of Γ-invariant cross ratios on the Gromov boundary $\partial_\infty\Gamma$. This draws an unexpected similarity between cubulations and the many other geometric structures that can be encoded in terms of boundary cross ratios: for instance, points of Teichmüller space, Hitchin representations, geodesic currents.

Statement of originality.

All of the material in this thesis is adapted from my joint work with J. Beyrer and M. Incerti-Medici. More precisely:

- Sections 5.1–5.2 and Chapters 3, 7 and 9 are joint work with J. Beyrer and appear in [BF19b];
- Sections 2.1.4–2.1.5, Section 5.3 and Chapters 4 and 8 are joint work with J. Beyrer and appear in [BF19a];
- Chapter 6 is joint work with J. Beyrer and M. Incerti-Medici; it appears in [BFIM19].

Except where otherwise indicated, cited, or commonly known, I declare that the contents of this thesis are, to the best of my knowledge, original and the result of my own collaborative work. I also remark that:

- Talia Fernós suggested looking at the regular boundary in relation to Theorem F;
- Anthony Genevois brought Bourdon's paper [Bou97] to my attention. Proposition 8.2.1 was inspired by Bourdon's Lemme 2.4.8;
- Mark Hagen suggested a similar proof of Proposition 2.1.8;
- Sam Shepherd supplied the current proof of Proposition 3.2.4 via Lemma 3.1.2;
- the idea of balanced and unbalanced hyperplanes in Section 2.1.5 stemmed from a conversation with Alessandro Sisto.

This thesis has not been submitted for a degree at another university.

Elia Fioravanti
Oxford, 7 April 2019

Contents

1 Introduction. 1

1.1 Marked-length-spectrum rigidity. 1

1.2 Cross ratios. 6

1.3 Geometric structures as cross ratios. 10

1.4 On the proof of Theorem E and Theorem G. 14

1.5 On the proof of Theorem H. 16

2 Preliminaries. 20

2.1 CAT(0) cube complexes. 20

2.1.1 Definitions and notation. 20

2.1.2 The Roller boundary ∂X. 22

2.1.3 Convexity and bridges. 26

2.1.4 Combinatorial geodesics vs CAT(0) geodesics. 28

2.1.5 Median barycentres. 29

2.2 Automorphisms of cube complexes. 30

2.2.1 Axes and minimal sets. 30

2.2.2 Actions on cube complexes. 32

2.2.3 The regular boundary $\partial_{\mathrm{reg}} X$. 35

3 Cross ratios on cube complexes. 37

3.1 Gromov products. 37

3.2 Definition and first properties. 38

3.3 Opposite points. 41

3.4 Extending op-preserving bijections. 43

3.5 Ample sets of regular points. 45

4 Contracting boundaries of cube complexes. 47
4.1 Contracting geodesics. 47
4.2 Roller boundaries vs contracting boundaries. 48
4.3 Contracting non-terminating ultrafilters. 52

5 From length spectra to boundary cross ratios. 54
5.1 Strongly contracting isometries. 54
5.2 Constructing a Möbius map. 57
5.3 Transferring cross ratios. 60

6 Extending Möbius maps I: the globally-defined case. 64
6.1 Extremal vertices and straight geodesics. 64
6.2 Straight points. 66
6.3 Skinny vertices. 68
6.4 The isomorphism and its uniqueness. 70

7 Extending Möbius maps II: cube complexes with no free faces. 73
7.1 Proof outline. 73
7.2 Existence of OR medians. 75
7.3 Extending partial isomorphisms. 80

8 Extending Möbius maps III: cubulations of hyperbolic groups. 84
8.1 Traces at infinity. 85
8.2 Towards hyperplane recognition. 89
8.3 Trust issues. 91
8.4 Traces vs cross ratios. 93
8.5 Concluding the proof. 98
8.6 Epilogue. 100

9 The space of cubical actions of a group. 103

Index 106

Bibliography 108

Chapter 1

Introduction.

1.1 Marked-length-spectrum rigidity.

As motivation for the entire thesis, we start by recalling the following notorious open problem.

Conjecture 1 (The Marked-Length-Spectrum Rigidity Conjecture, Problem 3.1 in [BK85])**.** Let M and N be closed Riemannian manifolds[1] with negative sectional curvature. If $\phi\colon \pi_1 M \to \pi_1 N$ is an isomorphism that preserves lengths of geodesic representatives, there exists a (unique) isometry $\Phi\colon M \to N$ with $\Phi_* = \phi$.

The 2-dimensional case was settled in the foundational work of Otal [Ota90]. In higher dimension, progress on the conjecture has been strikingly limited. An exception is the case when one of the two manifolds is locally symmetric, which was solved by Hamenstädt [Ham99].

Even for more general spaces, only a handful of length-spectrum rigidity results has been obtained so far. To the best of our knowledge, all progress has been restricted to 2-dimensional spaces (graphs [CM87], non-positively curved Riemannian surfaces [Cro90, CFF92], surfaces with singularities [HP97, BL18], certain non-positively curved square complexes [CM13]) and particularly rigid Finsler manifolds (2-step nilmanifolds [Ebe94], strictly convex, real projective manifolds with the Hilbert metric [Kim01b]).

[1]The original formulation of the conjecture has the additional requirement that M and N be diffeomorphic. We will drop this assumption, as this does not seem to affect what is known on the conjecture. We remark that, since M and N are aspherical and have isomorphic fundamental groups, they are homotopy equivalent. In dimension at least 5, this already implies that M and N are *homeomorphic* [FJ93]; however, there are examples of homeomorphic negatively curved manifolds that are not *diffeomorphic* [FJ89].

Before we go any further, we would like to rephrase Conjecture 1 in the language of geometric group theory. Consider for a moment an arbitrary metric space (X, d) and an (isometric) group action $\Gamma \curvearrowright X$. To each element $g \in \Gamma$, we can associate its *translation length*:

$$\ell_X(g) = \inf_{x \in X} d(x, gx).$$

The resulting function $\ell_X \colon \Gamma \to \mathbb{R}_{\geq 0}$ is usually known as *length function* or *(marked) length spectrum*.

When Γ is the fundamental group of a Riemannian manifold M acting on the universal cover $X = \widetilde{M}$ by deck transformations, the translation length $\ell_X(g)$ coincides with the length of a shortest closed geodesic in the free homotopy class of $g \in \pi_1 M$. Thus, Conjecture 1 is equivalent to the following:

Conjecture 2. Let X and Y be simply connected Riemannian manifolds with negative sectional curvature. Let a group[2] Γ act properly discontinuously and cocompactly on X and Y. If the functions ℓ_X and ℓ_Y coincide, there exists a unique Γ-equivariant isometry $\widetilde{\Phi} \colon X \to Y$.

We see no reason why one should wonder about the above problem only in the setting of Riemannian manifolds, or only in negative curvature:

Question 3. Let X and Y be CAT(0) spaces with the *geodesic extension property*[3]. Let a group Γ act properly discontinuously and cocompactly on X and Y with $\ell_X = \ell_Y$. Does there exist a Γ-equivariant isometry $\widetilde{\Phi} \colon X \to Y$?

The geodesic extension property is automatically satisfied by all Riemannian manifolds by the Hopf-Rinow theorem. Without this requirement, it would be easy to construct counterexamples to Question 3, for instance by glueing in a few loose edges around our spaces (cf. Example 8.6.1 and Figure 8.1).

While tackling Question 3 in full generality is likely to be a hopeless endeavour, it should be possible to say something on specific classes of CAT(0) spaces. The goal of this thesis is to address Question 3 in the case when X and Y are CAT(0) *cube complexes*.

Cube complexes were originally introduced by Gromov [Gro87] and have quickly become ubiquitous within geometric group theory. They are better understood than

[2] If we want a genuine equivalence with Conjecture 1, we should also require Γ to be torsion-free. The formulation above, however, is closer to the setting that we will consider in the rest of thesis.

[3] Definition II.5.7 in [BH99].

arbitrary spaces of non-positive curvature, due to the distinctive features of their geometry. For instance, unlike arbitrary Hadamard manifolds, CAT(0) cube complexes always contain many totally geodesic, codimension-one hypersurfaces.

Still, cubical geometry is by no means trivial, as it displays a vast array of phenomena and it encodes the geometry of a large number of groups. Examples of *cubulated* groups — i.e. those groups that act properly discontinuously and cocompactly on a CAT(0) cube complex — include right-angled Artin groups, hyperbolic or right-angled Coxeter groups [NR03], hyperbolic free-by-cyclic groups [HW15, HW16], hyperbolic 3-manifold groups [BW12], most non-geometric 3-manifold groups [PW14, HP15, PW18], many arithmetic lattices in $SO(n,1)$ [BHW11, HW12], small cancellation groups [Wis04], random groups at low density [OW11], etc.

In our setting, it is more appropriate to endow CAT(0) cube complexes with their ℓ^1 *metric*, also known as *combinatorial metric*. This is bi-Lipschitz equivalent to the CAT(0) metric, but it has the advantage of satisfying all the good properties of median geometry [Rol98, CDH10, Fio17]. This will be crucial to many of our results and we will go into more detail on what fails for the CAT(0) metric later in the introduction.

When a group Γ acts on a CAT(0) cube complex X by cubical automorphisms, the (combinatorial) length function ℓ_X only takes integer values. Indeed, it suffices to compute the infimum in the definition of ℓ_X over vertices of the barycentric subdivision X' (cf. Proposition 2.2.1); these always have integer displacement with respect to the combinatorial metric.

The following two theorems are our main results[4] in relation to Question 3.

Theorem A. *Let X and Y be irreducible* CAT(0) *cube complexes with no free faces; suppose moreover that* $\mathrm{Aut}(X)$ *and* $\mathrm{Aut}(Y)$ *contain uniform lattices. Let Γ be a group*[5] *and let $\Gamma \curvearrowright X$ and $\Gamma \curvearrowright Y$ be essential, non-elementary actions. The two ℓ^1 length functions coincide if and only if X and Y are Γ-equivariantly isomorphic.*

Theorem B. *Let a Gromov hyperbolic group*[6] *Γ act properly*[7] *and cocompactly on essential, hyperplane-essential* CAT(0) *cube complexes X and Y. The two actions have the same ℓ^1 length function if and only if X and Y are Γ-equivariantly isomorphic.*

[4]Also see Theorem M for cube complexes — such as non-hyperbolic Davis complexes — where neither of the two theorems may apply.

[5]We are not requiring Γ to be countable, nor — in case Γ has a topology — that actions have continuous orbits.

[6]To avoid ambiguities, we specify that our Gromov hyperbolic groups are always assumed to be finitely generated.

[7]For countable groups, we use the word 'properly' as a synonym of 'properly discontinuously'.

For precise definitions, we refer the reader to Section 2.2.2. Here we only provide some justification for our hypotheses:

- Regarding hyperbolicity of Γ in Theorem B, note that the groups involved in Conjectures 1 and 2 are all Gromov hyperbolic.

- Having no free faces (cf. Theorem A) is equivalent to the geodesic extension property for the CAT(0) metric [8]. This is a natural assumption as our motivation comes from Riemannian manifolds.

- Cube complexes with no free faces are automatically essential and hyperplane-essential (cf. Theorem B). Being essential and hyperplane-essential, however, is a much weaker requirement.

  Every cocompact group action on a CAT(0) cube complex X can be collapsed[9] to an action on an *essential*, *hyperplane-essential* CAT(0) cube complex $X_\bullet$. This procedure preserves most additional properties that the original action may have had. In particular, the collapsing map $X \twoheadrightarrow X_\bullet$ has uniformly bounded fibres and it is an equivariant quasi-isometry.

- In the setting of Theorem A, every cocompact action is essential. Moreover, every properly discontinuous, cocompact action of a non-virtually-cyclic group is non-elementary (Lemma 2.2.5). Of course, Theorem A also applies to many non-proper and non-cocompact actions.

- Irreducibility cannot be dropped from Theorem A, as the induced action on the barycentric subdivision $\Gamma \curvearrowright X'$ and the diagonal action $\Gamma \curvearrowright X \times X$ always have the same length function $2 \cdot \ell_X$.

Theorem A considers more general actions and more general groups than Theorem B, at the price of imposing stronger assumptions on the cube complexes. The interest of Theorem B lies in the fact that this is the *optimal* result for cubulations of hyperbolic groups. Examples 8.6.1 and 8.6.2 show that essentiality and hyperplane-essentiality — unlike having no free faces — are also *necessary* conditions for any form of length-spectrum rigidity to hold.

In this regard, note that cubulations of hyperbolic groups normally have many free faces. A good example is provided by surface groups $\Gamma = \pi_1 S$: if a cubulation of Γ has cyclic hyperplane-stabilisers and $\dim \geq 3$, then it must have free faces. On the other

[8] See Definition II.5.9 and Proposition II.5.10 in [BH99].

[9] This is proved in Section 3 of [CS11] and Theorem A of [HT18].

hand, one can easily construct plenty of essential, hyperplane-essential cubulations of arbitrarily high dimension by applying Sageev's construction [Sag95, Sag97, BW12] to sufficiently large families of closed curves on S.

The same construction provides a plethora of pairwise non-isomorphic cubulations for many other groups. When $G = F_n$, it suffices to add some punctures to the closed surface S. If instead G is the fundamental group of a closed hyperbolic 3-manifold, one can replace curves by quasi-Fuchsian surfaces [KM12]. Note that all these examples are *automatically* essential and hyperplane-essential, even without resorting to the collapsing procedure described earlier.

Remarks. Although we have preferred to state Theorem A and Theorem B for *cube* complexes, they hold, more generally, for *cuboid* complexes. In such complexes, we allow edges to take arbitrary (positive) real lengths. The price to pay is that we can merely conclude that X and Y are Γ-equivariantly *isometric* in this context, i.e. said isometries will in general not map vertices to vertices.

All results of this thesis admit straightforward generalisations to the context of CAT(0) cuboid complexes. With few exceptions, proofs require no modification. In order to give a simplified treatment, we have decided to only consider the setting of *cube* complexes in this document. The reader will find more details on the cuboidal case in our papers [BF19b, BF19a].

Theorem A can be used to compactify certain spaces of cubulations. The picture that one should have in mind is that of Culler and Vogtmann's Outer Space [CV86], where every point is an action of the free group F_n on a simplicial tree. This space embeds into the space of projectivised length functions $\mathbb{P}(\mathbb{R}^{F_n})$ and can be compactified by adding length functions of certain non-proper actions on real trees [CM87, CL95].

Given a group Γ and an integer $D \geq 0$, let us denote by $\mathrm{Cub}_D(\Gamma)$ the collection of all actions with unbounded orbits of Γ on CAT(0) cube complexes with countably many vertices and dimension at most D. We identify actions that are Γ-equivariantly isomorphic. Each of these actions must admit elements with positive translation length by a result of Sageev (see Theorem 5.1 in [Sag95]), so we can consider the composition

$$\mathrm{Cub}_D(\Gamma) \xrightarrow{\ell} \mathbb{R}^\Gamma \setminus \{0\} \longrightarrow \mathbb{P}(\mathbb{R}^\Gamma),$$

where the first arrow simply assigns to every action its ℓ^1 length function. Theorems A and B can be seen as stating that this composition is injective on a large subset of $\mathrm{Cub}_D(\Gamma)$. Endowing $\mathbb{R}^\Gamma$ with the product topology and $\mathbb{P}(\mathbb{R}^\Gamma)$ with the quotient topology, we will prove the following.

Theorem C. *Let Γ be finitely generated. For every $D \geq 0$, the image of the map $\mathrm{Cub}_D(\Gamma) \to \mathbb{P}(\mathbb{R}^\Gamma)$ is relatively compact.*

We draw the attention of the reader to the generality of Theorem C, since the actions in $\mathrm{Cub}_D(\Gamma)$ are only required to have unbounded orbits. Every point of the closure of the image of $\mathrm{Cub}_D(\Gamma)$ in $\mathbb{P}(\mathbb{R}^\Gamma)$ is the length function of an action of Γ on a finite rank median space with unbounded orbits (Proposition 9.0.2 and Remark 9.0.3).

Theorems A and C can in particular be applied to the *Salvetti blow-ups* constructed in [CSV17] in order to study the group of untwisted outer automorphisms $U(\Gamma) \leq \mathrm{Out}(\Gamma)$ of a right-angled Artin group Γ. The following is an immediate consequence:

Corollary D. *Let Γ be a right-angled Artin group that does not split as a direct product. The Charney-Stambaugh-Vogtmann Outer Space for $U(\Gamma)$ [CSV17] continuously injects into $\mathbb{P}(\mathbb{R}^\Gamma)$ with relatively compact image.*

The 2-dimensional case was obtained in [CM13, Vij15]. We would expect the map in Corollary D to be a topological embedding, but this is not immediate and we have not investigated the matter in this thesis.

Our main tool in proving Theorems A and B is a new notion of *cross ratio* on boundaries of CAT(0) cube complexes. In this, we draw inspiration from the pioneering work of Otal on the 2-dimensional case of Conjecture 1 [Ota90, Ota92].

Before we say anything on proofs, we would like to provide some background on cross ratios and their relation to marked-length-spectrum rigidity.

1.2 Cross ratios.

Cross ratios arise naturally in the study of boundaries at infinity of negatively curved spaces. A classical example is provided by the standard projective cross ratio[10] on $\partial_\infty \mathbb{H}^2 \simeq \mathbb{RP}^1$. More generally, given a CAT(-1) space X and four pairwise-distinct points x, y, z, w in the Gromov boundary $\partial_\infty X$, their cross ratio is defined as:

$$\mathrm{cr}(x, y, z, w) = (x \cdot z)_p + (y \cdot w)_p - (x \cdot w)_p - (y \cdot z)_p,$$

[10] One should rather take the *logarithm* of the *absolute value* of the projective cross ratio if this is to match the definition given below.

where $(- \cdot -)_p$ denotes Gromov products[11] with respect to a basepoint $p \in X$. The value of $\mathrm{cr}(x, y, z, w)$ is independent of the choice of p and it is preserved by the diagonal action of the isometry group of X on $(\partial_\infty X)^4$.

Cross ratios have proved a valuable tool in relation to strong rigidity results [Ota90, Bou96, Bou97]. Our interest in them is motivated by a well-known connection with the Marked-Length-Spectrum Rigidity Conjecture. Indeed, it turns out that cross ratios of certain 4-tuples of boundary points can be expressed in terms of length functions (cf. Théorème 2.2 in [Ota92] and Section 5 in [Kim01a]). As a consequence, we have the following:

Theorem (Theorem 5.1 in [Bis15])**.** *Let X and Y be simply connected Riemannian manifolds with negative sectional curvature. Let a torsion-free group Γ act properly discontinuously and cocompactly on X and Y and let $f\colon \partial_\infty X \to \partial_\infty Y$ be the unique Γ-equivariant homeomorphism. We have $\ell_X = \ell_Y$ if and only if f is cross-ratio preserving.*

Conjectures 1 and 2 now become equivalent to the Riemannian case of the following conjecture, which seems to be folklore in the field.

Conjecture 4. Let X and Y be CAT(-1) spaces with the geodesic extension property, endowed with properly discontinuous, cocompact actions of a group Γ. Every Γ-equivariant, cross-ratio preserving homeomorphism $f\colon \partial_\infty X \to \partial_\infty Y$ extends to a Γ-equivariant isometry $F\colon X \to Y$.

Progress on Conjecture 4 has been just as limited as progress on the Marked-Length-Spectrum Rigidity Conjecture. Nevertheless, it is interesting to remark that the homeomorphism f is known to extend to a $(1, \log 2)$-quasi-isometry [Bis15].

One can similarly define cross ratios on *contracting boundaries*[12] of CAT(0) spaces and on Gromov boundaries of Gromov hyperbolic spaces. In this more general context, unfortunately, the connection between Conjecture 1 and Conjecture 4 normally breaks down. As we shall see, this is not the case for CAT(0) cube complexes endowed with their combinatorial metric (Theorem F).

Quite surprisingly — and this is perhaps the most innovative feature of our work — we will also be able to answer positively several analogues of Conjecture 4 in the cubical context; see Theorems E, G and H. Theorems A and B will follow from these results.

[11] See e.g. Definition III.H.3.15 in [BH99].

[12] Contracting boundaries of CAT(0) spaces were introduced in [CS15].

Before we go any further, it is interesting to remark that the "coarse analogue" of Conjecture 4 is well-known to have a positive answer. If X is a Gromov hyperbolic or CAT(0) space, its *quasi-isometry* type is fully determined by the cross ratio, respectively, on the Gromov or contracting boundary [Pau96, CCM18].

Let us now consider a CAT(0) cube complex X endowed with its combinatorial metric d. Given vertices $x, y, z, w \in X$, it is natural to define their *cross ratio* as:

$$\mathrm{cr}(x, y, z, w) = d(x, w) + d(y, z) - d(x, z) - d(y, w),$$

by analogy with the context of CAT(-1) spaces. In our setting, this can be rewritten as:

$$\mathrm{cr}(x, y, z, w) = \#\mathscr{W}(x, z|y, w) - \#\mathscr{W}(x, w|y, z) \in \mathbb{Z},$$

where $\mathscr{W}(x, z|y, w)$ is the collection of hyperplanes separating x and z from y and w. The latter formula *continuously* extends the cross ratio to the horofunction boundary of (X, d), aka the *Roller boundary* ∂X (Proposition 3.2.4).

Unlike the visual boundary[13] $\partial_\infty X$ — which is the horofunction boundary associated to the CAT(0) metric on X — the Roller boundary ∂X is always totally disconnected. It is important to observe this, since —as mentioned — the cross ratio on ∂X is both continuous and $\mathbb{Z}$-valued.

With these definitions, we can prove the following analogue of Conjecture 4.

Theorem E. *Let X and Y be* CAT(0) *cube complexes with no extremal vertices and not isomorphic to $\mathbb{R}$. Every cross-ratio preserving bijection $f\colon \partial X \to \partial Y$ uniquely extends to a cubical isomorphism $F\colon X \to Y$.*

Note that the theorem does not require finite dimensional or locally finite cube complexes, nor any group action. We need to rule out extremal vertices[14] or one could modify a bounded portion of the cube complex X without affecting its boundary ∂X. In relation to Conjecture 4, we observe that, when the CAT(0) metrics have the geodesic extension property, there are no extremal vertices (Remark 6.1.4). However, having no extremal vertices is a much weaker assumption.

Unfortunately, Theorem E is of no help in the proof of Theorems A and B. The point is that we still have not said anything on how to relate length functions and cross ratios in the setting of CAT(0) cube complexes. This turns out to be a much more

[13]For the visual boundary of a CAT(0) space and the Gromov boundary of a Gromov hyperbolic space, we refer the reader, respectively, to Chapters II.8 and III.H.3 in [BH99]. We denote both boundaries by $\partial_\infty X$, as these coincide whenever both defined.

[14]See Definition 6.1.1.

delicate matter than in CAT(-1) spaces: Actions with the same length spectrum will only give rise to cross-ratio preserving maps between certain *subsets* of the Roller boundaries. We address this matter now.

We will be interested in the subset $\partial_{\mathrm{reg}} X \subseteq \partial X$ known as *regular boundary*. This was introduced by [Fer18, FLM18] in the study of random walks on cube complexes. As we will describe later (Theorem 4.2), the regular boundary shares a large subset with the contracting boundary $\partial_c X$ of Charney and Sultan [CS15].

The following result is the first step in the proofs of both Theorem A and Theorem B. This should clarify why Theorem E is not a good answer to the cubical analogue of Conjecture 4.

Theorem F. *Let X and Y be locally finite, finite dimensional, irreducible* CAT(0) *cube complexes. Given essential, non-elementary actions $\Gamma \curvearrowright X$ and $\Gamma \curvearrowright Y$ with the same ℓ^1 length function, there exists a Γ-equivariant, cross-ratio preserving bijection $f\colon \mathcal{A} \to \mathcal{B}$ between nonempty, Γ-invariant subsets $\mathcal{A} \subseteq \partial_{\mathrm{reg}} X$ and $\mathcal{B} \subseteq \partial_{\mathrm{reg}} Y$.*

Note that the hypotheses of Theorem F are satisfied both in the setting of Theorem A and in that of Theorem B. After Theorem F, however, the proofs of Theorems A and B part ways. We will need two different results on the extension of cross-ratio preserving boundary maps.

In the case of Theorem A, the proof is completed with the following:

Theorem G. *Let X and Y be irreducible* CAT(0) *cube complexes with no free faces and endowed with essential, non-elementary actions of a group Γ. Suppose that* $\mathrm{Aut}(X)$ *and* $\mathrm{Aut}(Y)$ *contain uniform lattices. Let $f\colon \mathcal{A} \to \mathcal{B}$ be a Γ-equivariant, cross-ratio preserving bijection, where $\mathcal{A} \subseteq \partial_{\mathrm{reg}} X$ and $\mathcal{B} \subseteq \partial_{\mathrm{reg}} Y$ are nonempty and Γ-invariant. Then, there exists a unique Γ-equivariant cubical isomorphism $F\colon X \to Y$ extending f.*

In the setting of Theorem B, the cube complexes are Gromov hyperbolic and the regular boundary[15] $\partial_{\mathrm{reg}} X \subseteq \partial X$ is naturally identified with a subset $\partial_{\mathrm{reg}} X \subseteq \partial_\infty X$ of the Gromov boundary. The latter subset is *co-meagre*[16], as it is exactly the complement of the union of the Gromov boundaries of the hyperplanes of X.

Now, the map f provided by Theorem F is just a restriction of the unique Γ-equivariant homeomorphism $f\colon \partial_\infty X \to \partial_\infty Y$. The proof of Theorem B is thus completed with the following:

[15]It is interesting to remark that, for hyperbolic cube complexes, the regular boundary also coincides with the collection of non-terminating ultrafilters introduced in [NS13].

[16]A set is *co-meagre* if its complement is a countable union of sets whose closures have empty interior. By Baire's theorem, co-meagre subsets of $\partial_\infty \Gamma$ are dense.

Theorem H. *Let a non-elementary Gromov hyperbolic group Γ act properly and cocompactly on essential, hyperplane-essential* CAT(0) *cube complexes X and Y. Let $f\colon \partial_\infty X \to \partial_\infty Y$ be the unique Γ-equivariant homeomorphism. Suppose that there exists a nonempty, Γ-invariant subset $\Omega \subseteq \partial_{\mathrm{reg}} X$ such that $f(\Omega) \subseteq \partial_{\mathrm{reg}} Y$ and cross ratios of elements of Ω^4 are preserved by f. There exists a unique Γ-equivariant isomorphism $F\colon X \to Y$ extending f.*

We believe it is interesting to compare Theorem H with two old results obtained in different settings. As mentioned earlier, Paulin showed that homeomorphisms of Gromov boundaries arise from *quasi-isometries* if and only if they *almost preserve* the boundary cross ratio (Theorem 1.2 in [Pau96]). Theorem H can thus be regarded as an exact version of this result when the spaces are cube complexes.

Paulin's arguments share some (superficial) similarities with our proof of Theorem G. Theorem H, on the other hand, is obtained through completely different techniques. In fact, a more fitting comparison can be drawn with Bourdon's Proposition 2.4.7 from [Bou97], whose statement strikingly resembles that of Theorem H.

This is not a coincidence, as Bourdon's buildings $I_{p,q}$ — and, more generally, all Fuchsian buildings without triangular chambers — can be given a natural structure of CAT(0) square complex; see e.g. Section 2.2 in [GM18]. The hyperplanes of the CAT(0) square complex correspond to Bourdon's *arbre-murs*, along with a choice of a preferred side. With these observations in mind, Bourdon's *birapport combinatoire* on $\partial_\infty I_{p,q} \simeq \partial_\infty \Gamma_{p,q}$ becomes a manifestation of our cross ratio. Proposition 2.4.7 in [Bou97] and part of Theorem 1.5 in [Xie06] are then special cases of Theorem H.

We will sketch the proofs of Theorems E, G and H later in the introduction, as these are probably the most interesting results of this thesis. Before that, however, we would like to describe some more consequences of these results.

1.3 Geometric structures as cross ratios.

So far, we have only discussed very specific instances of cross ratios, for instance those on Roller boundaries of cube complexes and on Gromov boundaries of CAT(-1) spaces. However, these objects can be given a unifying abstract definition.

Given a group Γ, the contracting boundary $\partial_c \Gamma$ was introduced by Cordes [Cor17]; when Γ is Gromov hyperbolic, this coincides with the Gromov boundary $\partial_\infty \Gamma$. We

denote by $\partial_c\Gamma^{(4)}$ the set of 4-tuples of pairwise-distinct points in the contracting boundary. A map $\mathbb{B}\colon \partial_c\Gamma^{(4)} \to \mathbb{R}$ is said to be an *(abstract) cross ratio* if the following are satisfied:

(i) $\mathbb{B}(x,y,z,w) = -\mathbb{B}(y,x,z,w)$;

(ii) $\mathbb{B}(x,y,z,w) = \mathbb{B}(z,w,x,y)$;

(iii) $\mathbb{B}(x,y,z,w) = \mathbb{B}(x,y,z,t) + \mathbb{B}(x,y,t,w)$;

(iv) $\mathbb{B}(x,y,z,w) + \mathbb{B}(y,z,x,w) + \mathbb{B}(z,x,y,w) = 0$.

We say that $\mathbb{B}$ is *invariant* if it is preserved by the diagonal action of Γ on $\partial_c\Gamma^{(4)}$. Similar notions appear e.g. in [Ota92, Ham97, Lab06, Lab07].

Abstract cross ratios provide a unified interpretation of many *geometric structures*, thus proving a valuable tool to study various spaces of representations.

For instance, when S is a closed hyperbolic surface and $\Gamma = \pi_1 S$, every point of the Teichmüller space of S yields identifications $\partial_\infty\Gamma \simeq \partial_\infty\mathbb{H}^2$ and the cross ratio on $\partial_\infty\mathbb{H}^2$ can be pulled back to an invariant cross ratio on $\partial_\infty\Gamma$. The latter uniquely determines the original point of Teichmüller space [Bon88]. More generally, the space of negatively curved metrics on S embeds into the space of invariant cross ratios on $\partial_\infty\Gamma$ [Ota90] and so does the space of geodesic currents.

Another setting where cross ratios play a central role is the study of representations of surface groups into higher-rank Lie groups. A striking result of Labourie identifies the space of Hitchin representations $\rho\colon \Gamma \to PSL_n\mathbb{R}$ with a space of Hölder-continuous invariant cross ratios on $\partial_\infty\Gamma$ [Lab07].

Cubulations can be viewed as yet another geometric structure that groups can be endowed with. As one more consequence of Theorems G and H, large spaces of cubulations of Γ naturally inject into the space of invariant $\mathbb{Z}$-valued cross ratios on $\partial_c\Gamma$.

Theorem I. *Let Γ be a finitely generated, non-virtually-cyclic group.*

(1) Every proper cocompact action of Γ on an irreducible CAT(0) *cube complex with no free faces X canonically determines an invariant cross ratio:*

$$\mathrm{cr}_X\colon \partial_c\Gamma^{(4)} \to \mathbb{Z}.$$

(2) Given another action $\Gamma \curvearrowright Y$ as above, the cross ratios cr_X and cr_Y coincide if and only if X and Y are Γ-equivariantly isomorphic.

It is worth pointing out that, under the hypotheses of Theorem I, the contracting boundary $\partial_c\Gamma$ is always nonempty and, in fact, infinite.

Theorem J. *Let Γ be a non-elementary Gromov hyperbolic group.*

(1) Every proper cocompact action of Γ on an essential CAT(0) *cube complex X canonically determines an invariant cross ratio*

$$\mathrm{cr}_X \colon \partial_\infty\Gamma^{(4)} \to \mathbb{Z}.$$

There exists a co-meagre subset $\mathcal{C} \subseteq \partial_\infty\Gamma$ such that cr_X is continuous at all points of $\mathcal{C}^{(4)} \subseteq \partial_\infty\Gamma^{(4)}$.

(2) Let X and Y be hyperplane-essential. If there exists a co-meagre subset $\mathcal{D} \subseteq \partial_\infty\Gamma$ such that the cross ratios cr_X and cr_Y coincide on $\mathcal{D}^{(4)} \subseteq \partial_\infty\Gamma^{(4)}$, then X and Y are Γ-equivariantly isomorphic.

In particular, two essential hyperplane-essential cubulations of a hyperbolic group yield the same boundary cross ratio if and only if they are Γ-equivariantly isomorphic. Examples 8.6.1 and 8.6.2 show that — in a very strong sense — neither essentiality nor hyperplane-essentiality can be dropped.

Identifying $\partial_\infty\Gamma \simeq \partial_\infty X$ in Theorem J, the cross ratio cr_X is continuous at all 4-tuples with coordinates in $\partial_{\mathrm{reg}}X \subseteq \partial_\infty X$ (Proposition 5.3.4). As mentioned earlier, this is the subset of points that do not lie in the Gromov boundary of any hyperplane of X. It is interesting to remark that $\mathcal{C}$ can be chosen to be independent of the cube complex X. By Baire's theorem, the union $\overline{\mathcal{C}}$ of all Gromov boundaries of infinite-index, quasi-convex subgroups of Γ is meagre in $\partial_\infty\Gamma$. Every point of the co-meagre set $\partial_\infty\Gamma \setminus \overline{\mathcal{C}}$ thus lies in $\partial_{\mathrm{reg}}X$ for all essential cubulations of Γ.

The reader might wonder why, instead, no form continuity is mentioned in Theorem I. As a matter of fact, cr_X is again continuous[17] at a "large" subset of $\partial_c\Gamma^{(4)}$, but it does not make sense to speak of *co-meagre* subsets in this context. Indeed, the entire contracting boundary is often itself meagre, even in the Cashen-Mackay topology. In fact, $\partial_c\Gamma$ is likely to be a Baire space if and only if Γ is hyperbolic. The latter observation follows from Theorem 7.6 in [CM18] when Γ is a toral relatively hyperbolic group. We thank Chris Cashen for pointing this out to us.

As mentioned, Theorems I and J follow from Theorems G and H, respectively. An ingredient is still missing, however, namely a study of the relationship between the Roller boundary ∂X and the contracting boundary $\partial_c X$.

[17]Both endowing $\partial_c\Gamma$ with the visual topology and with the topologies of [Cor17, CM18].

Fixing a basepoint $p \in X$, every point of ∂X is represented by a combinatorial ray based at p. We denote by $\partial_{\mathrm{cu}} X \subseteq \partial X$ the subset of points that are represented by *contracting* combinatorial rays. We endow $\partial_{\mathrm{cu}} X$ with the restriction of the (totally disconnected) topology of ∂X. We moreover denote by $\partial_c^{\mathrm{vis}} X$ the space obtained by endowing the contracting boundary $\partial_c X$ with the restriction of the visual topology on the visual boundary $\partial_\infty X$.

Responding to a suggestion in the introduction of [CS15], we prove:

Theorem K. *Let X be a uniformly locally finite* CAT(0) *cube complex.*

(1) There exists a natural continuous surjection $\Phi\colon \partial_{\mathrm{cu}} X \longrightarrow \partial_c^{\mathrm{vis}} X$ with finite fibres. Collapsing its fibres, Φ descends to a homeomorphism.

(2) If X is hyperbolic, we have $\partial_{\mathrm{cu}} X = \partial X$ and $\partial_c^{\mathrm{vis}} X = \partial_\infty X$.

The reader will find additional details on the map Φ in Section 4.2, especially in Remark 4.2.4 and Theorem 4.2.6. We stress that — whenever flats are present — it is not possible to represent the *entire* visual boundary $\partial_\infty X$ as a quotient of a subset of the Roller boundary.

Now, Theorems I and J are obtained by picking canonical barycentres for the fibres of the map Φ. These are built through a new construction of barycentres for bounded cube complexes, which we describe in Section 2.1.5.

We conclude this section with the following question, which is naturally brought to mind by Theorem J.

Question. Let Γ be a non-elementary hyperbolic group. Is it possible to provide a set of conditions completely characterising which invariant, $\mathbb{Z}$-valued cross ratios on $\partial_\infty \Gamma$ arise from cubulations of Γ?

Theorem 1.1 in [Lab07] is a result of this type in the context of Hitchin representations. A complete answer to the above question might provide a new procedure to cubulate groups.

In this regard, note that $\partial_\infty \Gamma$ is endowed with a continuous, invariant, $\mathbb{R}$-*valued* cross ratio whenever Γ acts properly and cocompactly on a CAT(-1) space. So it would also be interesting to determine under what circumstances an invariant $\mathbb{R}$-valued cross ratio can be discretised to an invariant $\mathbb{Z}$-valued cross ratio. Of course, one should be very careful when making speculations, as, for instance, uniform lattices in $SU(n,1)$ and $Sp(n,1)$ are not cubulable [DP16, NR97].

1.4 On the proof of Theorem E and Theorem G.

In both cases, we have a cross-ratio preserving map $f\colon \partial X \dashrightarrow \partial Y$ defined on (subsets of) the Roller boundaries. Our aim is to extend f to an isomorphism $F\colon X \to Y$.

This is one of the parts of this thesis where it becomes absolutely crucial to endow our cube complexes with their ℓ^1 metrics. Indeed, the corresponding horofunction compactification $X \cup \partial X$, is endowed with a structure of *median algebra* [Rol98, NS13, Fio17], unlike the *visual compactification* $X \cup \partial_\infty X$. This means that to three points $x, y, z \in \partial X$ we can associate a canonical barycentre $m(x, y, z) \in X$.

We will essentially aim to represent each vertex $v \in X$ as $m(x, y, z)$ for points $x, y, z \in \partial X$, and then set $F(m(x, y, z)) = m(f(x), f(y), f(z))$. There are two main issues with this approach:

(1) It is not a priori clear that F is well-defined. There might exist two triples $(x, y, z), (x', y', z') \in (\partial X)^3$ for which the medians $m(x, y, z)$ and $m(x', y', z')$ coincide, while $m(f(x), f(y), f(z))$ and $m(f(x'), f(y'), f(z'))$ are different.

(2) In Theorem G, the map f is only defined on the subset $\mathcal{A} \subseteq \partial_{\mathrm{reg}} X$, so we need to be able to pick x, y, z within $\partial_{\mathrm{reg}} X$. In many cases, however, there exist vertices of X that cannot be represented as medians of regular points. An example is provided by the vertices that we add to any cube complex when we barycentrically subdivide it.

A naive solution to issue (1) might aim to recognise purely in terms of cross ratios whether the condition "$m(x, y, z) = m(x', y', z')$" holds. This proves to be a delicate matter, as the cross ratios of the six points alone will in general not suffice.

The problem is illustrated in Figure 1.1 in the simple case with $x = x'$ and $y = y'$. In both cube complexes, all 4-tuples obtained by permuting x, y, z and z' have vanishing cross ratio, although we have $m(x, y, z) = m(x, y, z')$ on the left and $m(x, y, z) \neq m(x, y, z')$ on the right.

This is a characteristic feature of cross ratios on cube complexes of dimension at least 3. Indeed, cross ratios obtained by permuting coordinates of a 4-tuple (x, y, z, w) will only provide information on the respective differences between the cardinalities of $\mathscr{W}(x, y|z, w)$, $\mathscr{W}(x, z|y, w)$, $\mathscr{W}(x, w|y, z)$, and never on these three quantities themselves. Note in this regard that these three sets of hyperplanes are pairwise transverse.

These problems can be circumvented by representing vertices v as medians of triples (x, y, z) that satisfy the additional condition $x \operatorname{op}_z y$ (Definition 3.3.2). This means that all geodesics with endpoints x and y must pass through v. This forces

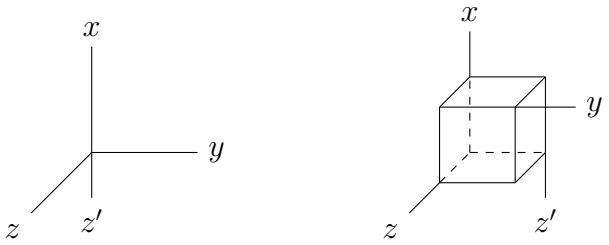


Figure 1.1: The cube complex X, on the left, is a tree with a single branch point and four boundary points x, y, z and z'. Pictured on the right is a portion of $Y \simeq \mathbb{R}^3$; the points x, y, z and z' lie in the Roller boundary $\partial\mathbb{R}^3$.

one of the three sets $\mathscr{W}(x, y|z, w)$, $\mathscr{W}(x, z|y, w)$, $\mathscr{W}(x, w|y, z)$ to be empty for every choice of w. In relation to Figure 1.1, we have $x \operatorname{op}_z y$ on the left, but not on the right. For more details, we refer the reader to Section 3.3.

It can be shown that most vertices v are of the form $v = m(x, y, z)$ for a triple (x, y, z) with $x \operatorname{op}_z y$ (Lemma 6.3.1). Moreover, such triples can be characterised in terms of cross ratios (Proposition 3.3.4 and Lemma 6.1.5). It follows that the map f preserves the property $x \operatorname{op}_z y$ and, from there, it is not hard to obtain Theorem E.

The proof of Theorem G, unfortunately, is a lot more involved. Issue (2) is now exacerbated by the fact that we need to represent vertices $v \in X$ as $v = m(x, y, z)$ not only for $x, y, z \in \partial_{\mathrm{reg}} X$, but also with $x \operatorname{op}_z y$. We say that v is an *OR median* if this is possible (Definition 7.1.3).

One of the most technical parts of the thesis is devoted to showing that, under the hypotheses of Theorem G, there always exists at least one OR median (Theorem 7.1.5). It is then possible to define an isometry F between the subsets of OR medians (Theorem 3.5.3). Finally, the proof of Theorem G is completed by extending this partial isometry to a global isomorphism. This is achieved by relying on the following result, which we believe of independent interest.

Theorem L. *Let X and Y be irreducible, locally finite* CAT(0) *cube complexes with no free faces, endowed with cocompact actions $G_1 \curvearrowright X$ and $G_2 \curvearrowright Y$. Let nonempty subsets $A \subseteq X$ and $B \subseteq Y$ be invariant, respectively, under G_1 and G_2. Every distance-preserving bijection $\phi\colon A \to B$ admits a unique extension to a cubical isomorphism $F\colon X \to Y$.*

It is interesting to remark that absence of free faces is not strictly necessary in our proof of Theorem L. What we really need is essentiality, hyperplane-essentiality and Lemma 7.2.3, which should be considerably weaker.

Thus, Theorems A and G fundamentally only require absence of free faces in order to show the existence of an OR median (Theorem 7.1.5 and Section 7.2). If the existence of OR medians is given a priori, the following variation on Theorem A can be useful.

Theorem M. *Let X and Y be locally finite, finite dimensional, irreducible, sector-heavy* CAT(0) *cube complexes. Let essential, non-elementary actions $\Gamma \curvearrowright X$ and $\Gamma \curvearrowright Y$ have the same ℓ^1 length function. There exists a Γ-equivariant isometry between the subsets of OR medians.*

Here, we say that X is *sector-heavy* if the intersection of any family of pairwise-transverse halfspaces contains a halfspace (Definition 2.2.9). This is automatically satisfied when X has no free faces (Proposition 2.2.10).

Theorem M applies for instance to all Davis complexes associated to (irreducible) right-angled Coxeter groups. Although these can have many free faces, they are always sector-heavy and all their vertices are OR medians.

In relation to Theorem M, we should mention that we are not aware of any (reasonable) cube complex where the subset of OR medians is empty. Thus, the following is in order:

Question. Let X be a locally finite, finite dimensional, irreducible CAT(0) cube complex admitting an essential, non-elementary action of a finitely generated group Γ. Does X always contain an OR median?

1.5 On the proof of Theorem H.

Despite the similarities in the statements, the proof of Theorem H follows a completely different outline to Theorems E and G. Rather than relying on the median structure of $X \cup \partial X$, we will make crucial use of hyperbolicity in order to recognise hyperplanes and halfspaces from their traces in the Gromov boundary $\partial_\infty X$.

As mentioned, some of our arguments were inspired by the work of Bourdon on Fuchsian buildings. However, unlike the strictly 2-dimensional setting of Fuchsian buildings, our cube complexes can have arbitrarily high dimension and this will seriously complicate proofs.

Let us denote by $\mathscr{W}(X)$ and $\mathscr{H}(X)$, respectively, the collections of all hyperplanes and all halfspaces of the cube complex X. The rough idea is that it should be possible

to reconstruct the structure of the *halfspace pocset*[18] $(\mathscr{H}(X), \subseteq, *)$ simply by looking at the Gromov boundary $\partial_\infty X$ and the cubical cross ratio cr.

Overlooking various complications, there are two (bipartite) steps:

(Ia) Every $\mathfrak{h} \in \mathscr{H}(X)$ is 'visible' in $\partial_\infty X$.

(Ib) Given $\mathfrak{h}, \mathfrak{k} \in \mathscr{H}(X)$, we have $\mathfrak{h} \subseteq \mathfrak{k}$ if and only if $\partial_\infty \mathfrak{h} \subseteq \partial_\infty \mathfrak{k}$. [19]

(IIa) For every $\mathfrak{w} \in \mathscr{W}(X)$, there exists $\mathfrak{w}' \in \mathscr{W}(Y)$ with $f(\partial_\infty \mathfrak{w}) = \partial_\infty \mathfrak{w}'$.

(IIb) For every $\mathfrak{h} \in \mathscr{H}(X)$, there exists $\mathfrak{h}' \in \mathscr{H}(Y)$ with $f(\partial_\infty \mathfrak{h}) = \partial_\infty \mathfrak{h}'$.

Steps (Ia) and (IIb) show that the Γ-equivariant homeomorphism $f\colon \partial_\infty X \to \partial_\infty Y$ induces a Γ-equivariant bijection $f_*\colon \mathscr{H}(X) \to \mathscr{H}(Y)$. By Step (Ib), the bijection f_* preserves inclusions and, by general theory of CAT(0) cube complexes, f_* must be induced by a Γ-equivariant isomorphism $F\colon X \to Y$.

Steps (Ia) and (Ib) are the key point where, respectively, *essentiality* and *hyperplane-essentiality* come into play. Example 8.6.1 shows that, if X is not essential, some halfspaces may be invisible in $\partial_\infty X$, i.e. Step (Ia) fails. Without hyperplane-essentiality, instead, $\partial_\infty X$ may not be able to tell whether two halfspaces are nested or not. This is exactly the problem in Example 8.6.2, where *transverse* halfspaces $\mathfrak{h}$ and $\mathfrak{k}$ have $\partial_\infty \mathfrak{h} = \partial_\infty \mathfrak{k}$.

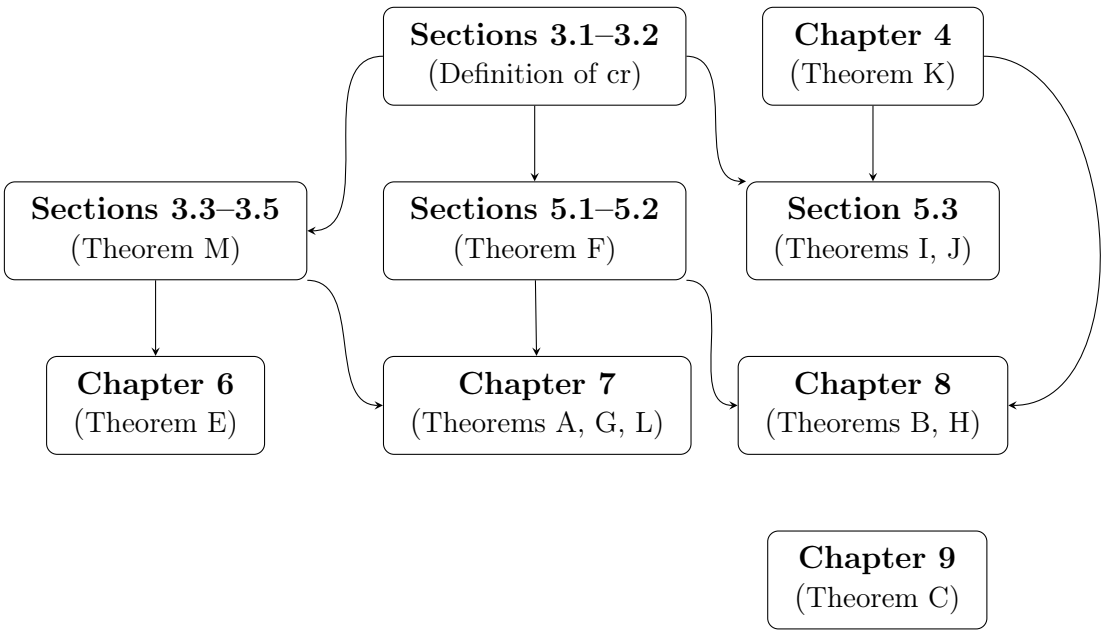
Regarding Step (IIa), it is not hard to use cr to characterise which pairs of points $\xi, \eta \in \partial_\infty X$ lie in the Gromov boundary of a common hyperplane (Proposition 8.2.1). This property is then preserved by f, which is all one needs if no two hyperplanes share asymptotic directions (e.g. in Fuchsian buildings). In general, we will require more elaborate arguments (Lemma 8.4.1 and Proposition 8.4.3) based on the fact that $\partial_\infty \Gamma$ cannot be covered by limit sets of infinite-index quasi-convex subgroups.

Finally, there is a deceiving similarity between the statements of Steps (IIa) and (IIb), but the proof of the latter is significantly more complicated. Given $\mathfrak{w} \in \mathscr{W}(X)$ bounding $\mathfrak{h} \in \mathscr{H}(X)$, the set $\partial_\infty \mathfrak{h} \setminus \partial_\infty \mathfrak{w}$ is a union of connected components of $\partial_\infty X \setminus \partial_\infty \mathfrak{w}$. However, $\partial_\infty X \setminus \partial_\infty \mathfrak{w}$ will in general have many more components than there are halfspaces bounded by $\mathfrak{w}$. The case to keep in mind is when $\Gamma = \pi_1 S$, for a closed oriented surface S, and the hyperplane-stabiliser $\Gamma_\mathfrak{w} < \Gamma$ is the fundamental group of a subsurface of S with at least 3 boundary components.

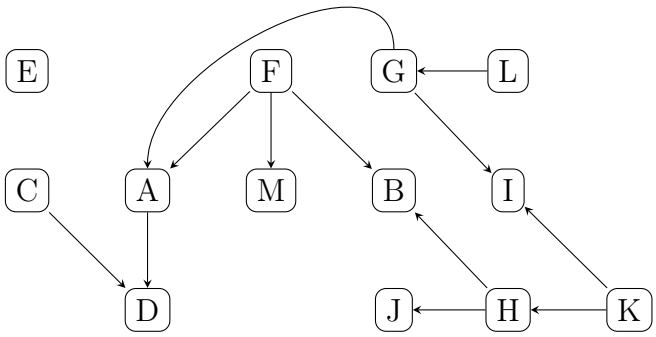
[18] This is just the set of halfspaces, endowed with the order relation given by inclusions, and with the order-reversing involution coming from taking complements.

[19] This is not true in general, but it is how one should think about things. It only fails when $\partial_\infty \mathfrak{h} = \partial_\infty \mathfrak{k}$ and $\mathfrak{k} \subsetneq \mathfrak{h}$, in which case $\mathfrak{k}$ and $\mathfrak{h}$ are at finite Hausdorff distance anyway.

Not all $\Gamma_{\mathfrak{w}}$-invariant partitions of the set of connected components of $\partial_\infty X \setminus \partial_\infty \mathfrak{w}$ arise from a halfspace of X. Thus, one cannot recover $\partial_\infty \mathfrak{h}$ from the knowledge of $\partial_\infty \mathfrak{w}$ purely through topological and dynamical arguments. We will instead rely again on the cross ratio cr in order to circumvent these issues. Step (IIb) will be finally completed in Theorem 8.5.3.



Above is a scheme of the material required in each chapter. Theorems are listed under the chapters where they are proved.
Proofs sometimes rely on results obtained later in the thesis. So we provide below a scheme of which theorems are required in each other's proofs. Only results appearing in the introduction are listed.



Chapter 2

Preliminaries.

2.1 CAT(0) cube complexes.

2.1.1 Definitions and notation.

We will assume a certain familiarity with basic properties of CAT(0) cube complexes. The reader can consult for instance [Sag14] and the first sections of [CN05, CS11, NS13, CFI16] for an introduction to the subject. Here, we merely fix notation and recall various standard facts that we will be relying on over the course of the thesis.

Let X be a simply connected cube complex satisfying Gromov's no-$\triangle$-condition; see 4.2.C in [Gro87] and Chapter II.5 in [BH99]. As has become customary, we will not distinguish between X and its 0-skeleton, denoting both simply by X. It will always be clear from the context how to interpret the notation, although — unless specified otherwise — points $v \in X$ will normally be understood to be vertices.

The Euclidean metrics on the cubes of X fit together to yield a CAT(0) metric on X, thus justifying the terminology *"*CAT(0) *cube complex"*. Such CAT(0) metric, however, will generally be of no interest to us[1] and we instead prefer to endow X with its ℓ^1 *metric* (aka *combinatorial metric*). This is the induced path metric obtained by endowing each cube $[0,1]^k \subseteq X$ with the restriction of the ℓ^1 metric of $\mathbb{R}^k$. The notation $d(-,-)$ will always refer to this latter metric.

Given vertices $v, w \in X$, their combinatorial distance $d(v,w)$ coincides with the length of a shortest path connecting v and w within the 1-skeleton; as such, it is always a natural number. In finite dimensional cube complexes, the CAT(0) metric and the combinatorial metric are bi-Lipschitz equivalent and complete. It is however important to remark that (X,d) is *uniquely geodesic* only when X is a tree.

[1] Sections 2.1.4 and 4.2 are an exception.

We will employ the more familiar and concise expression "CAT(0) cube complex" with the meaning of "simply connected cube complex satisfying Gromov's no-$\triangle$-condition and endowed with its combinatorial metric d". Unless specified otherwise, all geodesics will be implicitly assumed to be combinatorial geodesics contained in the 1-skeleton. We refer to bi-infinite geodesics simply as *lines*. In some situations, especially in Sections 2.1.4 and 4.2, we will also need to consider geodesics with respect to the CAT(0) metric. In this case, we will use the terminology "combinatorial geodesic/segment/ray/line" and "CAT(0) geodesic/segment/ray/line".

We denote by X' the *barycentric subdivision* (aka *first cubical subdivision*) of X. This is the CAT(0) cube complex obtained by adding a vertex $v(c)$ at the centre of each cube $c \subseteq X$; vertices $v(c)$ and $v(c')$ are joined by an edge of X' if c is a codimension-one face of c' or vice versa. Each k-cube of X gives rise to 2^k k-cubes of X'. We assign all edges of X' length 1, so the natural inclusion $X \hookrightarrow X'$ is a homothety of a factor 2.

Let $\mathscr{W}(X)$ and $\mathscr{H}(X)$ be, respectively, the sets of hyperplanes and of halfspaces of X. We simply write $\mathscr{W}$ and $\mathscr{H}$ when there is no need to specify the cube complex. Let $\mathfrak{h}^*$ denote the complement of the halfspace $\mathfrak{h}$. Endowing $\mathscr{H}$ with the poset structure given by inclusions $\subseteq$, the involution $*$ is order-reversing. The triple $(\mathscr{H}, \subseteq, *)$ is thus a *po***c***set*, in the sense of [Sag14].

Three hyperplanes of X form a *facing triple* if they are pairwise disjoint and none of the three separates the other two. Three halfspaces form a *facing triple* if they are pairwise disjoint. Two halfspaces $\mathfrak{h}$ and $\mathfrak{k}$ are *nested* if either $\mathfrak{h} \subseteq \mathfrak{k}$ or $\mathfrak{k} \subseteq \mathfrak{h}$.

Two distinct hyperplanes are *transverse* if they cross. Similarly, we say that two halfspaces — or a halfspace and a hyperplane — are transverse if the corresponding hyperplanes are. Halfspaces $\mathfrak{h}$ and $\mathfrak{k}$ are transverse if and only if all four intersections $\mathfrak{h} \cap \mathfrak{k}$, $\mathfrak{h}^* \cap \mathfrak{k}$, $\mathfrak{h} \cap \mathfrak{k}^*$ and $\mathfrak{h}^* \cap \mathfrak{k}^*$ are nonempty. Finally, we say that two subsets $\mathcal{U}, \mathcal{V} \subseteq \mathscr{W}$ are *transverse* if every element of $\mathcal{U}$ is transverse to every element of $\mathcal{V}$.

Every hyperplane $\mathfrak{w}$ — and more generally every intersection $\mathfrak{w}_1 \cap ... \cap \mathfrak{w}_k$ of pairwise-transverse hyperplanes $\mathfrak{w}_1, ..., \mathfrak{w}_k$ — can itself be regarded as a CAT(0) cube complex. Its cells are precisely the intersections $\mathfrak{w}_1 \cap ... \cap \mathfrak{w}_k \cap c$, where $c \subseteq X$ is a cube. Alternatively, $\mathfrak{w}_1 \cap ... \cap \mathfrak{w}_k$ can be viewed as a subcomplex of the barycentric subdivision of X.

Given $\mathfrak{w} \in \mathscr{W}(X)$, the set of hyperplanes of the CAT(0) cube complex $\mathfrak{w}$ is naturally identified with the set of hyperplanes of X that are transverse to $\mathfrak{w}$. We thus denote by $\mathscr{W}(\mathfrak{w})$ this subset of $\mathscr{W}(X)$. We also denote by $C(\mathfrak{w})$ the *carrier* of $\mathfrak{w}$, i.e. its neighbourhood of radius $\frac{1}{2}$ in X; this is a subcomplex of X.

If $e \subseteq X$ is an edge, we write $\mathfrak{w}(e) \in \mathscr{W}(X)$ for the hyperplane crossed by e. A hyperplane $\mathfrak{w}$ is *adjacent* to a vertex $v \in X$ if $\mathfrak{w} = \mathfrak{w}(e)$ for an edge e incident to v. The set of hyperplanes adjacent to v will be denoted by $\mathscr{W}_v$.

The *link* of v (denoted $\mathrm{lk}_X(v)$ or just $\mathrm{lk}(v)$) is the graph having $\mathscr{W}_v$ as vertex set and edges joining two hyperplanes exactly when they are transverse. Equivalently, vertices of $\mathrm{lk}(v)$ are edges of X incident to v, two edges being connected by an edge of $\mathrm{lk}(v)$ if and only if they span a square in X. Cubes based at v are in one-to-one correspondence with complete subgraphs of $\mathrm{lk}(v)$; we refer to the latter as *cliques*. The *degree* of v is $\deg(v) := \#\mathscr{W}_v = \#\mathrm{lk}^0(v)$.

Given vertices $e, f \in \mathrm{lk}(v)$, we write $\delta(e, f)$ for the minimum number of edges of $\mathrm{lk}(v)$ in a path from e to f. If no such path exists, we set $\delta(e, f) = +\infty$. When $\mathrm{lk}(v)$ is connected, δ is a metric on its 0-skeleton. We will sometimes refer to vertices of $\mathrm{lk}(v)$ simply as *elements of* $\mathrm{lk}(v)$, in order to avoid confusion with vertices of X.

We will generally confuse all geodesics (and quasi-geodesics) with their images as subsets of X. Every geodesic $\gamma \subseteq X$ can be viewed as a collection of edges; distinct edges $e, e' \subseteq \gamma$ must yield distinct hyperplanes $\mathfrak{w}(e)$ and $\mathfrak{w}(e')$. If $\gamma \subseteq X$ is an (oriented) geodesic, we write $\gamma(0)$ for its initial vertex and $\gamma(n)$ for its n-th vertex.

Denote by $\mathscr{W}(\gamma)$ the collection of hyperplanes crossed by (the edges of) γ. If two geodesics γ and γ' share an endpoint $v \in X$, their union $\gamma \cup \gamma'$ is again a geodesic if and only if $\mathscr{W}(\gamma) \cap \mathscr{W}(\gamma') = \emptyset$. The following is then a simple observation.

Lemma 2.1.1. *Given $v \in X$ and rays $r_1, r_2 \subseteq X$ based at v, consider the sets $\mathscr{W}_i = \mathscr{W}(r_i) \cap \mathscr{W}_v$. The union $r_1 \cup r_2$ is a line if and only if $\mathscr{W}_1 \cap \mathscr{W}_2 = \emptyset$.*

Proof. If $r_1 \cup r_2$ is not a geodesic, there exists $\mathfrak{w} \in \mathscr{W}(r_1) \cap \mathscr{W}(r_2)$. If $\mathfrak{w}$ is not adjacent to v, let $\mathfrak{u}$ be a hyperplane closest to v among those that separate v from $\mathfrak{w}$; otherwise, let us set $\mathfrak{u} = \mathfrak{w}$. In any case, $\mathfrak{u}$ is adjacent to v and lies in both $\mathscr{W}(r_i)$. Thus, we have $\mathfrak{u} \in \mathscr{W}_1 \cap \mathscr{W}_2$. □

2.1.2 The Roller boundary ∂X.

Let X be a CAT(0) cube complex. Given a vertex $v \in X$, we denote by $\sigma_v \subseteq \mathscr{H}$ the set of all halfspaces containing v. It satisfies the following properties:

(1) given any two halfspaces $\mathfrak{h}, \mathfrak{k} \in \sigma_v$, we have $\mathfrak{h} \cap \mathfrak{k} \neq \emptyset$;

(2) for any hyperplane $\mathfrak{w} \in \mathscr{W}$, a side of $\mathfrak{w}$ lies in σ_v;

(3) every descending chain of halfspaces in σ_v is finite.

Subsets $\sigma \subseteq \mathscr{H}$ satisfying (1) – (3) are known as *DCC ultrafilters* (cf. [Sag14]). We refer to a set $\sigma \subseteq \mathscr{H}$ satisfying only (1) and (2) simply as an *ultrafilter*.

Let $\iota\colon X \to 2^{\mathscr{H}}$ denote the map that takes each vertex v to the set σ_v. Its image $\iota(X)$ is precisely the collection of all *DCC ultrafilters*. Endowing $2^{\mathscr{H}}$ with the product topology, we can consider the closure $\overline{\iota(X)}$, which coincides with the set of all *ultrafilters*. Equipped with the subspace topology, this is a totally disconnected, compact, Hausdorff space known as the *Roller compactification* of X [BCG$^+$09, NS13]; we denote it by $\overline{X}$. The *Roller boundary* ∂X is defined as the difference $\overline{X} \setminus X$.

The inclusion $\iota\colon X \to \overline{X}$ is always continuous[2]. If, moreover, X is locally finite, then ι is a topological embedding, X is open in $\overline{X}$ and ∂X is compact. We prefer to imagine ∂X as a set of points at infinity, rather than a set of ultrafilters. We will therefore write $x \in \partial X$ for points in the Roller boundary and employ the notation $\sigma_x \subseteq \mathscr{H}$ to refer to the ultrafilter representing x.

Given a (combinatorial) ray $r \subseteq X$ and a hyperplane $\mathfrak{w} \in \mathscr{W}$, there exists a unique side $\mathfrak{h} \in \mathscr{H}$ of $\mathfrak{w}$ such that $r \setminus \mathfrak{h}$ is bounded. The collection of all such halfspaces forms an ultrafilter and we denote by $r(+\infty) \in \partial X$ the corresponding point; we refer to $r(+\infty)$ as the *endpoint at infinity* of r. Fixing a base vertex $v \in X$, every point of ∂X is of the form $r(+\infty)$ for a ray r based at v. We obtain a bijection between points of ∂X and rays based at v, where we need to identify the rays r_1 and r_2 if $\mathscr{W}(r_1) = \mathscr{W}(r_2)$. See Proposition A.2 in [Gen16b] for details.

For a vertex $v \in X$ and a halfspace $\mathfrak{h} \in \mathscr{H}$, we have $v \in \mathfrak{h}$ if and only if $\mathfrak{h} \in \sigma_v$. We can extend the halfspace $\mathfrak{h} \subseteq X$ to a subset $\overline{\mathfrak{h}} \subseteq \overline{X}$ by declaring that a point $x \in \partial X$ lies in $\overline{\mathfrak{h}}$ if and only if $\mathfrak{h} \in \sigma_x$. Note that this yields a partition $\overline{X} = \overline{\mathfrak{h}} \sqcup \overline{\mathfrak{h}^*}$; our notation is justified by the observation that $\overline{\mathfrak{h}}$ is simply the closure of $\mathfrak{h}$ in $\overline{X}$. For simplicity, we will often omit the overline symbol outside of the current section, thus not distinguishing between a halfspace $\mathfrak{h} \subseteq X$ and its extension $\overline{\mathfrak{h}} \subseteq \overline{X}$.

Given subsets $A, B \subseteq \overline{X}$, we adopt the notation:

$$\begin{aligned}
\mathscr{H}(A|B) &= \{\mathfrak{h} \in \mathscr{H} \mid B \subseteq \overline{\mathfrak{h}},\ A \subseteq \overline{\mathfrak{h}^*}\},\\
\mathscr{W}(A|B) &= \{\mathfrak{w} \in \mathscr{W} \mid \mathfrak{w} \text{ bounds an element of } \mathscr{H}(A|B)\},\\
\mathscr{W}(A) &= \bigcup_{x,y \in A} \mathscr{W}(x|y).
\end{aligned}$$

We will simplify expressions like $\mathscr{W}(\{x\}|\{y,z\})$ to $\mathscr{W}(x|y,z)$. It is immediate from the definitions that $\mathscr{W}(x|y) \neq \emptyset$ if and only if $x, y \in \overline{X}$ are distinct. If $\mathfrak{w} \in \mathscr{W}(A|B)$, we say that $\mathfrak{w}$ *separates* A and B. To avoid possible ambiguities, we adopt the

[2]Note that ι is only defined on the 0-skeleton, which has the discrete topology.

convention that hyperplanes $\mathfrak{w}$ are not contained in either of their sides; in particular, $\mathfrak{w} \notin \mathscr{W}(\mathfrak{w}|A)$ for every $A \subseteq \overline{X}$.

Lemma 2.1.2. *Given points $x, y, z, w \in \overline{X}$, the sets $\mathscr{W}(x, y|z, w)$, $\mathscr{W}(x, z|y, w)$ and $\mathscr{W}(x, w|y, z)$ are pairwise transverse.*

Proof. Consider $\mathfrak{h} \in \mathscr{H}(x, y|z, w)$ and $\mathfrak{k} \in \mathscr{H}(x, z|y, w)$. Since we have $x \in \mathfrak{h}^* \cap \mathfrak{k}^*$, $y \in \mathfrak{h}^* \cap \mathfrak{k}$, $z \in \mathfrak{h} \cap \mathfrak{k}^*$ and $w \in \mathfrak{h} \cap \mathfrak{k}$, we conclude that $\mathfrak{h}$ and $\mathfrak{k}$ are transverse. Hence $\mathscr{W}(x, y|z, w)$ and $\mathscr{W}(x, z|y, w)$ are transverse; the same argument shows that these sets are also transverse to $\mathscr{W}(x, w|y, z)$. □

We will denote by $\partial_\infty X$ the *visual boundary*[3] of X, in order to distinguish it from the Roller boundary ∂X. It is interesting to remark that $\partial_\infty X$ coincides with the horofunction boundary associated to the CAT(0) metric on X. On the other hand, the identity of X naturally extends to a homeomorphism between ∂X and the horofunction boundary of (X, d)[4]. In case X is Gromov hyperbolic, we will denote by $\partial_\infty X$ also the *Gromov boundary* of (X, d), as this is naturally homeomorphic to the visual boundary.

The characterisation of ∂X in terms of ultrafilters has the advantage of providing a natural structure of *median algebra* on $\overline{X}$, as we now explain. Given any three points $x, y, z \in \overline{X}$, the set

$$(\sigma_x \cap \sigma_y) \cup (\sigma_y \cap \sigma_z) \cup (\sigma_z \cap \sigma_x) = (\sigma_x \cup \sigma_y) \cap (\sigma_y \cup \sigma_z) \cap (\sigma_z \cup \sigma_x)$$

is the collection of all those halfspaces that contain at least two of the three points. It is immediate to observe that this is an ultrafilter. We denote by $m(x, y, z)$ the point of $\overline{X}$ that it represents and refer to it as the *median* of x, y, z. The map $m\colon \overline{X}^3 \to \overline{X}$ is continuous and satisfies certain symmetries that make $(\overline{X}, m)$ into a median algebra[5].

Given $x, y \in \overline{X}$, the *interval* between x and y is the set:

$$I(x, y) = \{z \in \overline{X} \mid \mathscr{W}(z|x, y) = \emptyset\} = \{z \in \overline{X} \mid m(x, y, z) = z\}.$$

Note that we always have $I(x, x) = \{x\}$. Given $x, y, z \in \overline{X}$, the median $m(x, y, z) \in \overline{X}$ is the unique point that lies in all three intervals $I(x, y)$, $I(y, z)$ and $I(z, x)$.

Given vertices $v, w \in X$, we have $I(v, w) \subseteq X$ and this is exactly the union of all geodesics with endpoints v and w. The median map $m\colon \overline{X}^3 \to \overline{X}$ maps X^3 into X, thus also endowing X with a structure of median algebra. The latter is well-behaved

[3] See Definition II.8.6 in [BH99].

[4] This is an unpublished result of U. Bader and D. Guralnik; see [CL11, FLM18] for a proof.

[5] The exact definition is irrelevant for our purposes and can be found in [Rol98, CDH10, Fio17].

with respect to the combinatorial metric d: Given any three vertices $v_1, v_2, v_3 \in X$, their median $m = m(v_1, v_2, v_3)$ is the only vertex of X that satisfies

$$d(v_i, v_j) = d(v_i, m) + d(m, v_j)$$

for all $1 \leq i < j \leq 3$. In other words, the metric space (X, d) is a *median space*[6].

This is a good moment to introduce two notions that will be crucial to the rest of thesis, especially to Chapters 6 and 7. For the time being, we will not discuss them any further.

Definition 2.1.3. A subset $\mathcal{A} \subseteq \partial X$ is *concise* if $I(x, y) \cap X \neq \emptyset$ for all distinct points $x, y \in \mathcal{A}$.

Definition 2.1.4. A subset $\mathcal{A} \subseteq \partial X$ is *ample* if, for every $k \geq 1$ and every choice of pairwise transverse halfspaces $\mathfrak{h}_1, ..., \mathfrak{h}_k$, the intersection $\mathcal{A} \cap \overline{\mathfrak{h}_1} \cap ... \cap \overline{\mathfrak{h}_k}$ is nonempty.

Given two vertices $v, w \in X$, we have $d(v, w) = \#\mathscr{W}(v|w)$. By analogy, we can define $d(x, y) = \#\mathscr{W}(x|y)$ for all points $x, y \in \overline{X}$. The resulting function $d\colon \overline{X} \times \overline{X} \to \mathbb{N} \cup \{+\infty\}$ satisfies all axioms of a metric, except that it can indeed take the value $+\infty$. We write $x \sim y$ if x and y satisfy $d(x, y) < +\infty$. This is an equivalence relation on $\overline{X}$; we refer to its equivalence classes as *components*.

Given $x \in \overline{X}$, we denote by $Z(x)$ the only component of $\overline{X}$ that contains the point x. When $x \in X$, we have $Z(x) = X$. For every component $Z \subseteq \overline{X}$, the pair (Z, d) is a metric space. Joining points of Z by an edge whenever they are at distance 1 and adding k-cubes whenever we see their 1-skeleta, we can give (Z, d) a structure of CAT(0) cube complex (see [Gur05, Fio17]).

We conclude by obtaining a few simple results which will be needed later on.

Lemma 2.1.5. *Suppose that* $\dim X < +\infty$. *Let* $\mathfrak{h}$ *and* $\mathfrak{k}$ *be disjoint halfspaces with* $\mathfrak{h} \neq \mathfrak{k}^*$. *Suppose that there exist points* $x \in \mathfrak{h} \cap \partial X$ *and* $y \in \mathfrak{k} \cap \partial X$ *with* $x \sim y$. *Then, there exists an infinite chain* $\mathfrak{j}_0 \supsetneq \mathfrak{j}_1 \supsetneq ...$ *of halfspaces transverse to both* $\mathfrak{h}$ *and* $\mathfrak{k}$.

Proof. Pick a point $p \in X$ with $p \in \mathfrak{h}^* \cap \mathfrak{k}^*$. The set $\mathscr{W}(p|x, y)$ is infinite, as $\mathscr{W}(p|x)$ is infinite and $\mathscr{W}(x|y)$ is finite. Since σ_p is a DCC ultrafilter, the sets $\mathscr{W}(p|\mathfrak{h})$ and $\mathscr{W}(p|\mathfrak{k})$ are finite and the set $A = \mathscr{W}(p|x, y) \setminus (\mathscr{W}(p|\mathfrak{h}) \cup \mathscr{W}(p|\mathfrak{k}))$ is infinite. Any hyperplane in A is transverse to both $\mathfrak{h}$ and $\mathfrak{k}$. Since any two elements of $\mathscr{H}(p|x, y)$ are either transverse or nested, the required chain is obtain by applying Ramsey's theorem to A. □

[6]We have only defined the median map on the 0-skeleton of X here. However, it should be noted that the ℓ^1 metric d endows the entire cube complex with a structure of median space.

Lemma 2.1.6. *Let $r \subseteq X$ be a ray and set $x = r(+\infty)$. Given a point $x' \in \partial X$ with $x' \sim x$, there exists a ray $r' \subseteq X$ satisfying $r'(0) = r(0)$, $r'(+\infty) = x'$ and such that the Hausdorff distance $d_{\mathrm{Haus}}(r, r')$ is at most $d(x, x')$.*

Proof. It suffices to consider the case when $d(x, x') = 1$. Let $\mathfrak{w}$ be the only hyperplane separating x and x'; let $\mathfrak{h}$ denote the side of $\mathfrak{w}$ containing x. Note that $\mathfrak{w}$ must be transverse to all but finitely many hyperplanes in $\mathscr{W}(r)$. Thus, r intersects the carrier $C(\mathfrak{w})$ in a sub-ray $\gamma \subseteq r$.

If $r(0) \in \mathfrak{h}$, the ray γ does not cross the hyperplane $\mathfrak{w}$. Let γ' be the ray such that $\mathscr{W}(\gamma(n), \gamma'(n)) = \{\mathfrak{w}\}$ for all $n \geq 0$. In this case, the ray r' is obtained by following r up to $\gamma(0)$, crossing $\mathfrak{w}$, and finally following γ' all the way to x'.

If instead $r(0) \in \mathfrak{h}^*$, there exists $k \geq 0$ such that $\mathscr{W}(\gamma(k), \gamma(k+1)) = \{\mathfrak{w}\}$. Let γ'' be the ray with $\mathscr{W}(\gamma''(n)|\gamma(n+k+1)) = \{\mathfrak{w}\}$ for all $n \geq 0$; in particular, $\gamma''(0) = \gamma(k)$ and $\gamma''(+\infty) = x'$. We construct r' by following r up to $\gamma(k)$ and then, rather than crossing $\mathfrak{w}$, following γ'' until x'. □

2.1.3 Convexity and bridges.

Let X be a CAT(0) cube complex. All results and definitions mentioned here can be found e.g. in [Rol98, Fio17], unless a different reference is provided.

A subset $C \subseteq \overline{X}$ is *convex* if $I(x, y) \subseteq C$ for all $x, y \in C$. Thus, a subset $C \subseteq X$ is convex if and only if all geodesics between points of C are contained in C. All components of $\overline{X}$ are convex. Given $u, v \in \overline{X}$, the interval $I(u, v) \subseteq \overline{X}$ is convex and so is every halfspace $\mathfrak{h} \subseteq X$. In fact, halfspaces are precisely those convex subsets of X that have convex complement.

Intersections of convex sets are convex. Given a subset $A \subseteq \overline{X}$, we denote by $\mathrm{Hull}(A)$ the smallest convex subset of $\overline{X}$ containing A. We remark that $\mathrm{Hull}(A) \cap X$ coincides with the intersection of all halfspaces $\mathfrak{h} \subseteq X$ with the property that A is contained in the extension $\overline{\mathfrak{h}} \subseteq \overline{X}$. If $A \subseteq X$, we have $\mathrm{Hull}(A) \subseteq X$, as X is a convex subset of $\overline{X}$. In this case, $\mathrm{Hull}(A)$ is itself a CAT(0) cube complex and a subcomplex of X; its hyperplane set is identified with the subset $\mathscr{W}(A) \subseteq \mathscr{W}(X)$.

If $C \subseteq \overline{X}$ is convex and closed, there exists a unique *gate-projection* $\pi_C \colon \overline{X} \to C$, namely a map with the property that $\pi_C(x) \in I(x, y)$ for every $x \in X$ and every $y \in C$. Gate-projections are 1-Lipschitz with respect to the extension of d to $\overline{X}$. When $u, v \in \overline{X}$ and $C = I(u, v)$, the gate-projection is given by $x \mapsto m(u, v, x)$.

If $C \subseteq X$ is convex, its closure $\overline{C} \subseteq \overline{X}$ is closed and convex. As such, it admits a gate-projection $\pi_{\overline{C}} \colon \overline{X} \to \overline{C}$; we will simply write $\pi_C \colon X \to C$ for the restriction

of $\pi_{\overline{C}}$ to X. Note that this is just the nearest-point projection with respect to the combinatorial metric d[7].

Given a closed convex subset $C \subseteq \overline{X}$ and two points $x, y \in \overline{X}$, we always have $\mathscr{W}(x|\pi_C(x)) = \mathscr{W}(x|C)$ and $\mathscr{W}(x|y) \cap \mathscr{W}(C) = \mathscr{W}(\pi_C(x)|\pi_C(y))$; see Proposition 2.3 in [Fio17]. If $C_1, ..., C_k \subseteq \overline{X}$ are convex and pairwise intersect, then $C_1 \cap ... \cap C_k \neq \emptyset$; this is known as Helly's lemma.

We now consider two convex subsets $C_1, C_2 \subseteq X$ and simply write $\pi_i \colon X \to C_i$ for the two gate-projections. We say that points $x_1 \in C_1$ and $x_2 \in C_2$ form a *pair of gates* if $\pi_1(x_2) = x_1$ and $\pi_2(x_1) = x_2$. Equivalently, we have $\mathscr{W}(x_1|x_2) = \mathscr{W}(C_1|C_2)$, or $d(x_1, x_2) = d(C_1, C_2)$. We denote by $S_1 \subseteq C_1$ and $S_2 \subseteq C_2$ the subsets of points that are part of a pair of gates. The *bridge* $B = B(C_1, C_2)$ is the set $\mathrm{Hull}(S_1 \cup S_2)$; we also refer to S_i as the *shore* of C_i.

The following appears e.g. in Section 2.G of [CFI16] or Section 2.2 of [Fio19].

Lemma 2.1.7. *(1) Bridges and shores are convex subcomplexes of X. We have $S_i = B \cap C_i$ and $S_1 = \pi_1(S_2) = \pi_1(C_2)$, $S_2 = \pi_2(S_1) = \pi_2(C_1)$. The restrictions $\pi_1|_{S_2} \colon S_2 \to S_1$ and $\pi_2|_{S_1} \colon S_1 \to S_2$ are cubical isomorphisms and they are inverse to each other.*

(2) The bridge B is a disjoint union of the intervals $I(x_1, \pi_2(x_1))$ as x_1 varies in S_1. Setting $I = I(x_1, \pi_2(x_1))$ for any choice of $x_1 \in S_1$, the two gate-projections π_{S_1} and π_I give rise to a cubical isomorphism $B \to S_1 \times I$.

(3) We have $\mathscr{W}(S_1) = \mathscr{W}(S_2) = \mathscr{W}(C_1) \cap \mathscr{W}(C_2)$ and $\mathscr{W}(B) = \mathscr{W}(S_1) \sqcup \mathscr{W}(C_1|C_2)$.

(4) For all $y_1 \in C_1$ and $y_2 \in C_2$, we have:

$$\begin{aligned} d(y_1, y_2) &= d(y_1, S_1) + d(\pi_B(y_1), \pi_B(y_2)) + d(y_2, S_2) \\ &= d(y_1, S_1) + d(\pi_{S_1}(y_1), \pi_{S_1}(y_2)) + d(C_1, C_2) + d(y_2, S_2). \end{aligned}$$

Of course, the role of S_1 can equally be played by S_2 in Lemma 2.1.7. We denote by $S = S(C_1, C_2)$ the isomorphism type of S_1 and S_2 and refer to it as the *abstract shore*. The gate-projections π_{S_i} yield the same map $\pi_S \colon B \to S$.

We say that two disjoint halfspaces $\mathfrak{h}, \mathfrak{k} \in \mathscr{H}(X)$ are *strongly separated* if the abstract shore $S(\mathfrak{h}, \mathfrak{k})$ consists of a single point. Equivalently, no hyperplane is transverse to both $\mathfrak{h}$ and $\mathfrak{k}$. We also say that the hyperplanes bounding $\mathfrak{h}$ and $\mathfrak{k}$ are *strongly separated*. Note that, in this case, there exists a unique vertex $v \in \mathfrak{h} \cap B(\mathfrak{h}, \mathfrak{k})$; for all $x \in \overline{\mathfrak{h}}$ and $y \in \overline{\mathfrak{k}}$, we have $m(x, v, y) = v$ and $v \in I(x, y)$.

[7]The metric d is not CAT(0), but there still exists a *unique* point of C that is closest to x.

2.1.4 Combinatorial geodesics vs CAT(0) geodesics.

The following result is probably well-known to experts, but a proof does not seem to appear in the literature. We provide it here for completeness.

Throughout this subsection, we will always specify whether geodesics are meant with respect to the CAT(0) metric on X, or rather with respect to the combinatorial metric d. We stress that Hausdorff distances, however, will always be calculated with respect to the combinatorial metric.

Proposition 2.1.8. *Let X be a D-dimensional* CAT(0) *cube complex. Every* CAT(0) *ray based at a vertex of X is at Hausdorff distance at most D from a combinatorial ray with the same origin.*

Given a combinatorial geodesic γ, the hyperplanes of $\mathscr{W}(\gamma)$ can be arranged in a (finite or infinite) sequence $(\mathfrak{w}_n)_{n\geq 0}$ according to the order in which they are crossed by γ. We denote this sequence by $\mathfrak{s}(\gamma)$.

Lemma 2.1.9. *Let $(\mathfrak{w}_n)$ be a sequence of pairwise distinct hyperplanes of X and let $p \in X$ be a vertex. There exists a combinatorial geodesic γ based at p with $\mathfrak{s}(\gamma) = (\mathfrak{w}_n)$ if and only if, for every $n \geq 0$, we have $\mathscr{W}(p|\mathfrak{w}_n) = \{\mathfrak{w}_0, ..., \mathfrak{w}_{n-1}\} \setminus \mathscr{W}(\mathfrak{w}_n)$.*

Proof. If there exists a combinatorial geodesic γ with $\gamma(0) = p$ and $\mathfrak{s}(\gamma) = (\mathfrak{w}_n)$, each point $\gamma(n)$ lies in the carrier $C(\mathfrak{w}_n)$. Thus $\mathscr{W}(p|\mathfrak{w}_n) \subseteq \mathscr{W}(p|\gamma(n))$ and every element of $\mathscr{W}(p|\gamma(n))$ either crosses $\mathfrak{w}_n$ or lies in $\mathscr{W}(p|\mathfrak{w}_n)$. We conclude that $\mathscr{W}(p|\mathfrak{w}_n) = \{\mathfrak{w}_0, ..., \mathfrak{w}_{n-1}\} \setminus \mathscr{W}(\mathfrak{w}_n)$ for all $n \geq 0$.

Assuming instead that the sequence $(\mathfrak{w}_n)$ satisfies the latter condition, we are going to construct points $p_n \in C(\mathfrak{w}_n)$ with $\mathscr{W}(p|p_n) = \{\mathfrak{w}_0, ..., \mathfrak{w}_{n-1}\}$. We then obtain γ by setting $\gamma(n) = p_n$. We proceed by induction on $n \geq 0$, observing that the case $n = 0$ immediately follows from $\mathscr{W}(p|\mathfrak{w}_0) = \emptyset$.

Given $p_n \in C(\mathfrak{w}_n)$ with $\mathscr{W}(p|p_n) = \{\mathfrak{w}_0, ..., \mathfrak{w}_{n-1}\}$, let $p_{n+1} \in X$ be the only vertex with $\mathscr{W}(p_n|p_{n+1}) = \{\mathfrak{w}_n\}$. As $\mathscr{W}(p|p_{n+1}) = \{\mathfrak{w}_0, ..., \mathfrak{w}_n\}$, we only need to show that p_{n+1} lies in the carrier $C(\mathfrak{w}_{n+1})$. If this failed, there would exist a hyperplane $\mathfrak{u} \in \mathscr{W}(p_{n+1}|\mathfrak{w}_{n+1})$ and we would have either $\mathfrak{u} \in \mathscr{W}(p, p_{n+1}|\mathfrak{w}_{n+1})$ or $\mathfrak{u} \in \mathscr{W}(p_{n+1}|p, \mathfrak{w}_{n+1})$. The former is forbidden by $\mathscr{W}(p|\mathfrak{w}_{n+1}) \subseteq \{\mathfrak{w}_0, ..., \mathfrak{w}_n\}$, whereas the latter would clash with the fact that, for $0 \leq i \leq n$, each $\mathfrak{w}_i$ either crosses $\mathfrak{w}_{n+1}$ or lies in $\mathscr{W}(p|\mathfrak{w}_{n+1})$. □

Proof of Proposition 2.1.8. Let ρ be a CAT(0) ray based at a vertex $p \in X$. For every hyperplane $\mathfrak{w} \in \mathscr{W}(X)$ there exists exactly one side of $\mathfrak{w}$ that has unbounded

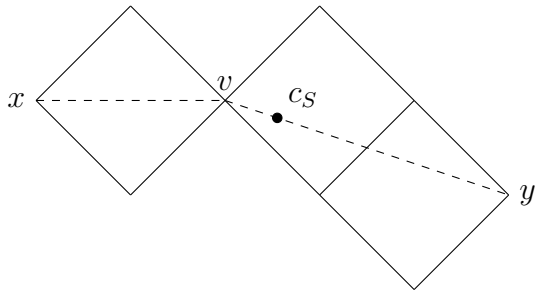


Figure 2.1: The pictured cube complex S consists of three squares. The CAT(0) geodesic from x to y is the unique longest geodesic in S. The CAT(0) distance between v and the barycentre c_S is $\frac{1}{2}(\sqrt{5}-\sqrt{2})$. This number is not of the form $\frac{1}{2^n}\sqrt{a^2+b^2}$ for any $a, b, n \in \mathbb{N}$, so c_S is not a vertex of any iterated cubical subdivision of S.

intersection with ρ. The collection of these halfspaces forms an ultrafilter $\sigma \subseteq \mathscr{H}(X)$ representing a point $x \in \partial X$.

A hyperplane is crossed by ρ if and only if it lies in the set $\mathscr{W}(p|x)$; thus, ρ is entirely contained in the subcomplex $I(p,x) \cap X$. Let $(\mathfrak{w}_n)_{n\geq 0}$ be an ordering of the elements of $\mathscr{W}(p|x)$, so that $m < n$ if $\mathfrak{w}_m$ is crossed by ρ before $\mathfrak{w}_n$. By Lemma 2.1.9, there exists a combinatorial ray r from p to x such that $\mathfrak{s}(\gamma) = (\mathfrak{w}_n)$. It remains to prove that $d_{\mathrm{Haus}}(r, \rho) \leq D$.

Let $u \in r$ and $v \in \rho$ be points (not necessarily vertices) that are not separated by any hyperplane of X. Note that for every point $u \in r$ there exists such a point $v \in \rho$ and vice versa. Let $\iota \colon I(p,x) \cap X \to \mathbb{R}^D$ be an ℓ^1-isometric cubical embedding; it exists for instance by Theorem 1.14 in [BCG$^+$09]. Under the map ι, preimages of convex sets are convex and, therefore, preimages of halfspaces are halfspaces. It follows that the points $\iota(u)$ and $\iota(v)$ are not separated by any hyperplane of $\mathbb{R}^D$, hence they lie in a translate of a unit cube of $\mathbb{R}^D$. Thus $d(u,v) = d(\iota(u), \iota(v)) \leq D$. □

2.1.5 Median barycentres.

Let S be a bounded CAT(0) cube complex. Considering the CAT(0) metric on S, there exists a unique barycentre $c_S \in S$. This is the centre of the unique smallest CAT(0) ball containing S; see e.g. Proposition II.2.7 in [BH99] or Proposition 3.73 in [DK17]. However, the point c_S is in general not a vertex of S, nor a vertex of any iterated cubical subdivision. This is illustrated in Figure 2.1.

In Section 5.3 we will need a different notion of barycentre, which we now introduce. It will always be a vertex of the first cubical subdivision S'.

In the discussion below, points of S are assumed to be vertices and we only consider the combinatorial metric, as in the rest of the thesis. Let $\mathfrak{h} \in \mathscr{H}(S)$ be a side of the

hyperplane $\mathfrak{w}$ and let $x \in \mathfrak{h}$ and $y \in \mathfrak{h}^*$ be vertices maximising the distance from $\mathfrak{w}$. We say that $\mathfrak{w}$ is *balanced* if $d(x, \mathfrak{w}) = d(y, \mathfrak{w})$ and *unbalanced* otherwise. If $d(x, \mathfrak{w}) > d(y, \mathfrak{w})$, we call $\mathfrak{h}$ *heavy* and $\mathfrak{h}^*$ *light*.

Lemma 2.1.10. *Given halfspaces $\mathfrak{h}, \mathfrak{k} \in \mathscr{H}(S)$ with $\mathfrak{h} \cap \mathfrak{k} = \emptyset$ and $\mathfrak{h} \neq \mathfrak{k}^*$, at least one of them is light.*

Proof. Let $\mathfrak{w}$ and $\mathfrak{u}$ be the hyperplanes associated to $\mathfrak{h}$ and $\mathfrak{k}$, respectively. Pick $x \in \mathfrak{h}$ and $y \in \mathfrak{k}$ maximising the distance from $\mathfrak{w}$ and $\mathfrak{u}$, respectively. Note that $\mathfrak{k} \subseteq \mathfrak{h}^*$ and $\mathfrak{h} \subseteq \mathfrak{k}^*$. If $\mathfrak{h}$ and $\mathfrak{k}$ were both not light, we would have

$$d(x, \mathfrak{w}) \geq d(y, \mathfrak{w}) \geq d(y, \mathfrak{u}) + d(\mathfrak{u}, \mathfrak{w}) > d(y, \mathfrak{u})$$

and similarly $d(y, \mathfrak{u}) > d(x, \mathfrak{w})$. This is a contradiction. □

In particular, any two heavy halfspaces intersect and any two balanced hyperplanes are transverse. By Helly's Lemma, the intersection of all heavy halfspaces is nonempty. It is a cube $c \subseteq S$ cut by all balanced hyperplanes.

The centre of c is a vertex m_S of the cubical subdivision S'. We refer to it as the *median barycentre* of S. Note that m_S is a vertex of S if and only if every hyperplane is unbalanced. For instance, $m_S = v$ in Figure 2.1.

We remark that, given bounded CAT(0) cube complexes S_1 and S_2 and a cubical isomorphism $F\colon S_1 \to S_2$, we have $F(m_{S_1}) = m_{S_2}$.

2.2 Automorphisms of cube complexes.

Throughout this section, let X be a CAT(0) cube complex endowed with its combinatorial metric d. A *(cubical) automorphism* is an automorphism of the cellular structure on X. Equivalently, it is an isometry $X \to X$ that takes vertices to vertices. We denote by $\mathrm{Aut}(X)$ the group of all cubical automorphisms.

2.2.1 Axes and minimal sets.

Let us consider an element $g \in \mathrm{Aut}(X)$. An *axis* for g is a $\langle g \rangle$-invariant (combinatorial) geodesic $\gamma \subseteq X$, none of whose points is fixed by g. The *minimal set* $\mathrm{Min}_X(g) \subseteq X$ is defined as either the union of all axes of g, or the collection of vertices that are fixed by g; we also write simply $\mathrm{Min}(g)$ when this causes no confusion.

We say that g acts *stably without inversions* (Definition 4.1 in [Hag07]) if no power of g preserves a hyperplane and swaps its two sides. Note that every $g \in \mathrm{Aut}(X)$

acts stably without inversions on the barycentric subdivision X' (Lemma 4.2 in *op. cit.*). We say that g acts *non-transversely* (cf. Definition 3.15 in [FFT18]) if, for every $\mathfrak{w} \in \mathscr{W}(\mathrm{Min}(g))$, the hyperplanes $g\mathfrak{w}$ and $\mathfrak{w}$ are not transverse.

The notation ℓ_X was defined in the introduction.

Proposition 2.2.1. *Let $g \in \mathrm{Aut}(X)$ act stably without inversions.*

(1) We have $\mathrm{Min}(g) \neq \emptyset$. Thus, either g fixes a vertex and $\ell_X(g) = 0$, or g admits an axis and $\ell_X(g) > 0$.

(2) We have $\mathrm{Min}(g) \subseteq \mathrm{Min}(g^n)$ and $\ell_X(g^n) = |n| \cdot \ell_X(g)$ for all $n \in \mathbb{Z}$.

(3) A vertex $v \in X$ satisfies $d(v, gv) = \ell_X(g)$ if and only if $v \in \mathrm{Min}(g)$.

(4) If g acts non-transversely, $\mathrm{Min}(g)$ is a convex subcomplex of X.

(5) If $\dim X = D < +\infty$, the automorphism $g^{D!}$ acts non-transversely.

Proof. Parts (1)–(3) follow from Corollary 6.2 in [Hag07]. By Lemma 3.10 and Proposition 3.17 in [FFT18], it suffices to prove part (4) in the case when g fixes a vertex of X (indeed, in the notation of *op. cit.*, this also implies that X_g^{fix} is a convex subcomplex of X_g^{ell}). Let us then assume that g fixes vertices $v, w \in X$. Given $\mathfrak{w} \in \mathscr{W}(v|w)$, we have $g\mathfrak{w} \in \mathscr{W}(gv|gw) = \mathscr{W}(v|w)$ and $d(v, \mathfrak{w}) = d(v, g\mathfrak{w})$. It follows that $\mathfrak{w}$ and $g\mathfrak{w}$ are either transverse or equal. Since g acts non-transversely, it must fix every element of $\mathscr{W}(v|w)$. It follows that $I(v, w) \subseteq \mathrm{Min}(g)$.

Finally, we prove part (5). Given $\mathfrak{w} \in \mathscr{W}(X)$, the hyperplanes $g^i\mathfrak{w}$ cannot be pairwise transverse for $0 \leq i \leq D$; pick i so that $\mathfrak{w}$ and $g^i\mathfrak{w}$ are not transverse. If $\mathfrak{w}$ crosses an axis of g, we can then choose a side $\mathfrak{h}$ of $\mathfrak{w}$ so that $g^i\mathfrak{h} \subsetneq \mathfrak{h}$. Since i divides $D!$, we have $g^{D!}\mathfrak{h} \subsetneq \mathfrak{h}$ and the hyperplanes $g^{D!}\mathfrak{w}$ and $\mathfrak{w}$ are disjoint in this case. The remaining elements of $\mathscr{W}(\mathrm{Min}(g))$ must separate two axes of g, or two fixed vertices. It suffices to consider the latter case, as, in the notation of [FFT18], we can always pass to X_g^{ell}. Now, if $v, w \in X$ are fixed by g and $n \in \mathbb{N}$, there are at most D hyperplanes of $\mathscr{W}(v|w)$ at distance exactly $n + \frac{1}{2}$ from v; indeed, they are pairwise transverse. Since, moreover, these hyperplanes are permuted by g, we conclude that $g^{D!}$ fixes every element of $\mathscr{W}(v|w)$. Hence $g^{D!}$ acts non-transversely. □

Given vertices $v, w \in X$ and an integer $k \geq 0$, we denote by $\mathscr{H}_k(v|w) \subseteq \mathscr{H}(v|w)$ the subset of halfspaces at distance $k + \frac{1}{2}$ from v. Note that the elements of $\mathscr{H}_k(v|w)$ are pairwise transverse, so $\#\mathscr{H}_k(v|w) \in \{0, 1, ..., D\}$, where $D = \dim X$. Moreover,

$$d(v, w) = \sum_{k \geq 0} \#\mathscr{H}_k(v|w).$$

Lemma 2.2.2. *Let* $\dim X = D < +\infty$ *and consider an automorphism* $g \in \mathrm{Aut}(X)$ *that fixes a unique vertex* $v \in X$*. For all* $w \in X$*, we have* $d(w, gw) \geq \frac{2}{D} \cdot d(v, w)$*.*

Proof. Given a vertex $w \in X$ and an integer $k \geq 0$, suppose for the sake of contradiction that $g \cdot \mathscr{H}_k(v|w) = \mathscr{H}_k(v|w)$ and that these sets are nonempty. The elements of $\mathscr{H}_k(v|w)$ are bounded by pairwise-transverse hyperplanes, which determine a cube $c \subseteq X$ preserved by g. It follows that g fixes the (unique) vertex of c that is farthest from v, contradicting uniqueness of v.

For all $k \geq 0$, the set $g \cdot \mathscr{H}_k(v|w) = \mathscr{H}_k(v|gw)$ is then either empty or different from $\mathscr{H}_k(v|w)$. Setting $n_k = \#\mathscr{H}_k(v|w) \leq D$ and $m = m(v, w, gw)$, we have:

$$\#\mathscr{H}_k(v|m) = \#\left(\mathscr{H}_k(v|w) \cap \mathscr{H}_k(v|gw)\right) \leq \tfrac{n_k - 1}{n_k} \cdot \#\mathscr{H}_k(v|w) \leq \tfrac{D-1}{D} \cdot \#\mathscr{H}_k(v|w).$$

The lemma now follows from two inequalities:

$$d(v|m) = \sum_{k \geq 0} \#\mathscr{H}_k(v|m) \leq \tfrac{D-1}{D} \cdot \sum_{k \geq 0} \#\mathscr{H}_k(v|w) = \tfrac{D-1}{D} \cdot d(v, w),$$

$$d(w, gw) = 2 \cdot d(w, m) = 2 \cdot (d(v, w) - d(v, m)) \geq 2 \cdot \left(1 - \tfrac{D-1}{D}\right) \cdot d(v, w). \qquad \square$$

Corollary 2.2.3. *Let* $\dim X = D < +\infty$ *and consider* $g \in \mathrm{Aut}(X)$ *acting stably without inversions and non-transversely. We have* $d(v, gv) \geq \frac{2}{D} \cdot d(v, \mathrm{Min}(g)) + \ell_X(g)$ *for every* $v \in X$*.*

Proof. By Proposition 2.2.1, the minimal set $\mathrm{Min}_X(g)$ is a convex subcomplex of X. Passing to the *restriction quotient* (in the sense of p. 860 of [CS11]) that kills all hyperplanes of $\mathrm{Min}_X(g)$, the corollary follows from Lemma 2.2.2. $\square$

2.2.2 Actions on cube complexes.

Throughout this subsection, let $\Gamma \curvearrowright X$ be an action by automorphisms of a group Γ. We do not require Γ to be countable, nor — in case Γ has a topology — that actions have continuous orbits.

We start by recording the following observation (see e.g. Exercise 1.6 in [Sag14]).

Lemma 2.2.4. *Let* $\mathfrak{w}_1, ..., \mathfrak{w}_k$ *be pairwise transverse hyperplanes and let us denote by* $\Gamma_i < \Gamma$ *the stabiliser of* $\mathfrak{w}_i$*. If the action* $\Gamma \curvearrowright X$ *is cocompact, then so is the action* $\Gamma_1 \cap ... \cap \Gamma_k \curvearrowright \mathfrak{w}_1 \cap ... \cap \mathfrak{w}_k$*.*

The cube complex X is *irreducible* if it cannot be split as a product of lower-dimensional cube complexes. Every finite dimensional cube complex admits a canonical decomposition as product of irreducible cube complexes (Proposition 2.6 in [CS11]); we refer to it as the *De Rham decomposition*.

The action $\Gamma \curvearrowright X$ is *non-elementary* if there are no finite orbits in $X \cup \partial_\infty X$; we recall that $\partial_\infty X$ denotes the *visual* boundary. We say that an element $g \in \Gamma$ *flips* a halfspace $\mathfrak{h} \in \mathscr{H}(X)$ if we have $\mathfrak{h}^* \cap g\mathfrak{h}^* = \emptyset$ and $g\mathfrak{h} \neq \mathfrak{h}^*$. The halfspace $\mathfrak{h}$ is *(*Γ*-)flippable* if such a group element exists and *(*Γ*-)unflippable* otherwise.

The cube complex X is *essential* if no halfspace is contained in a metric neighbourhood of the corresponding hyperplane. The action $\Gamma \curvearrowright X$ is *essential* if no Γ-orbit is contained in a metric neighbourhood of a halfspace. Thus, a cocompact action $\Gamma \curvearrowright X$ is essential if and only if X is essential.

Lemma 2.2.5. *Let X be irreducible and essential; let Γ be a non-virtually-cyclic group. Every properly discontinuous, cocompact action $\Gamma \curvearrowright X$ is non-elementary.*

Proof. Since Γ acts properly discontinuously, it is a countable group. If Γ had a finite orbit within X, it would have to be a finite group. Suppose for the sake of contradiction that Γ has a finite orbit in $\partial_\infty X$. Passing to a finite index subgroup, we assume that Γ fixes a point $\xi \in \partial_\infty X$; let r be a CAT(0)-ray representing ξ.

Let $\mathfrak{w}$ be a hyperplane crossed by r at a point p. Observe that $d(r(t), \mathfrak{w})$ diverges for $t \to +\infty$. Otherwise r would be asymptotic to a different CAT(0)-ray, based at p and entirely contained in $\mathfrak{w}$, which is forbidden by the CAT(0) condition. We conclude that the side of $\mathfrak{w}$ that contains $r(0)$ is unflippable. On the other hand, Theorem 4.7 in [CS11] shows that every element of $\mathscr{H}(X)$ is Γ-flippable, a contradiction. □

Proposition 2.2.6. *Let X be finite dimensional and let $\Gamma \curvearrowright X$ be essential and non-elementary.*

1. *Every halfspace is Γ-flippable.*

2. *If $\mathfrak{h} \subseteq \mathfrak{k}$ are halfspaces, there exists $g \in \Gamma$ with $g\mathfrak{k} \subsetneq \mathfrak{h} \subseteq \mathfrak{k}$.*

Proof. These are the Flipping and Double Skewering Lemmas from [CS11]. □

Remark 2.2.7. Essential, Gromov hyperbolic CAT(0) cube complexes are irreducible. Indeed, essentiality guarantees that X has no bounded factors in its De Rham decomposition, while hyperbolicity implies that there is at most one unbounded factor.

We say that X is *hyperplane-essential* if all its hyperplanes are essential cube complexes. An action $\Gamma \curvearrowright X$ is *hyperplane-essential* if each hyperplane-stabiliser acts essentially on the corresponding hyperplane. By Lemma 2.2.4, a cocompact action $\Gamma \curvearrowright X$ is hyperplane-essential if and only if X is hyperplane-essential.

We say that X is *cocompact* if the action $\mathrm{Aut}(X) \curvearrowright X$ is cocompact. The following is basically Lemma 5.5 in [NS13], although the proof in *loc. cit.* has the unstated assumption that group actions be properly discontinuous. A complete proof at this level of generality will appear in [HW].

Proposition 2.2.8. *Let X be cocompact, locally finite, essential, hyperplane-essential and irreducible. For any two transverse halfspaces $\mathfrak{h}_1$ and $\mathfrak{h}_2$, there exists a halfspace $\mathfrak{k} \subseteq \mathfrak{h}_1 \cap \mathfrak{h}_2$.*

A *free face* in X is a non-maximal cube $c \subseteq X$ that is contained in a unique maximal cube. When X has no free faces, Proposition 2.2.10 below provides a stronger version of the previous result.

Definition 2.2.9. We say that X is *sector-heavy* if, for every $k \geq 2$ and every choice of pairwise transverse halfspaces $\mathfrak{h}_1, ..., \mathfrak{h}_k$, there exists a halfspace $\mathfrak{k} \subseteq \mathfrak{h}_1 \cap ... \cap \mathfrak{h}_k$.

Proposition 2.2.10. *If X is cocompact, locally finite, irreducible and has no free faces, then X is sector-heavy.*

Proof. Since X is finite dimensional and without free faces, X is essential. Observe moreover that any intersection of pairwise-transverse hyperplanes of X is itself a cube complex with no free faces. We conclude that X is also hyperplane-essential and the case $k = 2$ follows from Proposition 2.2.8. For the general case, we proceed by induction on $k \geq 3$.

Let $\mathfrak{h}_1, ..., \mathfrak{h}_k$ be pairwise transverse halfspaces. By the inductive hypothesis, there exists a halfspace $\mathfrak{j} \subseteq \mathfrak{h}_1 \cap ... \cap \mathfrak{h}_{k-1}$; let $\mathfrak{w}$ be the hyperplane corresponding to $\mathfrak{j}$. If $\mathfrak{w} \subseteq \mathfrak{h}_k$, either we have $\mathfrak{j} \subseteq \mathfrak{h}_k$ and we are done, or $\mathfrak{h}_k^* \subseteq \mathfrak{j}$; in the latter case, $\mathfrak{j}$ intersects $\mathfrak{h}_1^*$, a contradiction. If $\mathfrak{w}$ is transverse to $\mathfrak{h}_k$, we also immediately conclude by invoking Proposition 2.2.8, which yields a halfspace $\mathfrak{k} \subseteq \mathfrak{j} \cap \mathfrak{h}_k$. We are left to consider the case when $\mathfrak{w} \subseteq \mathfrak{h}_k^*$.

Let $\Gamma < \mathrm{Aut}(X)$ be the subgroup that preserves each of the halfspaces $\mathfrak{h}_1, ..., \mathfrak{h}_{k-1}$ and denote by $\mathfrak{w}_i$ the hyperplane corresponding to $\mathfrak{h}_i$. Since X has no free faces, the cube complex $\mathfrak{w}_1 \cap ... \cap \mathfrak{w}_{k-1}$ is essential. The action $\Gamma \curvearrowright \mathfrak{w}_1 \cap ... \cap \mathfrak{w}_{k-1}$ is cocompact by Lemma 2.2.4 and therefore essential. Now, applying Proposition 3.2 in [CS11] to the hyperplane $\mathfrak{w}_1 \cap ... \cap \mathfrak{w}_k$ of the cube complex $\mathfrak{w}_1 \cap ... \cap \mathfrak{w}_{k-1}$, we obtain an element

$g \in \Gamma$ with $\mathfrak{h}_k \subsetneq g\mathfrak{h}_k$. For large $n \geq 1$, either $\mathfrak{w}$ is transverse to $g^n\mathfrak{h}_k$ or $\mathfrak{w} \subseteq g^n\mathfrak{h}_k$. The discussion above then shows that, in both cases, there exists a halfspace

$$\mathfrak{j}' \subseteq \mathfrak{h}_1 \cap ... \cap \mathfrak{h}_{k-1} \cap g^n\mathfrak{h}_k = g^n(\mathfrak{h}_1 \cap ... \cap \mathfrak{h}_k).$$

We conclude by taking $\mathfrak{k} = g^{-n}\mathfrak{j}'$. □

2.2.3 The regular boundary $\partial_{\text{reg}}X$.

Let X be irreducible throughout this subsection.

The following notion first appeared as Definition 7.3 in [Fer18] and Definition 5.8 in [FLM18]. Also see Proposition 7.5 in [Fer18] for our formulation.

Definition 2.2.11. A point $x \in \partial X$ is *regular* if, for every $\mathfrak{h} \in \sigma_x$, there exists $\mathfrak{k} \in \sigma_x$ such that $\mathfrak{k}$ and $\mathfrak{h}^*$ are strongly separated. Equivalently, σ_x contains a descending chain $\mathfrak{h}_0 \supsetneq \mathfrak{h}_1 \supsetneq ...$ such that $\mathfrak{h}_n^*$ and $\mathfrak{h}_{n+1}$ are strongly separated for every $n \geq 0$. We refer to the latter as a *strongly separated chain* and denote by $\partial_{\text{reg}}X \subseteq \partial X$ the subset of regular points.

We record a few observations for later use.

Lemma 2.2.12. *[Lemma 5.12 in [FLM18]] Let $\mathfrak{h}_0 \supsetneq \mathfrak{h}_1 \supsetneq ...$ be a strongly separated chain. There exists a unique point $x \in \overline{X}$ that lies in each $\mathfrak{h}_n$. Moreover, x is regular.*

Lemma 2.2.13. *Let $\Gamma \curvearrowright X$ be essential and non-elementary. For every finite subset $\mathcal{F} \subseteq \partial_{\text{reg}}X$, there exists $g \in \Gamma$ with $g\mathcal{F} \cap \mathcal{F} = \emptyset$.*

Proof. The orbit $\Gamma \cdot x$ is infinite for every $x \in \mathcal{F}$; this follows for instance from Proposition 5.2 in [Fio18] and the fact that Γ has no finite orbits in the visual boundary of X. Lemma 2.2.12 shows that any strongly separated chain containing a point $y \in (\Gamma \cdot x) \setminus \mathcal{F}$ must include a halfspace $\mathfrak{h}$ with $\mathcal{F} \subseteq \mathfrak{h}^*$. Part (1) of Proposition 2.2.6 now yields $g \in \Gamma$ with $g\mathfrak{h}^* \cap \mathfrak{h}^* = \emptyset$, so in particular $g\mathcal{F}$ and $\mathcal{F}$ are disjoint. □

Remark 2.2.14. Lemma 5.14 in [FLM18] shows that the median $m(x, y, z)$ lies in X whenever the three points $z \in \overline{X}$ and $x, y \in \partial_{\text{reg}}X$ are pairwise distinct. In particular, every subset of $\partial_{\text{reg}}X$ is concise (see Definition 2.1.3).

With reference to the proof sketch of Theorem H in Section 1.5 of the introduction, the following is the formulation of Steps (Ia) and (Ib) that we will rely on. Note that hyperbolicity is not required here.

Proposition 2.2.15. *Let $\Gamma \curvearrowright X$ be essential and non-elementary. Consider two halfspaces $\mathfrak{h}_1, \mathfrak{h}_2$ and a nonempty Γ-invariant subset $\mathcal{A} \subseteq \partial_{\mathrm{reg}} X$.*

(1) The intersections $\mathfrak{h}_i \cap \mathcal{A}$ and $\mathfrak{h}_i^* \cap \mathcal{A}$ are always nonempty.

Assume now that X is also locally finite, cocompact and hyperplane-essential.

(2) The halfspaces $\mathfrak{h}_1$ and $\mathfrak{h}_2$ are transverse if and only if the set $\mathcal{A}$ intersects each of the four sectors $\mathfrak{h}_1 \cap \mathfrak{h}_2$, $\mathfrak{h}_1^* \cap \mathfrak{h}_2$, $\mathfrak{h}_1 \cap \mathfrak{h}_2^*$ and $\mathfrak{h}_1^* \cap \mathfrak{h}_2^*$.

(3) If we have $\mathfrak{h}_1 \cap \mathcal{A} \subsetneq \mathfrak{h}_2 \cap \mathcal{A}$, then $\mathfrak{h}_1 \subsetneq \mathfrak{h}_2$.

Proof. We start by showing that $\mathcal{A}$ intersects every halfspace $\mathfrak{h}$. Let us pick $x \in \mathcal{A}$ and a strongly separated chain $\mathfrak{k}_0 \supsetneq \mathfrak{k}_1 \supsetneq ...$ containing x. The hyperplane bounding $\mathfrak{h}$ must be contained in $\mathfrak{k}_n^*$ for all sufficiently large values of n. We thus have either $x \in \mathfrak{k}_n \subseteq \mathfrak{h}$ or $\mathfrak{k}_n \cap \mathfrak{h} = \emptyset$ for some $n \geq 0$. In the latter case, Proposition 2.2.6 yields an element $g \in \Gamma$ with $g\mathfrak{k}_n \subseteq \mathfrak{h}$. Hence $gx \in \mathfrak{h} \cap \mathcal{A}$ and this proves part (1).

If $\mathfrak{h}_1$ and $\mathfrak{h}_2$ are not transverse, one of the four intersections $\mathfrak{h}_1 \cap \mathfrak{h}_2$, $\mathfrak{h}_1^* \cap \mathfrak{h}_2$, $\mathfrak{h}_1 \cap \mathfrak{h}_2^*$, $\mathfrak{h}_1^* \cap \mathfrak{h}_2^*$ is empty and, in particular, it cannot contain any point of $\mathcal{A}$. If $\mathfrak{h}_1$ and $\mathfrak{h}_2$ *are* transverse, on the other hand, all four intersections contain a halfspace by Proposition 2.2.8. It then follows from part (1) that they all intersect $\mathcal{A}$ and this proves part (2).

We conclude by addressing part (3). If $\mathfrak{h}_1 \cap \mathcal{A} \subsetneq \mathfrak{h}_2 \cap \mathcal{A}$, part (2) shows that $\mathfrak{h}_1$ and $\mathfrak{h}_2$ cannot be transverse. We then have either $\mathfrak{h}_1 \subsetneq \mathfrak{h}_2$, or $\mathfrak{h}_2 \subseteq \mathfrak{h}_1$, or $\mathfrak{h}_2^* \subseteq \mathfrak{h}_1$ or $\mathfrak{h}_2 \subseteq \mathfrak{h}_1^*$. In the second case, we would have $\mathfrak{h}_1 \cap \mathcal{A} \subsetneq \mathfrak{h}_2 \cap \mathcal{A} \subseteq \mathfrak{h}_1 \cap \mathcal{A}$, in the third case $\mathfrak{h}_1 \cap \mathfrak{h}_2^* \cap \mathcal{A} \supseteq \mathfrak{h}_2^* \cap \mathcal{A} \neq \emptyset$ and in the fourth $\mathfrak{h}_1 \cap \mathfrak{h}_2^* \cap \mathcal{A} \supseteq \mathfrak{h}_1 \cap \mathcal{A} \neq \emptyset$. These all lead to a contradiction and we must therefore have $\mathfrak{h}_1 \subsetneq \mathfrak{h}_2$. □

In relation to part (3) of Proposition 2.2.15, note however that $\mathfrak{h}_1 \cap \mathcal{A} \subseteq \mathfrak{h}_2 \cap \mathcal{A}$ *does not* imply $\mathfrak{h}_1 \subseteq \mathfrak{h}_2$, as we might actually have $\mathfrak{h}_1 \cap \mathcal{A} = \mathfrak{h}_2 \cap \mathcal{A}$ and $\mathfrak{h}_2 \subsetneq \mathfrak{h}_1$ in this case.

As an immediate consequence of part (1) of Proposition 2.2.15, we have:

Corollary 2.2.16. *Let $\Gamma \curvearrowright X$ be essential and non-elementary. If X is sector-heavy, every nonempty Γ-invariant subset $\mathcal{A} \subseteq \partial_{\mathrm{reg}} X$ is ample.*

Chapter 3

Cross ratios on cube complexes.

We introduce the cross ratio on the Roller boundary ∂X in Sections 3.1 and 3.2. Then, in Sections 3.3, 3.4 and 3.5, we obtain some preliminary results on extensions of cross-ratio preserving maps between (subsets of) Roller boundaries. Corollary 3.4.4 and Theorem 3.5.3 will be crucial ingredients of Chapters 6 and 7, respectively.

At the end of the chapter, we anticipate a proof of Theorem M relying on Theorem F.

3.1 Gromov products.

Let X be a CAT(0) cube complex endowed with its combinatorial metric d. Given a base vertex $v \in X$, the *Gromov product* of $x, y \in \overline{X}$ is defined as:

$$(x \cdot y)_v := \#\mathscr{W}(v|x,y) = d\big(v, m(v,x,y)\big) \in \mathbb{N} \cup \{+\infty\}.$$

When $x, y \in X$, the above quantity coincides with the usual Gromov product:

$$(x \cdot y)_v = \tfrac{1}{2} \cdot \big[d(v,x) + d(v,y) - d(x,y)\big].$$

Note that $(x \cdot y)_v = +\infty$ if and only if $m(v,x,y) \in \partial X$.

Lemma 3.1.1. *Consider points $x, y, z \in \overline{X}$ and a vertex $v \in X$.*

(1) We have $m(x,y,z) \in X$ if and only if each of the three intervals $I(x,y)$, $I(y,z)$, $I(z,x)$ intersects X.

(2) We have $(x \cdot y)_v < +\infty$ if and only if $I(x,y)$ intersects X.

Proof. Note that the medians $m(x,y,z)$ and $m(x,y,v)$ lie in $I(x,y)$. If $I(x,y) \subseteq \partial X$, we then have $m(x,y,z) \in \partial X$, $m(x,y,v) \in \partial X$ and $(x \cdot y)_v = +\infty$. This shows one

implication of both (1) and (2). If $I(x,y)$, $I(y,z)$, $I(z,x)$ all intersect X, Helly's Lemma guarantees that $I(x,y) \cap I(y,z) \cap I(z,x) \cap X$ is nonempty, i.e. $m(x,y,z)$ lies in X. Applying this observation to the case $z = v$, we see that $I(x,y) \cap X \neq \emptyset$ is enough to guarantee that $m(x,y,v) \in X$; hence $(x \cdot y)_v < +\infty$. □

Thus, if $\mathcal{A} \subseteq \partial X$ is a concise subset, we have $(x \cdot y)_v < +\infty$ for all distinct points $x, y \in \mathcal{A}$ and all $v \in X$.

Lemma 3.1.2. *If X is locally finite, the function $(- \cdot -)_v \colon \overline{X}^2 \to \mathbb{N} \cup \{+\infty\}$ is continuous for every vertex $v \in X$.*

Proof. Consider points $x, y \in \overline{X}$. Given any finite subset $F \subseteq \mathscr{H}(v|x,y)$, let $U \subseteq \overline{X}$ be the set of points represented by ultrafilters containing F. Note that U is open, that x and y lie in U and that $(x' \cdot y')_v \geq \#F$ for all $x', y' \in U$. In particular, if $(x \cdot y)_v = +\infty$, this shows that the Gromov product is continuous at the pair (x,y).

Suppose instead that $(x \cdot y)_v < +\infty$. The median $m = m(v,x,y)$ lies in X by Lemma 3.1.1. Let V_x denote the set of points $x' \in \overline{X}$ with the property that $\mathscr{W}(v|x') \cap \mathscr{W}_m = \mathscr{W}(v|x) \cap \mathscr{W}_m$. Similarly, we define V_y as the set of points $y' \in \overline{X}$ with $\mathscr{W}(v|y') \cap \mathscr{W}_m = \mathscr{W}(v|y) \cap \mathscr{W}_m$. Note that $x \in V_x$ and $y \in V_y$. Since X is locally finite, the set $\mathscr{W}_m$ is finite; hence V_x and V_y are open in $\overline{X}$. Given $(x', y') \in V_x \times V_y$, no element of $\mathscr{W}_m$ separates x and x' or y and y'. It follows that no element of $\mathscr{W}_m$ separates m and $m(v,x',y')$; hence $m(v,x',y') = m$ and $(x' \cdot y')_v = (x \cdot y)_v$. □

3.2 Definition and first properties.

Let Ω be a set and $\mathscr{A} \subseteq \Omega^4$ a subset closed under permutations of the four coordinates. A map $\mathbb{B} \colon \mathscr{A} \to \mathbb{R} \cup \{\pm\infty\}$ is an *(abstract) cross ratio* on Ω if it satisfies the following conditions for all 4-tuples (x,y,z,w), (x,y,z,t) and (x,y,t,w) lying in $\mathscr{A}$:

(i) $\mathbb{B}(x,y,z,w) = -\mathbb{B}(y,x,z,w)$;

(ii) $\mathbb{B}(x,y,z,w) = \mathbb{B}(z,w,x,y)$;

(iii) $\mathbb{B}(x,y,z,w) = \mathbb{B}(x,y,z,t) + \mathbb{B}(x,y,t,w)$;

(iv) $\mathbb{B}(x,y,z,w) + \mathbb{B}(y,z,x,w) + \mathbb{B}(z,x,y,w) = 0$.

This definition should be compared e.g. with those in [Ota92, Ham97, Lab06].

Back to the CAT(0) cube complex X, let us fix a base vertex $v \in X$. We consider the subset $\mathscr{A} \subseteq (\overline{X})^4$ of 4-tuples (x,y,z,w) such that at most one of the three

quantities $(x \cdot y)_v + (z \cdot w)_v$, $(x \cdot z)_v + (y \cdot w)_v$ and $(x \cdot w)_v + (y \cdot z)_v$ is infinite. By part (2) of Lemma 3.1.1, the set $\mathscr{A}$ does not depend on the choice of v. By Lemma 3.1.2, it is an open subset of $(\overline{X})^4$. Note moreover that, by Remark 2.2.14 and Lemma 3.1.1, every 4-tuple of pairwise distinct points of $\partial_{\mathrm{reg}} X$ lies in $\mathscr{A}$.

It is easy to check that the map $\mathrm{cr}_v \colon \mathscr{A} \to \mathbb{Z} \cup \{\pm\infty\}$ defined by

$$\mathrm{cr}_v(x, y, z, w) = (x \cdot z)_v + (y \cdot w)_v - (x \cdot w)_v - (y \cdot z)_v$$

is well-defined and a cross ratio on $\overline{X}$. When x, y, z and w lie in X, we have

$$\mathrm{cr}_v(x, y, z, w) = d(x, w) + d(y, z) - d(x, z) - d(y, w),$$

which shows independence of the choice of the basepoint v. In general, basepoint-independence follows from the following result:

Proposition 3.2.1. *For every $v \in X$ and every $(x, y, z, w) \in \mathscr{A}$, we have*

$$\mathrm{cr}_v(x, y, z, w) = \#\mathscr{W}(x, z|y, w) - \#\mathscr{W}(x, w|y, z).$$

Proof. We show that every hyperplane $\mathfrak{w} \in \mathscr{W}$ gives the same contribution to both sides of the equality. Note that

$$\mathrm{cr}_v(x, y, z, w) = \#\mathscr{W}(v|x, z) + \#\mathscr{W}(v|y, w) - \#\mathscr{W}(v|x, w) - \#\mathscr{W}(v|y, z).$$

Every $\mathfrak{w} \in \mathscr{W}(x, z|y, w)$ contributes to either $\mathscr{W}(v|x, z)$ or $\mathscr{W}(v|y, w)$ by $+1$, without affecting $\mathscr{W}(v|x, w)$ and $\mathscr{W}(v|y, z)$. Similarly, every hyperplane $\mathfrak{w} \in \mathscr{W}(x, w|y, z)$ decreases $-\mathscr{W}(v|x, w) - \mathscr{W}(v|y, z)$ by 1 and leaves $\mathscr{W}(v|x, z)$ and $\mathscr{W}(v|y, w)$ invariant. Thus, it suffices to check that hyperplanes $\mathfrak{w} \notin \mathscr{W}(x, z|y, w) \sqcup \mathscr{W}(x, w|y, z)$ do not affect $\mathrm{cr}_v(x, y, z, w)$.

This is clear if all four points x, y, z and w lie on the same side of $\mathfrak{w}$, or if $\mathfrak{w} \in \mathscr{W}(x, y|z, w)$. The remaining case is when exactly three of the four points lie on one side of $\mathfrak{w}$. Performing a sequence of moves $(x \leftrightarrow y, z \leftrightarrow w)$ and $(x \leftrightarrow z, y \leftrightarrow w)$, which leave $\mathrm{cr}_v(x, y, z, w)$ invariant, we reduce to the case when $\mathfrak{w} \in \mathscr{W}(x|y, z, w)$. If v is not on the same side of $\mathfrak{w}$ as x, the hyperplane $\mathfrak{w}$ does not contribute to any summand of $\mathrm{cr}_v(x, y, z, w)$. Otherwise $\mathfrak{w} \in \mathscr{W}(x, v|y, z, w)$; in this case the only contributions to $\mathrm{cr}_v(x, y, z, w)$ arise from $\mathscr{W}(v|y, w)$ and $\mathscr{W}(v|y, z)$ and they cancel each other. □

We remark that the right-hand side of the equality in Proposition 3.2.1 is in general a cross ratio on a set strictly larger than $\mathscr{A}$.

Corollary 3.2.2. *The map* $\mathrm{cr}_v\colon \mathscr{A} \to \mathbb{Z} \cup \{\pm\infty\}$ *is independent of the choice of* v*. All automorphisms of* X *preserve* cr_v*.*

Definition 3.2.3. We will write $\mathrm{cr}\colon \mathscr{A} \to \mathbb{Z} \cup \{\pm\infty\}$ from now on and refer to it as *the cross ratio* on $\overline{X}$ (or ∂X).

Despite taking values in a discrete set, the cross ratio cr is generally continuous. Note to this regard that $\overline{X}$ and ∂X are totally disconnected. Endowing $\mathscr{A} \subseteq (\overline{X})^4$ with the subspace topology, the following is immediate from Lemma 3.1.2.

Proposition 3.2.4. *If* X *is locally finite, the cross ratio* cr *is continuous.*

It is often useful to simultaneously record information about the cross ratios of all 4-tuples obtained by permuting the coordinates of (x, y, z, w). This is the purpose of Definition 3.2.5 below. Note that identities (i) and (ii) imply that $|\mathrm{cr}(x, y, z, w)|$ is invariant under a subgroup of order 8 of the symmetric group on 4 elements. Thus, we only need to record $24/8 = 3$ 'meaningful' values for every subset $\{x, y, z, w\}$. These values are precisely the three cross ratios appearing in identity (iv), so they are not independent.

Given $a_1, b_1, c_1, a_2, b_2, c_2 \in \mathbb{N} \cup \{+\infty\}$, we say that the triples (a_1, b_1, c_1) and (a_2, b_2, c_2) are *equivalent* if there exists $n \geq 0$ such that

$$a_i = a_j + n; \qquad b_i = b_j + n; \qquad c_i = c_j + n,$$

where $\{i, j\} = \{1, 2\}$. We denote the equivalence class of (a, b, c) by $[\![a : b : c]\!]$. Note that $[\![+\infty : +\infty : +\infty]\!]$ is the only class consisting of a single triple; every other equivalence class has a unique representative with at least one zero entry. We also remark that all the triples in a given equivalence class have the same infinite entries.

Definition 3.2.5. Given $x, y, z, w \in \overline{X}$ and $v \in X$, the *cross ratio triple* $\mathrm{crt}_v(x, y, z, w)$ is the equivalence class:

$$\Big[\!\Big[(x \cdot y)_v + (z \cdot w)_v : (x \cdot z)_v + (y \cdot w)_v : (x \cdot w)_v + (y \cdot z)_v\Big]\!\Big].$$

This is the same as $[\![\#\mathscr{W}(x, y|z, w) : \#\mathscr{W}(x, z|y, w) : \#\mathscr{W}(x, w|y, z)]\!]$ whenever (x, y, z, w) lies in $\mathscr{A}$, as a consequence of Proposition 3.2.1.

Note that crt_v is always independent of the choice of v. This follows from Proposition 3.2.1 when $(x, y, z, w) \in \mathscr{A}$ and is clear otherwise. We thus simply write crt.

All entries of a cross ratio triple are nonnegative. The three cross ratios in identity (iv) above are recovered by taking the difference of two entries of the triple. We

will employ asterisks $*$ when we do not want to specify a coordinate of $\mathrm{crt}(x,y,z,w)$. We thus write "$\mathrm{crt}(x,y,z,w) = [\![* : 0 : 1]\!]$" rather than "$\mathrm{crt}(x,y,z,w) = [\![a : 0 : 1]\!]$ and $a \in \mathbb{N} \cup \{+\infty\}$".

Let now Y be another CAT(0) cube complex. We write $\mathscr{A}(X)$ rather than just $\mathscr{A}$ when it is necessary to specify the cube complex under consideration. The following makes precise the notion of 'cross-ratio preserving map', which we have used in the introduction.

Definition 3.2.6. Let $\mathcal{A} \subseteq \partial X$ and $\mathcal{B} \subseteq \partial Y$ be subsets. A map $f\colon \mathcal{A} \to \mathcal{B}$ is *Möbius* if $f^4(\mathscr{A}(X) \cap \mathcal{A}^4) \subseteq \mathscr{A}(Y)$ and, for each $(x,y,z,w) \in \mathscr{A}(X) \cap \mathcal{A}^4$, we have

$$\mathrm{cr}(f(x), f(y), f(z), f(w)) = \mathrm{cr}(x,y,z,w).$$

The latter happens if and only if $\mathrm{crt}(f(x), f(y), f(z), f(w)) = \mathrm{crt}(x,y,z,w)$ for all 4-tuples $(x,y,z,w) \in \mathscr{A}(X) \cap \mathcal{A}^4$.

A bijection $f\colon \mathcal{A} \to \mathcal{B}$ is Möbius if and only if its inverse $f^{-1}\colon \mathcal{B} \to \mathcal{A}$ is.

3.3 Opposite points.

As described in the introduction, *opposite points* (Definition 3.3.2) will be crucial to avoid the issue depicted in Figure 1.1. Here, we briefly introduce this notion.

Given $x_1, x_2 \in \partial X$, we consider the CAT(0) cube complex $T = X \cap I(x_1, x_2)$.

Lemma 3.3.1. *The following conditions are equivalent for a vertex $v \in T$:*

(1) $\mathrm{lk}_T(v)$ *has exactly two connected components;*

(2) $\mathrm{lk}_T(v)$ *is a union of two disjoint cliques;*

(3) $I(x_1, x_2) = I(x_1, v) \cup I(v, x_2)$.

In this case, the connected components of $\mathrm{lk}_T(v)$ *are given by the sets of edges crossing the hyperplanes of* $\mathscr{W}(v|x_1)$ *and* $\mathscr{W}(v|x_2)$.

Proof. For each i, the hyperplanes in $\mathscr{W}(v|x_i) \cap \mathscr{W}_v$ are pairwise transverse and originate a clique in $\mathrm{lk}(v)$. As $\mathscr{W}(T) = \mathscr{W}(x_1|x_2) = \mathscr{W}(x_1|v) \sqcup \mathscr{W}(v|x_2)$, either $\mathrm{lk}_T(v)$ is connected, or v satisfies (1) and the two connected components of $\mathrm{lk}_T(v)$ are given precisely by $\mathscr{W}(v|x_1)$ and $\mathscr{W}(v|x_2)$. This shows the equivalence of (1) and (2).

Regarding (2) $\Rightarrow$ (3), if a point $w \in I(x_1, x_2)$ did not lie in $I(x_1, v) \cup I(v, x_2)$, the set $\mathscr{W}(v|w)$ would intersect both $\mathscr{W}(v|x_1)$ and $\mathscr{W}(v|x_2)$. For $i = 1, 2$, pick hyperplanes

$\mathfrak{w}_i \in \mathscr{W}(v|w) \cap \mathscr{W}(v|x_i)$; choosing $\mathfrak{w}_i$ closest to v, we can assume that $\mathfrak{w}_i \in \mathscr{W}_v$. Let $\mathfrak{h}_i$ be the the side of $\mathfrak{w}_i$ that contains v. Since v satisfies condition (2), we must have $\mathfrak{h}_1^* \cap \mathfrak{h}_2^* = \emptyset$; on the other hand, w lies in both $\mathfrak{h}_1^*$ and $\mathfrak{h}_2^*$. This is a contradiction.

Finally, if $\mathfrak{w}_1 \in \mathscr{W}(v|x_1)$ and $\mathfrak{w}_2 \in \mathscr{W}(v|x_2)$ are adjacent to v and transverse, there exists a point $w \in I(x_1, x_2)$ with $\mathscr{W}(v|w) = \{\mathfrak{w}_1, \mathfrak{w}_2\}$. In this case, we have $w \notin I(x_1, v) \cup I(v, x_2)$. This shows (3) $\Rightarrow$ (1). □

We refer to a vertex $v \in T$ satisfying the equivalent conditions in Lemma 3.3.1 as a *cut point* for T.

Definition 3.3.2. Given $x, y, z \in \overline{X}$, we say that x and y are *opposite* with respect to z (written $x \operatorname{op}_z y$) if the median $m = m(x, y, z)$ lies in X and it is a cut point for $I(x, y) \cap X$. Equivalently, $I(x, y) = I(x, m) \cup I(m, y)$.

The terminology above is inspired by the case when z coincides with the median $m(x, y, z) \in X$. We are however mostly interested in the general situation where the points x, y and z lie in ∂X.

Remark 3.3.3. Consider points $x, y, z \in \overline{X}$ with $x \operatorname{op}_z y$ and $m = m(x, y, z)$. Given any $w \in \partial X$, we either have $(x \cdot w)_m = 0$ or $(y \cdot w)_m = 0$. Indeed, the gate-projection $m(x, y, w)$ falls either in $I(y, m)$ or in $I(x, m)$.

The rest of this section is devoted to showing that Möbius bijections take triples of opposite points to triples with the same property. A similar result under slightly different hypotheses will be Lemma 6.1.5.

Proposition 3.3.4. *Let X be locally finite. Let $\mathcal{A} \subseteq \partial X$ be ample and concise. Given pairwise-distinct points $x_1, x_2, y \in \mathcal{A}$, the condition $x_1 \operatorname{op}_y x_2$ fails if and only if there exists $z \in \mathcal{A}$ such that $\operatorname{crt}(x_1, x_2, y, z) = [\![a : b : c]\!]$ with $a < \min\{b, c\}$ and $\max\{b, c\} < +\infty$.*

Proof. Since $\mathcal{A}$ is concise, the point $m = m(x_1, x_2, y)$ lies in X by Lemma 3.1.1. We have:

$$\operatorname{crt}_m(x_1, x_2, y, z) = [\![(y \cdot z)_m : (x_2 \cdot z)_m : (x_1 \cdot z)_m]\!].$$

If $x_1 \operatorname{op}_y x_2$ and $z \in \partial X$, Remark 3.3.3 yields $\min\{(x_2 \cdot z)_m, (x_1 \cdot z)_m\} = 0$ and we cannot have $(y \cdot z)_m < 0$.

Suppose now instead that x_1 and x_2 are not opposite with respect to y. There exist transverse hyperplanes $\mathfrak{w}_i \in \mathscr{W}(m|x_i) \cap \mathscr{W}_m$. Denoting by $m' \in X$ the point with $\mathscr{W}(m|m') = \{\mathfrak{w}_1, \mathfrak{w}_2\}$, the hyperplanes $\mathfrak{w}_1$ and $\mathfrak{w}_2$ lie in $\mathscr{W}(m'|y) \cap \mathscr{W}_{m'}$. This

set is finite, as X is locally finite. Let $\mathfrak{h}_1, ..., \mathfrak{h}_k$ be the sides containing m' of the hyperplanes in $\mathscr{W}(m'|y) \cap \mathscr{W}_{m'}$.

As $\mathcal{A}$ is ample, there exists a point $z \in \mathcal{A} \cap \mathfrak{h}_1 \cap ... \cap \mathfrak{h}_k$. This property and Lemma 2.1.1 imply that $(y \cdot z)_m = 0$. The fact that $\mathcal{A}$ is concise guarantees that the Gromov products $(x_1 \cdot z)_m$ and $(x_2 \cdot z)_m$ are finite. Observing that $\mathfrak{w}_i \in \mathscr{W}(m|x_i, z)$, we moreover have $(x_i \cdot z)_m \geq 1$. Since $\mathrm{crt}(x_1, x_2, y, z) = [\![0 : (x_2 \cdot z)_m : (x_1 \cdot z)_m]\!]$, this concludes the proof. □

Consider now another CAT(0) cube complex Y.

Definition 3.3.5. Let $\mathcal{A} \subseteq \partial X$ and $\mathcal{B} \subseteq \partial Y$ be subsets. A bijection $f \colon \mathcal{A} \to \mathcal{B}$ is *op-preserving* if we have $x \operatorname{op}_z y \Leftrightarrow f(x) \operatorname{op}_{f(z)} f(y)$ for all $x, y, z \in \partial X$.

Proposition 3.3.4 immediately yields:

Corollary 3.3.6. *Let X and Y be locally finite. If $\mathcal{A} \subseteq \partial X$ and $\mathcal{B} \subseteq \partial Y$ are ample and concise, every Möbius bijection $f \colon \mathcal{A} \to \mathcal{B}$ is op-preserving.*

A closely related result is Corollary 6.1.6.

3.4 Extending op-preserving bijections.

Let X and Y be CAT(0) cube complexes. Throughout this section, we consider two subsets $\mathcal{A} \subseteq \partial X$, $\mathcal{B} \subseteq \partial Y$ and an op-preserving, Möbius bijection $f \colon \mathcal{A} \to \mathcal{B}$.

We saw in Corollary 3.3.6 that all Möbius bijections are op-preserving if $\mathcal{A}$, $\mathcal{B}$ are ample and concise and X, Y are locally finite. However, we make **no assumptions** on $\mathcal{A}$ and $\mathcal{B}$ in this section. Indeed, we will need to apply our results to $\mathcal{A} = \partial X$ in Chapter 6, which is not concise. We also make no assumptions on X and Y.

We are about to describe how triples of opposite points can be used to extend f into the cube complexes X and Y. This will be a crucial ingredient of Chapters 6 and 7. We start with a few simple results culminating in Corollary 3.4.4.

In order to avoid cumbersome formulas, we employ the following notation for points $x, y, z, w \in \partial X$ and vertices $v \in Y$:

$$(x \cdot y)^f_v := (f(x) \cdot f(y))_v, \qquad m^f(x, y, z) := m(f(x), f(y), f(z)),$$
$$x \operatorname{op}^f_z y \Leftrightarrow f(x) \operatorname{op}_{f(z)} f(y), \qquad \mathrm{crt}^f(x, y, z, w) := \mathrm{crt}(f(x), f(y), f(z), f(w)).$$

We also consider points $x_1, x_2, x, y_1, y_2, y \in \partial X$ with $x_1 \operatorname{op}_x x_2$ and $y_1 \operatorname{op}_y y_2$. We set:

$$m_x = m(x_1, x_2, x), \qquad m_y = m(y_1, y_2, y),$$
$$m'_x = m^f(x_1, x_2, x), \qquad m'_y = m^f(y_1, y_2, y).$$

Lemma 3.4.1. *Given $u \in \partial X$ with $(x_1, x_2, x, u) \in \mathscr{A}(X)$, we have:*

$$(x_1 \cdot u)_{m_x} = (x_1 \cdot u)^f_{m'_x}, \quad (x_2 \cdot u)_{m_x} = (x_2 \cdot u)^f_{m'_x}, \quad (x \cdot u)_{m_x} = (x \cdot u)^f_{m'_x}.$$

Proof. Observe that:

$$\mathrm{crt}_{m_x}(x_1, x_2, x, u) = [\![(x \cdot u)_{m_x} : (x_2 \cdot u)_{m_x} : (x_1 \cdot u)_{m_x}]\!],$$

$$\mathrm{crt}^f_{m'_x}(x_1, x_2, x, u) = [\![(x \cdot u)^f_{m'_x} : (x_2 \cdot u)^f_{m'_x} : (x_1 \cdot u)^f_{m'_x}]\!].$$

Since $x_1 \,\mathrm{op}_x\, x_2$, Remark 3.3.3 shows that either $(x_1 \cdot u)_{m_x} = 0$ or $(x_2 \cdot u)_{m_x} = 0$. Since f is op-preserving, we have $x_1 \,\mathrm{op}^f_x\, x_2$ and also one among $(x_1 \cdot u)^f_{m'_x}$ and $(x_2 \cdot u)^f_{m'_x}$ must vanish. The equality $\mathrm{crt}(x_1, x_2, x, u) = \mathrm{crt}^f(x_1, x_2, x, u)$ then implies that we have $(x_1 \cdot u)_{m_x} = (x_1 \cdot u)^f_{m'_x}$, $(x_2 \cdot u)_{m_x} = (x_2 \cdot u)^f_{m'_x}$ and $(x \cdot u)_{m_x} = (x \cdot u)^f_{m'_x}$. □

Lemma 3.4.2. *Let $u, v \in \partial X$ be two points such that the 4-tuples (x_1, x_2, u, v), (x_1, x_2, x, u) and (x_1, x_2, x, v) all lie in $\mathscr{A}(X)$. Then $(u \cdot v)_{m_x} = (u \cdot v)^f_{m'_x}$.*

Proof. We have $\mathrm{crt}(x_1, x_2, u, v) = \mathrm{crt}^f(x_1, x_2, u, v)$. Equating the cross ratio triples $\mathrm{crt}_{m_x}(x_1, x_2, u, v)$ and $\mathrm{crt}^f_{m'_x}(x_1, x_2, u, v)$, we obtain

$$[\![(u \cdot v)_{m_x} : b : c]\!] = [\![(u \cdot v)^f_{m'_x} : b' : c']\!],$$

where Lemma 3.4.1 yields $b = b'$ and $c = c'$. Hence $(u \cdot v)_{m_x} = (u \cdot v)^f_{m'_x}$. □

Proposition 3.4.3. *We have: $d(m_x, m_y) = (y_1 \cdot y_2)_{m_x} + |(y_1 \cdot y)_{m_x} - (y_2 \cdot y)_{m_x}|$.*

Proof. Set $v = m(y_1, y_2, m_x)$. As v is the gate-projection of m_x to the interval $I(y_1, y_2)$, we have $d(m_x, m_y) = d(m_x, v) + d(v, m_y)$, where $d(m_x, v) = (y_1 \cdot y_2)_{m_x}$. Up to exchanging y_1 and y_2, we can assume that v lies within $I(m_y, y_2)$. Since no element of $\mathscr{W}(v|y_2) = \mathscr{W}(m_x, y_1|y_2)$ separates m_x and y, it follows that the set $\mathscr{W}(m_x, y_1|y_2, y)$ is empty. We conclude that $(y_2 \cdot y)_{m_x} = \#\mathscr{W}(m_x|y_1, y_2, y)$. On the other hand, observing that $\mathscr{W}(v|m_y) = \mathscr{W}(m_x, y_2|y_1, y)$, we have:

$$\mathscr{W}(m_x|y_1, y) = \mathscr{W}(m_x|y_1, y_2, y) \sqcup \mathscr{W}(v|m_y)$$

and $(y_1 \cdot y)_{m_x} = (y_2 \cdot y)_{m_x} + d(v, m_y)$. □

Lemma 3.4.2 and Proposition 3.4.3 immediately yield the following.

Corollary 3.4.4. *Let $x_1, x_2, x, y_1, y_2, y \in \partial X$ be such that $x_1 \,\mathrm{op}_x\, x_2$ and $y_1 \,\mathrm{op}_y\, y_2$. Suppose that (x_1, x_2, u, v) and (x_1, x_2, x, u) lie in $\mathscr{A}(X)$ whenever u and v are distinct elements of the set $\{y_1, y_2, y\}$. Then, we have $d(m_x, m_y) = d(m'_x, m'_y)$. In particular, m'_x and m'_y coincide if and only if m_x and m_y do.*

3.5 Ample sets of regular points.

Let X and Y be locally finite CAT(0) cube complexes. In this section, we apply Corollary 3.4.4 to a Möbius bijection $f\colon \mathcal{A} \to \mathcal{B}$ between nonempty, ample sets of regular points $\mathcal{A} \subseteq \partial_{\mathrm{reg}} X$ and $\mathcal{B} \subseteq \partial_{\mathrm{reg}} Y$. Note that $\mathcal{A}$ and $\mathcal{B}$ are concise by Remark 2.2.14 and f is op-preserving by Corollary 3.3.6.

The main result of the section will be Theorem 3.5.3, which is a key ingredient of Chapter 7. We start by considering the set:

$$\mathfrak{M}_{\mathcal{A}}(X) = \{m(x,y,z) \mid x,y,z \in \mathcal{A},\ x \operatorname{op}_z y\} \subseteq X.$$

The subset $\mathfrak{M}_{\mathcal{B}}(Y) \subseteq Y$ is defined similarly. When $\mathcal{A} = \partial_{\mathrm{reg}} X$ and $\mathcal{B} = \partial_{\mathrm{reg}} Y$, we simply write $\mathfrak{M}(X)$ and $\mathfrak{M}(Y)$.

Since $\mathcal{A}$ is concise, a 4-tuple $(x,y,z,w) \in \mathcal{A}^4$ lies in $\mathscr{A}(X)$ if and only if no three of the four points x, y, z and w coincide. Applying Corollary 3.4.4 to f and f^{-1}, we immediately obtain maps $\phi\colon \mathfrak{M}_{\mathcal{A}}(X) \to \mathfrak{M}_{\mathcal{B}}(Y)$ and $\psi\colon \mathfrak{M}_{\mathcal{B}}(Y) \to \mathfrak{M}_{\mathcal{A}}(X)$ that preserve distances and satisfy:

$$\phi(m(x,y,z)) = m^f(x,y,z), \qquad \psi(m(x,y,z)) = m^{f^{-1}}(x,y,z),$$

whenever $x \operatorname{op}_z y$. Since f is op-preserving, the compositions $\phi \circ \psi$ and $\psi \circ \phi$ are the identity. We have shown:

Corollary 3.5.1. *The map $\phi\colon \mathfrak{M}_{\mathcal{A}}(X) \to \mathfrak{M}_{\mathcal{B}}(Y)$ is a distance-preserving bijection.*

We will also need the following simple observation.

Lemma 3.5.2. *The sets $\mathfrak{M}_{\mathcal{A}}(X)$ and $\mathfrak{M}(X)$ coincide.*

Proof. Since $\mathcal{A} \subseteq \partial_{\mathrm{reg}} X$, it is clear that $\mathfrak{M}_{\mathcal{A}}(X) \subseteq \mathfrak{M}(X)$. Let us consider a vertex $v \in \mathfrak{M}(X)$ and $x,y,z \in \partial_{\mathrm{reg}} X$ with $m(x,y,z) = v$ and $x \operatorname{op}_z y$.

Consider a strongly separated chain $\mathfrak{k}_0 \supsetneq \mathfrak{k}_1 \supsetneq \ldots$ containing x. Given $\mathfrak{w} \in \mathscr{W}_v$, the halfspace $\mathfrak{k}_n$ is disjoint from $\mathfrak{w}$ and on the same side of $\mathfrak{w}$ as x for all sufficiently large values of n. Since $\mathscr{W}_v$ is finite, there exists a halfspace $\mathfrak{h}_x \in \sigma_x$ that lies on the same side as x with respect to all hyperplanes in $\mathscr{W}_v$. We similarly define halfspaces $\mathfrak{h}_y$ and $\mathfrak{h}_z$. Since $\mathcal{A}$ is ample, we can pick points $x' \in \mathcal{A} \cap \mathfrak{h}_x$, $y' \in \mathcal{A} \cap \mathfrak{h}_y$ and $z' \in \mathcal{A} \cap \mathfrak{h}_z$. These satisfy $x' \operatorname{op}_{z'} y'$ and $m(x',y',z') = v$. □

Theorem 3.5.3. *(1) The Möbius bijection $f\colon \mathcal{A} \to \mathcal{B}$ determines a distance-preserving bijection $\phi\colon \mathfrak{M}(X) \to \mathfrak{M}(Y)$.*

(2) If f is equivariant with respect to actions of a group Γ, so is ϕ.

(3) The map ϕ extends f in the following sense: if vertices $v_n \in \mathfrak{M}(X)$ converge to a point $w \in \mathcal{A}$, we have $\phi(v_n) \to f(w)$.

(4) If $F\colon X \to Y$ is an isomorphism that restricts to f on $\mathcal{A}$, then F and ϕ coincide on $\mathfrak{M}(X)$.

Proof. Part (1) follows from Corollary 3.5.1 and Lemma 3.5.2, whereas parts (2) and (4) are immediate from our construction. We are left to show that $\phi(v_n) \to f(w)$ whenever vertices $v_n \in \mathfrak{M}(X)$ converge to a point $w \in \mathcal{A}$.

Say $v_n = m(x_n, y_n, z_n)$ with $x_n, y_n, z_n \in \mathcal{A}$ and $x_n \operatorname{op}_{z_n} y_n$. Let us write a_n, b_n and c_n for the values $(x_n \cdot w)_{v_0}$, $(y_n \cdot w)_{v_0}$ and $(z_n \cdot w)_{v_0}$ ordered so that $a_n \le b_n \le c_n$. Consider a strongly separated chain $\mathfrak{k}_0 \supsetneq \mathfrak{k}_1 \supsetneq \ldots$ containing w. Note that we have $v_n \to w$ if and only if, for every $k \ge 0$, there exists $N(k) \ge 0$ such that $v_n \in \mathfrak{k}_k$ for all $n \ge N(k)$; the latter happens if and only if at least two elements of $\{x_n, y_n, z_n\}$ lie in $\mathfrak{k}_k$. We conclude that $v_n \to w$ if and only if $b_n \to +\infty$. Lemma 3.4.2 now shows that $v_n \to w$ if and only if the points $\phi(v_n) = m^f(x_n, y_n, z_n)$ converge to $f(w)$. □

Note that Theorem M now immediately follows from Theorem F, Corollary 2.2.16 and Theorem 3.5.3. The *"sets of OR medians"* in the statement of Theorem M are precisely $\mathfrak{M}(X)$ and $\mathfrak{M}(Y)$ (cf. Definition 7.1.3).

Chapter 4

Contracting boundaries of cube complexes.

This chapter is mainly concerned with the proof of Theorem K (see Theorem 4.2.6 and Remark 4.2.4). Throughout it, we consider a finite dimensional, locally finite CAT(0) cube complex X.

4.1 Contracting geodesics.

Definition 4.1.1. Let Y be a proper metric space. Given a closed subset $A \subseteq Y$, we denote by $\pi_A \colon Y \to 2^A$ the nearest-point projection to A. If $B \subseteq Y$, we write $\pi_A(B)$ instead of $\bigcup_{b\in B} \pi_A(b)$. A closed subset $A \subseteq Y$ is *contracting* if there exists $D > 0$ such that every metric ball B disjoint from A satisfies $\operatorname{diam}(\pi_A(B)) \leq D$.

Given a function $M \colon [1, +\infty) \to [0, +\infty)$, a quasi-geodesic $\gamma \subseteq Y$ is *M-Morse* if, for every $C > 0$ and every (C, C)-quasi-geodesic η with endpoints on γ, the entire η is contained in the (open) $M(C)$-neighbourhood of γ. We say that γ is *Morse* if it is M-Morse for some function M.

We refer the reader to [ACGH17] for a detailed discussion of contracting subsets and the Morse property in general metric spaces.

Recall now that the cube complex X is endowed with its combinatorial metric. All geodesics will be combinatorial in this section.

Definition 4.1.2. We say that a geodesic $\gamma \subseteq X$ is *C-lean* if there do not exist transverse subsets $\mathcal{U} \subseteq \mathscr{W}(\gamma)$ and $\mathcal{V} \subseteq \mathscr{W}(X)$ such that $\min\{\#\mathcal{U}, \#\mathcal{V}\} > C$ and such that $\mathcal{U} \sqcup \mathcal{V}$ does not contain facing triples. We say that γ is *lean* if it is C-lean for some $C \geq 0$.

For the following, see Corollary 3.7 in [Gen16b] and Lemma 4.6 in [Gen17].

Theorem 4.1.3. *For a ray $\gamma \subseteq X$, we have: lean $\Leftrightarrow$ contracting $\Leftrightarrow$ Morse.*

Given rays γ and γ' at finite Hausdorff distance, it is clear from definitions that γ satisfies the above conditions if and only if γ' does.

Lemma 4.1.4. *Let α and γ be rays in X with $\alpha(+\infty) \sim \gamma(+\infty)$ in ∂X. If γ is contracting, then α is at finite Hausdorff distance from γ, hence contracting.*

Proof. By Lemma 2.1.6, it suffices to consider the case when $\alpha(+\infty) = \gamma(+\infty)$. By Theorem 4.1.3, there exists $C > 0$ such that γ is C-lean. Set $p = \gamma(0)$, $q = \alpha(0)$, $x = \alpha(+\infty) = \gamma(+\infty)$ and $I = I(x, p)$. Since $\mathscr{W}(q|I) = \mathscr{W}(q|x, p)$ is a finite subset of $\mathscr{W}(q|x) = \mathscr{W}(\alpha)$, the intersection $\alpha \cap I$ is a sub-ray of α. We conclude by showing that $d(u, \gamma) \leq 2C$ for every point $u \in I$.

Pick a point $v \in \gamma$ with $d(u, p) = d(v, p)$ and set $m = m(p, u, v)$. Note that $\mathscr{W}(m|u)$ and $\mathscr{W}(m|v)$ are contained in $\mathscr{W}(x|p) = \mathscr{W}(\gamma)$. Every halfspace $\mathfrak{h} \in \mathscr{H}(m|u)$ is transverse to every halfspace $\mathfrak{k} \in \mathscr{H}(m|v)$; indeed, we have $m \in \mathfrak{h}^* \cap \mathfrak{k}^*$, $u \in \mathfrak{h} \cap \mathfrak{k}^*$, $v \in \mathfrak{h}^* \cap \mathfrak{k}$ and $x \in \mathfrak{h} \cap \mathfrak{k}$. Moreover, the sets $\mathscr{W}(m|u)$ and $\mathscr{W}(m|v)$ have the same size and contain no facing triples. We conclude that $\#\mathscr{W}(m|u) = \#\mathscr{W}(m|v) \leq C$, since γ is C-lean. This shows that $d(u, v) \leq 2C$. □

4.2 Roller boundaries vs contracting boundaries.

Unlike the rest of the thesis, this section employs both the combinatorial and CAT(0) metrics on X; we will specify each time whether geodesics are meant with respect to the former or latter.

The *contracting boundary* $\partial_c X$ was introduced in [CS15]. Disregarding topologies for the moment, $\partial_c X$ is the subset of the visual boundary $\partial_\infty X$ that consists of points represented by contracting CAT(0) rays. In order to relate the contracting boundary $\partial_c X$ and the Roller boundary ∂X, we introduce the following (cf. Lemma 4.1.4):

Definition 4.2.1. A point $x \in \partial X$ is *contracting* (or a *contracting ultrafilter*) if one (equivalently, all) combinatorial ray representing x is contracting. We denote the set of contracting ultrafilters by $\partial_{\mathrm{cu}} X \subseteq \partial X$.

We stress that our definition of contracting point is not equivalent to the one in Remark 6.7 of [FLM18]; in fact, our notion is weaker, as it does not imply that $x \in \partial_{\mathrm{reg}} X$.

In general, the inclusion $\partial_{\mathrm{cu}}X \subseteq \partial X$ is strict. If however X is Gromov hyperbolic, every combinatorial ray in X is contracting (see e.g. Theorem 3.3 in [Gen16a]) and we have $\partial_{\mathrm{cu}}X = \partial X$.

Lemma 4.1.4 shows that the set $\partial_{\mathrm{cu}}X$ is a union of $\sim$-equivalence classes. The following result provides more information.

Lemma 4.2.2. *(1) Every component of $\partial_{\mathrm{cu}}X$ is bounded.*

(2) Points $x \in \partial_{\mathrm{cu}}X$ and $y \in \partial X$ lie in the same component if and only if they satisfy $I(x,y) \cap X = \emptyset$.

Proof. Consider a point $x \in \partial_{\mathrm{cu}}X$ and a C-lean combinatorial ray r with $x = r(+\infty)$; set $p = r(0)$. We simultaneously prove both parts of the lemma by showing that, for $y \in \partial X$, the condition $(x \cdot y)_p = +\infty$ implies $d(x,y) \leq C$ (cf. Lemma 3.1.1).

Suppose for the sake of contradiction that $\mathscr{W}(p|x,y)$ is infinite and $\mathscr{W}(x|y)$ contains a finite subset $\mathcal{U}$ with $\#\mathcal{U} > C$. Given a hyperplane $\mathfrak{w} \in \mathscr{W}(x|y)$ and a halfspace $\mathfrak{h} \in \mathscr{H}(p|x,y)$, either $\mathfrak{w} \subseteq \mathfrak{h}$ or $\mathfrak{w}$ and $\mathfrak{h}$ are transverse. Fixing $\mathfrak{w}$, there are at most $d(p,\mathfrak{w}) < +\infty$ halfspaces $\mathfrak{h} \in \mathscr{H}(p|\mathfrak{w})$. Thus, almost every hyperplane in $\mathscr{W}(p|x,y)$ is transverse to all elements of $\mathcal{U}$. As $\mathscr{W}(p|x,y)$ and $\mathcal{U}$ contain no facing triples, this violates C-leanness of r. □

Lemma 4.2.3. *Consider a point $x \in \partial_{\mathrm{cu}}X$. There exists an infinite descending chain of halfspaces $\mathfrak{h}_0 \supsetneq \mathfrak{h}_1 \supsetneq \ldots$ such that $\bigcap \mathfrak{h}_n = Z(x)$ and such that the shores $S(\mathfrak{h}_n^*, \mathfrak{h}_{n+1})$ are finite cube complexes of uniformly bounded diameter.*

Proof. Let r be a contracting combinatorial ray with $x = r(+\infty)$. Theorem 3.9 in [Gen16b] yields an infinite chain of halfspaces $\mathfrak{h}_0 \supsetneq \mathfrak{h}_1 \supsetneq \ldots$ such that the shores $S_n = S(\mathfrak{h}_n^*, \mathfrak{h}_{n+1})$ have uniformly bounded diameter and such that $x \in \mathfrak{h}_n$ for every $n \geq 0$. Since shores embed as subcomplexes of X, they are locally finite. Boundedness then implies that each S_n is finite.

Now, observe that $\mathfrak{h}_n^* \cap Z(x) = \emptyset$ for all $n \geq 0$. Otherwise, there would exist an integer $k \geq 0$ and a point $y \in \mathfrak{h}_k^* \cap Z(x)$. We would then have $y \in \mathfrak{h}_n^* \cap Z(x)$ for all $n \geq k$, violating the fact that $d(x,y) < +\infty$. This shows that $Z(x)$ is contained in each $\mathfrak{h}_n$. Given a point $z \in \bigcap \mathfrak{h}_n$, we have $I(x,z) \subseteq \bigcap \mathfrak{h}_n$ and hence $I(x,z) \cap X = \emptyset$. Lemma 4.2.2 then shows that $z \in Z(x)$. This proves that $\bigcap \mathfrak{h}_n = Z(x)$ and concludes the proof. □

Remark 4.2.4. If X is *uniformly* locally finite, part (1) of Lemma 4.2.2 can actually be promoted to say that components of $\partial_{\mathrm{cu}}X$ are finite. Indeed, let x and $\mathfrak{h}_0 \supsetneq \mathfrak{h}_1 \supsetneq \ldots$

be as in the statement of Lemma 4.2.3. Given $y \in Z(x)$, every $\mathfrak{w} \in \mathscr{W}(x|y)$ must be transverse to all but finitely many of the $\mathfrak{h}_n$. In particular, $\mathfrak{w}$ must be a hyperplane of almost all shores $S(\mathfrak{h}_n^*, \mathfrak{h}_{n+1})$. The latter are uniformly finite, as they embed as uniformly bounded subcomplexes of the uniformly locally finite cube complex X. We conclude that only finitely many hyperplanes of X can separate two points of $Z(x)$, i.e. that $Z(x)$ is finite.

Since X is finite dimensional, its combinatorial and CAT(0) metrics are quasi-isometric. In particular, the notion of Morse quasi-geodesic is independent of our choice of one of the two metrics.

We remark that Morse quasi-geodesics in complete CAT(0) spaces always stay within bounded distance of contracting CAT(0) geodesics (see e.g. Lemma 2.5, Theorem 2.9 and the proof of Corollary 2.10 in [CS15]). Along with Theorem 4.1.3 and Proposition 2.1.8, this observation yields the following.

Corollary 4.2.5. *Every contracting combinatorial ray is at finite Hausdorff distance from a contracting* CAT(0) *ray. Every contracting* CAT(0) *ray is at finite Hausdorff distance from a contracting combinatorial ray.*

Every point $x \in \partial_{\mathrm{cu}} X$ is represented by combinatorial rays in X. These rays are all contracting by Lemma 4.1.4 and Corollary 4.2.5 shows that they are at finite Hausdorff distance from a unique family of pairwise-asymptotic contracting CAT(0) rays. This yields a map

$$\Phi \colon \partial_{\mathrm{cu}} X \longrightarrow \partial_c X.$$

The next result describes a few properties of Φ. We endow $\partial_{\mathrm{cu}} X$ with the restriction of the topology of ∂X. We write $\partial_c^{\mathrm{vis}} X$ to refer to the contracting boundary $\partial_c X$ endowed with the restriction of the visual topology on $\partial_\infty X$. Although this is not one of the standard topologies on $\partial_c X$ ([CS15, CM18]), it is all that we will need in most of the thesis.

Theorem 4.2.6. *The map $\Phi \colon \partial_{\mathrm{cu}} X \to \partial_c^{\mathrm{vis}} X$ is a continuous surjection, whose fibres are precisely the $\sim$-equivalence classes in $\partial_{\mathrm{cu}} X$. Moreover, Φ descends to a homeomorphism $\overline{\Phi} \colon \partial_{\mathrm{cu}} X / \sim \to \partial_c^{\mathrm{vis}} X$.*

Proof. Surjectivity is immediate from Corollary 4.2.5. Lemma 4.1.4 shows that Φ is constant on $\sim$-equivalence classes. On the other hand, if γ and γ' are combinatorial rays with $\gamma(+\infty) \not\sim \gamma'(\infty)$, it is clear that the distance $d(\gamma(n), \gamma'(n))$ diverges as n goes to infinity. Thus, the fibres of Φ are precisely the components of $\partial_{\mathrm{cu}} X$.

We now prove continuity of Φ. Fix a vertex $v \in X$ and let $D = \dim X$. Given points $x, y \in \partial_{\mathrm{cu}} X$, let ρ_x and ρ_y be the CAT(0) rays from v to $\Phi(x)$ and $\Phi(y)$, respectively. Let r_x and r_y be combinatorial rays based at v and satisfying $d_{\mathrm{Haus}}(r_x, \rho_x) \leq D$, $d_{\mathrm{Haus}}(r_y, \rho_y) \leq D$, as provided by Proposition 2.1.8. The points $x' = r_x(+\infty)$ and $y' = r_y(+\infty)$ lie in the components $Z(x)$ and $Z(y)$, respectively. Theorem 4.1.3 moreover shows that r_x and r_y are C-lean for some constant $C > 0$.

Given an integer $n \geq 0$, we define an open neighbourhood $U_n(x)$ of x in $\partial_{\mathrm{cu}} X$ as follows. Pick halfspaces $\mathfrak{h}_3 \subseteq \mathfrak{h}_2 \subseteq \mathfrak{h}_1$ with $Z(x) \subseteq \mathfrak{h}_3$, $d(v, \mathfrak{h}_1) > n$, $\#\mathscr{W}(\mathfrak{h}_1^* | \mathfrak{h}_2) > C$ and $\#\mathscr{W}(\mathfrak{h}_2^* | \mathfrak{h}_3) > C$. Their existence is ensured by Lemma 4.2.3. The neighbourhood $U_n(x)$ is then the subset of $\partial_{\mathrm{cu}} X$ consisting of points that lie within $\mathfrak{h}_3$.

Claim. For all $y \in U_n(x)$, we have $d(r_x(n), r_y(n)) \leq 5C$.

As $d_{\mathrm{Haus}}(r_x, \rho_x) \leq D$ and $d_{\mathrm{Haus}}(r_y, \rho_y) \leq D$, the claim shows that, given any open neighbourhood V of $\Phi(x)$ in $\partial_c^{\mathrm{vis}} X$, we must have $\Phi(U_n(x)) \subseteq V$ for all sufficiently large n. We thus complete the proof of continuity of Φ by proving the claim.

Proof of Claim. Let the halfspaces $\mathfrak{h}_1$, $\mathfrak{h}_2$ and $\mathfrak{h}_3$ be as above. Since $y \in \mathfrak{h}_3$, Lemma 2.1.5 implies that $x', y' \in \mathfrak{h}_2$. Indeed, as $\#\mathscr{W}(\mathfrak{h}_2^* | \mathfrak{h}_3) > C$ and r_x is C-lean, no infinite chain of halfspaces can be transverse to both $\mathfrak{h}_2$ and $\mathfrak{h}_3$.

We set $q_x = r_x(n)$, $q_y = r_y(n)$ and $m = m(v, q_x, q_y)$. Note that the points q_x, q_y and m lie in $\mathfrak{h}_1^*$ as $d(v, \mathfrak{h}_1) > n$. Let $A_x \subseteq \mathscr{H}(m | q_x)$ and $A_y \subseteq \mathscr{H}(m | q_y)$ be the subsets of halfspaces containing $\mathfrak{h}_2$; set $a_x = \# A_x$ and $a_y = \# A_y$.

Each $\mathfrak{k} \in \mathscr{H}(m | q_x) \setminus A_x$ satisfies $\mathfrak{k}^* \cap \mathfrak{h}_2 \neq \emptyset$, but also $m \in \mathfrak{k}^* \cap \mathfrak{h}_2^*$, $q_x \in \mathfrak{k} \cap \mathfrak{h}_2^*$ and $x' \in \mathfrak{k} \cap \mathfrak{h}_2$. We conclude that each halfspace in $\mathscr{H}(m | q_x) \setminus A_x$ is transverse to $\mathfrak{h}_2$ and, similarly, to $\mathfrak{h}_1$. Since $\mathscr{W}(\mathfrak{h}_1^* | \mathfrak{h}_2) \subseteq \mathscr{W}(r_x)$ and $\#\mathscr{W}(\mathfrak{h}_1^* | \mathfrak{h}_2) > C$, the fact that r_x is C-lean implies that $d(m, q_x) - a_x \leq C$. Similarly, we obtain $d(m, q_y) - a_y \leq C$.

The sets A_x and A_y are transverse. As $A_x \subseteq \mathscr{W}(r_x)$, leanness of r_x implies that $\min\{a_x, a_y\} \leq C$. Since $d(m, q_x) = d(m, q_y)$ by construction, we have

$$|a_x - a_y| = |(d(m, q_x) - a_x) - (d(m, q_y) - a_y)| \leq C.$$

Hence $\max\{a_x, a_y\} \leq 2C$ and

$$d(q_x, q_y) = d(m, q_x) + d(m, q_y) \leq (a_x + C) + (a_y + C) \leq 5C. \quad \square$$

Finally, in order to prove that $\overline{\Phi}$ is a homeomorphism, let us first obtain the following property: Given $x, x_n \in \partial_{\mathrm{cu}} X$ and $y \in \partial X$ with $\Phi(x_n) \to \Phi(x)$ and $x_n \to y$, we must have $y \in Z(x)$.

Fix a basepoint $v \in X$. Let ρ_n and ρ be CAT(0) rays from v to x_n and x, respectively. Proposition 2.1.8 yields combinatorial rays r_n and r based at v, with $d_{\text{Haus}}(r_n, \rho_n) \leq D$ and $d_{\text{Haus}}(r, \rho) \leq D$. By Theorem 4.1.3, there exists $C \geq 0$ such that r is C-lean. Let $\mathfrak{h}_0 \supsetneq \mathfrak{h}_1 \supsetneq \ldots$ be a chain in $\sigma_x \setminus \sigma_v$ with $\bigcap \mathfrak{h}_n = Z(x)$, as provided by Lemma 4.2.3. Passing to a subchain if necessary, we can assume that $\#\mathscr{W}(\mathfrak{h}_k^*|\mathfrak{h}_{k+1}) > C$ for all $k \geq 0$

Given any $k \geq 0$, there exists $N(k) \geq 0$ such that r_n enters $\mathfrak{h}_k$ for all $n \geq N(k)$. Indeed, the rays ρ_n converge to ρ uniformly on compact sets and r_n, r are uniformly Hausdorff-close to ρ_n, ρ. We conclude that the points $z_n := r_n(+\infty)$ lie in $\mathfrak{h}_k$ for all $n \geq N(k)$. Since $x_n \in Z(z_n)$, we must also have $x_n \in \mathfrak{h}_{k-1}$. This follows from Lemma 2.1.5, as, by leanness of r, there is no infinite chain of halfspaces transverse to both $\mathfrak{h}_k$ and $\mathfrak{h}_{k-1}$. We conclude that $y = \lim x_n$ lies in $\bigcap \mathfrak{h}_n = Z(x)$, as required.

Now, since $\partial_\infty X$ is metrisable, so is $\partial_c^{\text{vis}} X$. The fact that $\overline{\Phi}$ is a homeomorphism will thus follow if we prove that $\overline{\Phi}^{-1}$ is sequentially continuous. Let us denote by $p\colon \partial_{\text{cu}} X \to \partial_{\text{cu}} X/\!\sim$ the quotient projection. Suppose for the sake of contradiction that points $x, x_n \in \partial_{\text{cu}} X$ are given so that $\Phi(x_n) \to \Phi(x)$, but $p(x_n) \not\to p(x)$. There exists an open neighbourhood V of $p(x)$ in $\partial_{\text{cu}} X/\!\sim$ such that no $p(x_n)$ lies in V; hence no x_n lies in $p^{-1}(V)$. Since ∂X is compact and metrisable, a subsequence x_{n_k} converges to a point $y \in \partial X$; as $p^{-1}(V)$ is open, we have $y \notin p^{-1}(V)$. However, the set $p^{-1}(V)$ is a union of $\sim$-equivalence classes and we have shown above that $y \sim x \in p^{-1}(V)$, a contradiction. □

Remark 4.2.7. Given points $x, x_n \in \partial_{\text{cu}} X$ satisfying $\Phi(x_n) \to \Phi(x)$ and $Z(x) = \{x\}$, we always have $x_n \to x$. Otherwise, compactness of ∂X would yield a subsequence x_{n_k} converging to a point $y \in \partial X$ different from x. We have shown during the proof of Theorem 4.2.6 (right after the claim) that $y \in Z(x) = \{x\}$.

When X is Gromov hyperbolic, we have $\partial_{\text{cu}} X = \partial X$ and $\partial_c^{\text{vis}} X$ coincides with the Gromov boundary $\partial_\infty X$. This case is worth highlighting:

Corollary 4.2.8. *If X is Gromov hyperbolic, the surjection $\Phi\colon \partial X \to \partial_\infty X$ descends to a homeomorphism $\overline{\Phi}\colon \partial X/\!\sim \to \partial_\infty X$.*

4.3 Contracting non-terminating ultrafilters.

Let X be irreducible throughout this section. The following definition is originally due to Nevo and Sageev (cf. Section 3.1 in [NS13]).

Definition 4.3.1. A point $x \in \partial X$ is *non-terminating* if $Z(x) = \{x\}$. Equivalently, the poset $(\sigma_x, \subseteq)$ does not have minimal elements.

We denote by $\partial_{\mathrm{nt}} X \subseteq \partial X$ the subset of non-terminating ultrafilters and set $\partial_{\mathrm{cnt}} X = \partial_{\mathrm{cu}} X \cap \partial_{\mathrm{nt}} X$. By Theorem 4.2.6, the restriction to $\partial_{\mathrm{cnt}} X$ of the map $\Phi\colon \partial_{\mathrm{cu}} X \to \partial_c X$ is injective. Therefore, we will identify the sets $\partial_{\mathrm{cnt}} X \subseteq \partial_{\mathrm{nt}} X$ and $\Phi(\partial_{\mathrm{cnt}} X) \subseteq \partial_c X$, even writing simply $\partial_{\mathrm{cnt}} X \subseteq \partial_c X$.

Remark 4.3.2. It follows from Theorem 4.2.6 and Remark 4.2.7 that $\partial_{\mathrm{cnt}} X$ inherits the same topology from $\partial_{\mathrm{cu}} X \subseteq \partial X$ and $\partial_c^{\mathrm{vis}} X$.

It will be useful to make the following observation.

Lemma 4.3.3. *We have* $\partial_{\mathrm{cnt}} X = \partial_{\mathrm{cu}} X \cap \partial_{\mathrm{reg}} X$.

Proof. We have $\partial_{\mathrm{cnt}} X \supseteq \partial_{\mathrm{cu}} X \cap \partial_{\mathrm{reg}} X$ since it is clear from Definition 2.2.11 that $\partial_{\mathrm{reg}} X \subseteq \partial_{\mathrm{nt}} X$. For the other inclusion, consider a point $x \in \partial_{\mathrm{cnt}} X$. Lemma 4.2.3 yields an infinite descending chain of halfspaces $\mathfrak{h}_0 \supsetneq \mathfrak{h}_1 \supsetneq \ldots$ such that the shores $S(\mathfrak{h}_n^*, \mathfrak{h}_{n+1})$ are finite and $\bigcap \mathfrak{h}_n = \{x\}$. In particular, for every $n \geq 0$, only finitely many hyperplanes are transverse to both $\mathfrak{h}_n$ and $\mathfrak{h}_{n+1}$. We are going to show that, for every $k \geq 0$, there exists $N > k$ such that the halfspaces $\mathfrak{h}_k^*$ and $\mathfrak{h}_N$ are strongly separated. This implies that $\mathfrak{h}_k^*$ and $\mathfrak{h}_n$ are strongly separated for every $n \geq N$, providing a strongly separated chain in σ_x and hence showing that x is regular.

Suppose for the sake of contradiction that, for some $k \geq 0$ and every $n > k$, there exists a hyperplane $\mathfrak{w}_n$ transverse to both $\mathfrak{h}_k$ and $\mathfrak{h}_n$. Note that each $\mathfrak{w}_n$ is in particular transverse to $\mathfrak{h}_k$ and $\mathfrak{h}_{k+1}$ and that there are only finitely many such hyperplanes. We conclude that $\mathfrak{w}_n = \mathfrak{w}$ for a hyperplane $\mathfrak{w}$ and infinitely many values of n. In particular, $\mathfrak{w}$ is transverse to all $\mathfrak{h}_n$ with $n \geq k$. It follows that there exists a point $y \in \partial X$ such that $y \in \mathfrak{h}_n$ for all $n \geq 0$ and $\mathfrak{w} \in \mathscr{W}(x|y)$, violating the fact that $\bigcap \mathfrak{h}_n = \{x\}$. □

Contracting non-terminating ultrafilters are plentiful, as they for instance arise in relation to *strongly contracting automorphisms* (see Definition 5.1.1, Proposition 5.1.3 and Lemma 5.2.1). When X is Gromov hyperbolic, $\partial_{\mathrm{cnt}} X$ is even co-meagre in $\partial_\infty X$ (Lemmas 8.0.1 and 8.1.5).

Chapter 5

From length spectra to boundary cross ratios.

The main result of this chapter is Theorem F, allowing us to translate length-spectrum rigidity questions into the problem of extending certain Möbius maps. We also give proofs of Theorems I and J.

In Section 5.1, we introduce *strongly contracting isometries* — a particular class of elements in automorphism groups of irreducible CAT(0) cube complexes. Their length functions are related to cross ratios via Theorem 5.1.6, which will act as our bridge between length spectra and cross ratios. The proof of Theorem F is then completed in Section 5.2. We will require a random-walk argument of [FLM18], in order to show that strongly contracting isometries are always 'sufficiently abundant'.

Finally, in Section 5.3, we consider the simplified context of properly discontinuous, cocompact actions. Assuming Theorems G and H, we anticipate a proof of Theorems I and J.

5.1 Strongly contracting isometries.

Let X be an irreducible, finite dimensional, locally finite CAT(0) cube complex throughout this section.

Definition 5.1.1. A cubical automorphism $g \in \mathrm{Aut}(X)$ is *strongly contracting* if there exist halfspaces $\mathfrak{h}_1$ and $\mathfrak{h}_2$ such that $g\mathfrak{h}_1 \subseteq \mathfrak{h}_2 \subseteq \mathfrak{h}_1$ and both pairs $(\mathfrak{h}_2, \mathfrak{h}_1^*)$ and $(g\mathfrak{h}_1, \mathfrak{h}_2^*)$ are strongly separated.

As motivation for our terminology, observe that strongly contracting automorphisms are contracting isometries (Lemma 6.2 in [CS11]).

Remark 5.1.2. Let $\mathfrak{h}_1$, $\mathfrak{h}_2$ and g be as in Definition 5.1.1. Given any $\mathfrak{w} \in \mathscr{W}(X)$, strong separation yields $n \in \mathbb{Z}$ and $i \in \{1,2\}$ such that $\mathfrak{w} \subseteq g^n\mathfrak{h}_i \cap g^{n+1}\mathfrak{h}_i^*$. In particular, for every $k \in \mathbb{Z} \setminus \{0\}$, the hyperplanes $\mathfrak{w}$ and $g^k\mathfrak{w}$ are distinct and not transverse. This shows that g acts stably without inversions and non-transversely, as defined in Section 2.2.1.

We now prove a sequence of results leading up to Theorem 5.1.6, which is the focal point of this section.

Proposition 5.1.3. *Let $g \in \mathrm{Aut}(X)$ be strongly contracting.*

(1) The automorphism g fixes exactly two points $g^\pm \in \partial_{\mathrm{cnt}}X$. They also lie in $\partial_{\mathrm{reg}}X$.

(2) For any $x \in \overline{X} \setminus \{g^\pm\}$, we have $g^n x \to g^+$ and $g^{-n}x \to g^-$ as $n \to +\infty$.

(3) Given $\mathfrak{j}_1, \mathfrak{j}_2 \in \mathscr{H}(g^-|g^+)$, there exists N such that $g^n\mathfrak{j}_1 \subseteq \mathfrak{j}_2$ for all $n \geq N$.

Proof. Let $\mathfrak{h}_1$ and $\mathfrak{h}_2$ be halfspaces as in Definition 5.1.1. By Lemma 2.2.12, there exists a unique point g^+ lying in $g^n\mathfrak{h}_1$ for all $n \geq 0$. Similarly, there exists a unique point g^- lying in $g^n\mathfrak{h}_1^*$ for all $n \leq 0$. By Remark 5.1.2 and Proposition 2.2.1, there exists a $\langle g \rangle$-invariant line $\gamma \subseteq X$. As γ must cross all hyperplanes bounding the halfspaces $g^n\mathfrak{h}_1$, the endpoints of γ must be $g^\pm$. Furthermore, Theorem 3.9 in [Gen16b] shows that γ is a contracting geodesic. We conclude that $g^\pm \in \partial_{\mathrm{cu}}X$, while it is clear that from their definition that $g^\pm \in \partial_{\mathrm{reg}}X$. Lemma 4.3.3 now yields $g^\pm \in \partial_{\mathrm{cnt}}X$.

It is clear that the points $g^\pm$ are fixed by g; their uniqueness will follow from part (2), thus completing part (1). In turn, part (2) follows from part (3), as every $x \in \overline{X} \setminus \{g^\pm\}$ lies in $g^n\mathfrak{h}_1 \cap g^{n+1}\mathfrak{h}_1^*$ for some $n \in \mathbb{Z}$.

Let thus $\mathfrak{w}_1$ and $\mathfrak{w}_2$ be the hyperplanes bounding $\mathfrak{j}_1$ and $\mathfrak{j}_2$. By Remark 5.1.2, there exist integers $n_1, n_2 \in \mathbb{Z}$ and indices $i_1, i_2 \in \{1,2\}$ such that $\mathfrak{w}_1 \subseteq g^{n_1}\mathfrak{h}_{i_1} \cap g^{n_1+1}\mathfrak{h}_{i_1}^*$ and $\mathfrak{w}_2 \subseteq g^{n_2}\mathfrak{h}_{i_2} \cap g^{n_2+1}\mathfrak{h}_{i_2}^*$. Since $\mathfrak{j}_1$ and $\mathfrak{j}_2$ lie in $\mathscr{H}(g^-|g^+)$, we have $\mathfrak{j}_1 \subseteq g^{n_1}\mathfrak{h}_{i_1}$ and $g^{n_2+1}\mathfrak{h}_{i_2} \subseteq \mathfrak{j}_2$. Hence $g^n\mathfrak{j}_1 \subseteq \mathfrak{j}_2$ for all $n \geq n_2 - n_1 + 2$. □

In relation to part (1) of Proposition 5.1.3, we will write $g_X^\pm \in \partial_{\mathrm{cnt}}X$ when it is necessary to specify the cube complex on which our groups are acting.

Proposition 5.1.4. *Let $g \in \mathrm{Aut}(X)$ be strongly contracting. A vertex $v \in X$ lies in the interval $I(g^-, g^+)$ if and only if $d(v, gv) = \ell_X(g)$.*

Proof. By Remark 5.1.2, the automorphism g acts stably without inversions and non-transversely. Proposition 2.2.1 then shows that $d(v, gv) = \ell_X(g)$ holds if and only if v lies in the convex subcomplex $\mathrm{Min}(g) \subseteq X$.

If a halfspace $\mathfrak{h}$ contains a semi-axis of g, then it also contains either g^+ or g^-. Hence $\mathfrak{h}$ contains an entire axis of g if and only if it contains both g^- and g^+ and this happens if and only if $\mathfrak{h}$ contains *all* axes of g. On the one hand, the intersection of all such halfspaces must coincide with the convex set $\mathrm{Min}(g)$; on the other, it must be $\mathrm{Hull}(\{g^\pm\}) \cap X = I(g^-, g^+) \cap X$. This concludes the proof. □

Proposition 5.1.5. *Let $g, h \in \mathrm{Aut}(X)$ be strongly contracting, with $\#\{g^\pm, h^\pm\} = 4$. For all sufficiently large values of $n > 0$, the automorphism $g^n h^n$ is strongly contracting. Furthermore, we have $(g^n h^n)^+ \to g^+$ and $(g^n h^n)^- \to h^-$ for $n \to +\infty$.*

Proof. Let $\mathfrak{h}_0 \supsetneq \mathfrak{h}_1 \supsetneq ...$ and $\mathfrak{k}_0 \supsetneq \mathfrak{k}_1 \supsetneq ...$ be strongly separated chains containing g^+ and h^-, respectively. Up to discarding finitely many halfspaces, we can assume that $\mathfrak{h}_0 \subseteq \mathfrak{k}_0^*$ and that g^- and h^+ lie in $\mathfrak{h}_0^* \cap \mathfrak{k}_0^*$. Since h^+ and g^- are also contained in strongly separated chains, we can find halfspaces $\mathfrak{j} \ni h^+$ and $\mathfrak{m} \ni g^-$ such that $\mathfrak{h}_0$, $\mathfrak{k}_0$, $\mathfrak{j}$ and $\mathfrak{m}$ are pairwise disjoint.

Given $k \geq 0$, part (3) of Proposition 5.1.3 provides $n_k \geq 0$ such that $h^n \mathfrak{k}_k^* \subseteq \mathfrak{j}$ and $g^n \mathfrak{m}^* \subseteq \mathfrak{h}_k$ for all $n \geq n_k$. Observe that we then have:

$$g^n h^n \mathfrak{k}_k^* \subseteq g^n \mathfrak{j} \subseteq g^n \mathfrak{m}^* \subseteq \mathfrak{h}_k.$$

In particular, $g^n h^n$ is strongly contracting for all $n \geq n_1$. In light of the construction of $(g^n h^n)^\pm$ in Proposition 5.1.3, this also shows that $(g^n h^n)^+$ lies in $\mathfrak{h}_k$ and $(g^n h^n)^-$ lies in $\mathfrak{k}_k$ for $n \geq n_k$. We conclude that $(g^n h^n)^+$ and $(g^n h^n)^-$ converge to g^+ and h^-, respectively, as k goes to infinity. □

The following should be compared to Théorème 2.2 in [Ota92] and Theorem 1 in Section 5 of [Kim01a] — both concerned with the context of simply connected, negatively curved Riemannian manifolds.

Theorem 5.1.6. *Let $g, h \in \mathrm{Aut}(X)$ be strongly contracting, with $\#\{g^\pm, h^\pm\} = 4$. There exists $N \geq 0$ such that every $n \geq N$ satisfies:*

$$\mathrm{cr}(g^-, h^-, g^+, h^+) = \ell_X(g^n) + \ell_X(h^n) - \ell_X(g^n h^n).$$

Proof. Let $\mathfrak{h}_0 \supsetneq \mathfrak{h}_1 \supsetneq ...$ and $\mathfrak{h}_0^* \supsetneq \mathfrak{h}_{-1}^* \supsetneq ...$ be strongly separated chains containing g^+ and g^-, respectively. Similarly define halfspaces $\mathfrak{k}_n$ for $h^\pm$. We can assume that $\mathfrak{h}_0^*$, $\mathfrak{h}_1$, $\mathfrak{k}_0^*$ and $\mathfrak{k}_1$ are pairwise disjoint.

Given $k \geq 1$, Proposition 5.1.5 yields $n_k \geq 0$ such that $(g^n h^n)^- \in \mathfrak{k}^*_{-k-1}$ and $(g^n h^n)^+ \in \mathfrak{h}_1$ for all $n \geq n_k$. Enlarging n_k if necessary, part (3) of Proposition 5.1.3 allows us to moreover assume that $g^n \mathfrak{h}_{-k-1} \subseteq \mathfrak{h}_k$ and $h^n \mathfrak{k}_{-k} \subseteq \mathfrak{k}_k$ for all $n \geq n_k$. Let $x_k \in \mathfrak{h}_{-k}$ be the only point in the bridge $B(\mathfrak{h}_{-k}, \mathfrak{h}^*_{-k-1})$ and $y_k \in \mathfrak{k}_{-k}$ the only point in the bridge $B(\mathfrak{k}_{-k}, \mathfrak{k}^*_{-k-1})$.

For $n \geq n_k$, the points h^- and $(g^n h^n)^-$ lie in $\mathfrak{k}^*_{-k-1}$, while h^+ and $(g^n h^n)^+$ lie in $\mathfrak{k}_{-k}$. Thus, the point y_k belongs to $I(h^-, h^+) \cap I((g^n h^n)^-, (g^n h^n)^+)$ and similarly $x_k \in I(g^-, g^+)$. Proposition 5.1.4 now implies that we have $\ell_X(g^n h^n) = d(y_k, g^n h^n y_k)$, $\ell_X(h^n) = d(y_k, h^n y_k)$ and $\ell_X(g^n) = d(x_k, g^n x_k)$.

Observe that $y_k \in \mathfrak{k}^*_{-k+1} \subseteq \mathfrak{k}^*_0 \subseteq \mathfrak{h}^*_1 \subseteq \mathfrak{h}^*_k$; hence $g^{-n} y_k \in g^{-n} \mathfrak{h}^*_k \subseteq \mathfrak{h}^*_{-k-1}$ since we are assuming that $g^n \mathfrak{h}_{-k-1} \subseteq \mathfrak{h}_k$. On the other hand, the point $h^n y_k$ lies in $h^n \mathfrak{k}_{-k} \subseteq \mathfrak{k}_k \subseteq \mathfrak{k}_1 \subseteq \mathfrak{h}_0 \subseteq \mathfrak{h}_{-k}$. Since x_k is a gate for $\mathfrak{h}_{-k}$ and $\mathfrak{h}^*_{-k-1}$, we conclude that $m(g^{-n} y_k, x_k, h^n y_k) = x_k$ for all $n \geq n_k$. This implies that $m(y_k, g^n x_k, g^n h^n y_k) = g^n x_k$. Hence, for all $n \geq n_k$:

$$d(y_k, g^n h^n y_k) = d(y_k, g^n x_k) + d(g^n x_k, g^n h^n y_k) = d(y_k, g^n x_k) + d(x_k, h^n y_k)$$

and we have:

$$\begin{aligned} \ell_X(g^n) + \ell_X(h^n) - \ell_X(g^n h^n) &= d(x_k, g^n x_k) + d(y_k, h^n y_k) - d(y_k, g^n h^n y_k) = \\ &= d(x_k, g^n x_k) + d(y_k, h^n y_k) - d(y_k, g^n x_k) - d(x_k, h^n y_k) = \\ &= \mathrm{cr}(x_k, y_k, g^n x_k, h^n y_k). \end{aligned}$$

Proposition 3.2.4 shows that $\mathrm{cr}(x_k, y_k, g^n x_k, h^n y_k) = \mathrm{cr}(g^-, h^-, g^+, h^+)$ for all sufficiently large $k \geq 0$ and every $n \geq n_k$. It follows that every sufficiently large $n \geq 0$ satisfies $\mathrm{cr}(g^-, h^-, g^+, h^+) = \ell_X(g^n) + \ell_X(h^n) - \ell_X(g^n h^n)$. □

5.2 Constructing a Möbius map.

Let X and Y be irreducible, finite dimensional, locally finite CAT(0) cube complexes. Let a group Γ act essentially and non-elementarily on X and Y.

In this section we prove Theorem F. In order to make notation less confusing, the constructed map will be denoted by ψ, rather than f as in the statement of the theorem. We start by finding strongly contracting automorphisms.

Lemma 5.2.1. *There exists $g \in \Gamma$ that simultaneously acts as a strongly contracting automorphism of X and Y.*

Before proving the lemma, a note is needed in case Γ is uncountable.

Remark 5.2.2. There exist countable subgroups $\Gamma_X, \Gamma_Y \leq \Gamma$ so that the actions $\Gamma_X \curvearrowright X$ and $\Gamma_Y \curvearrowright Y$ are essential and non-elementary. We can thus replace Γ with the countable subgroup $\langle \Gamma_X, \Gamma_Y \rangle$ in Lemma 5.2.1.

As to the existence of Γ_X, observe that X has only countably many hyperplanes. Thus, essentiality of the action $\Gamma \curvearrowright X$ is witnessed by countably many elements of Γ, which generate a countable subgroup $\Gamma_1 \leq \Gamma$ acting essentially on X. Furthermore, there exists a countable subgroup $\Gamma_2 \leq \Gamma$ such that $\Gamma_2 \curvearrowright X$ is non-elementary. We can even take Γ_2 to be a finitely generated free group and its existence follows e.g. from Propositions 6.4(2) and 5.2 in [Fio18]. Finally, we set $\Gamma_X = \langle \Gamma_1, \Gamma_2 \rangle$.

Proof of Lemma 5.2.1. By Remark 5.2.2, we can assume that Γ is countable. Set $\Omega = \Gamma^{\mathbb{N}}$ and let μ be a probability measure on Γ whose support generates Γ as a semigroup. Denote by $Z_n(\omega)$ the random walk on Γ induced by μ, where $\omega \in \Omega$, and let $\mathbb{P}$ be the corresponding measure on Ω. Lemmas 11.3 and 11.4 in [FLM18] and the proof of Theorem 11.5 in *loc. cit.* show that

$$\lim_{n \to +\infty} \tfrac{1}{n} \cdot \#\{k \leq n \mid Z_k(\omega) \text{ is strongly contracting in } X\} = 1$$

for $\mathbb{P}$-almost every $\omega \in \Omega$. Applying the same reasoning to the action $\Gamma \curvearrowright Y$, we conclude that there exists an element $g \in \Gamma$ that is strongly contracting in both X and Y. □

In the rest of the section, let us fix an element $g \in \Gamma$ that is simultaneously strongly contracting in X and Y. We denote by $g_X^+ \in \partial_{\mathrm{reg}} X$ and $g_Y^+ \in \partial_{\mathrm{reg}} Y$ the points introduced in Proposition 5.1.3. Observe that hgh^{-1} is strongly contracting for each $h \in \Gamma$ and $(hgh^{-1})_X^+ = h \cdot g_X^+$.

Proposition 5.2.3. *Assume that the length functions ℓ_X and ℓ_Y coincide. An element $k \in \Gamma$ fixes $g_X^+ \in \partial_{\mathrm{reg}} X$ if and only if it fixes $g_Y^+ \in \partial_{\mathrm{reg}} Y$.*

Proof. Consider the sets $U_1 = \{g_X^\pm\} \cup \{k \cdot g_X^\pm\}$ and $U_2 = \{g_Y^\pm\} \cup \{k \cdot g_Y^\pm\}$. Lemma 2.2.13 guarantees the existence of $\gamma_1, \gamma_2 \in \Gamma$ with $\gamma_1 U_1 \cap U_1 = \emptyset$ and $\gamma_2 U_2 \cap (U_2 \cup \gamma_1^{-1} U_2) = \emptyset$. If $\gamma_2 g_X^+ \notin U_1$, we set $\gamma = \gamma_2$, otherwise $\gamma = \gamma_1 \gamma_2$. In any case, we have $\gamma g_X^+ \notin U_1$ and $\gamma g_Y^+ \notin U_2$. The element $h_1 = \gamma g \gamma^{-1}$ is strongly contracting in both X and Y; moreover, $(h_1)_X^+ \notin U_1$ and $(h_1)_Y^+ \notin U_2$. Taking $h_2 = h_1^n g h_1^{-n}$ for a sufficiently large integer $n \geq 0$, part (2) of Proposition 5.1.3 guarantees that $\{(h_2)_X^\pm\} \cap U_1 = \emptyset$ and $\{(h_2)_Y^\pm\} \cap U_2 = \emptyset$.

Now, we consider the elements $a_n = h_2^{-n}(kg^{-n}k^{-1})$ for all $n \geq 0$. Proposition 5.1.5 shows that $(a_n)_X^+ \to (h_2)_X^-$ and $(a_n)_X^- \to k \cdot g_X^+$. Proposition 3.2.4 then implies that

$$\operatorname{cr}\left((a_n)_X^-, g_X^-, (a_n)_X^+, g_X^+\right) \longrightarrow \operatorname{cr}\left(k \cdot g_X^+, g_X^-, (h_2)_X^-, g_X^+\right),$$

as the points g_X^-, g_X^+ and $(h_2)_X^-$ are pairwise distinct. Observe that the limit is $-\infty$ if and only if k fixes g_X^+. The same observations hold within Y.

Our choice of h_2 implies that $(a_n)_X^+ \neq (h_2)_X^-$ and $(a_n)_X^- \neq k \cdot g_X^+$ for all sufficiently large $n \geq 0$. In particular, the sequences $(a_n)_X^+$ and $(a_n)_X^-$ are not eventually constant and the four points $(a_n)_X^-$, g_X^-, $(a_n)_X^+$ and g_X^+ are pairwise distinct for all large $n \geq 0$. As the same holds in Y, Theorem 5.1.6 and the assumption on length functions yield:

$$\operatorname{cr}\left((a_n)_X^-, g_X^-, (a_n)_X^+, g_X^+\right) = \operatorname{cr}\left((a_n)_Y^-, g_Y^-, (a_n)_Y^+, g_Y^+\right)$$

for all large n. We conclude that the limits coincide:

$$\operatorname{cr}\left(k \cdot g_X^+, g_X^-, (h_2)_X^-, g_X^+\right) = \operatorname{cr}\left(k \cdot g_Y^+, g_Y^-, (h_2)_Y^-, g_Y^+\right).$$

Hence $k \cdot g_X^+ = g_X^+$ if and only if $k \cdot g_Y^+ = g_Y^+$. □

We now consider the orbits $\mathcal{A} = \Gamma \cdot g_X^+ \subseteq \partial_{\mathrm{reg}} X$ and $\mathcal{B} = \Gamma \cdot g_Y^+ \subseteq \partial_{\mathrm{reg}} Y$. By Proposition 5.2.3, there exists a unique Γ-equivariant bijection $\psi\colon \mathcal{A} \to \mathcal{B}$ with $\psi(g_X^+) = g_Y^+$. Observe that every point of $\mathcal{A}$ is of the form k_X^+ for some element $k \in \Gamma$ conjugate to g; moreover, $\psi(k_X^+) = k_Y^+$.

Remark 5.2.4. If there exist conjugates $a, b \in \Gamma$ of g such that $a_X^- = b_X^-$, but $a_X^+ \neq b_X^+$, we replace g with g^{-1} in the construction of the set $\mathcal{A}$. This way:

- either $a_X^- \neq b_X^-$ for all conjugates a, b of g that satisfy $a_X^+ \neq b_X^+$,
- or, for every $x \in \mathcal{A}$, there exist conjugates $a, b \in \Gamma$ of g such that $x = a_X^+ = b_X^+$ and $a_X^- \neq b_X^-$.

The need for this property will become apparent in the proof of Lemma 5.2.5.

Lemma 5.2.5. *If $x, y \in \mathcal{A}$ are distinct, there exist elements $g_n \in \Gamma$ satisfying:*

$$(g_n)_X^- \to x, \quad (g_n)_Y^- \to \psi(x); \qquad (g_n)_X^+ \to y, \quad (g_n)_Y^+ \to \psi(y).$$

Proof. Let $a, b \in \Gamma$ be conjugates of g such that $x = a_X^+$ and $y = b_X^+$; we also have $\psi(x) = a_Y^+$ and $\psi(y) = b_Y^+$. Replacing a if necessary, Remark 5.2.4 enables us to ensure that $a_X^- \neq b_X^-$. If $a_X^- = b_X^+$, we can take all g_n to coincide with a^{-1}; similarly, if $b_X^- = a_X^+$, we set $g_n = b$. If none of these equalities is verified, we have $\#\{a_X^\pm, b_X^\pm\} = 4$ and, by Proposition 5.1.5, it suffices to take $g_n = b^n a^{-n}$. □

Proposition 5.2.6. *If ℓ_X and ℓ_Y coincide, the map ψ is Möbius.*

Proof. It suffices to show that ψ preserves cross ratios of 4-tuples (x, y, z, w) with pairwise-distinct coordinates. All other 4-tuples have cross ratio 0 or $\pm\infty$, depending on which coordinates coincide. Lemma 5.2.5 yields elements $g_n, h_n \in \Gamma$ with:

$$(g_n)_X^- \to x, \quad (g_n)_Y^- \to \psi(x); \qquad (h_n)_X^- \to y, \quad (h_n)_Y^- \to \psi(y);$$

$$(g_n)_X^+ \to z, \quad (g_n)_Y^+ \to \psi(z); \qquad (h_n)_X^+ \to w, \quad (h_n)_Y^+ \to \psi(w).$$

Theorem 5.1.6 implies that:

$$\mathrm{cr}((g_n)_Y^-, (h_n)_Y^-, (g_n)_Y^+, (h_n)_Y^+) = \mathrm{cr}((g_n)_X^-, (h_n)_X^-, (g_n)_X^+, (h_n)_X^+)$$

for all $n \geq 0$. Proposition 3.2.4 now shows that the two sides of the above equality converge to $\mathrm{cr}(\psi(x), \psi(y), \psi(z), \psi(w))$ and $\mathrm{cr}(x, y, z, w)$. □

This concludes the proof of Theorem F.

Remark 5.2.7. When group actions are properly discontinuous and cocompact, the discussion in this section can be considerably simplified. For instance, Proposition 5.2.3 does not require the assumption on length functions (cf. Lemma 5.3.1) and Remark 5.2.4 also becomes unnecessary.

5.3 Transferring cross ratios.

In this section, we rely on Theorems G and H in order to anticipate a proof of Theorems I and J. We also recall a few well-known facts on the dynamics of groups acting on their boundary. The latter will be heavily needed in Chapter 8.

Given a finitely generated group Γ, its contracting boundary $\partial_c\Gamma$ was introduced in [Cor17, CM18]; contrary to these references, we will not endow $\partial_c\Gamma$ with any topology.

Fixing a word metric on Γ, we say that $g \in \Gamma$ is *Morse* if $n \mapsto g^n$ is a Morse quasi-geodesic in Γ. This notion is independent of the chosen word metric. Every Morse element $g \in \Gamma$ fixes exactly two points $g^{\pm\infty} \in \partial_c\Gamma$ (see e.g. Definition 9.1, Theorem 2.2 and Theorem 9.4 in [CM18]).

Given a proper cocompact action $\Gamma \curvearrowright X$, the Milnor-Schwarz lemma shows that orbit maps $\Gamma \to X$ are Γ-equivariant quasi-isometries. They are at finite distance from each other and all yield the same Γ-equivariant bijection $o_X \colon \partial_c\Gamma \to \partial_c X$. This follows from either Proposition 4.2 in [Cor17] or Corollary 6.2 in [CM18].

Now assume that X is an irreducible CAT(0) cube complex and that $g \in \Gamma$ is strongly contracting. We can consider the points $g^{\pm} \in \partial_{\mathrm{cnt}} X \subseteq \partial_c X$ introduced in Proposition 5.1.3; here we are identifying $\partial_{\mathrm{cnt}} X$ with a subset of $\partial_c X$ as described in Section 4.3. In this case, g is a Morse element in Γ by Theorem 4.1.3. In particular, we have $\partial_c \Gamma \neq \emptyset$.

Lemma 5.3.1. *Let $\Gamma \curvearrowright X$ be proper cocompact and X irreducible. For every strongly contracting element $g \in \Gamma$, we have $o_X(g^{+\infty}) = g^+$ and $o_X(g^{-\infty}) = g^-$.*

Proof. The discussion above already shows that $o_X(\{g^{\pm\infty}\}) = \{g^{\pm}\}$ as these are the only two Γ-fixed points in $\partial_c \Gamma$ and $\partial_c X$, respectively. Let $\partial_c^{\mathcal{FQ}} X$ denote the contracting boundary of X, endowed with the Cashen-Mackay topology (Definition 5.4 in [CM18]). By Theorem 9.4 and Corollary 6.2 in *op. cit.*, the element g acts on $\partial_c^{\mathcal{FQ}} X$ with north-south dynamics. Note that g also acts on $\partial_c^{\mathrm{vis}} X$ with north-south dynamics, by Proposition 5.1.3 and Theorem 4.2.6. Since the identity map $\partial_c^{\mathcal{FQ}} X \to \partial_c^{\mathrm{vis}} X$ is continuous (Section 7 in [CM18]), we conclude that the attracting/repelling fixed points of g are the same for the two topologies. □

If Γ is Gromov hyperbolic, every infinite-order element $g \in \Gamma$ is Morse and acts on the Gromov boundary $\partial_\infty \Gamma$ with north-south dynamics.

Let now Γ act properly and cocompactly on irreducible CAT(0) cube complexes X and Y. In the following discussion, we do not require Γ to be hyperbolic, but we do assume that Γ is not virtually cyclic. In particular, $\partial_c \Gamma$ is infinite.

We would like to exploit the map Φ introduced in Section 4.2 to transfer the cross ratio on ∂X to a Γ-invariant cross ratio on $\partial_c X$ and $\partial_c \Gamma$. To this end, we will need the following notion.

Definition 5.3.2. A subset $\mathcal{A} \subseteq \partial_{\mathrm{cu}} X$ is a *section* (of the map Φ) if $\mathcal{A}$ intersects each fibre of Φ at exactly one point. In particular, $\partial_{\mathrm{cnt}} X \subseteq \mathcal{A}$.

Remark 5.3.3. If $\mathcal{A}$ is a section of Φ and $x, y \in \mathcal{A}$ are distinct points, Lemma 4.2.2 shows that x and y have finite Gromov product. In particular, if $x, y, z, w \in \mathcal{A}$ are pairwise distinct, the median $m(x, y, z)$ lies in X and the cross ratio $\mathrm{cr}(x, y, z, w)$ is well-defined and finite (cf. Lemma 3.1.1).

Note that it is always possible to find a Γ-invariant section $\mathcal{A} \subseteq \partial_{\mathrm{cu}} X$, up to subdividing X. Indeed, since every component of $\partial_{\mathrm{cu}} X$ is finite by Remark 4.2.4, we can consider the first cubical subdivision X' and pick the median barycentre (cf. Section 2.1.5) of each component of $\partial_{\mathrm{cu}} X'$. The resulting subset $\mathcal{A} \subseteq \partial_{\mathrm{cu}} X'$ is a

Γ-invariant section of $\Phi' \colon \partial_{\mathrm{cu}} X' \to \partial_c X' \simeq \partial_c X$. We assign length 1 to every edge of X'; in particular, the inclusion $X \hookrightarrow X'$ is a homothety doubling distances.

Restricting the cross ratio on $\partial X'$ to the set $\mathcal{A}$ and identifying $\mathcal{A} \simeq \partial_c \Gamma$ via the composition $o_X^{-1} \circ \Phi'$, we obtain an invariant cross ratio

$$\mathrm{cr}_X \colon \partial_c \Gamma^{(4)} \longrightarrow \mathbb{Z}.$$

By Remark 5.3.3, the value $\mathrm{cr}_X(x, y, z, w)$ is well-defined and finite as soon as the points $x, y, z, w \in \partial_c \Gamma$ are pairwise distinct. We can extend cr_X to 4-tuples (x, y, z, w) with $\#\{x, y, z, w\} \leq 3$ as long as no three of the four points coincide. If $x \neq y$ and $x \neq z$, we set $\mathrm{cr}_X(x, x, y, z) = 0$ and $\mathrm{cr}_X(x, y, x, z) = +\infty$, while the other values of cr_X can be recovered using its symmetries. We refer to cr_X as the *cubical cross ratio* on $\partial_c \Gamma$ associated to the cubulation $\Gamma \curvearrowright X$.

We can endow $\partial_c \Gamma$ with the pull-back of the topology of $\partial_c^{\mathrm{vis}} X$, or with one of the topologies of [Cor17, CM18], but cr_X will rarely be continuous on its entire domain. Indeed, cr_X takes integer values, whereas $\partial_c \Gamma$ can be connected (for instance when Γ is a one-ended hyperbolic group). Nevertheless, we have the following:

Proposition 5.3.4. *Endowing $\partial_c \Gamma$ with any of the three topologies above, cr_X is continuous at every 4-tuple with coordinates in $\partial_{\mathrm{cnt}} X \subseteq \partial_c X$.*

Proof. It suffices to consider $\partial_c \Gamma$ endowed with the pull-back of the topology of $\partial_c^{\mathrm{vis}} X$, as this is the coarsest of the three. In this case, the result follows from Proposition 3.2.4 and Remark 4.3.2. □

Proof of Theorem I. Part (1) is immediate from the previous discussion. Let now $\Gamma \curvearrowright X$ and $\Gamma \curvearrowright Y$ be two cubulations satisfying the hypotheses of part (2). Note that X and Y are essential, as they have no free faces, and the two actions are non-elementary by Lemma 2.2.5. Lemma 5.2.1 yields an element $g \in \Gamma$ that acts as a strongly contracting automorphism on both X and Y. The Γ-equivariant bijection $f = o_Y \circ o_X^{-1} \colon \partial_c X \to \partial_c Y$ takes $g_X^+ \in \partial_{\mathrm{cnt}} X$ to $g_Y^+ \in \partial_{\mathrm{cnt}} Y$ by Lemma 5.3.1. Setting $\mathcal{A} = \Gamma \cdot g_X^+ \subseteq \partial_{\mathrm{cnt}} X$, we have $f(\mathcal{A}) = \Gamma \cdot g_Y^+ \subseteq \partial_{\mathrm{cnt}} Y$ and cross ratios of these points are preserved. We moreover have $\mathcal{A} \subseteq \partial_{\mathrm{reg}} X$ and $f(\mathcal{A}) \subseteq \partial_{\mathrm{reg}} Y$ by Lemma 4.3.3. We conclude by Theorem G □

Proof of Theorem J. We begin by observing that X and Y are irreducible by Remark 2.2.7. Part (1) then follows from the above discussion and Proposition 5.3.4, if we show that $\mathcal{C} := \partial_{\mathrm{nt}} X$ is co-meagre in $\partial_\infty X$. This will be proved in Corollary 8.1.6.

Regarding part (2), suppose that X is also hyperplane-essential, let $\Gamma \curvearrowright Y$ be another cubulation with the same properties and consider the Γ-equivariant homeomorphism $f = o_Y \circ o_X^{-1} \colon \partial_\infty X \to \partial_\infty Y$. Suppose that cr_X and cr_Y coincide on $\mathcal{D}^{(4)} \subseteq \partial_\infty \Gamma^{(4)}$, for a co-meagre subset $\mathcal{D} \subseteq \partial_\infty \Gamma$. The set $\mathcal{D} \cap \partial_{\mathrm{nt}} X \cap f^{-1}(\partial_{\mathrm{nt}} Y)$ is co-meagre and so is the intersection of all its (countably many[1]) Γ-translates, which we denote by Ω. By Baire's theorem, Ω is nonempty and we conclude by applying Theorem H. $\square$

[1]Note that here Γ is finitely generated, hence countable.

Chapter 6

Extending Möbius maps I: the globally-defined case.

This chapter is devoted to the proof of Theorem E. Throughout, we consider CAT(0) cube complexes X, Y and a Möbius bijection $f\colon \partial X \to \partial Y$.

6.1 Extremal vertices and straight geodesics.

In this section, we introduce the following two notions.

Definition 6.1.1. A vertex $v \in X$ is *extremal* if there exists an edge $e \subseteq X$ incident to v such that every other edge incident to v spans a square with e.

Equivalently, $\mathrm{lk}(v)$ is a cone over one of its vertices. This happens if and only if a neighbourhood of v splits as $[0,1) \times N$ for a subcomplex $N \subseteq X$.

Definition 6.1.2. A geodesic $\gamma \subseteq X$ is *straight* if no two hyperplanes in $\mathscr{W}(\gamma)$ are transverse.

Our interest in cube complexes with no extremal vertices is motivated by the following straightforward observation.

Lemma 6.1.3. *If X has no extremal vertices, every edge can be extended to a straight (bi-infinite) line.*

We say that X is *complete* if there is no infinite ascending chain of cubes (cf. [Lea13]). Finite dimensional cube complexes are always complete.

Remark 6.1.4. If X is complete and it has no free faces, then X has no extremal vertices. Indeed, consider a vertex $v \in X$ contained in an edge e. Since X is complete,

there exists a maximal cube $c_1 \subseteq X$ containing e. Let $c \subseteq c_1$ be the face such that $v \in c$ and $c_1 \simeq c \times e$. Since c is not a free face, there exists a maximal cube $c_2 \subseteq X$ with $c_1 \cap c_2 = c$. Let $e' \subseteq c_2$ be an edge such that $v \in e'$ and $e' \not\subseteq c$. The edges e and e' do not span a square or e' and c_1 would span a cube properly containing c_1. Hence v is not an extremal vertex.

Nevertheless, the reader will realise that CAT(0) cube complexes with no extremal vertices are much more common than cube complexes with no free faces. For instance, the (universal cover of) the Davis complex [Dav08] associated to a right-angled Coxeter group G often has free faces, but it only has extremal vertices when $G \simeq \mathbb{Z}/2\mathbb{Z} \times H$ for some parabolic subgroup $H < G$.

We will need the following slight tweak of Proposition 3.3.4.

Lemma 6.1.5. *Let X have no extremal vertices. Given points $x_1, x_2, y \in \partial X$ with $m(x_1, x_2, y) \in X$, the condition $x_1 \operatorname{op}_y x_2$ fails if and only if there exists a point $z \in \partial X$ such that $\operatorname{crt}(x_1, x_2, y, z) = [\![a : b : c]\!]$ and $a < \min\{b, c\} < +\infty$.*

Proof. Setting $m = m(x_1, x_2, y)$, we have:

$$\operatorname{crt}_m(x_1, x_2, y, z) = [\![(y \cdot z)_m : (x_2 \cdot z)_m : (x_1 \cdot z)_m]\!].$$

If $x_1 \operatorname{op}_y x_2$ and $z \in \partial X$, Remark 3.3.3 yields $\min\{(x_2 \cdot z)_m, (x_1 \cdot z)_m\} = 0$ and we cannot have $(y \cdot z)_m < 0$.

Consider instead the case when x_1 and x_2 are not opposite with respect to y. There exist transverse hyperplanes $\mathfrak{w}_i \in \mathscr{W}(m|x_i)$ adjacent to m. Denote by $m' \in X$ the vertex with $\mathscr{W}(m|m') = \{\mathfrak{w}_1, \mathfrak{w}_2\}$. By Lemma 6.1.3, there exists a straight ray r such that $r(1) = m'$ and $\mathscr{W}(r(0)|r(1)) = \{\mathfrak{w}_1\}$; we set $z = r(+\infty)$. Let γ and γ_2 be rays based at $r(0)$ satisfying $\gamma(1) = m$, $\gamma(+\infty) = y$ and $\gamma_2(+\infty) = x_2$. As $r(0)$, y and x_2 are all on the same side of $\mathfrak{w}_1$, Lemma 2.1.1 implies that the unions $\gamma \cup r$ and $\gamma_2 \cup r$ are lines. Hence $(y \cdot z)_m = 0$ and $(x_2 \cdot z)_m = 1$. Since $\mathfrak{w}_1 \in \mathscr{W}(m|x_1, z)$, we also have $(x_1 \cdot z)_m \geq 1$. We conclude that $\operatorname{crt}(x_1, x_2, y, z) = [\![0 : 1 : c]\!]$ with $c \geq 1$. □

Recall that we are considering a Möbius bijection $f \colon \partial X \to \partial Y$. Again, the following should be compared to Corollary 3.3.6, though more care is required here.

Corollary 6.1.6. *Suppose that X and Y have no extremal vertices. Given points $x, y, z \in \partial X$, we have $x \operatorname{op}_z y$ if and only if $f(x) \operatorname{op}_{f(z)} f(y)$.*

Proof. Fix a base vertex $v \in X$ and assume that $x \operatorname{op}_z y$. In particular, the median $m(x, y, z)$ lies in X. By Lemma 3.1.1, the latter is equivalent to the Gromov products $(x \cdot y)_v$, $(y \cdot z)_v$ and $(z \cdot x)_v$ being all finite. In other words, the 4-tuples (x, x, y, y), (y, y, z, z) and (z, z, x, x) all lie in $\mathscr{A}(X)$. As f is Möbius, it takes these 4-tuples into $\mathscr{A}(Y)$ and we must have $m(f(x), f(y), f(z)) \in Y$.

Now, if we did not have $f(x) \operatorname{op}_{f(z)} f(y)$, Lemma 6.1.5 would yield $w \in \partial Y$ with $\operatorname{crt}(f(x), f(y), f(z), w) = [\![a : b : c]\!]$ and $a < \min\{b, c\} < +\infty$. In particular $(f(x), f(y), f(z), w) \in \mathscr{A}(X)$ and hence $\operatorname{crt}(x, y, z, f^{-1}(w)) = [\![a : b : c]\!]$, contradicting Lemma 6.1.5. This shows that $x \operatorname{op}_z y \Rightarrow f(x) \operatorname{op}_{f(z)} f(y)$ and the converse implication follows by considering $f^{-1} \colon B \to A$. □

6.2 Straight points.

Assume that X and Y have no extremal vertices throughout this section. Corollary 6.1.6 then guarantees that our Möbius bijection f is op-preserving.

In order to apply the results of Section 3.4, we need to ensure that the hypotheses of Corollary 3.4.4 are satisfied. To this purpose, will consider the following class of boundary points.

Definition 6.2.1. A point $x \in \partial X$ is *straight* if there exists a straight ray $r \subseteq X$ with $r(+\infty) = x$; equivalently, x is an endpoint of a straight line. We denote by $\partial_s X \subseteq \partial X$ the set of straight boundary points.

Observe that two points $x, y \in \partial X$ are endpoints of a straight line γ if and only if the interval $I(x, y) \cap X$ is isomorphic to $\mathbb{R}$. Indeed, $I(x, y) \cap X$ coincides with γ in this case. The following result characterises such situations in a similar way to Lemma 6.1.5.

Lemma 6.2.2. *Two points $x, y \in \partial X$ are endpoints of a straight line if and only if both the following are verified:*

(1) $I(x, y) \cap X \neq \emptyset$;

(2) for no $z, w \in \partial X$ we have $\operatorname{crt}(x, y, z, w) = [\![a : b : c]\!]$ and $a < \min\{b, c\} < +\infty$.

Proof. We begin by assuming that x and y are endpoints of a straight line γ. Condition (1) is clearly satisfied and we are now going to prove condition (2).

Given $z, w \in \partial X$ with $m(x, y, z) \in X$, we have $x \operatorname{op}_z y$ and Lemma 6.1.5 shows that we cannot have $\operatorname{crt}(x, y, z, w) = [\![a : b : c]\!]$ and $a < \min\{b, c\} < +\infty$. The same

happens if $m(x,y,w) \in X$, as can be observed by simply swapping z and w. We are left to examine the situation where $m(x,y,z)$ and $m(x,y,w)$ lie in the boundary; note that both medians must then belong to the set $\{x,y\}$. We can assume that $m(x,y,z) \neq m(x,y,w)$, as otherwise the first coordinate of $\mathrm{crt}(x,y,z,w)$ is infinite. If $m(x,y,z) = x$ and $m(x,y,w) = y$, we have $\mathrm{crt}(x,y,z,w) = [\![* : \infty : 0]\!]$; otherwise, $\mathrm{crt}(x,y,z,w) = [\![* : 0 : \infty]\!]$. In all cases $\mathrm{crt}(x,y,z,w)$ is not of the form $[\![a : b : c]\!]$ with $a < \min\{b,c\} < +\infty$ and condition (2) is satisfied.

We now assume that $I(x,y) \cap X \neq \emptyset$, but x and y are not endpoints of a straight line. We will show that condition (2) fails. The intersection $I(x,y) \cap X$ cannot be one-dimensional or it would be isomorphic to $\mathbb{R}$. Hence $I(x,y) \cap X$ contains a square s; denote by $\mathfrak{w}$ and $\mathfrak{w}'$ its hyperplanes. Let v_x and v_y be the vertices of s such that $\{\mathfrak{w}, \mathfrak{w}'\}$ is disjoint from $\mathscr{W}(x|v_x)$ and $\mathscr{W}(y|v_y)$. Lemma 6.1.3 shows that there exist straight rays r_x and r_y, based at v_x and v_y respectively, such that their first crossed hyperplane is $\mathfrak{w}$. We set $z = r_x(+\infty)$ and $w = r_y(+\infty)$. Lemma 2.1.1 implies that $(z \cdot w)_{v_x} = 0$, $(x \cdot z)_{v_x} = 0$ and $(y \cdot w)_{v_x} = 1$. Moreover, $(y \cdot z)_{v_x} \geq 1$ as $\mathfrak{w}$ separates v_x from y and z. We conclude that $\mathrm{crt}(x,y,z,w) = [\![0 : 1 : c]\!]$ with $c = (x \cdot w)_{v_x} + (y \cdot z)_{v_x} \geq 1$. □

Proposition 6.2.3. *(1) We have $x \in \partial_s X$ if and only if $f(x) \in \partial_s Y$.*

(2) If $x, y \in \partial X$ are endpoints of a straight line, so are the points $f(x)$ and $f(y)$.

Proof. As a boundary point is straight if and only if it is an endpoint of a straight line, part (1) follows from part (2). If $x, y \in \partial_s X$ are endpoints of a straight line $\gamma \subseteq X$, we have $I(x,y) \cap X \neq \emptyset$. By part (2) of Lemma 3.1.1, this is equivalent to the fact that (x,x,y,y) lies in $\mathscr{A}(X)$. We conclude that $I(f(x), f(y)) \cap Y \neq \emptyset$.

Now, if $f(x)$ and $f(y)$ were not endpoints of a straight line, Lemma 6.2.2 would yield points $z, w \in \partial Y$ with $\mathrm{crt}(f(x), f(y), z, w) = [\![a : b : c]\!]$ and $a < \min\{b,c\} < +\infty$. However, $\mathrm{crt}(x, y, f^{-1}(z), f^{-1}(w))$ would then have the same form, contradicting Lemma 6.2.2. □

The next result is our main motivation for considering straight points.

Lemma 6.2.4. *Consider $x \in \partial_s X$ and a vertex $v \in X$. Given $y, z \in \partial_s X$ with $(y \cdot z)_v < +\infty$, at least one of the Gromov products $(x \cdot y)_v$, $(x \cdot z)_v$ is finite.*

Proof. Let r_x, r_y and r_z be straight rays representing x, y and z, respectively. As the symmetric difference $\mathscr{W}(r_x) \triangle \mathscr{W}(v|x)$ is contained in $\mathscr{W}(r_x(0)|v)$, the intersection $\mathcal{U}_x = \mathscr{W}(r_x) \cap \mathscr{W}(v|x)$ is cofinite in $\mathscr{W}(v|x)$ and does not contain transverse hyperplanes. The same holds for $\mathcal{U}_y = \mathscr{W}(r_y) \cap \mathscr{W}(v|y)$ and $\mathcal{U}_z = \mathscr{W}(r_z) \cap \mathscr{W}(v|z)$.

If we had $(x \cdot y)_v = (x \cdot z)_v = +\infty$, the set $\mathscr{W}(v|x)$ would have infinite intersection with both $\mathscr{W}(v|y)$ and $\mathscr{W}(v|z)$. In particular, both $\mathcal{U}_x \cap \mathcal{U}_y$ and $\mathcal{U}_x \cap \mathcal{U}_z$ would be infinite. As any hyperplane separating two elements of $\mathscr{W}(r_y)$ must lie in $\mathscr{W}(r_y)$, the intersections $\mathcal{U}_x \cap \mathcal{U}_y$ and $\mathcal{U}_x \cap \mathcal{U}_z$ would then be cofinite in $\mathcal{U}_x$. Hence $\mathcal{U}_y \cap \mathcal{U}_z$ would be infinite, contradicting the fact that $(y \cdot z)_v < +\infty$. □

We employ the notation $m^f(\cdot,\cdot,\cdot)$ as in Section 3.4.

Corollary 6.2.5. *If $x_1, x_2, x, y_1, y_2, y \in \partial_s X$ and $x_1 \operatorname{op}_x x_2$, $y_1 \operatorname{op}_y y_2$, we have:*

$$d(m(x_1, x_2, x), m(y_1, y_2, y)) = d(m^f(x_1, x_2, x), m^f(y_1, y_2, y)).$$

Proof. We only need to verify the hypotheses of Corollary 3.4.4. To this end, let u and v be distinct elements of the set $\{y_1, y_2, y\}$ and fix a basepoint $p \in X$. As the Gromov products $(x_1 \cdot x_2)_p$, $(x_1 \cdot x)_p$ and $(x_2 \cdot x)_p$ are all finite, Lemma 6.2.4 shows that $(x_1, x_2, x, u) \in \mathscr{A}(X)$. If (x_1, x_2, u, v) does not lie in $\mathscr{A}(X)$, we can assume, up to permuting the points, that either $(x_1 \cdot u)_p = (x_2 \cdot u)_p = +\infty$ or $(x_1 \cdot u)_p = (x_1 \cdot v)_p = +\infty$. As $(x_1 \cdot x_2)_p$ and $(u \cdot v)_p$ are finite, both situations are ruled out by Lemma 6.2.4. □

6.3 Skinny vertices.

Let again X have no extremal vertices and let us assume that $\dim X \geq 1$. We now address the problem of which vertices $v \in X$ can be represented as median of triples such as those in Corollary 6.2.5.

We say that a vertex $v \in X$ is *skinny* if $\deg(v) = 2$. Skinny vertices are exactly those that are 'invisible from the boundary', as we now describe.

Lemma 6.3.1. *For a vertex $v \in X$, the following are equivalent:*

(1) v is not skinny;

(2) there exist $x, y, z \in \partial_s X$ such that $m(x, y, z) = v$ and $x \operatorname{op}_z y$.

Proof. If there exist $x, y, z \in \partial X$ with $m(x, y, z) = v$, the rays from v to x, y and z must begin with three pairwise distinct edges. Hence $\deg(v) \geq 3$, which shows (2) ⇒ (1). Regarding the implication (1) ⇒ (2), we have $\deg(v) \geq 3$ as soon as v is not skinny. Indeed, $\deg(v) = 0$ can only happen if X is a single point and $\deg(v) = 1$ would violate the assumption that X has no extremal vertices. Let e_1, e_2 and e_3 be pairwise distinct edges incident to v; as v is not extremal, we can assume that $\mathfrak{w}(e_1)$ and $\mathfrak{w}(e_2)$ are not transverse. Lemma 6.1.3 allows us to extend each e_i to a straight

ray r_i. By Lemma 2.1.1, each union $r_i \cup r_j$ is a line if $i \neq j$. Setting $x = r_1(+\infty)$, $y = r_2(+\infty)$ and $z = r_3(+\infty)$, we thus have $m(x, y, z) = v$. As $r_1 \cup r_2$ is straight, we also have $x \operatorname{op}_z y$. □

Denote by $\mathscr{V} \subseteq X$ the set of skinny vertices. Let $\mathscr{S} \subseteq X$ be the union of all edges intersecting $\mathscr{V}$. Let $\mathscr{F} \subseteq X$ be the full subcomplex with vertex set $X \setminus \mathscr{V}$. We call $\mathscr{S}$ and $\mathscr{F}$ the *skinny* and *fat parts* of X, respectively. We remark that $X = \mathscr{F} \cup \mathscr{S}$ and that every vertex in $\mathscr{S} \setminus \mathscr{F}$ is skinny. We will employ the notation $\mathscr{V}(X)$, $\mathscr{S}(X)$ and $\mathscr{F}(X)$ when it is necessary to specify the cube complex under consideration.

If $X \not\approx \mathbb{R}$, each connected component of $\mathscr{S}$ is either a straight ray or a straight segment; we refer to these as *skinny rays* and *skinny segments*. Every skinny ray intersects $\mathscr{F}$ at a single vertex; given $v \in \mathscr{F}$, we denote by $\mathscr{R}(v) \subseteq \partial_s X$ the set of endpoints at infinity of skinny rays based at v.

Lemma 6.3.2. *A vertex $v \in X$ and a point $x \in \partial X$ are endpoints of a skinny ray if and only if there exist $y, z \in \partial_s X$ with the following properties:*

(1) $m(x, y, z) = v$ and $x \operatorname{op}_z y$;

(2) for every $w \in \partial_s X \setminus \{x\}$ we have $\operatorname{crt}(x, y, z, w) = [\![* : * : 0]\!]$.

Proof. Suppose that v and x are endpoints of a skinny ray r. Lemma 6.1.3 allows us to extend r to a straight line γ; let y be the endpoint of γ other than x. As $\deg(v) \geq 3$ by definition, we can construct a straight ray r' based at v and disjoint from γ. Setting $z = r'(+\infty)$, we have $m(x, y, z) = v$ and $x \operatorname{op}_z y$. If $w \in \partial X$ and $w \neq x$, we must have $m(x, y, w) \notin I(x, v) \setminus \{v\}$; Remark 3.3.3 then shows that $m(x, y, w) \in I(v, y)$ and $(x \cdot w)_v = 0$. As $(y \cdot z)_v$ also vanishes, $\operatorname{crt}(x, y, z, w)$ is of the form $[\![* : * : 0]\!]$ as required.

Conversely, suppose that y and z are given satisfying condition (1). If $I(x, v) \cap X$ is not a skinny ray, it contains a vertex $u \neq v$ with $\deg(u) \geq 3$. Let $\mathfrak{w}_x \in \mathscr{W}(u|x)$ and $\mathfrak{w}_v \in \mathscr{W}(u|v)$ be hyperplanes adjacent to u. Let e be an edge incident to u and not crossing $\mathfrak{w}_x$ or $\mathfrak{w}_v$; let γ be a straight ray extending e and set $w = \gamma(+\infty)$. We have $\mathfrak{w}_x \in \mathscr{W}(w|x)$ so $w \neq x$. Similarly, we have $\mathfrak{w}_v \in \mathscr{W}(w|v)$; this shows that $(x \cdot w)_v > 0$ and, along with $x \operatorname{op}_v y$, it guarantees that $m(x, y, w) \in I(x, v)$ and $(y \cdot w)_v = 0$. Thus condition (2) fails, as $\operatorname{crt}(x, y, z, w) = [\![* : 0 : c]\!]$ with $c = (x \cdot w)_v > 0$. □

6.4 The isomorphism and its uniqueness.

In this section, we complete the proof of Theorem E. Let X and Y have no extremal vertices and suppose that they are neither single points, nor isomorphic to $\mathbb{R}$.

Let $v \in X$ be a non-skinny vertex. Lemma 6.3.1 provides $x_1, x_2, x \in \partial_s X$ with $m(x_1, x_2, x) = v$ and $x_1 \operatorname{op}_x x_2$. We define a map $F\colon \mathscr{F}(X) \to \mathscr{F}(Y)$ by setting $F(v) = m^f(x_1, x_2, x)$. Note that F is well-defined and distance-preserving by Corollary 6.2.5. Applying the same construction to the inverse $f^{-1}\colon \partial Y \to \partial X$, we obtain $H\colon \mathscr{F}(Y) \to \mathscr{F}(X)$. By Propositions 6.1.6 and 6.2.3, the compositions $F \circ H$ and $H \circ F$ are the identity. We conclude that F is surjective and, in fact, an isometric bijection of fat parts.

Theorem 6.4.1. *The map F extends to a cubical isomorphism $F\colon X \to Y$.*

Proof. Observe that two vertices $v_1, v_2 \in \mathscr{F}(X)$ are endpoints of a skinny segment if and only if no $v_3 \in \mathscr{F}(X) \setminus \{v_1, v_2\}$ satisfies $d(v_1, v_2) = d(v_1, v_3) + d(v_3, v_2)$. Thus v_1 and v_2 are endpoints of a skinny segment of length ℓ if and only if $F(v_1)$ and $F(v_2)$ are. We can therefore isometrically extend F over all skinny segments in X.

We are left to deal with skinny rays. We conclude the proof by showing that $f(\mathscr{R}(v)) = \mathscr{R}(F(v))$ for all $v \in \mathscr{F}(X)$. The inclusion $f(\mathscr{R}(v)) \subseteq \mathscr{R}(F(v))$ suffices, as we can then apply the same argument to f^{-1}. Consider $x \in \mathscr{R}(v)$ and let $y, z \in \partial_s X$ be the points provided by Lemma 6.3.2. Since $x \in \partial_s X$, we have $m^f(x, y, z) = F(v)$ and, by Proposition 6.1.6, also $x \operatorname{op}^f_z y$. Given $w \in \partial_s Y \setminus \{f(x)\}$, the point $f^{-1}(w)$ lies in the set $\partial_s X \setminus \{x\}$ by Proposition 6.2.3. Lemma 6.3.2 shows that $\operatorname{crt}(x, y, z, f^{-1}(w))$ is of the form $[\![* : * : 0]\!]$ and, by Lemma 6.2.4, the 4-tuple $(x, y, z, f^{-1}(w))$ lies in $\mathscr{A}(X)$. Hence $\operatorname{crt}(f(x), f(y), f(z), w) = [\![* : * : 0]\!]$ for all $w \in \partial_s Y \setminus \{f(x)\}$. Lemma 6.3.2 finally implies that $f(x) \in \mathscr{R}(F(v))$. □

Now, the isomorphism $F\colon X \to Y$ extends to an isomorphism of median algebras $\overline{F}\colon \overline{X} \to \overline{Y}$. We conclude the proof of Theorem E via the following result.

Theorem 6.4.2. *The map F is the only cubical isomorphism with $\overline{F}|_{\partial X} = f$.*

Proof. The uniqueness of F is clear from our construction. We need to prove that $\overline{F}(x) = f(x)$ for every $x \in \partial X$. First, we suppose that $x \in \partial_s X$.

Let γ be a straight line with an endpoint at x; denote by y the other endpoint of γ. We can assume that x is not endpoint of a skinny ray, as $\overline{F}$ and f clearly coincide on those. Thus, there exist vertices $v_n \in \gamma$ with $\deg(v_n) \geq 3$ and $v_n \to x$; we can moreover assume that $v_{n+1} \in I(v_n, x) \setminus \{v_n\}$.

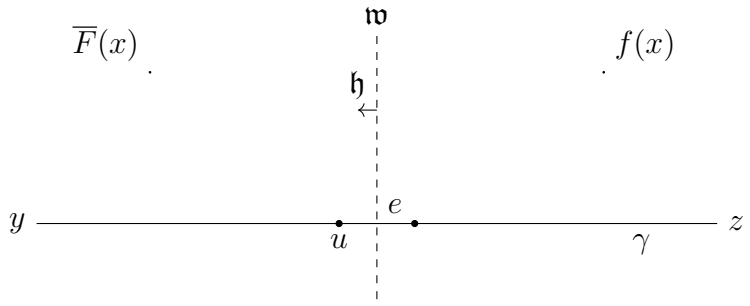


Figure 6.1: The case when $x \in \partial X \setminus \partial_s X$.

Let e_n be an edge with $e_n \cap \gamma = \{v_n\}$. Extending e_n to a straight ray, we construct $z_n \in \partial_s X$ with $m(x, y, z_n) = v_n$ and $x \operatorname{op}_{z_n} y$. Note that $\operatorname{crt}_{v_0}(x, y, z_0, z_n) = [\![* : 0 : c_n]\!]$, where $c_n = (z_n \cdot x)_{v_0} = d(v_0, v_n)$ is strictly increasing. By construction, $\overline{F}(x)$ is the limit of the sequence $F(v_n) = m^f(x, y, z_n)$. By Proposition 6.2.3, there exists a straight line $\gamma' \subseteq Y$ with endpoints $f(x)$ and $f(y)$. Now, the fact that $c_n \to +\infty$ implies that the points $m^f(x, y, z_n) \in \gamma'$ converge to $f(x)$. Hence $f(x) = \overline{F}(x)$.

We are left to handle points $x \in \partial X \setminus \partial_s X$. Suppose for the sake of contradiction that $\overline{F}(x)$ and $f(x)$ are separated by a hyperplane $\mathfrak{w}$. Consider an edge e crossing $\mathfrak{w}$ and extend e to a straight line γ; let u be the endpoint of e on the same side of $\mathfrak{w}$ as $\overline{F}(x)$. Name y and z the endpoints of γ so that $\mathfrak{w} \in \mathscr{W}(\overline{F}(x), y | f(x), z)$. The situation is portrayed in Figure 6.1. We are going to construct a point $w \in \partial_s X$ such that $\min\{(\overline{F}(x) \cdot y)_u, (\overline{F}(x) \cdot w)_u\} < +\infty$ and $m(y, z, w) \in X$. We first show how to use w to conclude the proof.

First, observe that we have $(\overline{F}(x) \cdot z)_u = 0$ and $(y \cdot z)_u = 0$. Our choice of w also implies that $(y \cdot w)_u$, $(z \cdot w)_u$ and at least one among $(\overline{F}(x) \cdot y)_u$ and $(\overline{F}(x) \cdot w)_u$ is finite, so $(\overline{F}(x), y, z, w) \in \mathscr{A}(X)$. As $\overline{F}$ and f coincide on $\{y, z, w\} \subseteq \partial_s X$, we have:

$$\operatorname{cr}\left(\overline{F}(x), w, y, z\right) = \operatorname{cr}\left(x, \overline{F}^{-1}(w), \overline{F}^{-1}(y), \overline{F}^{-1}(z)\right) = \\ = \operatorname{cr}\left(x, f^{-1}(w), f^{-1}(y), f^{-1}(z)\right) = \operatorname{cr}\left(f(x), w, y, z\right).$$

However, we have $(\overline{F}(x) \cdot y)_u \geq 0 = (f(x) \cdot y)_u$ and $(\overline{F}(x) \cdot z)_u = 0 < (f(x) \cdot z)_u$. Hence $\operatorname{cr}\left(\overline{F}(x), w, y, z\right) > \operatorname{cr}\left(f(x), w, y, z\right)$, a contradiction.

Regarding the construction of the point w, observe that γ contains at least two vertices v, v' of degree at least 3; this is because $\overline{F}(x)$ and $f(x)$ are not straight and project to different points of γ. We can assume that $v' \in I(v, z)$. Let e_y and e_z be the only edges at v that lie in $I(y, v)$ and $I(v, z)$, respectively; cf. Figure 6.2. Let moreover $\mathcal{E}$ be set of edges ϵ at v with $\mathfrak{w}(\epsilon) \in \mathscr{W}(v | \overline{F}(x))$. We distinguish three cases.

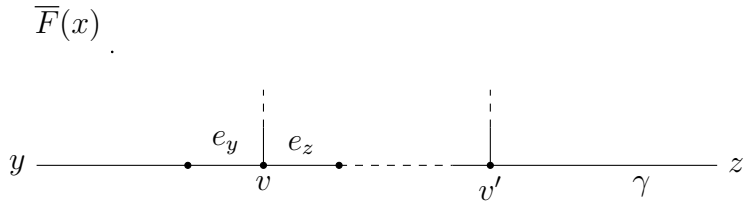


Figure 6.2: The general setup for the construction of the point w.

Case 1: either $(\overline{F}(x)\cdot y)_v < +\infty$ or $\#\mathcal{E} = 1$. It suffices to pick any edge ϵ incident to v and distinct from e_y and e_z; extending ϵ to a straight ray we obtain $w \in \partial_s X$ with $m(y,z,w) = v$. If $(\overline{F}(x)\cdot y)_v < +\infty$, we are done. If $(\overline{F}(x)\cdot y)_v = +\infty$ and $\#\mathcal{E} = 1$, we have $\mathcal{E} = \{e_y\}$; hence $(\overline{F}(x)\cdot w)_v = 0$ by Lemma 2.1.1.

Case 2: $(\overline{F}(x)\cdot y)_v = +\infty$ and no edge in $\mathcal{E}$ spans a square with e_z. Replacing v with v', we end up again in the situation where $\#\mathcal{E} = 1$, which was handled in the previous case.

Case 3: $(\overline{F}(x)\cdot y)_v = +\infty$, $\#\mathcal{E} \geq 2$ and some $\epsilon \in \mathcal{E}$ spans a square with e_z. Since v is not an extremal vertex, there exists an edge ϵ' at v that does not span a square with ϵ. In particular, $\epsilon' \neq e_z$ and $\epsilon' \notin \mathcal{E}$, which also ensures that $\epsilon' \neq e_y$. We extend ϵ to a straight ray r and set $w = r(+\infty)$. Lemma 2.1.1 implies that $m(y,z,w) = v$ and $(\overline{F}(x)\cdot w)_v = 0$. □

Chapter 7

Extending Möbius maps II: cube complexes with no free faces.

Let X and Y be CAT(0) cube complexes. This chapter is devoted to the proof of Theorem A. Having obtained Theorem F in Chapter 5, it suffices to prove Theorem G and this will be our sole focus.

The proof of Theorem G will be described in Section 7.1 and is divided into three steps. We have already carried out the middle one in Section 3.5. The other two steps will occupy, respectively, Sections 7.2 and 7.3. The third step, in Section 7.3, is actually the proof of Theorem L from the introduction.

7.1 Proof outline.

We begin by introducing a few key notions; cf. Definition 6.1.2.

Definition 7.1.1. An oriented geodesic $\alpha \subseteq X$ is *straight-looking (SL)* if no hyperplane in $\mathscr{W}(\alpha)$ is transverse to the first hyperplane crossed by α. A ray $r \subseteq X$ is *straight-looking regular (SLR)* if it is SL and $r(+\infty) \in \partial_{\mathrm{reg}} X$.

To avoid confusion with vertices of X, we will speak of *elements* of a link $\mathrm{lk}(v)$ with the implicit assumption that they are actually vertices of said link.

Definition 7.1.2. An oriented edge $e \subseteq X$ is *SLR-extendable* if it is the initial edge of an SLR ray. Given $v \in X$, we say that an element $e \in \mathrm{lk}(v)$ is *SLR-extendable* if the corresponding edge of X is SLR-extendable (when oriented away from v).

Definition 7.1.3. A vertex $v \in X$ is an *OR median* (OR stands for *'opposite regular'*) if there exist $x, y, z \in \partial_{\mathrm{reg}} X$ with $m(x, y, z) = v$ and $x \operatorname{op}_z y$. Recall from Section 3.5 that we denote by $\mathfrak{M}(X) \subseteq X$ the subset of OR medians.

Not every vertex of X is an OR median in general, even when $\partial_{\mathrm{reg}} X$ is infinite. For instance, if X is obtained by barycentrically subdividing a cube complex X_0, every OR median in X is a vertex of X_0. In fact, it is not even clear whether 'nicer' cube complexes should always contain such medians.

The following remark will be our main source of OR medians. Recall from Section 2.1.1 that $\mathrm{lk}(v)$ is endowed with a metric δ (possibly attaining the value $+\infty$).

Remark 7.1.4. If $\mathrm{lk}(v)$ contains pairwise-distinct, SLR-extendable elements e_1, e_2, e_3 with $\delta(e_1, e_2) \geq 2$, then v is an OR median.

Theorem G will be proved by piecing together three results. The first is the following, which will be obtained in Section 7.2.

Theorem 7.1.5. *Let X be irreducible, with no free faces and such that $\mathrm{Aut}(X)$ contains a uniform lattice. If X is neither a single point nor $\mathbb{R}$, then $\mathfrak{M}(X) \neq \emptyset$.*

The second is Theorem 3.5.3, which was obtained in Section 3.5. Finally, the third step is Theorem L from the introduction. Section 7.3 will be devoted to a proof of the latter. Using these three results:

Proof of Theorem G. Let $\mathcal{A}$, $\mathcal{B}$ and Γ be as in the statement of the theorem. Since the Γ-actions are non-elementary, X and Y are neither single points nor $\mathbb{R}$. Theorem 7.1.5 then guarantees that $\mathfrak{M}(X)$ and $\mathfrak{M}(Y)$ are nonempty. Proposition 2.2.10 and Corollary 2.2.16 show that $\mathcal{A}$ and $\mathcal{B}$ are ample. Theorem 3.5.3 now yields a Γ-equivariant distance-preserving bijection $\phi\colon \mathfrak{M}(X) \to \mathfrak{M}(Y)$. Since the action $\mathrm{Aut}(X) \curvearrowright X$ is cocompact and preserves $\mathfrak{M}(X)$, Theorem L shows that ϕ uniquely extends to a full isomorphism $F\colon X \to Y$.

Note that the uniqueness of F guarantees its Γ-equivariance. Indeed, given $g \in \Gamma$, the maps $g \circ F$ and $F \circ g$ are, respectively, extensions of $g \circ \phi$ and $\phi \circ g$. Since the latter coincide by equivariance of ϕ, we have $g \circ F = F \circ g$.

Since $\mathfrak{M}(X)$ is invariant under the essential action $\Gamma \curvearrowright X$, every point of $\partial_{\mathrm{reg}} X$ is limit of a sequence in $\mathfrak{M}(X)$. Part (3) of Theorem 3.5.3 thus guarantees that f is precisely the boundary extension of F. Finally, F is the unique isomorphism extending f by part (4) of Theorem 3.5.3 and the uniqueness part in Theorem L. $\square$

7.2 Existence of OR medians.

Let X be irreducible, locally finite, with no free faces and $\dim X \geq 1$. We consider a group G acting cocompactly on X.

Later in this section (from Proposition 7.2.7 onwards), we will require G to act properly discontinuously on X. Equivalently, there will be a homomorphism $G \to \mathrm{Aut}(X)$ with finite kernel and image a uniform lattice.

Recall that $\mathscr{W}_v$ denotes the set of hyperplanes adjacent to a vertex $v \in X$. It will also be useful to consider the projection $\pi_v \colon \overline{X} \to 2^{\mathscr{W}_v}$ defined by $\pi_v(x) = \mathscr{W}(v|x) \cap \mathscr{W}_v$. The proof of Theorem 7.1.5 will be based on the following dichotomy.

Lemma 7.2.1. *Either X contains a straight SLR ray, or there exist a point $x \in \partial_{\mathrm{reg}} X$, a vertex $v \in X$ and a hyperplane $\mathfrak{w} \in \mathscr{W}_v$ such that $\#\pi_v(x) \geq 2$ and no element of $\pi_v(x)$ is transverse to $\mathfrak{w}$.*

Proof. Pick a hyperplane $\mathfrak{u}$ and a point $x \in \partial_{\mathrm{reg}} X$. Let $\overline{x}$ denote the gate-projection of x to the carrier of $\mathfrak{u}$; by regularity of x, the point $\overline{x}$ lies inside X. Let $w \in X$ be the only vertex with $\mathscr{W}(w|\overline{x}) = \{\mathfrak{u}\}$ and let r be a ray from w to x. The ray r is SLR and, if it is straight, we are done.

Otherwise, a vertex $v = r(n)$ must satisfy $\#\pi_v(x) \geq 2$; we choose v so as to minimise $n \geq 1$. Let $\mathfrak{w}$ be the only hyperplane separating v and $r(n-1)$. If $\mathfrak{w}$ were transverse to some $\mathfrak{v} \in \pi_v(x)$, then $\mathfrak{w}$ and $\mathfrak{v}$ would both be adjacent to $r(n-1)$, violating minimality of n. This concludes the proof. □

It is in the second case of Lemma 7.2.1 that it will be particularly hard to prove Theorem 7.1.5. We will construct OR medians based on Remark 7.1.4, so we will need some results that allow us to piece together geodesics in order to assemble SLR rays. We now prove a sequence of results with this purpose, culminating in Theorem 7.2.9.

Lemma 7.2.2. *Given an element $g \in X$ and a vertex $v \in X$, there exists $n \geq 1$ such that every hyperplane in $\mathscr{W}_v \cap \mathscr{W}_{g^n v}$ is preserved by g^n.*

Proof. Let us write $\mathscr{W}_v = P \sqcup Q$, where P is the subset of hyperplanes preserved by some power of g. Possibly replacing g with a power, we can assume that g fixes every element of P. If a hyperplane $\mathfrak{w}$ lies in $\mathscr{W}_v \cap \mathscr{W}_{g^m v}$ for infinitely many $m > 0$, we must have $g^{-m}\mathfrak{w} \in \mathscr{W}_v$ for the same integers m; since $\mathscr{W}_v$ is finite, we then have $\mathfrak{w} \in P$. We conclude that, if n is sufficiently large, we have $\mathscr{W}_v \cap \mathscr{W}_{g^n v} \subseteq P$. □

Given an element $e \in \mathrm{lk}(v)$, we denote by $\mathfrak{w}(e)$ the hyperplane dual to the edge determined by e. Given a subset $F \subseteq \mathrm{lk}(v)$, we write $\mathfrak{w}(F) = \{\mathfrak{w}(e) \mid e \in F\}$. Note that each $g \in G$ takes e to an element $ge \in \mathrm{lk}(gv)$ and F to a subset $gF \subseteq \mathrm{lk}(gv)$; we have $\mathfrak{w}(ge) = g\mathfrak{w}(e)$ and $\mathfrak{w}(gF) = g\mathfrak{w}(F)$.

Lemma 7.2.3. *Given a vertex $v \in X$ and $e \in \mathrm{lk}(v)$, there exists $g \in G$ such that:*

(1) $\pi_v(gv) = \{\mathfrak{w}(e)\}$;

(2) every hyperplane in $\mathscr{W}_v \cap \mathscr{W}_{gv}$ is preserved by g.

Proof. Let $e_1 = e, e_2, ..., e_k$ form a maximal clique in $\mathrm{lk}(v)$. Consider the CAT(0) cube complex $Z = \mathfrak{w}(e_2) \cap ... \cap \mathfrak{w}(e_k)$ and let $H < G$ be the stabiliser of all halfspaces bounded by the $\mathfrak{w}(e_i)$ with $i \geq 2$. Note that v projects to a vertex $\overline{v} \in Z$ and e projects to an isolated point $\overline{e} \in \mathrm{lk}_Z(\overline{v})$. The action $H \curvearrowright Z$ is cocompact by Lemma 2.2.4 and, since Z has no free faces, it is essential. By Proposition 3.2 in [CS11], there exists an element $g \in H$ such that $\mathfrak{w}(\overline{e}) \in \mathscr{W}(\overline{v}|g^n\overline{v})$ for all $n \geq 1$; as $\overline{e}$ is isolated, we have $\pi_{\overline{v}}(g^n\overline{v}) = \{\mathfrak{w}(\overline{e})\}$.

Note that every hyperplane of X separating v and $g^n v$ must be transverse to all $\mathfrak{w}(e_i)$ with $i \geq 2$; hence $\pi_v(g^n v) = \{\mathfrak{w}(e)\}$ for all $n \geq 1$. Finally, Lemma 7.2.2 allows us to assume that g also satisfies condition (2). □

We collect here two useful remarks for the subsequent discussion.

Remark 7.2.4. Let $v, w \in X$ be vertices with $\pi_v(w) = \{\mathfrak{v}\}$. For a hyperplane $\mathfrak{w} \in \mathscr{W}_w$ with $\mathfrak{w} \neq \mathfrak{v}$, we have $\mathfrak{w} \in \mathscr{W}_v$ if and only if $\mathfrak{w}$ is transverse to $\mathfrak{v}$. Indeed, if $\mathfrak{w}$ and $\mathfrak{v}$ are not transverse, $\mathfrak{v}$ separates v and $\mathfrak{w}$; hence $\mathfrak{w} \notin \mathscr{W}_v$. On the other hand, if $\mathfrak{w}$ is not adjacent to v, there exists $\mathfrak{u} \in \mathscr{W}_v$ separating v and $\mathfrak{w}$. Since w lies in the carrier of $\mathfrak{w}$, we then have $\mathfrak{u} \in \mathscr{W}(v|w)$; hence $\mathfrak{u} = \mathfrak{v}$ and $\mathfrak{w}$ is not transverse to $\mathfrak{v}$.

Remark 7.2.5. Let $v \in X$ be a vertex and $c \subseteq \mathrm{lk}(v)$ a clique. We can always find a maximal clique $c' \subseteq \mathrm{lk}(v)$ with $c \cap c' = \emptyset$. In order to see this, let us denote the elements of c by $e_1, ..., e_k$. First, observe that there exists $e_1' \in \mathrm{lk}(v)$ such that $e_1', e_2, ..., e_k$ is a clique and $\delta(e_1, e_1') = 2$. Indeed, we can extend c to a maximal clique $\tilde{c}_1$ and, since X has no free faces, the clique $\tilde{c}_1 \setminus \{e_1\}$ must also be contained in a maximal clique $\tilde{c}_2 \neq \tilde{c}_1$. We can then take any point of $\tilde{c}_2 \setminus \tilde{c}_1$ as e_1'.

We repeat this procedure to obtain a clique $e_1', e_2', e_3, ..., e_k$ with $\delta(e_2, e_2') = 2$ and so on. In the end, we obtain a clique $e_1', ..., e_k'$ with $\delta(e_i, e_i') = 2$ for every i. The latter property implies that any maximal clique containing $\{e_1', ..., e_k'\}$ is disjoint from c.

A subgraph $L \subseteq \mathrm{lk}(v)$ is *irreducible* if it does not split as a join of two subgraphs. In general, we can always decompose $\mathrm{lk}(v) = L_1 * ... * L_s$ for pairwise-disjoint, irreducible subgraphs $L_i \subseteq \mathrm{lk}(v)$. The L_i are unique up to permutations and we refer to them as the *irreducible components* of $\mathrm{lk}(v)$.

Proposition 7.2.6. *Given an irreducible component $L \subseteq \mathrm{lk}(v)$ and a element $e \in L$, there exists $g \in G$ with $\pi_v(gv) = \{\mathfrak{w}(e)\}$ and $\mathscr{W}_v \cap \mathscr{W}_{gv} \cap \mathfrak{w}(L) = \emptyset$.*

Proof. Among elements that satisfy conditions (1) and (2) of Lemma 7.2.3, let us pick $g \in G$ minimising the cardinality of the set $\mathfrak{W} := \mathscr{W}_v \cap \mathscr{W}_{gv} \cap \mathfrak{w}(L)$. Suppose for the sake of contradiction that $\mathfrak{W}$ is nonempty.

Let $\mathcal{W} \subseteq L$ be the subset with $\mathfrak{w}(\mathcal{W}) = \mathfrak{W}$. Observe that $L \setminus \mathcal{W}$ is also nonempty, as it contains e. Since L does not split as a join of $\mathcal{W}$ and $L \setminus \mathcal{W}$, there exist $f \in L \setminus \mathcal{W}$ and $f' \in \mathcal{W}$ with $\delta(f, f') \geq 2$. Note that $\mathfrak{w}(gf)$ is not transverse to $\mathfrak{w}(e)$, or Remark 7.2.4 would yield $\mathfrak{w}(gf) \in \mathscr{W}_v \cap \mathscr{W}_{gv}$ and $\mathfrak{w}(gf) = g^{-1}\mathfrak{w}(gf) = \mathfrak{w}(f)$, violating $f \notin \mathcal{W}$. Note moreover that the hyperplane $\mathfrak{w}(f') = \mathfrak{w}(gf') \in \mathscr{W}_v \cap \mathscr{W}_{gv}$ is transverse to every element of $\mathscr{W}(v|gv)$. In particular $\mathfrak{w}(gf) \notin \pi_{gv}(v)$, as $\delta(gf, gf') = \delta(f, f') \geq 2$.

Now, Lemma 7.2.3 provides $h \in G$ with $\mathscr{W}(gv|hgv) \cap \mathscr{W}_{gv} = \{\mathfrak{w}(gf)\}$ and such that every hyperplane in $\mathscr{W}_{gv} \cap \mathscr{W}_{hgv}$ is preserved by h. As $\mathfrak{w}(gf)$ and $\mathfrak{w}(e)$ are not transverse and $\mathfrak{w}(gf) \notin \pi_{gv}(v)$, we can concatenate geodesics from v to gv and from gv to hgv in order to obtain an SL geodesic from v to hgv. It follows that $\pi_v(hgv) = \{\mathfrak{w}(e)\}$. Every hyperplane in $\mathscr{W}_v \cap \mathscr{W}_{hgv}$ must lie in $\mathscr{W}_v \cap \mathscr{W}_{gv} \cap \mathscr{W}_{hgv}$ and is therefore preserved by h, g and hg. Thus, hg satisfies conditions (1) and (2) of Lemma 7.2.3. By minimality of $\#\mathcal{W}$, the inclusion $\mathscr{W}_v \cap \mathscr{W}_{hgv} \cap \mathfrak{w}(L) \subseteq \mathfrak{W}$ must be an equality. Hence $\mathfrak{w}(f') \in \mathscr{W}_v \cap \mathscr{W}_{hgv}$ and $\mathfrak{w}(f') = \mathfrak{w}(gf')$ lies in $\mathscr{W}_{gv} \cap \mathscr{W}_{hgv}$, contradicting the fact that $\delta(gf, gf') = \delta(f, f') \geq 2$. □

In the rest of the section, we also require $G \curvearrowright X$ to be properly discontinuous.

Proposition 7.2.7. *Let $v \in X$ be a vertex and $c^-, c^+ \subseteq \mathrm{lk}(v)$ maximal cliques. There exists a strongly contracting $g \in G$ with $\pi_v(g^n v) = \mathfrak{w}(c^-)$ and $\pi_{g^n v}(v) = \mathfrak{w}(g^n c^+)$ for all $n \geq 1$, unless possibly if $X = \mathbb{R}$ and $c^- = c^+$.*

Proof. By Proposition 2.2.10, the cube complex X is sector-heavy. Let $\mathfrak{h}_+$ be a halfspace contained in the intersection of all sides of elements of $\mathfrak{w}(c^+)$ that do not contain v. Similarly, take $\mathfrak{h}_-$ in the intersection of all sides of elements of $\mathfrak{w}(c^-)$ that do not contain v. We only need to find a strongly contracting element $g \in G$ such that $g^n v \in \mathfrak{h}_-$ and $g^{-n} v \in \mathfrak{h}_+$ for all $n \geq 1$. For this, it suffices to construct a strongly

contracting element $g \in G$ with $g^+ \in \mathfrak{h}_-$ and $g^- \in \mathfrak{h}_+$, passing then to a large power if necessary.

The case $X = \mathbb{R}$ is immediate. If $X \neq \mathbb{R}$, the action $G \curvearrowright X$ is non-elementary by Lemma 2.2.5. Since X has no free faces, $G \curvearrowright X$ is also essential. Lemma 5.2.1 provides a strongly contracting element $h \in G$. Applying Corollary 2.2.16 to the set $G \cdot h^+ \subseteq \partial_{\mathrm{reg}} X$, we obtain $g_1, g_2 \in G$ with $g_1 \cdot h^+ \in \mathfrak{h}_-$, $g_2 \cdot h^+ \in \mathfrak{h}_+$ and $g_1 \cdot h^+ \neq g_2 \cdot h^+$. Considering strongly separated chains containing the points $g_1 \cdot h^+$ and $g_2 \cdot h^+$, we find strongly separated halfspaces $\mathfrak{k}_\pm$ with $\mathfrak{k}_+ \subseteq \mathfrak{h}_+$ and $\mathfrak{k}_- \subseteq \mathfrak{h}_-$. Part (2) of Proposition 2.2.6 now provides $g \in G$ with $g\mathfrak{k}_+^* \subseteq \mathfrak{k}_-$ and g^2 is the required strongly contracting element. □

The following is a small but essential improvement on Proposition 7.2.6.

Lemma 7.2.8. *Suppose that $X \neq \mathbb{R}$ and let $v \in X$ be a vertex with irreducible link. For every element $e \in \mathrm{lk}(v)$, there exists $g \in G$ satisfying $\pi_v(g^n v) = \{\mathfrak{w}(e)\}$ and $\mathscr{W}_v \cap \mathscr{W}_{g^n v} = \emptyset$ for all $n \geq 1$.*

Proof. We start by constructing an element $g \in G$ with $\pi_v(gv) = \{\mathfrak{w}(e)\}$, $\mathscr{W}_v \cap \mathscr{W}_{gv} = \emptyset$ and $\mathfrak{w}(ge) \notin \pi_{gv}(v)$. Proposition 7.2.6 already provides $h \in G$ with $\pi_v(hv) = \{\mathfrak{w}(e)\}$ and $\mathscr{W}_v \cap \mathscr{W}_{hv} = \emptyset$. Let $c \subseteq \mathrm{lk}(v)$ be the clique with $\pi_{hv}(v) = \mathfrak{w}(hc)$. If $e \notin c$, we can take $g = h$; suppose instead that $e \in c$.

Remark 7.2.5 yields a maximal clique $c' \subseteq \mathrm{lk}(v)$ with $c \cap c' = \emptyset$. Since $X \neq \mathbb{R}$, Proposition 7.2.7 provides an element $k \in G$ with $\pi_{hv}(kv) = \mathfrak{w}(hc')$ and $\pi_{kv}(hv) = \mathfrak{w}(kc')$. By Remark 7.2.4, no element of $\mathfrak{w}(hc')$ can be transverse to $\mathfrak{w}(e)$ or it would lie in $\mathscr{W}_v \cap \mathscr{W}_{hv} = \emptyset$. Note moreover that $\pi_{hv}(kv) = \mathfrak{w}(hc')$ is disjoint from $\pi_{hv}(v) = \mathfrak{w}(hc)$. It follows that a geodesic from v to kv is SL and passes through hv; in particular, we have $\pi_v(kv) = \{\mathfrak{w}(e)\}$ and $\mathscr{W}_v \cap \mathscr{W}_{kv} \subseteq \mathscr{W}_v \cap \mathscr{W}_{hv} = \emptyset$. Since the clique c' is maximal, we have $\pi_{kv}(v) = \pi_{kv}(hv) = \mathfrak{w}(kc')$, which does not contain $\mathfrak{w}(ke)$. In conclusion, we can take $g = k$.

We now show by induction on $n \geq 1$ that $\pi_v(g^n v) = \{\mathfrak{w}(e)\}$, $\mathscr{W}_v \cap \mathscr{W}_{g^n v} = \emptyset$, $\mathfrak{w}(g^n e) \notin \pi_{g^n v}(v)$ and, moreover, that a geodesic from v to $g^n v$ passes through all $g^i v$ with $1 \leq i \leq n-1$. This will conclude the proof.

The case $n = 1$ is immediate from our choice of g; let us assume that the above hold for some $n \geq 1$. Since $\mathfrak{w}(g^n e) \notin \pi_{g^n v}(v)$, a geodesic from v to $g^{n+1} v$ is obtained by concatenating any geodesic from v to $g^n v$ with any geodesic from $g^n v$ to $g^{n+1} v$. By convexity of hyperplane carriers, we then have $\mathscr{W}_v \cap \mathscr{W}_{g^{n+1} v} \subseteq \mathscr{W}_v \cap \mathscr{W}_{g^n v} = \emptyset$, $\mathscr{W}_v \cap \mathscr{W}(gv | g^{n+1} v) \subseteq \mathscr{W}_v \cap \mathscr{W}_{gv} = \emptyset$, $\mathscr{W}_{g^{n+1} v} \cap \mathscr{W}(v | g^n v) \subseteq \mathscr{W}_{g^{n+1} v} \cap \mathscr{W}_{g^n v} = \emptyset$. The last two inclusions guarantee that we have $\pi_v(g^{n+1} v) = \pi_v(gv) = \{\mathfrak{w}(e)\}$ and also

$\pi_{g^{n+1}v}(v) = \pi_{g^{n+1}v}(g^n v) \not\ni \mathfrak{w}(g^{n+1}e)$. This concludes the inductive step and the entire proof. □

The following will be our key tool in proving Theorem 7.1.5.

Theorem 7.2.9. *Consider an irreducible component $L \subseteq \mathrm{lk}(v)$ with $\#L \geq 3$. Given an element $e \in L$ and a maximal clique $c \subseteq L$, there exists $g \in G$ satisfying conditions (1) and (2) of Lemma 7.2.3 and, in addition:*

(3) $\mathscr{W}_v \cap \mathscr{W}_{gv} \cap \mathfrak{w}(L) = \emptyset$;

(4) $\pi_{gv}(v) = \mathfrak{w}(gc)$.

Proof. Let $\mathrm{lk}(v) = L_1 * ... * L_s$ be the decomposition into irreducible components, with $L_1 = L$. Choose maximal cliques $\kappa_i \subseteq L_i$ for each $i \geq 2$. Any two hyperplanes in $\mathfrak{U} := \mathfrak{w}(\kappa_2) \cup ... \cup \mathfrak{w}(\kappa_s)$ are transverse; let Z denote their intersection. The vertex $v \in X$ projects to a vertex $\overline{v} \in Z$ and $\mathrm{lk}_Z(\overline{v})$ is naturally identified with L_1. Let $H < G$ be the subgroup that stabilises every side of every hyperplane in $\mathfrak{U}$. By Lemma 2.2.4, the action $H \curvearrowright Z$ is proper and cocompact. Note that Z has no free faces. As $\mathrm{lk}_Z(\overline{v})$ is irreducible, with at least three vertices, Z is irreducible and $Z \neq \mathbb{R}$.

Observe that e and c project to an element $\overline{e} \in \mathrm{lk}_Z(\overline{v})$ and a maximal clique $\overline{c} \subseteq \mathrm{lk}_Z(\overline{v})$. Applying Lemma 7.2.8 to the action $H \curvearrowright Z$, we find $h_1 \in H$ with $\pi_{\overline{v}}(h_1\overline{v}) = \{\mathfrak{w}(\overline{e})\}$ and $\mathscr{W}_{\overline{v}} \cap \mathscr{W}_{h_1\overline{v}} = \emptyset$. By Lemma 7.2.2, we can moreover assume that h_1 preserves every hyperplane of X lying in $\mathscr{W}_v \cap \mathscr{W}_{h_1 v}$.

Let $\overline{c}_1 \subseteq \mathrm{lk}_Z(\overline{v})$ be the clique with $\pi_{h_1\overline{v}}(\overline{v}) = \{\mathfrak{w}(h_1\overline{c}_1)\}$. As Z has no free faces, Remark 7.2.5 provides a maximal clique $\overline{c}_2 \subseteq \mathrm{lk}_Z(\overline{v})$ with $\overline{c}_1 \cap \overline{c}_2 = \emptyset$. Proposition 7.2.7 and Lemma 7.2.2 yield an element $h_2 \in H$ satisfying $\pi_{h_1\overline{v}}(h_2h_1\overline{v}) = \{\mathfrak{w}(h_1\overline{c}_2)\}$, $\pi_{h_2h_1\overline{v}}(h_1\overline{v}) = \{\mathfrak{w}(h_2h_1\overline{c})\}$ and such that h_2 preserves every hyperplane of X lying in $\mathscr{W}_{h_1 v} \cap \mathscr{W}_{h_2h_1 v}$. Since $\pi_{h_1\overline{v}}(\overline{v})$ and $\pi_{h_1\overline{v}}(h_2h_1\overline{v})$ are disjoint, a geodesic from $\overline{v}$ to $h_2h_1\overline{v}$ passes through $h_1\overline{v}$. As $\mathscr{W}_{\overline{v}} \cap \mathscr{W}_{h_1\overline{v}} = \emptyset$, Remark 7.2.4 shows that every such geodesic is SL. In particular, we have $\pi_{\overline{v}}(h_2h_1\overline{v}) = \{\mathfrak{w}(\overline{e})\}$, $\mathscr{W}_{\overline{v}} \cap \mathscr{W}_{h_2h_1\overline{v}} \subseteq \mathscr{W}_{\overline{v}} \cap \mathscr{W}_{h_1\overline{v}} = \emptyset$ and, since $\overline{c}$ is maximal, also $\pi_{h_2h_1\overline{v}}(\overline{v}) = \pi_{h_2h_1\overline{v}}(h_1\overline{v}) = \mathfrak{w}(h_2h_1\overline{c})$.

Let us now set $g = h_2h_1$. Note that a hyperplane of X lies in $\mathscr{W}(v|gv)$ if and only if it is transverse to every element of $\mathfrak{U}$ and originates a hyperplane of Z that lies in $\mathscr{W}(\overline{v}|g\overline{v})$. It follows that g satisfies conditions (1) and (4).

Observing that a geodesic from v to gv passes through the point h_1v, we have $\mathscr{W}_v \cap \mathscr{W}_{gv} \subseteq \mathscr{W}_v \cap \mathscr{W}_{h_1 v} \cap \mathscr{W}_{h_2h_1 v}$ and it follows that g satisfies condition (2). Finally, each element of $\mathscr{W}_v \cap \mathscr{W}_{gv} \cap \mathfrak{w}(L_1)$ is transverse to every hyperplane in $\mathfrak{U}$ and therefore

originates a hyperplane of Z lying in $\mathscr{W}_{\overline{v}} \cap \mathscr{W}_{g\overline{v}}$. As the latter is empty, we conclude that g also satisfies condition (3). □

Proof of Theorem 7.1.5. According to Lemma 7.2.1, there are two cases to consider. First, suppose that there exists a straight SLR ray r and set $x = r(+\infty)$. Let $\mathfrak{u}$ be the first hyperplane crossed by r and let $\mathfrak{h}$ its side containing $r(0)$; let y be any regular point in $\mathfrak{h}$. Let $\mathfrak{k}$ be a halfspace strongly separated from $\mathfrak{h}$, with $x \in \mathfrak{k}$. For every regular point $z \in \mathfrak{k}$ with $z \neq x$, we have $x \operatorname{op}_z y$. If no such z exists, a sub-ray of r consists only of vertices that have degree 2 in X, hence $X = \mathbb{R}$ by cocompactness.

If no straight SLR ray exists, let $x \in \partial_{\mathrm{reg}} X$, $v \in X$ and $\mathfrak{w} \in \mathscr{W}_v$ be as provided by Lemma 7.2.1. There exist $f \in \mathrm{lk}(v)$ and a clique $c' \subseteq \mathrm{lk}(v)$ with $\mathfrak{w} = \mathfrak{w}(f)$ and $\pi_v(x) = \mathfrak{w}(c')$. The set $c' \cup \{f\}$ is entirely contained in a single irreducible component $L \subseteq \mathrm{lk}(v)$ and we have $\#L \geq 3$. We are going to show that every $e \in L$ is SLR-extendable. We will then conclude by taking SLR rays extending f and two elements of c', in order to define, respectively, points $x, y, z \in \partial_{\mathrm{reg}} X$ with $m(x, y, z) = v$ and $x \operatorname{op}_z y$.

Remark 7.2.5 provides a maximal clique $c \subseteq L$ such that $c \cap c' = \emptyset$. Given $e \in L$, let $g \in G$ be an element satisfying conditions (1)–(4) of Theorem 7.2.9. Since $\pi_{gv}(gx) = \mathfrak{w}(gc')$ is disjoint from $\pi_{gv}(v) = \mathfrak{w}(gc)$, a ray ρ from v to gx passes through gv. If some $\mathfrak{u} \in \mathfrak{w}(gc')$ were transverse to $\mathfrak{w}(e)$, Remark 7.2.4 would yield $\mathfrak{u} \in \mathscr{W}_v \cap \mathscr{W}_{gv}$ and we would have $\mathfrak{u} = g^{-1}\mathfrak{u} \in \mathfrak{w}(c') \subseteq \mathfrak{w}(L)$, violating the assumption that $\mathscr{W}_v \cap \mathscr{W}_{gv} \cap \mathfrak{w}(L) = \emptyset$. Hence, no element of $\mathfrak{w}(gc')$ is transverse to $\mathfrak{w}(e)$ and ρ is an SLR ray extending e. □

7.3 Extending partial isomorphisms.

Let X and Y be irreducible, locally finite and with no free faces. Let $\phi\colon A \to B$ be a distance-preserving bijection between nonempty sets of vertices $A \subseteq X$ and $B \subseteq Y$. We also assume that A and B are preserved, respectively, by cocompact actions $G_1 \curvearrowright X$ and $G_2 \curvearrowright Y$. This section is devoted to the proof of Theorem L.

We will obtain a cubical isomorphism between X and Y by progressively extending ϕ to vertices adjacent to A. In order to do so, we will need to recognise halfspaces adjacent to A by their intersections with A. This is the purpose of the next few results, up to Lemma 7.3.5.

Lemma 7.3.1. *Given halfspaces $\mathfrak{h}_1, \mathfrak{h}_2 \in \mathscr{H}(X)$ satisfying $\mathfrak{h}_1 \cap A = \mathfrak{h}_2 \cap A$, we have $\mathfrak{h}_1 \subseteq \mathfrak{h}_2$ or $\mathfrak{h}_2 \subseteq \mathfrak{h}_1$.*

Proof. The cocompact action $G_1 \curvearrowright X$ leaves A invariant and, since X has no free faces, it is essential and hyperplane-essential. The halfspaces $\mathfrak{h}_1$ and $\mathfrak{h}_2$ are not transverse, or Proposition 2.2.8 would provide a halfspace contained in $\mathfrak{h}_1 \cap \mathfrak{h}_2^*$ and, by essentiality of $G_1 \curvearrowright X$, we would have $\mathfrak{h}_1 \cap \mathfrak{h}_2^* \cap A \neq \emptyset$. We also cannot have $\mathfrak{h}_1 \subseteq \mathfrak{h}_2^*$ or $\mathfrak{h}_2^* \subseteq \mathfrak{h}_1$, as $\mathfrak{h}_1 \cap \mathfrak{h}_2 \supseteq \mathfrak{h}_1 \cap A \neq \emptyset$ and $\mathfrak{h}_1^* \cap \mathfrak{h}_2^* \supseteq \mathfrak{h}_1^* \cap A \neq \emptyset$. We conclude that either $\mathfrak{h}_1 \subseteq \mathfrak{h}_2$ or $\mathfrak{h}_2 \subseteq \mathfrak{h}_1$. □

Given a vertex $v \in A$, let $\mathcal{H}_v$ denote the collection of subsets $H \subseteq A$ with the property that, for every $x \in A$:

$$x \in H \;\Leftrightarrow\; \forall y \in H,\ d(x,y) < d(x,v) + d(v,y). \qquad (*)$$

Note that $d(x,y) < d(x,v) + d(p,v)$ is equivalent to $\pi_v(x) \cap \pi_v(y) \neq \emptyset$. If $\mathfrak{w} \in \mathscr{W}_v$ and $\mathfrak{h}$ is the side of $\mathfrak{w}$ that does not contain v, we set $H_{\mathfrak{w}} := \mathfrak{h} \cap A$.

Remark 7.3.2. For every $v \in A$ and $\mathfrak{w} \in \mathscr{W}_v$, there exists $x \in H_{\mathfrak{w}}$ with $\pi_v(x) = \{\mathfrak{w}\}$. Since A is G_1-invariant, this follows from Lemma 7.2.3 applied to the cocompact action $G_1 \curvearrowright X$.

Lemma 7.3.3. *For every $\mathfrak{w} \in \mathscr{W}_v$, we have $H_{\mathfrak{w}} \in \mathcal{H}_v$. Distinct elements of $\mathscr{W}_v$ yield distinct elements of $\mathcal{H}_v$.*

Proof. It suffices to show that $H_{\mathfrak{w}} \in \mathcal{H}_v$, as the second part of the lemma follows from Lemma 7.3.1. Since $\mathfrak{w} \in \pi_v(y)$ for all $y \in H_{\mathfrak{w}}$, the intersection $\pi_v(x) \cap \pi_v(y)$ is nonempty for all $x, y \in H_{\mathfrak{w}}$. For the other implication of $(*)$, Remark 7.3.2 provides $y_0 \in H_{\mathfrak{w}}$ with $\pi_v(y_0) = \{\mathfrak{w}\}$. If $x \in A$ and $\pi_v(x) \cap \pi_v(y_0) \neq \emptyset$, it follows that $\mathfrak{w} \in \pi_v(x)$, i.e. $x \in H_{\mathfrak{w}}$. □

Remark 7.3.4. If an element $H \in \mathcal{H}_v$ contains a point y_0 with $\#\pi_v(y_0) = 1$, we must have $H = H_{\mathfrak{u}}$ for the only hyperplane $\mathfrak{u} \in \pi_v(y_0)$. Indeed, property $(*)$ shows that $\mathfrak{u} \in \pi_v(y)$ for all $y \in H$. Hence, a point $x \in A$ lies in H if and only if $\mathfrak{u} \in \pi_v(x)$, i.e. if and only if $x \in H_{\mathfrak{u}}$.

Lemma 7.3.5. *An element $H \in \mathcal{H}_v$ is of the form $H = H_{\mathfrak{w}}$ for some hyperplane $\mathfrak{w} \in \mathscr{W}_v$ if and only if there exists $x \in H$ that belongs to no other element of $\mathcal{H}_v$.*

Proof. If, given $x \in A$, two hyperplanes $\mathfrak{u}$ and $\mathfrak{v}$ lie in $\pi_v(x)$, we have $x \in H_{\mathfrak{u}} \cap H_{\mathfrak{v}}$. By Lemma 7.3.3 and Remark 7.3.4, we conclude that x lies in a single element of $\mathcal{H}_v$ if and only if $\#\pi_v(x) = 1$. By Remarks 7.3.2 and 7.3.4, an element $H \in \mathcal{H}_v$ can contain such a point if and only if it is of the form $H = H_{\mathfrak{w}}$ for a hyperplane $\mathfrak{w} \in \mathscr{W}_v$. □

We will also employ the notation $\mathcal{H}_v$ and $H_{\mathfrak{w}}$ for vertices $v \in Y$ and hyperplanes $\mathfrak{w} \in \mathscr{W}(Y)$, of course replacing A with B in the definitions.

In the following discussion, we consider two edges $e_1, e_2 \subseteq X$, so that e_i has endpoints $v_i \in A$ and $w_i \in X$. We set $\mathfrak{w}_i = \mathfrak{w}(e_i)$ and denote by $\mathfrak{h}_i \in \mathscr{H}(X)$ the side of $\mathfrak{w}_i$ containing w_i. As ϕ preserves distances, it induces a bijection $\mathcal{H}_{v_i} \to \mathcal{H}_{\phi(v_i)}$. By Lemma 7.3.5, there exists a hyperplane $\phi_*(\mathfrak{w}_i) \in \mathscr{W}_{\phi(v_i)}$ with $\phi(H_{\mathfrak{w}_i}) = H_{\phi_*(\mathfrak{w}_i)}$. We define $\phi_*(\mathfrak{h}_i) \in \mathscr{H}(Y)$ as the side of $\phi_*(\mathfrak{w}_i)$ that satisfies:

$$\phi(\mathfrak{h}_i \cap A) = \phi(H_{\mathfrak{w}_i}) = H_{\phi_*(\mathfrak{w}_i)} = \phi_*(\mathfrak{h}_i) \cap B.$$

Finally, let $\widetilde{\phi}(w_i) \in \phi_*(\mathfrak{h}_i)$ denote the vertex of Y that is separated from $\phi(v_i)$ only by $\phi_*(\mathfrak{w}_i)$. We will show in a moment that this definition of $\widetilde{\phi}(w_i)$ is independent of the choice of e_i and v_i.

Proposition 7.3.6. *(1) We have $\mathfrak{h}_1 \cap A = \mathfrak{h}_2 \cap A$ if and only if $\mathfrak{h}_1 = \mathfrak{h}_2$.*

(2) If $\mathfrak{h}_1 \cap A = \mathfrak{h}_2^* \cap A$, exactly one of the following happens:

(a) $\mathfrak{h}_2^* \subsetneq \mathfrak{h}_1$ and $d(x, v_2) \leq d(x, v_1) + d(v_1, v_2) - 4$ for all $x \in \mathfrak{h}_1 \cap A$;

(b) $\mathfrak{h}_1 = \mathfrak{h}_2^*$ and $d(x_0, v_2) = d(x, v_1) + d(v_1, v_2) - 2$ for some point $x_0 \in \mathfrak{h}_1 \cap A$.

Proof. If we had $\mathfrak{h}_1 \cap A = \mathfrak{h}_2 \cap A$ and $\mathfrak{h}_1 \neq \mathfrak{h}_2$, Lemma 7.3.1 would allow us to assume that $\mathfrak{h}_1 \subsetneq \mathfrak{h}_2$. Since v_1 is adjacent to $\mathfrak{w}_1$, we would then have $v_1 \in \mathfrak{h}_2$. On the other hand, $v_1 \in A$ and $v_1 \notin \mathfrak{h}_1$, a contradiction.

We now address part (2). First, suppose that $\mathfrak{h}_1 = \mathfrak{h}_2^*$ and consider a point $x_0 \in A$ with $\pi_{v_1}(x_0) = \{\mathfrak{w}_1\}$, as provided by Remark 7.3.2. Note that $x_0 \in \mathfrak{h}_1 \cap A$. Since $v_2 \in \mathfrak{h}_2^* = \mathfrak{h}_1$ lies in the carrier of $\mathfrak{w}_2 = \mathfrak{w}_1$, every element of $\mathscr{W}(v_2|w_1)$ is transverse to $\mathfrak{w}_1$, whereas no element of $\mathscr{W}(w_1|x_0)$ is. It follows that $\mathscr{W}(v_2|w_1)$ and $\mathscr{W}(w_1|x_0)$ are disjoint, hence $m(v_1, v_2, x_0) = w_1$. Since $d(v_1, w_1) = 1$, we conclude that $d(x_0, v_2) = d(x, v_1) + d(v_1, v_2) - 2$.

Assume now instead that $\mathfrak{h}_1 \cap A = \mathfrak{h}_2^* \cap A$, but $\mathfrak{h}_1 \neq \mathfrak{h}_2^*$. By Lemma 7.3.1, we have either $\mathfrak{h}_2^* \subsetneq \mathfrak{h}_1$ or $\mathfrak{h}_1 \subsetneq \mathfrak{h}_2^*$. In the latter case, since v_2 is adjacent to $\mathfrak{w}_2 \subseteq \mathfrak{h}_1^*$, we would have $v_2 \in \mathfrak{h}_1^*$; so $v_2 \in \mathfrak{h}_2^* \setminus \mathfrak{h}_1$, a contradiction. We conclude that $\mathfrak{h}_2^* \subsetneq \mathfrak{h}_1$. Note that $v_2 \in \mathfrak{h}_2^* \cap A = \mathfrak{h}_1 \cap A$. Thus, we have $m = m(v_1, v_2, x) \in \mathfrak{h}_1 \cap \mathfrak{h}_2^*$ for every $x \in \mathfrak{h}_1 \cap A$. Since $v_1 \in \mathfrak{h}_1^* \cap A = \mathfrak{h}_2 \cap A$, it follows that $\mathfrak{w}_1, \mathfrak{w}_2$ lie in $\mathscr{W}(v_1|m)$ and $d(v_1, m) \geq 2$. Hence $d(x, v_2) \leq d(x, v_1) + d(v_1, v_2) - 4$. □

Recalling that ϕ preserves distances and $\phi(\mathfrak{h}_i \cap A) = \phi_*(\mathfrak{h}_i) \cap B$, this yields:

Corollary 7.3.7. *We have $\mathfrak{h}_1 = \mathfrak{h}_2$ if and only if $\phi_*(\mathfrak{h}_1) = \phi_*(\mathfrak{h}_2)$. Furthermore, $\mathfrak{h}_1 = \mathfrak{h}_2^*$ if and only if $\phi_*(\mathfrak{h}_1) = \phi_*(\mathfrak{h}_2)^*$.*

Corollary 7.3.8. *The point $\widetilde{\phi}(w_i)$ is independent of the choice of e_i and v_i. If $w_i \in A$, we have $\widetilde{\phi}(w_i) = \phi(w_i)$.*

Proof. It is clear that $\widetilde{\phi}(w_i) = \phi(w_i)$ when $w_i \in A$. For the other part of the lemma, observe that:

$$v_1 \neq v_2 \text{ and } w_1 = w_2 \iff d(v_1, v_2) = 2,\ v_2 \in \mathfrak{h}_1,\ v_1 \in \mathfrak{h}_2 \text{ and } \mathfrak{h}_1 \neq \mathfrak{h}_2^*.$$

Corollary 7.3.7 then shows that we have $w_1 = w_2 \Leftrightarrow \widetilde{\phi}(w_1) = \widetilde{\phi}(w_2)$. □

Let $\widetilde{A} \subseteq X$ denote the set of vertices that either lie in A or are connected to a vertex of A by an edge of X; the set $\widetilde{B} \subseteq Y$ is defined similarly. Corollary 7.3.8 yields a well-defined map $\widetilde{\phi}\colon \widetilde{A} \to \widetilde{B}$ extending $\phi\colon A \to B$. We can apply the same construction to the inverse $\psi = \phi^{-1}\colon B \to A$ and obtain an extension $\widetilde{\psi}\colon \widetilde{B} \to \widetilde{A}$. It is clear that the two compositions $\widetilde{\psi} \circ \widetilde{\phi}$ and $\widetilde{\phi} \circ \widetilde{\psi}$ are the identity. Note moreover that the sets $\widetilde{A}$ and $\widetilde{B}$ are still preserved by G_1 and G_2, respectively.

We are left to show that the map $\widetilde{\phi}$ is distance-preserving. Observe that:

$$d(w_1, v_2) = \begin{cases} d(v_1, v_2) + 1 & \text{if } v_2 \in \mathfrak{h}_1^*, \\ d(v_1, v_2) - 1 & \text{if } v_2 \in \mathfrak{h}_1; \end{cases}$$

$$d(w_1, w_2) = \begin{cases} d(v_1, v_2) + 2 & \text{if } v_1 \in \mathfrak{h}_2^* \text{ and } v_2 \in \mathfrak{h}_1^* \text{ and } \mathfrak{h}_1 \neq \mathfrak{h}_2, \\ d(v_1, v_2) - 2 & \text{if } v_1 \in \mathfrak{h}_2 \text{ and } v_2 \in \mathfrak{h}_1 \text{ and } \mathfrak{h}_1 \neq \mathfrak{h}_2^*, \\ d(v_1, v_2) & \text{otherwise, i.e. if } \begin{cases} \mathfrak{h}_1 \in \{\mathfrak{h}_2, \mathfrak{h}_2^*\}, \text{ or} \\ v_1 \in \mathfrak{h}_2 \text{ and } v_2 \in \mathfrak{h}_1^*, \text{ or} \\ v_1 \in \mathfrak{h}_2^* \text{ and } v_2 \in \mathfrak{h}_1. \end{cases} \end{cases}$$

Since v_1 and v_2 lie in A, we have $v_i \in \mathfrak{h}_j \Leftrightarrow \phi(v_i) \in \phi_*(\mathfrak{h}_j)$ for $i, j \in \{1, 2\}$. Along with the above equations, Corollary 7.3.7 immediately yields:

Corollary 7.3.9. *The bijection $\widetilde{\phi}\colon \widetilde{A} \to \widetilde{B}$ is distance-preserving.*

We can now repeatedly apply Corollary 7.3.9 to extend the domain of ϕ. Since A is invariant under the cocompact action $G_1 \curvearrowright X$, this process results in an extension to the entire X in finitely many steps.

This completes the proof of Theorem L, as uniqueness is clear from the construction. Note that we have only used the no-free-faces assumption in the form of Lemma 7.2.3, which is a priori much weaker.

Chapter 8

Extending Möbius maps III: cubulations of hyperbolic groups.

Sections 8.1–8.5 of this chapter are devoted to the proof of Theorem H. Then, in Section 8.6, we will deduce Theorem B and discuss why essentiality and hyperplane-essentiality cannot be removed from its statement.

Referring to the sketch of proof of Theorem H given in the introduction, we only need to carry out Steps (IIa) and (IIb), as Step (I) is provided by Proposition 2.2.15. The crucial results are thus Proposition 8.4.3 in Section 8.4 and Theorem 8.5.3 in Section 8.5. Before that, Sections 8.1, 8.2 and 8.3 develop all necessary ingredients.

Throughout the chapter, we fix a (finitely generated) Gromov hyperbolic group Γ, but we do not a priori assume that Γ is non-elementary (cf. Lemma 8.1.12). For the convenience of the reader, we recall here a couple of standard lemmas which do not require any action on a cube complex.

Lemma 8.0.1. *Let $H < \Gamma$ be an infinite-index quasi-convex subgroup. The limit set $\partial_\infty H \subseteq \partial_\infty \Gamma$ is nowhere-dense.*

Proof. By Lemma 2.9 in [GMRS98], there exists a point $\xi \in \partial_\infty \Gamma \setminus \partial_\infty H$. We can assume that H is non-elementary, or the lemma is trivial. Given any infinite-order element $h \in H$, the orbit $\langle h \rangle \cdot \xi$ accumulates on the point $h^{+\infty} \in \partial_\infty H$. Since the action $H \curvearrowright \partial_\infty H$ is minimal, the subset $H \cdot h^{+\infty}$ is dense in $\partial_\infty H$. Thus $\partial_\infty H \subseteq \overline{H \cdot \xi}$, while $H \cdot \xi \subseteq \partial_\infty \Gamma \setminus \partial_\infty H$. Since $\partial_\infty H$ is closed in $\partial_\infty \Gamma$, this concludes the proof. □

Lemma 8.0.2. *Let H and K be quasi-convex subgroups of Γ. If H and K have the same limit set, they are commensurable.*

Proof. Let $L \leq \Gamma$ denote the stabiliser of the limit set $\Lambda = \partial_\infty H = \partial_\infty K$. It is clear that L contains both H and K, so $\Lambda \subseteq \partial_\infty L$. Given a sequence $g_n \in L$ converging to

a point $\xi \in \partial_\infty L$, we have $g_n^{+\infty} \to \xi$. If ξ did not lie in Λ, we would have $g_n^{+\infty} \notin \Lambda$ for large n and north-south dynamics of g_n would violate the fact that Λ is preserved by L. We conclude that $\Lambda = \partial_\infty L$ and Lemma 2.9 in [GMRS98] shows that H and K have finite index in L. Hence $H \cap K$ has finite index in both H and K. □

8.1 Traces at infinity.

Let now Γ act properly and cocompactly on a CAT(0) cube complex X. We denote by $o_X \colon \partial_\infty \Gamma \to \partial_\infty X$ the only Γ-equivariant homeomorphism and by $\Phi \colon \partial X \to \partial_\infty X$ the map from Corollary 4.2.8. As usual, we will only endow X with its *combinatorial* metric. Note that every (combinatorial) ray $r \subseteq X$ determines a point of ∂X, but, by hyperbolicity, it also determines a point of $\partial_\infty X$.

The carrier $C(\mathfrak{w})$ of each hyperplane $\mathfrak{w} \in \mathscr{W}(X)$ is a convex subcomplex of X; in particular, it is itself hyperbolic and we can consider its boundary $\partial_\infty \mathfrak{w} \subseteq \partial_\infty X$. A point $\xi \in \partial_\infty X$ lies in $\partial_\infty \mathfrak{w}$ if and only if one/all rays representing ξ are at finite Hausdorff distance from $\mathfrak{w}$.

Definition 8.1.1. We refer to $\partial_\infty \mathfrak{w} \subseteq \partial_\infty X$ as the *trace at infinity* of $\mathfrak{w}$.

Denoting by $\Gamma_\mathfrak{w} \leq \Gamma$ the stabiliser of $\mathfrak{w}$, Lemma 2.2.4 guarantees that $\Gamma_\mathfrak{w}$ acts properly and cocompactly on the carrier $C(\mathfrak{w})$. In particular, the homeomorphism $o_X^{-1} \colon \partial_\infty X \to \partial_\infty \Gamma$ takes $\partial_\infty \mathfrak{w}$ to the limit set $\partial_\infty \Gamma_\mathfrak{w} \subseteq \partial_\infty \Gamma$.

Lemma 8.1.2. *Consider a hyperplane $\mathfrak{w} \in \mathscr{W}(X)$ and its two sides $\mathfrak{h}$ and $\mathfrak{h}^*$.*

(1) The subsets $\partial_\infty \mathfrak{h}$ and $\partial_\infty \mathfrak{h}^*$ are closed in $\partial_\infty X$. We have $\partial_\infty X = \partial_\infty \mathfrak{h} \cup \partial_\infty \mathfrak{h}^*$ and $\partial_\infty \mathfrak{h} \cap \partial_\infty \mathfrak{h}^* = \partial_\infty \mathfrak{w}$.

(2) If X is essential, the sets $\partial_\infty \mathfrak{h} \setminus \partial_\infty \mathfrak{w}$ and $\partial_\infty \mathfrak{h}^* \setminus \partial_\infty \mathfrak{w}$ are nonempty.

Proof. Since $\mathfrak{h}$ and $\mathfrak{h}^*$ are closed convex subsets of X, their boundaries are well-defined closed subsets of $\partial_\infty X$. Every ray in X intersects either $\mathfrak{h}$ or $\mathfrak{h}^*$ in a sub-ray, so every point of $\partial_\infty X$ lies in either $\partial_\infty \mathfrak{h}$ or $\partial_\infty \mathfrak{h}^*$. It is clear that $\partial_\infty \mathfrak{w} \subseteq \partial_\infty \mathfrak{h} \cap \partial_\infty \mathfrak{h}^*$. Conversely, if $\xi \in \partial_\infty \mathfrak{h} \cap \partial_\infty \mathfrak{h}^*$, we can consider rays r and r^* representing ξ and entirely contained, respectively, in $\mathfrak{h}$ and $\mathfrak{h}^*$. The function $t \mapsto d(r(t), r^*(t))$ must be uniformly bounded and, observing that $d(r(t), C(\mathfrak{w})) < d(r(t), \mathfrak{h}^*) \leq d(r(t), r^*(t))$, we see that r is contained in a metric neighbourhood of $C(\mathfrak{w})$. Hence $\xi \in \partial_\infty \mathfrak{w}$.

Finally, regarding part (2), consider a sequence of points $x_n \in \mathfrak{h}$ whose distances from $\mathfrak{h}^*$ diverge. Lemma 2.2.4 ensures that x_n can be chosen so that their gate-projections $\overline{x}_n$ to $C(\mathfrak{w})$ lie in a compact set. The Arzelà-Ascoli theorem guarantees

that, up to passing to a subsequence, the geodesics joining $\overline{x}_n$ and x_n converge to a ray that is contained in $\mathfrak{h}$ but in no metric neighbourhood of $\mathfrak{h}^*$. The corresponding point of $\partial_\infty X$ lies in $\partial_\infty \mathfrak{h} \setminus \partial_\infty \mathfrak{w}$. Similarly, $\partial_\infty \mathfrak{h}^* \setminus \partial_\infty \mathfrak{w}$ is also nonempty. □

Remark 8.1.3. In the setting of Lemma 8.1.2, it is interesting to note that $\partial_\infty \mathfrak{h} \setminus \partial_\infty \mathfrak{w}$ is a union of connected components of $\partial_\infty X \setminus \partial_\infty \mathfrak{w}$. This is because the partition $\partial_\infty X \setminus \partial_\infty \mathfrak{w} = (\partial_\infty \mathfrak{h} \setminus \partial_\infty \mathfrak{w}) \sqcup (\partial_\infty \mathfrak{h}^* \setminus \partial_\infty \mathfrak{w})$ consists of two disjoint closed subsets.

Remark 8.1.4. By Lemma 2.2.4, we have $\partial_\infty \mathfrak{w} = \emptyset$ if and only if $\mathfrak{w}$ is compact. If X is hyperplane-essential, this means that $\mathfrak{w}$ consists of a single point, i.e. $\mathfrak{w}$ is dual to an edge that disconnects X.

We remind the reader that, due to Corollary 4.2.8 and Remark 4.3.2, we feel entitled to implicitly identify the sets $\partial_{\mathrm{nt}} X \subseteq \partial X$ and $\Phi(\partial_{\mathrm{nt}} X) \subseteq \partial_\infty X$.

Lemma 8.1.5. *A point $\xi \in \partial_\infty X$ lies in $\partial_{\mathrm{nt}} X$ if and only if there does not exist any $\mathfrak{w} \in \mathscr{W}(X)$ with $\xi \in \partial_\infty \mathfrak{w}$.*

Proof. By part (1) of Lemma 8.1.2, a point $\xi \in \partial_\infty X$ lies in some trace at infinity $\partial_\infty \mathfrak{w}$ if and only if there exists a halfspace $\mathfrak{h} \in \mathscr{H}(X)$ with $\xi \in \partial_\infty \mathfrak{h} \cap \partial_\infty \mathfrak{h}^*$. Equivalently, there exist two rays representing ξ, one contained in $\mathfrak{h}$ and the other contained in $\mathfrak{h}^*$. Thus, ξ lies in a trace at infinity if and only if there exist distinct points $x, y \in \partial X$ with $\Phi(x) = \Phi(y) = \xi$. On the other hand, a point $\xi \in \partial_\infty X$ is non-terminating if and only if $\Phi^{-1}(\xi)$ is a singleton (cf. Theorem 4.2.6 and Definition 4.3.1). □

Corollary 8.1.6. *If X is essential, the subset $\partial_{\mathrm{nt}} X \subseteq \partial_\infty X$ is co-meagre.*

Proof. Since X is essential, all hyperplane stabilisers have infinite index in Γ. The union of the Gromov boundaries of the hyperplanes of X is then meagre in $\partial_\infty X$, by Lemmas 8.0.1 and 2.2.4. By Lemma 8.1.5, this union is precisely $\partial_\infty X \setminus \partial_{\mathrm{nt}} X$. □

In the rest of the section, we collect various facts that will be needed in the proof of Step (IIa). The key ingredient is Proposition 8.1.10.

Given hyperplanes $\mathfrak{u}, \mathfrak{w} \in \mathscr{W}(X)$, we denote by $B(\mathfrak{u}, \mathfrak{w}) \subseteq X$ the bridge associated to their carriers. When $\mathfrak{u}$ and $\mathfrak{w}$ are transverse, $B(\mathfrak{u}, \mathfrak{w})$ is simply the intersection of the two carriers.

Proposition 8.1.7. *Given $\mathfrak{u}, \mathfrak{w} \in \mathscr{W}(X)$, the group $\Gamma_{\mathfrak{u}} \cap \Gamma_{\mathfrak{w}}$ acts properly and cocompactly on $B(\mathfrak{u}, \mathfrak{w})$. Furthermore, $\partial_\infty B(\mathfrak{u}, \mathfrak{w}) = \partial_\infty \mathfrak{u} \cap \partial_\infty \mathfrak{w}$.*

Proof. Let S denote the intersection of $B = B(\mathfrak{u}, \mathfrak{w})$ with the carrier of $\mathfrak{u}$ and consider its stabiliser $\Gamma_S \leq \Gamma$. The action $\Gamma_S \curvearrowright S$ is cocompact by Proposition 2.7 in [HS18]. There exists a constant $D > 0$ such that S is contained in the D-neighbourhood of both $\mathfrak{u}$ and $\mathfrak{w}$. The finitely many hyperplanes that contain S in their D-neighbourhood are permuted by Γ_S and it follows that $\Gamma_\mathfrak{u} \cap \Gamma_\mathfrak{w}$ sits in Γ_S as a finite-index subgroup. Hence $\Gamma_\mathfrak{u} \cap \Gamma_\mathfrak{w}$ acts cocompactly on S and it also acts cocompactly on the bridge B, which is at finite Hausdorff distance from S. Finally, as $B \subseteq X$ is a subcomplex and $\Gamma \curvearrowright X$ is proper, the action $\Gamma_\mathfrak{u} \cap \Gamma_\mathfrak{w} \curvearrowright B$ is proper.

It is clear that $\partial_\infty B = \partial_\infty S \subseteq \partial_\infty \mathfrak{u}$. The same argument applied to the other shore of B shows that $\partial_\infty B \subseteq \partial_\infty \mathfrak{u} \cap \partial_\infty \mathfrak{w}$. Now, consider a point $\xi \in \partial_\infty \mathfrak{u} \setminus \partial_\infty S$ and a geodesic $\gamma \subseteq C(\mathfrak{u})$ representing ξ; note that the function $t \mapsto d(\gamma(t), S)$ must diverge with t. As soon as $\gamma(t) \notin S$, part (4) of Lemma 2.1.7 implies that

$$d(\gamma(t), C(\mathfrak{w})) = d(\gamma(t), S) + d(S, C(\mathfrak{w})).$$

In particular, the function $t \mapsto d(\gamma(t), C(\mathfrak{w}))$ must also diverge and $\xi \notin \partial_\infty \mathfrak{w}$. This shows that we have $\partial_\infty \mathfrak{u} \cap \partial_\infty \mathfrak{w} \subseteq \partial_\infty S = \partial_\infty B$, concluding the proof. □

We now introduce a preorder on $\mathscr{W}(X)$. We write $\mathfrak{u}_1 \preceq \mathfrak{u}_2$ if $\partial_\infty \mathfrak{u}_1 \subseteq \partial_\infty \mathfrak{u}_2$ and $\mathfrak{u}_1 \sim \mathfrak{u}_2$ if $\partial_\infty \mathfrak{u}_1 = \partial_\infty \mathfrak{u}_2$. The latter is an equivalence relation.

Recall that $\mathscr{W}(\mathfrak{w})$ denotes the set of hyperplanes transverse to $\mathfrak{w} \in \mathscr{W}(X)$.

Proposition 8.1.8. *Let X be hyperplane-essential.*

(1) We have $\mathfrak{u}_1 \preceq \mathfrak{u}_2$ if and only if $\mathscr{W}(\mathfrak{u}_1) \subseteq \mathscr{W}(\mathfrak{u}_2)$. Equivalently, $\mathfrak{u}_1$ and $\mathfrak{u}_2$ are not transverse and $\mathfrak{u}_1$ stays at bounded distance from $\mathfrak{u}_2$.

(2) If there exists $\mathfrak{w}$ with $\mathfrak{u}_1 \preceq \mathfrak{w}$ and $\mathfrak{u}_2 \preceq \mathfrak{w}$, then $\mathfrak{u}_1$ and $\mathfrak{u}_2$ are not transverse.

(3) There exists $N = N(X) \geq 0$ such that, given any hyperplane $\mathfrak{u}$ with $\partial_\infty \mathfrak{u} \neq \emptyset$, there exist at most N hyperplanes $\mathfrak{w}$ with $\mathfrak{u} \preceq \mathfrak{w}$.

Proof. We start by proving part (1). Let B denote the bridge $B(\mathfrak{u}_1, \mathfrak{u}_2)$ and let S_1 be the shore $B \cap C(\mathfrak{u}_1)$. If $\mathscr{W}(\mathfrak{u}_1) \subseteq \mathscr{W}(\mathfrak{u}_2)$, we have $\mathfrak{u}_2 \notin \mathscr{W}(\mathfrak{u}_1)$ and the carrier $C(\mathfrak{u}_1)$ is at finite Hausdorff distance from B. Thus, $\mathfrak{u}_1$ is contained in a metric neighbourhood of $\mathfrak{u}_2$ and this implies that $\mathfrak{u}_1 \preceq \mathfrak{u}_2$.

Conversely, suppose that $\mathfrak{u}_1 \preceq \mathfrak{u}_2$. Proposition 8.1.7 implies $\partial_\infty \mathfrak{u}_1 = \partial_\infty S_1$. Since S_1 is a convex subcomplex of $C(\mathfrak{u}_1)$, we will also denote by S_1 its projection to a convex subcomplex of the CAT(0) cube complex $\mathfrak{u}_1$. Note that $S_1 = \mathfrak{u}_1$, or there would exist a halfspace $\mathfrak{k}$ of $\mathfrak{u}_1$ with $S_1 \subseteq \mathfrak{k}$. In this case, applying part (2) of Lemma 8.1.2 to $\mathfrak{u}_1$,

we would obtain a point of $\partial_\infty \mathfrak{k}^* \setminus \partial_\infty \mathfrak{k} \subseteq \partial_\infty \mathfrak{u}_1 \setminus \partial_\infty S_1$, a contradiction. The fact that $S_1 = \mathfrak{u}_1$ yields $\mathscr{W}(\mathfrak{u}_1) = \mathscr{W}(S_1) = \mathscr{W}(\mathfrak{u}_1) \cap \mathscr{W}(\mathfrak{u}_2) \subseteq \mathscr{W}(\mathfrak{u}_2)$, completing part (1).

If $\mathfrak{u}_1$ and $\mathfrak{u}_2$ are transverse and $\mathfrak{u}_1 \preceq \mathfrak{w}$, it follows from part (1) that $\mathfrak{u}_2$ and $\mathfrak{w}$ are transverse. In particular $\mathfrak{u}_2 \not\preceq \mathfrak{w}$ and this proves part (2).

We finally address part (3). By hyperbolicity of X, there exists a constant $D = D(X) > 0$ such that, for any two hyperplanes $\mathfrak{w}_1, \mathfrak{w}_2$ with $\partial_\infty B(\mathfrak{w}_1, \mathfrak{w}_2) \neq \emptyset$, we have $d(C(\mathfrak{w}_1), C(\mathfrak{w}_2)) \leq D$. If $\mathfrak{u} \preceq \mathfrak{w}$, Proposition 8.1.7 gives $\partial_\infty B(\mathfrak{u}, \mathfrak{w}) = \partial_\infty \mathfrak{u} \neq \emptyset$, hence $d(C(\mathfrak{u}), C(\mathfrak{w})) \leq D$. Part (1) now shows that $C(\mathfrak{u})$ is contained in the $(D+1)$-neighbourhood of $C(\mathfrak{w})$. Since X admits a proper cocompact action, it is uniformly locally finite and there exist only uniformly finitely many such hyperplanes $\mathfrak{w}$. □

Definition 8.1.9. Given a hyperplane $\mathfrak{w}$, a point $\xi \in \partial_\infty \mathfrak{w}$ is *generic* if the only hyperplanes $\mathfrak{u}$ with $\xi \in \partial_\infty \mathfrak{u}$ are those that satisfy $\mathfrak{w} \preceq \mathfrak{u}$.

Proposition 8.1.10. *For every $\mathfrak{w} \in \mathscr{W}(X)$, the subset of generic points is dense in $\partial_\infty \mathfrak{w}$. In particular, generic points exist.*

Proof. Let $\mathscr{B}$ be the family of subsets of $\partial_\infty \mathfrak{w}$ of the form $\partial_\infty \mathfrak{w} \cap \partial_\infty \mathfrak{u}$, with $\mathfrak{u} \in \mathscr{W}(X)$ and $\mathfrak{w} \not\preceq \mathfrak{u}$. By definition, a point $\xi \in \partial_\infty \mathfrak{w}$ is generic if and only if it does not lie in the union of the elements of $\mathscr{B}$. Since $\mathscr{W}(X)$ is countable, so is $\mathscr{B}$. We can thus conclude by Baire's theorem, if we show that every $\mathfrak{B} \in \mathscr{B}$ is nowhere-dense in $\partial_\infty \mathfrak{w}$.

If $\mathfrak{B} = \partial_\infty \mathfrak{w} \cap \partial_\infty \mathfrak{u}$, set $H = \Gamma_\mathfrak{w} \cap \Gamma_\mathfrak{u}$. Proposition 8.1.7 shows that the homeomorphism $o_X \colon \partial_\infty \Gamma \to \partial_\infty X$ takes $\partial_\infty H \subseteq \partial_\infty \Gamma_\mathfrak{w}$ to $\mathfrak{B} \subseteq \partial_\infty \mathfrak{w}$. Since $\mathfrak{w} \not\preceq \mathfrak{u}$, the difference $\partial_\infty \mathfrak{w} \setminus \mathfrak{B}$ is nonempty and H must have infinite index in $\Gamma_\mathfrak{w}$. As H is quasi-convex, Lemma 8.0.1 shows that $\mathfrak{B} \subseteq \partial_\infty \mathfrak{w}$ is nowhere-dense. □

Remark 8.1.11. Assume that X is hyperplane-essential and consider a generic point $\xi \in \partial_\infty \mathfrak{w}$. Viewing $\mathfrak{w}$ itself as a CAT(0) cube complex, we have $\xi \in \partial_{\mathrm{nt}} \mathfrak{w}$. This follows from Lemma 8.1.5 since, if $\mathfrak{u} \in \mathscr{W}(\mathfrak{w})$, part (1) of Proposition 8.1.8 shows that $\mathfrak{w} \not\preceq \mathfrak{u}$ and hence $\xi \notin \partial_\infty \mathfrak{u}$.

This is a good point to make the following simple observation, which will only be needed in the proof of Proposition 8.4.5 later on.

Lemma 8.1.12. *Let X be essential and suppose that $\Gamma = \langle g \rangle \simeq \mathbb{Z}$. Then:*

(1) every hyperplane of X is compact;

(2) for every $\mathfrak{h} \in \mathscr{H}(X)$, there exists $N > 0$ such that $g^N \mathfrak{h}$ and $\mathfrak{h}^*$ are strongly separated.

Proof. Since $\partial_\infty X$ contains only two points ξ and η, part (2) of Lemma 8.1.2 shows that $\partial_\infty \mathfrak{w} = \emptyset$ for every $\mathfrak{w} \in \mathscr{W}(X)$. Lemma 2.2.4 then yields part (1).

Given any halfspace $\mathfrak{h}$, we have $\partial_\infty \mathfrak{h} \cap \partial_\infty \mathfrak{h}^* = \emptyset$. Up to swapping ξ and η, we have $\Phi^{-1}(\xi) \subseteq \mathfrak{h}$ and $\Phi^{-1}(\eta) \subseteq \mathfrak{h}^*$. Squaring g if necessary, we can assume that g fixes ξ and η. Since X admits a cocompact action, it is finite dimensional and there exists $n > 0$ such that $g^n\mathfrak{h}$ and $\mathfrak{h}$ are not transverse. Observe that $g^n\mathfrak{h} \cap \mathfrak{h}$ and $g^n\mathfrak{h}^* \cap \mathfrak{h}^*$ are both nonempty, as they contain $\Phi^{-1}(\xi)$ and $\Phi^{-1}(\eta)$, respectively. Replacing g with its inverse if necessary, we can assume that $g^n\mathfrak{h} \subseteq \mathfrak{h}$. As the hyperplane bounding $\mathfrak{h}$ is compact, there exist only finitely many hyperplanes transverse to $\mathfrak{h}$; they are all compact. Choosing a sufficiently large $k > 0$, we can thus ensure that none of them is transverse to $g^{kn}\mathfrak{h} \subsetneq \mathfrak{h}$, hence $g^{kn}\mathfrak{h}$ and $\mathfrak{h}^*$ are strongly separated. □

8.2 Towards hyperplane recognition.

This subsection is devoted to Propositions 8.2.1 and 8.2.5. The latter, in particular, will be our main tool in overcoming the difficulties, described in the introduction, regarding the passage from Step (IIa) to Step (IIb).

Let again $\Gamma \curvearrowright X$ be proper and cocompact and let $o_X \colon \partial_\infty \Gamma \to \partial_\infty X$ denote the only Γ-equivariant homeomorphism.

Proposition 8.2.1. *Consider two distinct points $\xi, \eta \in \partial_\infty X$ and four sequences $x_n, y_n, z_n, w_n \in \partial_{\mathrm{nt}} X$, where x_n and z_n converge to ξ while y_n and w_n converge to η. There exists $N \geq 0$ such that, for every $n \geq N$ and every $\mathfrak{w} \in \mathscr{W}(x_n, y_n | z_n, w_n)$, the points ξ and η lie in $\partial_\infty \mathfrak{w}$.*

Proof. Consider a geodesic α_n whose endpoints in ∂X are x_n and y_n; similarly, let β_n be a geodesic with endpoints z_n and w_n. We also pick a geodesic γ whose endpoints in $\partial_\infty X$ are precisely ξ and η. The convergence of the four sequences yields a constant $D > 0$ such that, for every point $p \in \gamma$, there exists $N \geq 0$ and points $a_n \in \alpha_n$, $b_n \in \beta_n$ with $d(a_n, p) \leq D$ and $d(b_n, p) \leq D$ for all $n \geq N$. Every hyperplane $\mathfrak{w}$ in $\mathscr{W}(x_n, y_n | z_n, w_n)$ separates α_n and β_n and must thus satisfy $d(p, \mathfrak{w}) \leq D$.

Since there are only finitely many hyperplanes at distance at most D from p, we are only left to show that every element of $\limsup \mathscr{W}(x_n, y_n | z_n, w_n)$ contains ξ and η in its trace at infinity. Let us consider a hyperplane $\mathfrak{w}$ that separates α_{n_k} and β_{n_k} for a diverging sequence of integers n_k. Up to passing to a further subsequence, the Arzelà-Ascoli theorem allows us to assume that the geodesics α_{n_k} converge locally uniformly to a geodesic α. The discussion above then shows that α is contained in a

metric neighbourhood of $C(\mathfrak{w})$ and has endpoints ξ and η in $\partial_\infty X$. Hence ξ and η lie in $\partial_\infty \mathfrak{w}$. □

Lemma 8.2.2. *For an infinite-order element $k \in \Gamma$ and a hyperplane $\mathfrak{w}$, the following are equivalent:*

(1) a non-trivial power of k preserves $\mathfrak{w}$;

(2) the points $o_X(k^{\pm\infty})$ lie in $\partial_\infty \mathfrak{w}$.

Fixing k, there are only finitely many hyperplanes satisfying these conditions.

Proof. Since X is hyperbolic, there exists a constant $D = D(X) > 0$ such that any two geodesic lines in X with the same endpoints in $\partial_\infty X$ are at Hausdorff distance at most D. If $\gamma \subseteq X$ is a geodesic with $o_X(k^{\pm\infty})$ as endpoints at infinity, every hyperplane satisfying condition (2) must contain γ in its D-neighbourhood. It follows that only finitely many hyperplanes of X satisfy condition (2). If $\mathfrak{w}$ is such a hyperplane, $k^n\mathfrak{w}$ also satisfies condition (2) for all $n > 0$, so we must have $k^n\mathfrak{w} = \mathfrak{w}$ for some $n > 0$. Conversely, if $k^n \in \Gamma_{\mathfrak{w}}$ for some $n > 0$, we have $k^{\pm\infty} = (k^n)^{\pm\infty} \in \partial_\infty \Gamma_{\mathfrak{w}}$. □

Definition 8.2.3. If $k \in \Gamma$ is infinite-order, we denote by $\mathscr{W}(k)$ the set of hyperplanes satisfying the equivalent conditions in Lemma 8.2.2. We say that k is *good* if it preserves every halfspace bounded by an element of $\mathscr{W}(k)$.

Remark 8.2.4. Every infinite-order element has a good power, as the set $\mathscr{W}(k)$ is always finite by Lemma 8.2.2.

We are interested in good elements because of the following result.

Proposition 8.2.5. *Consider distinct points $x, y \in \partial_{\mathrm{nt}} X$ and a good infinite-order element $k \in \Gamma$ that fixes neither of them. There exists $N \geq 0$ such that, for every $n \geq N$, we have $\mathscr{W}(x|y) \cap \mathscr{W}(k) = \mathscr{W}(k^n x, k^{-n}x|k^n y, k^{-n}y)$.*

Proof. Let us set $\mathscr{W}_n = \mathscr{W}(k^n x, k^{-n}x|k^n y, k^{-n}y)$ for the sake of simplicity. Since $\langle k \rangle$ acts with north-south dynamics on $\partial_\infty X$, the points $k^n x$ and $k^n y$ converge to $o_X(k^{+\infty})$, while $k^{-n}x$ and $k^{-n}y$ tend to $o_X(k^{-\infty})$ for $n \to +\infty$. Proposition 8.2.1 then yields an integer $N \geq 0$ such that $\mathscr{W}_n \subseteq \mathscr{W}(k)$ for all $n \geq N$. On the other hand, it is clear from the fact that k is good that $\mathscr{W}_n \cap \mathscr{W}(k) = \mathscr{W}(x|y) \cap \mathscr{W}(k)$. □

8.3 Trust issues.

Let Γ act properly and cocompactly on CAT(0) cube complexes X and Y. We fix two subsets $\mathcal{A} \subseteq \partial_{\mathrm{nt}} X$ and $\mathcal{B} \subseteq \partial_{\mathrm{nt}} Y$. Given $\mathcal{U} \subseteq \mathscr{W}(X)$, we employ the notation $\mathcal{U}(A|B) = \mathscr{W}(A|B) \cap \mathcal{U}$ for all subsets $A, B \subseteq \overline{X}$ and:

$$\mathrm{cr}_{\mathcal{U}}(x, y, z, w) = \#\mathcal{U}(x, z|y, w) - \#\mathcal{U}(x, w|y, z)$$

for all pairwise distinct points $x, y, z, w \in \partial_{\mathrm{nt}} X$. Given $\mathcal{V} \subseteq \mathscr{W}(Y)$, the same notation applies to subsets of $\overline{Y}$ and points of $\partial_{\mathrm{nt}} Y$.

As mentioned in the introduction, a key complication in the proof of Theorem H is the fact that $\mathrm{cr}_{\mathcal{U}}(x, y, z, w)$ does not provide any direct information on $\#\mathcal{U}(x, z|y, w)$ and $\#\mathcal{U}(x, w|y, z)$, only on their difference. The following notion is devised to address this problem.

Definition 8.3.1. We say that a 4-tuple $(a, b, c, d) \in (\partial_{\mathrm{nt}} X)^4$ is *$\mathcal{U}$-trustworthy* if at least one among the sets $\mathcal{U}(a, b|c, d)$, $\mathcal{U}(a, c|b, d)$ and $\mathcal{U}(a, d|b, c)$ is empty. When $\mathcal{U} = \mathscr{W}(X)$, we just say that (a, b, c, d) is *trustworthy*.

Lemma 8.3.2. *Given subsets $\mathcal{U} \subseteq \mathscr{W}(X)$ and $\mathcal{V} \subseteq \mathscr{W}(Y)$, consider pairwise distinct points $x_1, x_2, x_3, x_4 \in \mathcal{A}$ and a bijection $f \colon \mathcal{A} \to \mathcal{B}$ satisfying:*

$$\mathrm{cr}_{\mathcal{U}}\left(x_{\sigma(1)}, x_{\sigma(2)}, x_{\sigma(3)}, x_{\sigma(4)}\right) = \mathrm{cr}_{\mathcal{V}}\left(f(x_{\sigma(1)}), f(x_{\sigma(2)}), f(x_{\sigma(3)}), f(x_{\sigma(4)})\right)$$

for every permutation $\sigma \in \mathfrak{S}_4$. If (x_1, x_2, x_3, x_4) and $(f(x_1), f(x_2), f(x_3), f(x_4))$ are, respectively, $\mathcal{U}$-trustworthy and $\mathcal{V}$-trustworthy, then, for every $\sigma \in \mathfrak{S}_4$, we have:

$$\#\mathcal{U}\left(x_{\sigma(1)}, x_{\sigma(2)}|x_{\sigma(3)}, x_{\sigma(4)}\right) = \#\mathcal{V}\left(f(x_{\sigma(1)}), f(x_{\sigma(2)})|f(x_{\sigma(3)}), f(x_{\sigma(4)})\right).$$

Proof. Since (x_1, x_2, x_3, x_4) is $\mathcal{U}$-trustworthy, we can permute the four points so that $\mathcal{U}(x_1, x_2|x_3, x_4) = \emptyset$. In this case $\mathrm{cr}_{\mathcal{U}}(x_1, x_3, x_4, x_2) \geq 0$ and $\mathrm{cr}_{\mathcal{U}}(x_1, x_4, x_3, x_2) \geq 0$. We then have $\mathrm{cr}_{\mathcal{V}}(f(x_1), f(x_3), f(x_4), f(x_2)) \geq 0$ and $\mathrm{cr}_{\mathcal{V}}(f(x_1), f(x_4), f(x_3), f(x_2)) \geq 0$, which imply that the cardinality $\#\mathcal{V}(f(x_1), f(x_2)|f(x_3), f(x_4))$ is at most as large as the smallest of $\#\mathcal{V}(f(x_1), f(x_3)|f(x_2), f(x_4))$ and $\#\mathcal{V}(f(x_1), f(x_4)|f(x_2), f(x_3))$.

Since the 4-tuple $(f(x_1), f(x_2), f(x_3), f(x_4))$ is $\mathcal{V}$-trustworthy, this means that $\mathcal{V}(f(x_1), f(x_2)|f(x_3), f(x_4)) = \emptyset$. We thus have:

$$\begin{aligned}\#\mathcal{U}(x_1, x_3|x_2, x_4) &= \mathrm{cr}_{\mathcal{U}}(x_1, x_4, x_3, x_2) \\ &= \mathrm{cr}_{\mathcal{V}}(f(x_1), f(x_4), f(x_3), f(x_2)) = \#\mathcal{V}(f(x_1), f(x_3)|f(x_2), f(x_4)).\end{aligned}$$

The equality of $\#\mathcal{U}(x_1, x_4|x_2, x_3)$ and $\#\mathcal{V}(f(x_1), f(x_4)|f(x_2), f(x_3))$ is similar. □

The following will be our main source of trustworthy 4-tuples.

Definition 8.3.3. Let ξ and η be distinct points of $\partial_\infty X$. The pair (ξ, η) is *trustworthy* if there do not exist transverse hyperplanes $\mathfrak{u}_1$ and $\mathfrak{u}_2$ with $\xi, \eta \in \partial_\infty \mathfrak{u}_1 \cap \partial_\infty \mathfrak{u}_2$.

Lemma 8.3.4. *Consider four sequences $a_n, b_n, c_n, d_n \in \partial_{\mathrm{nt}} X$ and two distinct points $\xi, \eta \in \partial_\infty X$, where a_n and c_n converge to ξ, while b_n and d_n converge to η. If (ξ, η) is trustworthy, then there exists N such that (a_n, b_n, c_n, d_n) is trustworthy for all $n \geq N$.*

Proof. If (a_n, b_n, c_n, d_n) is not trustworthy for infinitely many values of n, we can pass to a subsequence in order to ensure that the sets $\mathscr{W}(a_n, b_n | c_n, d_n)$ and $\mathscr{W}(a_n, d_n | b_n, c_n)$ are all nonempty. Proposition 8.2.1 provides N such that, for all $n \geq N$, every element of the union $\mathscr{W}(a_n, b_n | c_n, d_n) \cup \mathscr{W}(a_n, d_n | b_n, c_n)$ contains ξ and η in its trace at infinity. Lemma 2.1.2, however, shows that the sets $\mathscr{W}(a_n, b_n | c_n, d_n)$ and $\mathscr{W}(a_n, d_n | b_n, c_n)$ are transverse, contradicting the fact that (ξ, η) is trustworthy. □

The next result applies for instance to the case when no three elements of $\mathcal{U}$ are transverse (see Lemma 2.1.2). It will only be needed in the proof of Theorem 8.5.3.

Lemma 8.3.5. *Suppose that, for a subset $\mathcal{U} \subseteq \mathscr{W}(X)$, every element of $\mathcal{A}^4$ is $\mathcal{U}$-trustworthy. Consider a partition $\mathcal{A} = \mathcal{P} \sqcup \mathcal{Q}$ with $\#\mathcal{P}, \#\mathcal{Q} \geq 2$ and such that $\mathcal{U}(x, y | z, w) \neq \emptyset$ for all $x, y \in \mathcal{P}$ and $z, w \in \mathcal{Q}$. Then, there exist $\overline{x}, \overline{y} \in \mathcal{P}$ and $\overline{z}, \overline{w} \in \mathcal{Q}$ such that $\mathcal{U}(\overline{x}, \overline{y} | \overline{z}, \overline{w}) = \mathcal{U}(\mathcal{P} | \mathcal{Q})$.*

Proof. Pick points $\overline{x}, \overline{y} \in \mathcal{P}$ and $\overline{z}, \overline{w} \in \mathcal{Q}$ so as to minimise the cardinality of $\mathcal{U}(\overline{x}, \overline{y} | \overline{z}, \overline{w})$; since $\#\mathcal{P}, \#\mathcal{Q} \geq 2$, this set is finite. It is clear that $\mathcal{U}(\overline{x}, \overline{y} | \overline{z}, \overline{w})$ contains $\mathcal{U}(\mathcal{P} | \mathcal{Q})$. Consider any hyperplane $\mathfrak{w} \in \mathcal{U}(\overline{x}, \overline{y} | \overline{z}, \overline{w})$ and let $\mathfrak{h}$ be its side containing $\overline{x}$ and $\overline{y}$. We are going to show that $\mathfrak{h} \cap \mathcal{A} \subseteq \mathcal{P}$ and the same argument will yield $\mathfrak{h}^* \cap \mathcal{A} \subseteq \mathcal{Q}$. This will conclude the proof as then $\mathfrak{h} \cap \mathcal{A} = \mathcal{P}$ and $\mathfrak{h}^* \cap \mathcal{A} = \mathcal{Q}$, which shows that $\mathfrak{w}$ separates $\mathcal{P}$ and $\mathcal{Q}$.

Suppose for the sake of contradiction that a point $u \in \mathfrak{h} \cap \mathcal{A}$ lies in $\mathcal{Q}$. By hypothesis, there exist hyperplanes $\mathfrak{u}_1 \in \mathcal{U}(\overline{x}, \overline{y} | u, \overline{z})$ and $\mathfrak{u}_2 \in \mathcal{U}(\overline{x}, \overline{y} | u, \overline{w})$. As $\mathcal{U}(\overline{x}, \overline{y} | u, \overline{z})$ and $\mathcal{U}(\overline{x}, \overline{y} | u, \overline{w})$ do not contain $\mathfrak{w}$ and cannot have fewer elements than $\mathcal{U}(\overline{x}, \overline{y} | \overline{z}, \overline{w})$, we must be able to choose $\mathfrak{u}_1$ and $\mathfrak{u}_2$ outside $\mathcal{U}(\overline{x}, \overline{y} | \overline{z}, \overline{w})$. In conclusion:

$$\mathfrak{u}_1 \in \mathcal{U}(\overline{x}, \overline{y}, \overline{w} | u, \overline{z}) \subseteq \mathcal{U}(\overline{y}, \overline{w} | u, \overline{z}), \quad \mathfrak{u}_2 \in \mathcal{U}(\overline{x}, \overline{y}, \overline{z} | u, \overline{w}) \subseteq \mathcal{U}(\overline{y}, \overline{z} | u, \overline{w}),$$

$$\mathfrak{w} \in \mathcal{U}(\overline{x}, \overline{y}, u | \overline{z}, \overline{w}) \subseteq \mathcal{U}(\overline{y}, u | \overline{z}, \overline{w}),$$

violating the fact that every 4-tuple, in particular $(u, \overline{y}, \overline{z}, \overline{w})$, is $\mathcal{U}$-trustworthy. □

8.4 Traces vs cross ratios.

We now assume that Γ is non-elementary. We consider proper cocompact actions of Γ on essential, hyperplane-essential CAT(0) cube complexes X and Y. These are irreducible by Remark 2.2.7. Writing $f := o_Y \circ o_X^{-1} \colon \partial_\infty X \to \partial_\infty Y$, let $\mathcal{A} \subseteq \partial_{\mathrm{nt}} X$ and $\mathcal{B} \subseteq \partial_{\mathrm{nt}} Y$ be nonempty Γ-invariant subsets with $f(\mathcal{A}) = \mathcal{B}$.

Assumption. Throughout this section, we fix subsets $\mathcal{U} \subseteq \mathscr{W}(X)$ and $\mathcal{V} \subseteq \mathscr{W}(Y)$ such that:

$$\mathrm{cr}_{\mathcal{U}}(x, y, z, w) = \mathrm{cr}_{\mathcal{V}}(f(x), f(y), f(z), f(w))$$

for all pairwise distinct points $x, y, z, w \in \mathcal{A}$. It will become clear in Section 8.5 how this relates to the proof of Theorem H.

The reader can think of the case where $\mathcal{U} = \mathscr{W}(X)$ and $\mathcal{V} = \mathscr{W}(Y)$, although we will need the full generality of the previous setup in Section 8.5.

Lemma 8.4.1. *Given a hyperplane $\mathfrak{u} \in \mathcal{U}$, a generic point $\xi \in \partial_\infty \mathfrak{u}$ and an arbitrary point $\eta \in \partial_\infty \mathfrak{u} \setminus \{\xi\}$, there exists a hyperplane $\mathfrak{v} \in \mathcal{V}$ with $f(\xi), f(\eta) \in \partial_\infty \mathfrak{v}$.*

Proof. Consider a (combinatorial) line $\gamma \subseteq C(\mathfrak{u})$ with endpoints ξ and η. Pick a hyperplane $\mathfrak{w}_0 \in \mathscr{W}(\gamma)$ and label the other elements of $\mathscr{W}(\gamma)$ as $\mathfrak{w}_n$, $n \in \mathbb{Z}$, according to the order in which they are crossed by γ and so that positive indices correspond to the half of γ ending at ξ. Let $\mathfrak{h}$ be any side of the hyperplane $\mathfrak{u}$ and let $\mathfrak{h}_n$ be the side of $\mathfrak{w}_n$ that contains the positive half of γ. For all $n > 0$, part (2) of Proposition 2.2.15 allows us to pick points of $\mathcal{A}$ as follows: $x_n \in \mathfrak{h} \cap \mathfrak{h}_n$, $y_n \in \mathfrak{h} \cap \mathfrak{h}^*_{-n}$, $z_n \in \mathfrak{h}^* \cap \mathfrak{h}_n$ and $w_n \in \mathfrak{h}^* \cap \mathfrak{h}^*_{-n}$. Note that every limit point of the x_n within ∂X must have infinite Gromov product with $\gamma(+\infty) \in \partial X$. Since $\Phi(\gamma(+\infty)) = \xi$, Lemma 4.2.2 and Corollary 4.2.8 show that $x_n \in \partial_{\mathrm{nt}} X$ converge to ξ. Similarly, z_n converge to ξ, while y_n and w_n converge to η.

By Remark 8.1.11 and Lemma 4.3.3, the generic point $\xi \in \partial_\infty \mathfrak{u}$ is represented by a regular point in the Roller boundary of the CAT(0) cube complex $\mathfrak{u}$. For all sufficiently large $n > 0$, it follows that $\mathfrak{w}_n \cap \mathfrak{u}$ and $\mathfrak{w}_{-n} \cap \mathfrak{u}$ are strongly separated as hyperplanes of the cube complex $\mathfrak{u}$. In other words, no hyperplane of X is transverse to $\mathfrak{u} \in \mathcal{U}(x_n, y_n | z_n, w_n)$ and $\mathfrak{w}_n, \mathfrak{w}_{-n} \in \mathscr{W}(x_n, z_n | y_n, w_n)$ at the same time. By Lemma 2.1.2, the sets $\mathscr{W}(x_n, w_n | y_n, z_n)$ are then all empty for large $n > 0$. It follows that $\mathrm{cr}_{\mathcal{U}}(x_n, z_n, y_n, w_n) > 0$, hence $\mathrm{cr}_{\mathcal{V}}(f(x_n), f(z_n), f(y_n), f(w_n)) > 0$ and $\mathcal{V}(f(x_n), f(y_n) | f(z_n), f(w_n)) \neq \emptyset$. As $f(x_n)$ and $f(z_n)$ converge to $f(\xi)$, while $f(y_n)$ and $f(w_n)$ converge to $f(\eta)$, Proposition 8.2.1 yields the required hyperplane $\mathfrak{v}$. $\square$

Definition 8.4.2. A hyperplane $\mathfrak{w} \in \mathcal{U}$ is *$\mathcal{U}$-boundary-maximal* if $\partial_\infty \mathfrak{w} \neq \emptyset$ and every hyperplane $\mathfrak{u} \in \mathcal{U}$ with $\mathfrak{w} \preceq \mathfrak{u}$ actually satisfies $\mathfrak{u} \sim \mathfrak{w}$. When $\mathcal{U} = \mathscr{W}(X)$, we simply speak of *boundary-maximal* hyperplanes.

Part (3) of Proposition 8.1.8 and Remark 8.1.4 show that boundary-maximal hyperplanes exist as soon as X is not a tree. In fact, for every hyperplane $\mathfrak{u}$ there exists a boundary-maximal hyperplane $\mathfrak{w}$ with $\mathfrak{u} \preceq \mathfrak{w}$.

The next result can be viewed as Step (IIa) from the introduction.

Proposition 8.4.3. *If $\mathfrak{u} \in \mathcal{U}$ is $\mathcal{U}$-boundary-maximal, there exists a $\mathcal{V}$-boundary-maximal hyperplane $\mathfrak{v} \in \mathcal{V}$ with $\partial_\infty \mathfrak{v} = f(\partial_\infty \mathfrak{u})$.*

Proof. Proposition 8.1.10 allows us to pick a generic point $\xi \in \partial_\infty \mathfrak{u}$. Note that $\mathfrak{u}$ is acted on properly and cocompactly by its stabiliser (Lemma 2.2.4), so $\partial_\infty \mathfrak{u} \setminus \{\xi\} \neq \emptyset$. In particular, Lemma 8.4.1 shows that there exists a hyperplane $\mathfrak{v} \in \mathcal{V}$ with $f(\xi) \in \partial_\infty \mathfrak{v}$. By part (3) of Proposition 8.1.8, we can take $\mathfrak{v}$ to be $\mathcal{V}$-boundary-maximal. Let $\zeta \in \partial_\infty \mathfrak{v} \setminus \{f(\xi)\}$ be generic.

Since $f(\xi)$ and ζ both lie in $\partial_\infty \mathfrak{v}$, Lemma 8.4.1 shows that ξ and $f^{-1}(\zeta)$ lie in the trace at infinity of some $\mathfrak{u}' \in \mathcal{U}$. Since ξ is generic, we have $\mathfrak{u} \preceq \mathfrak{u}'$ and, as $\mathfrak{u}$ is $\mathcal{U}$-boundary-maximal, we conclude that $\partial_\infty \mathfrak{u} = \partial_\infty \mathfrak{u}'$. In particular $f^{-1}(\zeta) \in \partial_\infty \mathfrak{u}$, showing that the closed subset $f(\partial_\infty \mathfrak{u})$ contains every generic point in $\partial_\infty \mathfrak{v} \setminus \{f(\xi)\}$. By Proposition 8.1.10, we have $\partial_\infty \mathfrak{v} \subseteq f(\partial_\infty \mathfrak{u})$ and we will now obtain the opposite inclusion with a similar argument.

Given generic points $\zeta \in \partial_\infty \mathfrak{v}$ and $\eta \in \partial_\infty \mathfrak{u} \setminus \{f^{-1}(\zeta)\}$, we apply Lemma 8.4.1 to the points $\eta, f^{-1}(\zeta) \in \partial_\infty \mathfrak{u}$. This shows that $f(\eta)$ and ζ lie in the trace at infinity of an element of $\mathcal{V}$, which, by genericity of ζ and $\mathcal{V}$-boundary-maximality of $\mathfrak{v}$, can be taken to coincide with $\mathfrak{v}$. Hence $f(\eta) \in \partial_\infty \mathfrak{v}$ and, by density of generic points, we have $f(\partial_\infty \mathfrak{u}) \subseteq \partial_\infty \mathfrak{v}$. □

We now prepare for Step (IIb), which will be completed in Theorem 8.5.3. Given a hyperplane $\mathfrak{w} \in \mathscr{W}(X)$ with $\partial_\infty \mathfrak{w} \neq \emptyset$, we denote by $\mathscr{T}_{\mathcal{U}}(\mathfrak{w})$ the set of hyperplanes $\mathfrak{u} \in \mathcal{U}$ satisfying $\mathfrak{u} \sim \mathfrak{w}$. We also write $\widetilde{\mathscr{T}_{\mathcal{U}}}(\mathfrak{w}) \subseteq \mathscr{H}(X)$ for the set of halfspaces bounded by elements of $\mathscr{T}_{\mathcal{U}}(\mathfrak{w})$. By Proposition 8.1.8, the set $\mathscr{T}_{\mathcal{U}}(\mathfrak{w})$ is finite and no two of its elements are transverse.

Definition 8.4.4. The *dual tree* $T(\mathfrak{w})$ is the restriction quotient (in the sense of p. 860 of [CS11]) determined by the subset $\mathscr{T}_{\mathcal{U}}(\mathfrak{w}) \subseteq \mathcal{U} \subseteq \mathscr{W}(X)$.

The observations above show that $T(\mathfrak{w})$ is a finite tree. Being a restriction quotient of X, it comes equipped with a projection $\pi_{\mathfrak{w}}\colon \overline{X} \to \overline{T(\mathfrak{w})} = T(\mathfrak{w})$. More precisely, $\pi_{\mathfrak{w}}$ takes the point $x \in \overline{X}$ to the point of $T(\mathfrak{w})$ represented by the ultrafilter $\sigma_x \cap \widetilde{\mathscr{T}_{\mathcal{U}}}(\mathfrak{w})$. We will be interested in the pseudo-metric $\delta_{\mathfrak{w}}$ defined on the set $\mathcal{A} \subseteq \partial_{\mathrm{nt}} X$ by the formula:

$$\delta_{\mathfrak{w}}(x,y) = d(\pi_{\mathfrak{w}}(x), \pi_{\mathfrak{w}}(y)) = \#(\mathscr{T}_{\mathcal{U}}(\mathfrak{w}) \cap \mathcal{U}(x|y)).$$

The same construction can be applied to hyperplanes $\mathfrak{w} \in \mathscr{W}(Y)$, yielding a subset $\mathscr{T}_{\mathcal{V}}(\mathfrak{w}) \subseteq \mathcal{V}$, a projection to a tree $\pi_{\mathfrak{w}}\colon \overline{Y} \to T(\mathfrak{w})$ and a pseudo-metric $\delta_{\mathfrak{w}}$ on $\mathcal{B}$.

The next result is the main ingredient of Step (IIb) (cf. Theorem 8.5.3). We will employ Definitions 8.2.3 and 8.3.3 in its proof.

Proposition 8.4.5. *Let $\mathfrak{u} \in \mathcal{U}$ and $\mathfrak{v} \in \mathcal{V}$ be, respectively, $\mathcal{U}$-boundary-maximal and $\mathcal{V}$-boundary-maximal, with $f(\partial_\infty \mathfrak{u}) = \partial_\infty \mathfrak{v}$. Given any two points $x, y \in \mathcal{A}$, we have $\delta_{\mathfrak{v}}(f(x), f(y)) = \delta_{\mathfrak{u}}(x,y)$.*

Proof. By Lemma 8.0.2, the subgroups $\Gamma_{\mathfrak{u}}$ and $\Gamma_{\mathfrak{v}}$ are commensurable. Let H be a common finite-index subgroup; in particular, H acts properly and cocompactly on both $\mathfrak{u}$ and $\mathfrak{v}$. Given an infinite-order element $h \in H$, we write $\mathcal{U}(h)$ and $\mathcal{V}(h)$ for the subsets of $\mathcal{U}$ and $\mathcal{V}$, respectively, corresponding to hyperplanes that contain $o_X(h^{\pm\infty})$ or $o_Y(h^{\pm\infty})$ in their trace at infinity. Note that we have $\mathscr{T}_{\mathcal{U}}(\mathfrak{u}) \subseteq \mathcal{U}(h)$ and $\mathscr{T}_{\mathcal{V}}(\mathfrak{v}) \subseteq \mathcal{V}(h)$ by Lemma 8.2.2.

Claim. *There exists an infinite-order element $h_0 \in H$ that is good in both X and Y, for which the pairs $(o_X(h_0^{-\infty}), o_X(h_0^{+\infty}))$ and $(o_Y(h_0^{-\infty}), o_Y(h_0^{+\infty}))$ are trustworthy, and for which the equalities $\mathscr{T}_{\mathcal{U}}(\mathfrak{u}) \cap \mathcal{U}(x|y) = \mathcal{U}(h_0) \cap \mathcal{U}(x|y)$ and $\mathscr{T}_{\mathcal{V}}(\mathfrak{v}) \cap \mathcal{V}(f(x)|f(y)) = \mathcal{V}(h_0) \cap \mathcal{V}(f(x)|f(y))$ are satisfied.*

The claim concludes the proof as follows. Lemma 8.3.4 shows that the 4-tuples $(h_0^n x, h_0^{-n} x | h_0^n y, h_0^{-n} y)$ and $(h_0^n f(x), h_0^{-n} f(x) | h_0^n f(y), h_0^{-n} f(y))$ are trustworthy for sufficiently large $n > 0$. Equivariance of f and Lemma 8.3.2 then imply that the sets $\mathcal{U}(h_0^n x, h_0^{-n} x | h_0^n y, h_0^{-n} y)$ and $\mathcal{V}(h_0^n f(x), h_0^{-n} f(x) | h_0^n f(y), h_0^{-n} f(y))$ have the same cardinality.

Since $x, y \in \partial_{\mathrm{nt}} X$, Lemma 8.1.5 shows that $x, y \notin \partial_\infty \mathfrak{u}$. On the other hand h_0 preserves $\mathfrak{u}$ and its two fixed points in $\partial_\infty X$ lie within $\partial_\infty \mathfrak{u}$ by Lemma 8.2.2. We conclude that h_0 does not fix x, y and, by Proposition 8.2.5, we finally obtain $\#(\mathcal{U}(h_0) \cap \mathcal{U}(x|y)) = \#(\mathcal{V}(h_0) \cap \mathcal{V}(f(x)|f(y)))$. By our choice of h_0, these two cardinalities are precisely $\delta_{\mathfrak{u}}(x,y)$ and $\delta_{\mathfrak{v}}(f(x), f(y))$.

Proof of claim. The points x, y are regular by Lemma 4.3.3. It follows that, among hyperplanes of X separating x and y, only finitely many are not strongly separated

from $\mathfrak{u}$. In light of Proposition 8.1.7, there are only finitely many hyperplanes $\mathfrak{u}'$ that separate x and y and satisfy $\partial_\infty \mathfrak{u}' \cap \partial_\infty \mathfrak{u} \neq \emptyset$; let us denote by $\mathfrak{u}_1, ..., \mathfrak{u}_r$ those that lie in $\mathcal{U} \setminus \mathscr{T}_{\mathcal{U}}(\mathfrak{u})$. Similarly let $\mathfrak{v}_1, ..., \mathfrak{v}_s$ be the elements of $\mathcal{V} \setminus \mathscr{T}_{\mathcal{V}}(\mathfrak{v})$ that separate $f(x)$ and $f(y)$ and satisfy $\partial_\infty \mathfrak{v}_i \cap \partial_\infty \mathfrak{v} \neq \emptyset$. By boundary-maximality of $\mathfrak{u}$ and $\mathfrak{v}$, we have $\partial_\infty \mathfrak{u}_i \cap \partial_\infty \mathfrak{u} \subsetneq \partial_\infty \mathfrak{u}$ and $\partial_\infty \mathfrak{v}_j \cap \partial_\infty \mathfrak{v} \subsetneq \partial_\infty \mathfrak{v}$ for all i and j. Lemma 8.0.1 and Proposition 8.1.7 then show that each of these subsets is nowhere-dense. Hence, the union $K \subseteq \partial_\infty H$ of all sets $o_X^{-1}(\partial_\infty \mathfrak{u}_i \cap \partial_\infty \mathfrak{u})$ and $o_Y^{-1}(\partial_\infty \mathfrak{v}_j \cap \partial_\infty \mathfrak{v})$ is nowhere-dense.

Observe that there exists $h \in H$ that acts as a strongly contracting automorphism on both $\mathfrak{u}$ and $\mathfrak{v}$. This follows from Lemma 5.2.1 if H is not virtually cyclic and from Lemma 8.1.12 otherwise. Since K is nowhere-dense, there exists a conjugate h_0 of h in H such that $h_0^{+\infty} \notin K$. Now, we have $o_X(h_0^{+\infty}) \notin \partial_\infty \mathfrak{u}_i$ and $o_Y(h_0^{+\infty}) \notin \partial_\infty \mathfrak{v}_j$ for all i and j, which implies that the inclusions $\mathscr{T}_{\mathcal{U}}(\mathfrak{u}) \cap \mathcal{U}(x|y) \subseteq \mathcal{U}(h_0) \cap \mathcal{U}(x|y)$ and $\mathscr{T}_{\mathcal{V}}(\mathfrak{v}) \cap \mathcal{V}(f(x)|f(y)) \subseteq \mathcal{V}(h_0) \cap \mathcal{V}(f(x)|f(y))$ are equalities.

Since h_0 acts as a strongly contracting automorphism on $\mathfrak{u}$, the points $o_X(h_0^{\pm\infty})$ lie in $\partial_{\mathrm{nt}}\mathfrak{u}$. Lemma 8.1.5 then shows that the pair $(o_X(h_0^{-\infty}), o_X(h_0^{+\infty}))$ is trustworthy. Similarly, we see that $(o_Y(h_0^{-\infty}), o_Y(h_0^{+\infty}))$ is trustworthy. Finally, as h_0 has infinite order, Remark 8.2.4 ensures that a power of h_0 is good. This concludes the proof of the claim and of the proposition. □

In the setting of Proposition 8.4.5, we get the commutative diagram:

$$\begin{array}{ccccc} \mathcal{A} & \xrightarrow{\pi_{\mathfrak{u}}} & \pi_{\mathfrak{u}}(\mathcal{A}) & \hookrightarrow & T(\mathfrak{u}) \\ \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle \psi} & & \vdots{\scriptstyle \Psi} \\ \mathcal{B} & \xrightarrow{\pi_{\mathfrak{v}}} & \pi_{\mathfrak{v}}(\mathcal{B}) & \hookrightarrow & T(\mathfrak{v}). \end{array}$$

Here $\pi_{\mathfrak{u}}(\mathcal{A})$ and $\pi_{\mathfrak{v}}(\mathcal{B})$ are precisely the quotient metric spaces associated to the pseudo-metric spaces $(\mathcal{A}, \delta_{\mathfrak{u}})$ and $(\mathcal{B}, \delta_{\mathfrak{v}})$. The map ψ is then provided by Proposition 8.4.5 and it is an isometry.

By part (1) of Proposition 2.2.15, we see that $\pi_{\mathfrak{u}}(\mathcal{A})$ and $\pi_{\mathfrak{v}}(\mathcal{B})$ contain all vertices of degree one in $T(\mathfrak{u})$ and $T(\mathfrak{v})$, respectively. The dashed arrow Ψ is then obtained through the following general fact about trees.

Lemma 8.4.6. *Let T_1 and T_2 be finite trees with all edges of length one. Let $V_i \subseteq T_i$ be sets of vertices containing all degree-one vertices of T_i. Every distance preserving bijection $\psi\colon V_1 \to V_2$ uniquely extends to an isomorphism $\Psi\colon T_1 \to T_2$.*

Proof. Let $V_i' \subseteq V_i$ denote the subsets of vertices of degree one. If $\#V_i' \leq 2$, the tree T_i is a segment and the lemma is clear. Let us therefore assume that V_1' and V_2' both contain at least three elements.

Extending every leaf of T_i to a ray, we embed T_i in a geodesically complete tree $\mathcal{T}_i$ with a natural bijection $\phi_i\colon V_i' \to \partial T_i$. Observing that the maps ϕ_i preserve cross ratios and ψ is an isometry, we conclude that the bijection $\phi_2\psi\phi_1^{-1}\colon \partial T_1 \to \partial T_2$ preserves cross ratios. It then follows from Theorem 4.3 in [BS17] that $\phi_2\psi\phi_1^{-1}$ admits a unique extension to an isometry $\Psi\colon \mathcal{T}_1 \to \mathcal{T}_2$.

Note that $m(x,y,z) = m(\phi_i(x),\phi_i(y),\phi_i(z))$ for all distinct points $x,y,z \in V_i'$. It follows that, given pairwise distinct points $x,y,z \in V_1'$, the map Ψ takes the median $m_1 := m(x,y,z)$ to $m_2 := m(\psi(x),\psi(y),\psi(z))$. Moreover, Ψ takes the ray from m_1 to $\phi_1(x)$ to the ray from m_2 to $\phi_2(\psi(x))$. Since

$$\begin{aligned} d(m_2,\Psi(x)) &= d(m_1,x) = \tfrac{1}{2}\cdot[d(x,y)+d(x,z)-d(y,z)] \\ &= \tfrac{1}{2}\cdot[d(\psi(x),\psi(y))+d(\psi(x),\psi(z))-d(\psi(y),\psi(z))] = d(m_2,\psi(x)), \end{aligned}$$

we conclude that $\Psi(x) = \psi(x)$. Thus Ψ and ψ coincide on V_1'. For every other point $w \in V_1$, we can find $x,y \in V_1'$ such that $d(x,y) = d(x,w) + d(w,y)$. It is then clear that Ψ and ψ coincide on the entire V_1. □

Corollary 8.4.7. *Let $\mathfrak{u}$ and $\mathfrak{v}$ be as in Proposition 8.4.5. For every $\mathfrak{h} \in \mathscr{H}(X)$:*

$$\#\{\mathfrak{j} \in \widetilde{\mathcal{T}_{\mathcal{U}}}(\mathfrak{u}) \mid \mathfrak{j}\cap\mathcal{A} = \mathfrak{h}\cap\mathcal{A}\} = \#\{\mathfrak{m} \in \widetilde{\mathcal{T}_{\mathcal{V}}}(\mathfrak{v}) \mid \mathfrak{m}\cap\mathcal{B} = f(\mathfrak{h}\cap\mathcal{A})\}.$$

Proof. We begin by observing that it suffices to prove the inequality

$$\#\{\mathfrak{j} \in \widetilde{\mathcal{T}_{\mathcal{U}}}(\mathfrak{u}) \mid \mathfrak{j}\cap\mathcal{A} = \mathfrak{h}\cap\mathcal{A}\} \leq \#\{\mathfrak{m} \in \widetilde{\mathcal{T}_{\mathcal{V}}}(\mathfrak{v}) \mid \mathfrak{m}\cap\mathcal{B} = f(\mathfrak{h}\cap\mathcal{A})\}. \qquad (*)$$

This already yields an equality if the right-hand side vanishes. Otherwise, there exists $\mathfrak{m}' \in \widetilde{\mathcal{T}_{\mathcal{V}}}(\mathfrak{v})$ with $\mathfrak{m}'\cap\mathcal{B} = f(\mathfrak{h}\cap\mathcal{A})$ and we can apply the same argument to f^{-1} and $\mathfrak{m}'$ to obtain the opposite inequality.

Now, if the left-hand side of $(*)$ vanishes, there is nothing to prove. Otherwise, there exists a halfspace $\mathfrak{j}' \in \widetilde{\mathcal{T}_{\mathcal{U}}}(\mathfrak{u})$ with $\mathfrak{j}'\cap\mathcal{A} = \mathfrak{h}\cap\mathcal{A}$ and our setup is not affected if we replace $\mathfrak{h}$ with $\mathfrak{j}'$. We can thus assume that $\mathfrak{h} \in \widetilde{\mathcal{T}_{\mathcal{U}}}(\mathfrak{u})$.

The projection $\pi_{\mathfrak{u}}(\mathfrak{h})$ is a halfspace of the tree $T(\mathfrak{u})$, with complement $\pi_{\mathfrak{u}}(\mathfrak{h}^*)$. Let C and C^* be the convex hulls, respectively, of $\pi_{\mathfrak{u}}(\mathfrak{h}\cap\mathcal{A})$ and $\pi_{\mathfrak{u}}(\mathfrak{h}^*\cap\mathcal{A})$. These are disjoint subtrees of $T(\mathfrak{u})$ and the union $C \sqcup C^*$ contains all degree-one vertices of $T(\mathfrak{u})$. The complement of $C \sqcup C^*$ must be an open arc $\alpha \subseteq T(\mathfrak{u})$ such that every vertex in α has degree two and lies outside $\pi_{\mathfrak{u}}(\mathcal{A})$. Thus, the hyperplanes associated to the set $\{\mathfrak{j} \in \widetilde{\mathcal{T}_{\mathcal{U}}}(\mathfrak{u}) \mid \mathfrak{j}\cap\mathcal{A} = \mathfrak{h}\cap\mathcal{A}\}$ are precisely the elements of $\mathcal{T}_{\mathcal{U}}(\mathfrak{u})$ that are dual to edges of α.

Observing that $\Psi\circ\pi_{\mathfrak{u}} = \pi_{\mathfrak{v}}\circ f$, we see that the sets $\Psi(C)$ and $\Psi(C^*)$ are the convex hulls, respectively, of $\pi_{\mathfrak{v}}(f(\mathfrak{h}\cap\mathcal{A}))$ and $\pi_{\mathfrak{v}}(f(\mathfrak{h}^*\cap\mathcal{A}))$ in $T(\mathfrak{v})$. Moreover, the arc $\Psi(\alpha)$ has the same length of α and its edges correspond to pairwise distinct elements of the set $\{\mathfrak{m}\in\widetilde{\mathscr{T}_{\mathcal{V}}}(\mathfrak{v}) \mid \mathfrak{m}\cap\mathcal{B} = f(\mathfrak{h}\cap\mathcal{A})\}$. This yields the desired inequality. □

8.5 Concluding the proof.

Let the group Γ, the cube complexes X and Y and the map f be as in the statement of Theorem H. We set $\mathcal{A} = \Omega$ and $\mathcal{B} = f(\Omega)$, so that we are in the general setup of Section 8.4. Here, however, we do not fix sets $\mathcal{U}$ and $\mathcal{V}$. Instead, we observe that, by the hypotheses of Theorem H, we have $\mathrm{cr}_X(x,y,z,w) = \mathrm{cr}_Y(f(x),f(y),f(z),f(w))$ for all pairwise distinct points $x,y,z,w\in\mathcal{A}$.

Given $\mathcal{U}\subseteq\mathscr{W}(X)$ and $\mathcal{V}\subseteq\mathscr{W}(Y)$, we denote by $\widetilde{\mathcal{U}}\subseteq\mathscr{H}(X)$ and $\widetilde{\mathcal{V}}\subseteq\mathscr{H}(Y)$ the collections of halfspaces bounded, respectively, by the elements of $\mathcal{U}$ and $\mathcal{V}$.

Definition 8.5.1. We say that subsets $\mathcal{U}\subseteq\mathscr{W}(X)$ and $\mathcal{V}\subseteq\mathscr{W}(Y)$ are *well-paired* if, for every $\mathfrak{h}\in\mathscr{H}(X)$ and $\mathfrak{k}\in\mathscr{H}(Y)$, we have:

$$\#\{\mathfrak{j}\in\widetilde{\mathcal{U}} \mid \mathfrak{j}\cap\mathcal{A} = \mathfrak{h}\cap\mathcal{A}\} = \#\{\mathfrak{m}\in\widetilde{\mathcal{V}} \mid \mathfrak{m}\cap\mathcal{B} = f(\mathfrak{h}\cap\mathcal{A})\},$$

$$\#\{\mathfrak{m}\in\widetilde{\mathcal{V}} \mid \mathfrak{m}\cap\mathcal{B} = \mathfrak{k}\cap\mathcal{B}\} = \#\{\mathfrak{j}\in\widetilde{\mathcal{U}} \mid \mathfrak{j}\cap\mathcal{A} = f^{-1}(\mathfrak{k}\cap\mathcal{B})\}.$$

In particular, if $\mathfrak{h}\in\widetilde{\mathcal{U}}$ and $\mathfrak{k}\in\widetilde{\mathcal{V}}$, the right-hand sides must be nonempty.

The following observation is immediate from definitions.

Lemma 8.5.2. *If the sets $\mathcal{U}\subseteq\mathscr{W}(X)$ and $\mathcal{V}\subseteq\mathscr{W}(Y)$ are well-paired and the points $x,y,z,w\in\mathcal{A}$ are pairwise distinct, we have:*

$$\mathcal{U}(x,y|z,w) = \mathcal{V}(f(x),f(y)|f(z),f(w)),$$

$$\mathrm{cr}_{\mathcal{U}}(x,y,z,w) = \mathrm{cr}_{\mathcal{V}}(f(x),f(y),f(z),f(w)).$$

The next result is the correct formulation of Step (IIb) from the introduction, as traces at infinity need to be counted "with multiplicity".

Theorem 8.5.3. *The sets $\mathscr{W}(X)$ and $\mathscr{W}(Y)$ are well-paired.*

Proof. We set $\mathcal{U}_0 = \mathscr{W}(X)$ and define inductively $\mathcal{U}_{k+1}\subsetneq\mathcal{U}_k$ as the subset of hyperplanes that are not $\mathcal{U}_k$-boundary-maximal. Let $\mathcal{U}_k^c = \mathscr{W}(X)\setminus\mathcal{U}_k$ and similarly define the sets $\mathcal{V}_k$ and $\mathcal{V}_k^c$ starting from $\mathcal{V}_0 = \mathscr{W}(Y)$.

We will now show by induction on k that the sets $\mathcal{U}_k^c$ and $\mathcal{V}_k^c$ are well-paired. The base case $k = 0$ is trivial, as $\mathcal{U}_0^c$ and $\mathcal{V}_0^c$ are empty. Assuming that $\mathcal{U}_k^c$ and $\mathcal{V}_k^c$ are well-paired for some $k \geq 0$, Lemma 8.5.2 yields

$$\mathrm{cr}_{\mathcal{U}_k^c}(x, y, z, w) = \mathrm{cr}_{\mathcal{V}_k^c}(f(x), f(y), f(z), f(w)),$$

whenever $x, y, z, w \in \mathcal{A}$ are pairwise distinct. On the other hand, observing that $\mathrm{cr}_X = \mathrm{cr}_{\mathcal{U}_k} + \mathrm{cr}_{\mathcal{U}_k^c}$ and $\mathrm{cr}_Y = \mathrm{cr}_{\mathcal{V}_k} + \mathrm{cr}_{\mathcal{V}_k^c}$, the fact that f takes cr_X to cr_Y implies that

$$\mathrm{cr}_{\mathcal{U}_k}(x, y, z, w) = \mathrm{cr}_{\mathcal{V}_k}(f(x), f(y), f(z), f(w)).$$

We are now in the setting of Section 8.4 and can apply Proposition 8.4.3. It follows that, for each $\mathfrak{u} \in \mathcal{U}_k \setminus \mathcal{U}_{k+1}$, there exists $\mathfrak{v} \in \mathcal{V}_k \setminus \mathcal{V}_{k+1}$ with $\partial_\infty \mathfrak{v} = f(\partial_\infty \mathfrak{u})$. Applying Corollary 8.4.7 to both f and f^{-1}, we see that the sets $\mathscr{T}_{\mathcal{U}_k}(\mathfrak{u})$ and $\mathscr{T}_{\mathcal{V}_k}(\mathfrak{v})$ are well-paired. Letting $\mathfrak{u}$ vary, these sets partition $\mathcal{U}_k \setminus \mathcal{U}_{k+1}$ and $\mathcal{V}_k \setminus \mathcal{V}_{k+1}$; we conclude that the latter are also well-paired. Observing that $\mathcal{U}_{k+1}^c = \mathcal{U}_k^c \sqcup (\mathcal{U}_k \setminus \mathcal{U}_{k+1})$ and $\mathcal{V}_{k+1}^c = \mathcal{V}_k^c \sqcup (\mathcal{V}_k \setminus \mathcal{V}_{k+1})$, we have finally shown that the sets $\mathcal{U}_{k+1}^c$ and $\mathcal{V}_{k+1}^c$ are well-paired, completing the proof of the inductive step.

Part (3) of Proposition 8.1.8 shows that, for sufficiently large values of k, the sets $\mathcal{U}_k$ and $\mathcal{V}_k$ are reduced to the subsets $\overline{\mathcal{U}} \subseteq \mathscr{W}(X)$ and $\overline{\mathcal{V}} \subseteq \mathscr{W}(Y)$ of hyperplanes with empty trace at infinity. We conclude the proof by showing that $\overline{\mathcal{U}}$ and $\overline{\mathcal{V}}$ are well-paired. Note that the arguments above already yield

$$\mathrm{cr}_{\overline{\mathcal{U}}}(x, y, z, w) = \mathrm{cr}_{\overline{\mathcal{V}}}(f(x), f(y), f(z), f(w))$$

for all pairwise distinct points $x, y, z, w \in \mathcal{A}$. By Remark 8.1.4, no two elements of $\overline{\mathcal{U}}$ or $\overline{\mathcal{V}}$ are transverse; in particular, Lemma 2.1.2 shows that every element of $\mathcal{A}^4$ is $\overline{\mathcal{U}}$-trustworthy and every element of $\mathcal{B}^4$ is $\overline{\mathcal{V}}$-trustworthy. Lemma 8.3.2 guarantees that

$$\#\overline{\mathcal{U}}(x, y|z, w) = \#\overline{\mathcal{V}}(f(x), f(y)|f(z), f(w)).$$

Thus, if $\mathfrak{h} \in \mathscr{H}(X)$ is bounded by an element of $\overline{\mathcal{U}}$, the set $\overline{\mathcal{V}}(x', y'|z', w')$ is nonempty for all $x', y' \in f(\mathfrak{h} \cap \mathcal{A})$ and $z', w' \in f(\mathfrak{h}^* \cap \mathcal{A})$. Lemma 8.3.5 then provides points $\overline{x}, \overline{y} \in f(\mathfrak{h} \cap \mathcal{A})$ and $\overline{z}, \overline{w} \in f(\mathfrak{h}^* \cap \mathcal{A})$ with

$$\begin{aligned}&\#\overline{\mathcal{V}}\big(f(\mathfrak{h} \cap \mathcal{A})\big|f(\mathfrak{h}^* \cap \mathcal{A})\big) = \#\overline{\mathcal{V}}\big(\overline{x}, \overline{y}\big|\overline{z}, \overline{w}\big)\\ =&\#\overline{\mathcal{U}}\big(f^{-1}(\overline{x}), f^{-1}(\overline{y})\big|f^{-1}(\overline{z}), f^{-1}(\overline{w})\big) \geq \#\overline{\mathcal{U}}\big(\mathfrak{h} \cap \mathcal{A}\big|\mathfrak{h}^* \cap \mathcal{A}\big) > 0.\end{aligned}$$

The opposite inequality is obtained with a similar argument and we conclude that $\#\overline{\mathcal{U}}(\mathfrak{h} \cap \mathcal{A}|\mathfrak{h}^* \cap \mathcal{A}) = \#\overline{\mathcal{V}}(f(\mathfrak{h} \cap \mathcal{A})|f(\mathfrak{h}^* \cap \mathcal{A}))$. This shows that $\overline{\mathcal{U}}$ and $\overline{\mathcal{V}}$ are well-paired, thus completing the proof of the theorem. □

Proof of Theorem H. We are going to show that $f\colon \partial_\infty X \to \partial_\infty Y$ induces a Γ-equivariant pocset isomorphism $f_*\colon (\mathscr{H}(X), \subseteq, *) \to (\mathscr{H}(Y), \subseteq, *)$. This will then yield a Γ-equivariant cubical isomorphism $F\colon X \to Y$ by Roller duality; see e.g. [Nic04, CN05, Sag14]. Uniqueness and the fact that F and f coincide on Ω will be clear from the construction.

We start by observing that the set $\mathcal{H}(\mathfrak{h}) := \{\mathfrak{j} \in \mathscr{H}(X) \mid \mathfrak{j} \cap \mathcal{A} = \mathfrak{h} \cap \mathcal{A}\}$ is totally ordered by inclusion for each $\mathfrak{h} \in \mathscr{H}(X)$. Indeed, given $\mathfrak{j}_1, \mathfrak{j}_2 \in \mathcal{H}(\mathfrak{h})$, the intersections $\mathfrak{j}_1 \cap \mathfrak{j}_2 \supseteq \mathfrak{h} \cap \mathcal{A}$ and $\mathfrak{j}_1^* \cap \mathfrak{j}_2^* \supseteq \mathfrak{h}^* \cap \mathcal{A}$ are nonempty by part (1) of Proposition 2.2.15. Moreover, $\mathfrak{j}_1$ and $\mathfrak{j}_2$ cannot be transverse, or part (2) of Proposition 2.2.15 would yield $\mathfrak{h} \cap \mathfrak{h}^* \supseteq \mathfrak{j}_1 \cap \mathfrak{j}_2^* \cap \mathcal{A} \neq \emptyset$. Hence $\mathfrak{j}_1 \subseteq \mathfrak{j}_2$ or $\mathfrak{j}_2 \subseteq \mathfrak{j}_1$.

Note that $\mathcal{H}(\mathfrak{h})$ is finite as, by part (2) of Lemma 4.2.2, the set $\#\mathscr{W}(x, y|z, w)$ is finite for all distinct points $x, y \in \mathfrak{h} \cap \mathcal{A}$, $z, w \in \mathfrak{h}^* \cap \mathcal{A}$. By Theorem 8.5.3, the set $f_*\mathcal{H}(\mathfrak{h}) := \{\mathfrak{m} \in \mathscr{H}(Y) \mid \mathfrak{m} \cap \mathcal{B} = f(\mathfrak{h} \cap \mathcal{A})\}$ is a chain of the same length as $\mathcal{H}(\mathfrak{h})$. Thus, there exists a unique order-preserving bijection between $\mathcal{H}(\mathfrak{h})$ and $f_*\mathcal{H}(\mathfrak{h})$ and this is exactly how we define f_* on $\mathcal{H}(\mathfrak{h})$.

Since the sets $\mathcal{H}(\mathfrak{h})$ partition $\mathscr{H}(X)$, we have defined a map $f_*\colon \mathscr{H}(X) \to \mathscr{H}(Y)$. This is a bijection, an inverse being provided by the same construction applied to f^{-1}. It is also clear that $f_*(\mathfrak{h}^*) = f_*(\mathfrak{h})^*$, as $\mathcal{H}(\mathfrak{h}^*)$ is exactly the set of complements of the elements of $\mathcal{H}(\mathfrak{h})$. We are thus only left to show that f_* preserves inclusions.

Consider $\mathfrak{h}, \mathfrak{k} \in \mathscr{H}(X)$ with $\mathfrak{h} \subseteq \mathfrak{k}$. We can assume that $\mathfrak{k} \notin \mathcal{H}(\mathfrak{h})$, as we already know that f is order-preserving on $\mathcal{H}(\mathfrak{h})$. Hence $\mathfrak{h} \cap \mathcal{A} \subsetneq \mathfrak{k} \cap \mathcal{A}$ and, by construction, $f_*(\mathfrak{h}) \cap \mathcal{B} = f(\mathfrak{h} \cap \mathcal{A}) \subsetneq f(\mathfrak{k} \cap \mathcal{A}) = f_*(\mathfrak{k}) \cap \mathcal{B}$. Part (3) of Proposition 2.2.15 finally yields $f_*(\mathfrak{h}) \subsetneq f_*(\mathfrak{k})$, concluding the proof. □

8.6 Epilogue.

We can finally use Theorem H to prove Theorem B.

Proof of Theorem B. If Γ is non-elementary, Theorem F provides a Γ-equivariant, Möbius bijection $f\colon \mathcal{A} \to \mathcal{B}$, where $\mathcal{A} \subseteq \partial_{\mathrm{reg}} X$ and $\mathcal{B} \subseteq \partial_{\mathrm{reg}} Y$ are nonempty Γ-invariant subsets. Recall from Section 5.2 that the map f (there called ψ) is just a restriction of the unique Γ-equivariant homeomorphism $o_Y \circ o_X^{-1}\colon \partial_\infty X \to \partial_\infty Y$. Since regular points are non-terminating, we conclude by Theorem H.

We are left to consider the case when Γ is virtually cyclic. If Γ is finite, X and Y must be single points, by essentiality; so let us assume that Γ is virtually isomorphic to $\mathbb{Z}$. By part (1) of Lemma 8.1.12, every hyperplane of X is compact and, since now X is hyperplane-essential, we must have $X \simeq \mathbb{R}$. The action $\Gamma \curvearrowright X$ factors

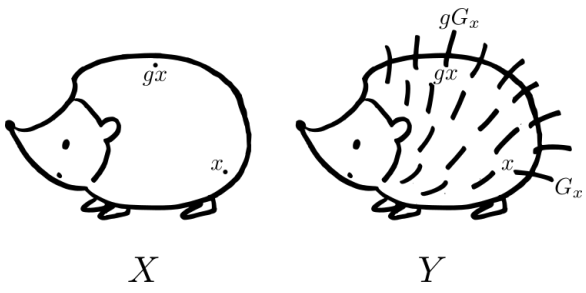


Figure 8.1: Giving a hedgehog back its spines.

through a faithful action of either $\mathbb{Z}$ or D_∞. In the former case, the only $g \in \Gamma$ with $\ell_X(g) = 0$ are those in the (finite) kernel of the action $\Gamma \curvearrowright X$. In the latter case, we have infinitely many $g \in \Gamma$ with $\ell_X(g) = 0$, for instance all reflections.

Since $\ell_X = \ell_Y$, the actions $\Gamma \curvearrowright X$ and $\Gamma \curvearrowright Y$ either both factor through a faithful action of $\mathbb{Z}$ or both factor through a faithful action of D_∞. In the former case, the two actions must coincide, as both the kernel and the $\mathbb{Z}$-action can be described in terms of length functions. In the latter, the two actions are Γ-equivariantly isomorphic, since actions $D_\infty \curvearrowright \mathbb{R}$ are determined, up to conjugacy, by the restriction to the maximal $\mathbb{Z}$ subgroup. □

As the next two examples demonstrate, there is no way of removing the essentiality and hyperplane-essentiality requirements from Theorems B and J.

Example 8.6.1. Let $\Gamma \curvearrowright X$ be any proper cocompact action on a CAT(0) cube complex. Fix a basepoint $x \in X$ and let $\Gamma_x \leq \Gamma$ denote its stabiliser. The disjoint union $Y = X \sqcup \Gamma/\Gamma_x$ is endowed with a natural Γ-action and we give it a structure of CAT(0) cube complex by adding edges connecting $g\Gamma_x$ and gx for every $g \in \Gamma$. The procedure is depicted in Figure 8.1 in a more general context.

The CAT(0) cube complex Y is irreducible and the action $\Gamma \curvearrowright Y$ is proper and cocompact. Note that Y is hyperplane-essential if and only if X is, but Y is never going to be essential, as all points of Γ/Γ_x are vertices of degree 1. It is easy to see that $\ell_X = \ell_Y$ and $\mathrm{cr}_X = \mathrm{cr}_Y$.

We now describe a general procedure that takes any finite dimensional CAT(0) cube complex X as input and gives out another CAT(0) cube complex $S(X)$ as output. We will refer to $S(X)$ as the *squarisation* of X. One can already get a good idea of the definition by looking at Figure 8.2.

Let $\mathscr{H}'$ denote the disjoint union of two copies of the pocset $(\mathscr{H}(X), \subseteq, *)$, labelling by $\mathfrak{h}_1$ and $\mathfrak{h}_2$ the two elements arising from $\mathfrak{h} \in \mathscr{H}(X)$. We turn $\mathscr{H}'$ into

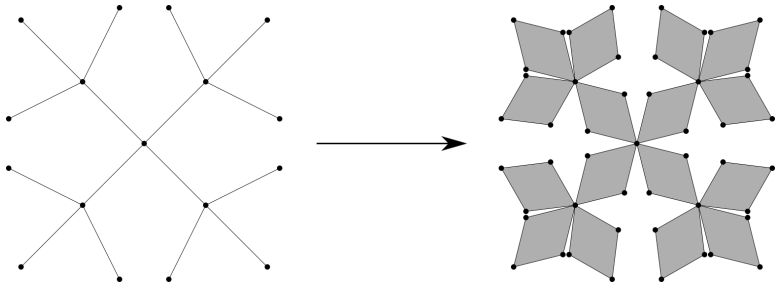


Figure 8.2: Part of a 4-regular tree T and its squarisation $S(T)$.

a pocset by declaring that, given $\mathfrak{h}, \mathfrak{k} \in \mathscr{H}(X)$ with $\mathfrak{h} \notin \{\mathfrak{k}, \mathfrak{k}^*\}$, we have $\mathfrak{h}_i \subseteq \mathfrak{k}_j$ if and only if $\mathfrak{h} \subseteq \mathfrak{k}$, no matter what the indices i and j are. On the other hand, the halfspaces $\mathfrak{h}_1$ and $\mathfrak{h}_2$ are transverse for all $\mathfrak{h} \in \mathscr{H}(X)$.

Now, $S(X)$ is obtained by applying Sageev's construction [Sag95, Sag14] to the pocset $(\mathscr{H}', \subseteq, *)$. Note that, if we had instead declared that $\mathfrak{h}_1$ and $\mathfrak{h}_2$ are nested, we would have obtained the first cubical subdivision X'.

Example 8.6.2. Let $\Gamma \curvearrowright X$ be any proper cocompact action on an essential, irreducible CAT(0) cube complex. The cube complexes X' and $S(X)$ are both essential, irreducible and naturally endowed with proper cocompact actions of Γ. It is easy to see that the actions $\Gamma \curvearrowright X'$ and $\Gamma \curvearrowright S(X)$ determine the same length function, namely the double of the length function associated to $\Gamma \curvearrowright X$. Moreover, the (comeagre) subset of $\partial_\infty \Gamma$ arising from non-terminating ultrafilters is the same for X, X' and $S(X)$ and there we have $\mathrm{cr}_{S(X)} = \mathrm{cr}_{X'} = 2 \cdot \mathrm{cr}_X$.

The failure of Theorems B and J is to be traced back to hyperplane-essentiality. Indeed, $S(X)$ is never hyperplane-essential. All its hyperplanes split as $S(\mathfrak{w}) \times [0, 1]$, where $\mathfrak{w}$ is the corresponding hyperplane of X.

Chapter 9

The space of cubical actions of a group.

In this chapter, we use ultralimits and median spaces to prove Theorem C. Our arguments are largely independent of the rest of the thesis. Let us fix a group Γ with a finite generating set $S \subseteq \Gamma$. We denote by $\|\cdot\|$ the word norm on Γ induced by S.

Suppose for a moment that Γ acts by isometries on a metric space (M, d). Given $x \in M$ and $g \in \Gamma$, we can consider the quantities:

$$\tau_M(x) = \max_{s \in S} d(x, sx), \qquad \tau_M = \inf_{x \in M} \tau_M(x), \qquad \ell_M(g) = \inf_{x \in M} d(x, gx).$$

Lemma 9.0.1. *For every $g \in \Gamma$ and $x \in M$, we have $d(x, gx) \leq \tau_M(x) \cdot \|g\|$. In particular, $\ell_M(g) \leq \tau_M \cdot \|g\|$.*

Proof. Writing $g = s_1 s_2 ... s_k$ with $s_i \in S$ and $k = \|g\|$, we have:

$$d(x, gx) \leq \sum_{i=0}^{k-1} d(s_1 ... s_i \cdot x, s_1 ... s_{i+1} \cdot x) \leq \sum_{i=1}^{k} d(x, s_i x) \leq k \cdot \tau_M(x). \qquad \square$$

If Γ acts on a CAT(0) cube complex X, we compute τ_X and ℓ_X with respect to the combinatorial metric, as usual. Given $g \in \Gamma$ and a point x in the interior of a cube $c \subseteq X$, it is not hard to see that there exists a vertex v of the barycentric subdivision $c' \subseteq X'$ with $d(v, gv) \leq d(x, gx)$. It follows that τ_X is always attained at a vertex of the barycentric subdivision X'.

We now consider a sequence of actions on CAT(0) cube complexes $\Gamma \curvearrowright X_n$ with $\dim X_n \leq D < +\infty$. For simplicity, we will write ℓ_n and τ_n with the meaning of ℓ_{X_n} and τ_{X_n}. Subdividing the complexes if necessary, we pick base vertices $o_n \in X_n$ with $\tau_n(o_n) = \tau_n$. We also fix a non-principal ultrafilter $\omega \subseteq 2^{\mathbb{N}}$.

Let (X_ω, o_ω) denote the ultralimit $\lim_\omega(\frac{1}{\tau_n} \cdot X_n, o_n)$. This is a complete, geodesic, median space. If moreover $\dim X_n \leq D < +\infty$ for all $n \geq 0$, the rank of X_ω is at most D (see e.g. Theorem 2.3 in [Bow13]). For every $g \in \Gamma$, Lemma 9.0.1 shows that:

$$d(o_\omega, go_\omega) = \lim_\omega \tfrac{1}{\tau_n} \cdot d(o_n, go_n) \leq \lim_\omega \tfrac{1}{\tau_n} \cdot \tau_n(o_n) \cdot \|g\| = \|g\| < +\infty.$$

It follows that we have an isometric action $\Gamma \curvearrowright X_\omega$. Let us write τ_ω and ℓ_ω rather than τ_{X_ω} and ℓ_{X_ω}.

Proposition 9.0.2. *We have $\tau_\omega = 1$ and $\ell_\omega(g^{D!}) = \lim_\omega \frac{1}{\tau_n} \cdot \ell_n(g^{D!})$ for all $g \in \Gamma$.*

Proof. It is clear that $\tau_\omega(o_\omega) = \lim_\omega \frac{1}{\tau_n} \cdot \tau_n(o_n) = 1$. On the other hand, for every $x_\omega = (x_n) \in X_\omega$, we have $\tau_\omega(x_\omega) = \lim_\omega \frac{1}{\tau_n} \cdot \tau_n(x_n) \geq 1$. Hence $\tau_\omega = \tau_\omega(o_\omega) = 1$.

Now, barycentrically subdividing the cube complexes if necessary, every $g \in \Gamma$ acts on every X_n stably without inversions. Furthermore, $g^{D!}$ acts non-transversely by Proposition 2.2.1. Let $p_n \in X_n$ be the vertex of $\mathrm{Min}_{X_n}(g^{D!})$ that is closest to o_n. By Corollary 2.2.3 and Lemma 9.0.1, we have:

$$d(o_n, p_n) \leq \tfrac{D}{2} \cdot d(o_n, g^{D!}o_n) \leq \tfrac{D}{2} \cdot \|g^{D!}\| \cdot \tau_n(o_n) = \tfrac{D}{2} \cdot \|g^{D!}\| \cdot \tau_n.$$

In particular, $\lim_\omega \frac{1}{\tau_n} \cdot d(o_n, p_n) < +\infty$ and $p_\omega = (p_n)$ is a point of X_ω. Now

$$d(p_\omega, g^{D!}p_\omega) = \lim_\omega \tfrac{1}{\tau_n} \cdot d(p_n, g^{D!}p_n) = \lim_\omega \tfrac{1}{\tau_n} \cdot \ell_n(g^{D!}),$$

and, similarly, $d(x_\omega, g^{D!}x_\omega) \geq d(p_\omega, g^{D!}p_\omega)$ for every $x_\omega \in X_\omega$. We conclude that $\ell_\omega(g^{D!}) = d(p_\omega, g^{D!}p_\omega) = \lim_\omega \frac{1}{\tau_n} \cdot \ell_n(g^{D!})$, completing the proof. □

Remark 9.0.3. Although these do not appear in the literature, it is reasonable to expect that analogues of the results of [Hag07] should hold for isometries of finite rank median spaces. In this case, we would have $\ell_\omega(g^{D!}) = D! \cdot \ell_\omega(g)$ in Proposition 9.0.2 and, in fact, $\ell_\omega(g) = \lim_\omega \frac{1}{\tau_n} \cdot \ell_n(g)$ for all $g \in \Gamma$.

An action $\Gamma \curvearrowright X$ has no global fixed point (equivalently: orbits are unbounded) if and only if $\tau_X \neq 0$. We define the *reduced length function* as $\overline{\ell}_X(g) := \ell_X(g)/\tau_X$.

Theorem 9.0.4. *If the actions $\Gamma \curvearrowright X_n$ have unbounded orbits, there exists and element $g \in \Gamma$ such that $\overline{\ell}_{X_n}(g) \not\to 0$.*

Proof. Suppose for the sake of contradiction that $\overline{\ell}_{X_n}(g) \to 0$ for all $g \in \Gamma$. The above discussion yields an action $\Gamma \curvearrowright X_\omega$, where X_ω is a complete, connected, finite rank, median space. By Proposition 9.0.2, we have $\tau_\omega = 1$ and $\ell_\omega(g^{D!}) = 0$ for all $g \in \Gamma$. In particular, $g^{D!}$ fixes a point of X_ω (this is clear from the proof of Proposition 9.0.2).

We conclude that the action $\langle g \rangle \curvearrowright X_\omega$ has bounded orbits for every $g \in \Gamma$. Since X_ω admits a Γ-invariant, complete CAT(0) metric (Theorem 1.1 in [Bow16]), Cartan's theorem yields a point of X_ω fixed by g (see e.g. Corollary II.2.8 in [BH99] or Theorem 3.74 in [DK17]). Now, every $g \in \Gamma$ fixes a point of X_ω and Theorem 3.1 in [Fio18] yields a global fixed point for $\Gamma \curvearrowright X_\omega$. This contradicts $\tau_\omega = 1$. □

As in the introduction, we can consider the composition:

$$\mathrm{Cub}_D(\Gamma) \xrightarrow{\bar{\ell}} \mathbb{R}^\Gamma \setminus \{0\} \longrightarrow \mathbb{P}(\mathbb{R}^\Gamma).$$

Note that we are now mapping $\mathrm{Cub}_D(\Gamma)$ into $\mathbb{R}^\Gamma \setminus \{0\}$ via *reduced* length functions, as this does not affect the projection to the projectivisation. Lemma 9.0.1 shows that reduced length functions all lie in the compact subspace $\prod_{g\in\Gamma}[0, \|g\|] \subseteq \mathbb{R}^\Gamma$ and, by Theorem 9.0.4, they do not accumulate on the zero function. The closure of the image of $\bar{\ell}\colon \mathrm{Cub}_D(\Gamma) \to \mathbb{R}^\Gamma \setminus \{0\}$ is then compact and projects to a compact subset of $\mathbb{P}(\mathbb{R}^\Gamma)$. This proves Theorem C.

Index

$(x \cdot y)_v$, 37
$(x \cdot y)_v^f$, 43
$B(C_1, C_2)$, 27
$C(\mathfrak{w})$, 21
$I(x, y)$, 24
$S(C_1, C_2)$, 27
$S(X)$, 101
$T(\mathfrak{w})$, 94
X', 21
$Z(x)$, 25
$\mathrm{cr}(x, y, z, w)$, 40
Φ, 50, 85
$\mathrm{crt}(x, y, z, w)$, 40
$\mathrm{crt}^f(x, y, z, w)$, 43
$\deg(v)$, 22
$\delta(e, f)$, 22
ℓ^1-metric, 20
$\ell_X(g)$, 2
$\mathrm{lk}(v)$, 22
$[\![a : b : c]\!]$, 40
$\mathcal{U}(A|B)$, 91
$\mathfrak{u}_1 \preceq \mathfrak{u}_2$, 87
$\mathfrak{u}_1 \sim \mathfrak{u}_2$, 87
$\mathrm{Cub}_D(\Gamma)$, 5
$\mathrm{cr}_{\mathcal{U}}$, 91
$\mathfrak{M}(X)$, 45
$\mathfrak{M}_{\mathcal{A}}(X)$, 45
$\mathfrak{w}(e)$, 22
$\mathscr{A}(X)$, 38
$\mathscr{F}(X)$, 69
$\mathscr{H}$, 21
$\mathscr{H}(A|B)$, 23
$\mathscr{R}(v)$, 69
$\mathscr{T}_{\mathcal{U}}(\mathfrak{w})$, 94
$\mathscr{W}$, 21
$\mathscr{W}(A|B)$, 23
$\mathscr{W}(A)$, 23
$\mathscr{W}(\gamma)$, 22
$\mathscr{W}(k)$, 90
$\mathscr{W}_v$, 22
op, 42
op^f, 43
$\overline{X}$, 23
$\overline{\ell}_X(g)$, 104
∂X, 23
$\partial_c X$, 48
$\partial_c \Gamma$, 60
$\partial_c^{\mathrm{vis}} X$, 50
$\partial_s X$, 66
$\partial_\infty X$, 24
$\partial_{\mathrm{cnt}} X$, 53
$\partial_{\mathrm{cu}} X$, 48
$\partial_{\mathrm{nt}} X$, 53
$\partial_{\mathrm{reg}} X$, 35
π_C, 26
$\pi_v(x)$, 75

σ_x, 22, 23
τ_M, 103
$\tau_M(x)$, 103
$g^{\pm\infty}$, 60
$g^{\pm}$, 55
$m(x, y, z)$, 24
$m^f(x, y, z)$, 43
o_X, 60, 85
$r(+\infty)$, 23
$x \sim y$, 25
Min(g), 30

abstract cross ratio, 38
abstract shore, 27
adjacent, 22
ample, 25
automorphism, 30
axis, 30

balanced, 30
barycentric subdivision, 21
boundary-maximal, 94
bridge, 27

carrier, 21
CAT(0) cube complex, 20
clique, 22
co-meagre, 9
cocompact, 33
combinatorial metric, 20
complete, 64
component, 25
concise, 25
contracting boundary, 48
contracting point/ultrafilter, 48
contracting set, 47
convex, 26
cross ratio, 40, 62
cross ratio triple, 40
cuboid complex, 5
cubulated, 3

DCC, 23
De Rham decomposition, 33
degree, 22

endpoint at infinity, 23
essential, 33
extremal, 64

facing triple, 21
fat part, 69
first cubical subdivision, 21
flip, 33
flippable, 33
free face, 34

gate-projection, 26
gates, 27
generic, 88
geodesic extension property, 2
good, 90
Gromov boundary, 24
Gromov product, 37

heavy, 30
hyperplane-essential, 33

interval, 24
irreducible, 33
irreducible component (of a link), 77

lean, 47
length function, 2
length spectrum, 2
light, 30

line, 21
link, 22

Möbius, 41
marked length spectrum, 2
median, 24
median algebra, 24
median barycentre, 30
median space, 25
minimal set, 30
Morse element, 60
Morse quasi-geodesic, 47

nested, 21
non-elementary, 33
non-terminating, 53
non-transversely, 31

op-preserving, 43
opposite, 42
OR median, 73

pocset, 21

reduced length function, 104
regular, 35
restriction quotient, 32
Roller boundary, 23
Roller compactification, 23

section, 61
sector-heavy, 34
separate, 23
shore, 27
skinny geodesic, 69
skinny part, 69
skinny vertex, 68
SL, 73
SLR, 73
SLR-extendable, 73
squarisation, 101
stably without inversions, 30
straight geodesic, 64
straight point, 66
straight-looking, 73
straight-looking regular, 73
strongly contracting, 54
strongly separated, 27
strongly separated chain, 35

trace at infinity, 85
translation length, 2
transverse, 21
trustworthy 4-tuple, 91
trustworthy pair, 92

ultrafilter, 23
unbalanced, 30
unflippable, 33

visual boundary, 24

well-paired, 98

Bibliography

[ACGH17] Goulnara N. Arzhantseva, Christopher H. Cashen, Dominik Gruber, and David Hume. Characterizations of Morse quasi-geodesics via superlinear divergence and sublinear contraction. *Doc. Math.*, 22:1193–1224, 2017.

[BCG+09] Jacek Brodzki, Sarah J. Campbell, Erik Guentner, Graham A. Niblo, and Nick J. Wright. Property A and CAT(0) cube complexes. *J. Funct. Anal.*, 256(5):1408–1431, 2009.

[BF19a] Jonas Beyrer and Elia Fioravanti. Cross ratios and cubulations of hyperbolic groups. *arXiv:1810.08087v3*, 2019.

[BF19b] Jonas Beyrer and Elia Fioravanti. Cross ratios on CAT(0) cube complexes and marked length-spectrum rigidity. *arXiv:1903.02447v2*, 2019.

[BFIM19] Jonas Beyrer, Elia Fioravanti, and Merlin Incerti-Medici. CAT(0) cube complexes are determined by their boundary cross ratio. *arXiv:1805.08478v3*, 2019.

[BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[BHW11] Nicolas Bergeron, Frédéric Haglund, and Daniel T. Wise. Hyperplane sections in arithmetic hyperbolic manifolds. *J. Lond. Math. Soc. (2)*, 83(2):431–448, 2011.

[Bis15] Kingshook Biswas. On Moebius and conformal maps between boundaries of CAT(-1) spaces. *Ann. Inst. Fourier (Grenoble)*, 65(3):1387–1422, 2015.

[BK85] Keith Burns and Anatole B. Katok. In collaboration with W. Ballman, M. Brin, P. Eberlein and R. Osserman. Manifolds with nonpositive curvature. *Ergodic Theory Dynam. Systems*, 5(2):307–317, 1985.

[BL18] Anja Bankovic and Christopher J. Leininger. Marked-length-spectral rigidity for flat metrics. *Trans. Amer. Math. Soc.*, 370(3):1867–1884, 2018.

[Bon88] Francis Bonahon. The geometry of Teichmüller space via geodesic currents. *Invent. Math.*, 92(1):139–162, 1988.

[Bou96] Marc Bourdon. Sur le birapport au bord des CAT(−1)-espaces. *Inst. Hautes Études Sci. Publ. Math.*, (83):95–104, 1996.

[Bou97] Marc Bourdon. Immeubles hyperboliques, dimension conforme et rigidité de Mostow. *Geom. Funct. Anal.*, 7(2):245–268, 1997.

[Bow13] Brian H. Bowditch. Coarse median spaces and groups. *Pacific J. Math.*, 261(1):53–93, 2013.

[Bow16] Brian H. Bowditch. Some properties of median metric spaces. *Groups Geom. Dyn.*, 10(1):279–317, 2016.

[BS17] Jonas Beyrer and Victor Schroeder. Trees and ultrametric Möbius structures. *p-Adic Numbers Ultrametric Anal. Appl.*, 9(4):247–256, 2017.

[BW12] Nicolas Bergeron and Daniel T. Wise. A boundary criterion for cubulation. *Amer. J. Math.*, 134(3):843–859, 2012.

[CCM18] Ruth Charney, Matthew Cordes, and Devin Murray. Quasi-Mobius Homeomorphisms of Morse boundaries. *arXiv:1801.05315v1*, 2018.

[CDH10] Indira Chatterji, Cornelia Druţu, and Frédéric Haglund. Kazhdan and Haagerup properties from the median viewpoint. *Adv. Math.*, 225(2):882–921, 2010.

[CFF92] C. Croke, A. Fathi, and J. Feldman. The marked length-spectrum of a surface of nonpositive curvature. *Topology*, 31(4):847–855, 1992.

[CFI16] Indira Chatterji, Talia Fernós, and Alessandra Iozzi. The median class and superrigidity of actions on CAT(0) cube complexes. *J. Topol.*, 9(2):349–400, 2016. With an appendix by Pierre-Emmanuel Caprace.

[CL95] Marshall M. Cohen and Martin Lustig. Very small group actions on **R**-trees and Dehn twist automorphisms. *Topology*, 34(3):575–617, 1995.

[CL11] Pierre-Emmanuel Caprace and Jean Lécureux. Combinatorial and group-theoretic compactifications of buildings. *Ann. Inst. Fourier (Grenoble)*, 61(2):619–672, 2011.

[CM87] Marc Culler and John W. Morgan. Group actions on **R**-trees. *Proc. London Math. Soc. (3)*, 55(3):571–604, 1987.

[CM13] Ruth Charney and Max Margolis. Length functions of 2-dimensional right-angled Artin groups. *Geom. Dedicata*, 166:31–45, 2013.

[CM18] Christopher H. Cashen and John M. Mackay. A metrizable topology on the contracting boundary of a group. *arXiv:1703.01482v3*, 2018.

[CN05] Indira Chatterji and Graham Niblo. From wall spaces to CAT(0) cube complexes. *Internat. J. Algebra Comput.*, 15(5-6):875–885, 2005.

[Cor17] Matthew Cordes. Morse boundaries of proper geodesic metric spaces. *Groups Geom. Dyn.*, 11(4):1281–1306, 2017.

[Cro90] Christopher B. Croke. Rigidity for surfaces of nonpositive curvature. *Comment. Math. Helv.*, 65(1):150–169, 1990.

[CS11] Pierre-Emmanuel Caprace and Michah Sageev. Rank rigidity for CAT(0) cube complexes. *Geom. Funct. Anal.*, 21(4):851–891, 2011.

[CS15] Ruth Charney and Harold Sultan. Contracting boundaries of CAT(0) spaces. *J. Topol.*, 8(1):93–117, 2015.

[CSV17] Ruth Charney, Nathaniel Stambaugh, and Karen Vogtmann. Outer space for untwisted automorphisms of right-angled Artin groups. *Geom. Topol.*, 21(2):1131–1178, 2017.

[CV86] Marc Culler and Karen Vogtmann. Moduli of graphs and automorphisms of free groups. *Invent. Math.*, 84(1):91–119, 1986.

[Dav08] Michael W. Davis. *The geometry and topology of Coxeter groups*, volume 32 of *London Mathematical Society Monographs Series*. Princeton University Press, Princeton, NJ, 2008.

[DK17] Cornelia Drutu and Michael Kapovich. Geometric group theory. In *AMS series "Colloquium Publications"*, 2017.

[DP16] Thomas Delzant and Pierre Py. Cubulable Kähler groups. *arXiv:1609.08474v1*, 2016.

[Ebe94] Patrick Eberlein. Geometry of 2-step nilpotent groups with a left invariant metric. *Ann. Sci. École Norm. Sup. (4)*, 27(5):611–660, 1994.

[Fer18] Talia Fernós. The Furstenberg–Poisson boundary and CAT(0) cube complexes. *Ergodic Theory Dynam. Systems*, 38(6):2180–2223, 2018.

[FFT18] Talia Fernós, Max Forester, and Jing Tao. Effective quasimorphisms on right-angled Artin groups. *arXiv:1602.05637v2*, 2018.

[Fio17] Elia Fioravanti. Roller boundaries for median spaces and algebras. *arXiv:1708.01005v2*, 2017.

[Fio18] Elia Fioravanti. The Tits alternative for finite rank median spaces. *Enseign. Math.*, 64(1-2):89–126, 2018.

[Fio19] Elia Fioravanti. Superrigidity of actions on finite rank median spaces. *Adv. Math.*, 352:1206–1252, 2019.

[FJ89] F. T. Farrell and L. E. Jones. Negatively curved manifolds with exotic smooth structures. *J. Amer. Math. Soc.*, 2(4):899–908, 1989.

[FJ93] F. T. Farrell and L. E. Jones. Topological rigidity for compact non-positively curved manifolds. In *Differential geometry: Riemannian geometry (Los Angeles, CA, 1990)*, volume 54 of *Proc. Sympos. Pure Math.*, pages 229–274. Amer. Math. Soc., Providence, RI, 1993.

[FLM18] Talia Fernós, Jean Lécureux, and Frédéric Mathéus. Random walks and boundaries of CAT(0) cubical complexes. *Comment. Math. Helv.*, 93(2):291–333, 2018.

[Gen16a] Anthony Genevois. Coning-off CAT(0) cube complexes. *arXiv:1603.06513v1*, 2016.

[Gen16b] Anthony Genevois. Contracting isometries of CAT(0) cube complexes and acylindrical hyperbolicity of diagram groups. *arXiv:1610.07791v1*, 2016.

[Gen17] Anthony Genevois. Hyperbolicities in CAT(0) cube complexes. *arXiv:1709.08843v1*, 2017.

[GM18] Anthony Genevois and Alexandre Martin. Automorphisms of graph products of groups from a geometric perspective. *arXiv:1809.08091v1*, 2018.

[GMRS98] Rita Gitik, Mahan Mitra, Eliyahu Rips, and Michah Sageev. Widths of subgroups. *Trans. Amer. Math. Soc.*, 350(1):321–329, 1998.

[Gro87] Misha Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[Gur05] Dan Guralnik. Coarse decompositions of boundaries for CAT(0) groups. *PhD Thesis, Technion, Haifa*, 2005.

[Hag07] Frédéric Haglund. Isometries of CAT(0) cube complexes are semi-simple. *arXiv:0705.3386v1*, 2007.

[Ham97] Ursula Hamenstädt. Cocycles, Hausdorff measures and cross ratios. *Ergodic Theory Dynam. Systems*, 17(5):1061–1081, 1997.

[Ham99] Ursula Hamenstädt. Cocycles, symplectic structures and intersection. *Geom. Funct. Anal.*, 9(1):90–140, 1999.

[HP97] Sa'ar Hersonsky and Frédéric Paulin. On the rigidity of discrete isometry groups of negatively curved spaces. *Comment. Math. Helv.*, 72(3):349–388, 1997.

[HP15] Mark F. Hagen and Piotr Przytycki. Cocompactly cubulated graph manifolds. *Israel J. Math.*, 207(1):377–394, 2015.

[HS18] Mark F. Hagen and Tim Susse. On hierarchical hyperbolicity of cubical groups. *arXiv:1609.01313v2*, 2018.

[HT18] Mark F. Hagen and Nicholas W. M. Touikan. Panel collapse and its applications. *arXiv:1712.06553v2*, 2018.

[HW] Mark F. Hagen and Henry Wilton. Guirardel cores for cube complexes. In preparation.

[HW12] Frédéric Haglund and Daniel T. Wise. A combination theorem for special cube complexes. *Ann. of Math. (2)*, 176(3):1427–1482, 2012.

[HW15] Mark F. Hagen and Daniel T. Wise. Cubulating hyperbolic free-by-cyclic groups: the general case. *Geom. Funct. Anal.*, 25(1):134–179, 2015.

[HW16] Mark F. Hagen and Daniel T. Wise. Cubulating hyperbolic free-by-cyclic groups: the irreducible case. *Duke Math. J.*, 165(9):1753–1813, 2016.

[Kim01a] Inkang Kim. Marked length rigidity of rank one symmetric spaces and their product. *Topology*, 40(6):1295–1323, 2001.

[Kim01b] Inkang Kim. Rigidity and deformation spaces of strictly convex real projective structures on compact manifolds. *J. Differential Geom.*, 58(2):189–218, 2001.

[KM12] Jeremy Kahn and Vladimir Markovic. Immersing almost geodesic surfaces in a closed hyperbolic three manifold. *Ann. of Math. (2)*, 175(3):1127–1190, 2012.

[Lab06] François Labourie. Anosov flows, surface groups and curves in projective space. *Invent. Math.*, 165(1):51–114, 2006.

[Lab07] François Labourie. Cross ratios, surface groups, $\mathrm{PSL}(n, \mathbf{R})$ and diffeomorphisms of the circle. *Publ. Math. Inst. Hautes Études Sci.*, (106):139–213, 2007.

[Lea13] Ian J. Leary. A metric Kan-Thurston theorem. *J. Topol.*, 6(1):251–284, 2013.

[Nic04] Bogdan Nica. Cubulating spaces with walls. *Algebr. Geom. Topol.*, 4:297–309, 2004.

[NR97] Graham A. Niblo and Lawrence Reeves. Groups acting on CAT(0) cube complexes. *Geom. Topol.*, 1:1–7, 1997.

[NR03] Graham A. Niblo and Lawrence D. Reeves. Coxeter groups act on CAT(0) cube complexes. *J. Group Theory*, 6(3):399–413, 2003.

[NS13] Amos Nevo and Michah Sageev. The Poisson boundary of CAT(0) cube complex groups. *Groups Geom. Dyn.*, 7(3):653–695, 2013.

[Ota90] Jean-Pierre Otal. Le spectre marqué des longueurs des surfaces à courbure négative. *Ann. of Math. (2)*, 131(1):151–162, 1990.

[Ota92] Jean-Pierre Otal. Sur la géometrie symplectique de l'espace des géodésiques d'une variété à courbure négative. *Rev. Mat. Iberoamericana*, 8(3):441–456, 1992.

[OW11] Yann Ollivier and Daniel T. Wise. Cubulating random groups at density less than 1/6. *Trans. Amer. Math. Soc.*, 363(9):4701–4733, 2011.

[Pau96] Frédéric Paulin. Un groupe hyperbolique est déterminé par son bord. *J. London Math. Soc. (2)*, 54(1):50–74, 1996.

[PW14] Piotr Przytycki and Daniel T. Wise. Graph manifolds with boundary are virtually special. *J. Topol.*, 7(2):419–435, 2014.

[PW18] Piotr Przytycki and Daniel T. Wise. Mixed 3-manifolds are virtually special. *J. Amer. Math. Soc.*, 31(2):319–347, 2018.

[Rol98] Martin A. Roller. Poc sets, median algebras and group actions. An extended study of Dunwoody's construction and Sageev's theorem. Preprint, University of Southampton, 1998.

[Sag95] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.

[Sag97] Michah Sageev. Codimension-1 subgroups and splittings of groups. *J. Algebra*, 189(2):377–389, 1997.

[Sag14] Michah Sageev. CAT(0) cube complexes and groups. In *Geometric group theory*, volume 21 of *IAS/Park City Math. Ser.*, pages 7–54. Amer. Math. Soc., Providence, RI, 2014.

[Vij15] Anna Vijayan. Compactifying the space of length functions of a right-angled Artin group. *arXiv:1510.05729v1*, 2015.

[Wis04] Daniel T. Wise. Cubulating small cancellation groups. *Geom. Funct. Anal.*, 14(1):150–214, 2004.

[Xie06] Xiangdong Xie. Quasi-isometric rigidity of Fuchsian buildings. *Topology*, 45(1):101–169, 2006.