



Cross ratios and cubulations of hyperbolic groups

Jonas Beyrer¹ · Elia Fioravanti¹

Received: 31 January 2020 / Revised: 31 January 2020 / Accepted: 25 November 2021 /
Published online: 11 December 2021
© The Author(s) 2021

Abstract

Many geometric structures associated to surface groups can be encoded in terms of invariant cross ratios on their circle at infinity; examples include points of Teichmüller space, Hitchin representations and geodesic currents. We add to this picture by studying cocompact cubulations of arbitrary Gromov hyperbolic groups G . Under weak assumptions, we show that the space of cubulations of G naturally injects into the space of G -invariant cross ratios on the Gromov boundary $\partial_\infty G$. A consequence of our results is that essential, hyperplane-essential, cocompact cubulations of hyperbolic groups are length-spectrum rigid, i.e. they are fully determined by their length function. This is the optimal length-spectrum rigidity result for cubulations of hyperbolic groups, as we demonstrate with some examples. In the hyperbolic setting, this constitutes a strong improvement on our previous work [4]. Along the way, we describe the relationship between the Roller boundary of a CAT(0) cube complex, its Gromov boundary and—in the non-hyperbolic case—the contracting boundary of Charney and Sultan. All our results hold for cube complexes with variable edge lengths.

Contents

1	Introduction	1548
2	Preliminaries	1555
2.1	CAT(0) cube complexes	1555
2.2	Combinatorial geodesics vs CAT(0) geodesics	1561
2.3	Median barycentres	1562
2.4	Cross ratios on cube complexes	1563
2.5	CAT(0) cuboid complexes	1564
3	The Morse property in cube complexes	1565
3.1	Contracting geodesics	1565
3.2	Roller boundaries vs contracting boundaries	1566

Communicated by Andreas Thom.

✉ Elia Fioravanti
fioravanti@mpim-bonn.mpg.de

¹ Max-Planck-Institut für Mathematik, Bonn, NRW, Germany

3.3 Contracting non-terminating ultrafilters	1570
4 Cubulations of hyperbolic groups	1572
4.1 Traces at infinity	1572
4.2 Towards hyperplane recognition	1576
4.3 Trust issues	1577
4.4 Traces vs cross ratios	1579
4.5 Concluding the proof	1584
5 Epilogue	1586
5.1 Cross ratios on contracting boundaries	1586
5.2 Marked length-spectrum rigidity	1588
References	1590

1 Introduction

Let G be a Gromov hyperbolic group. We denote by $\partial_\infty G^{(4)}$ the set of 4-tuples of pairwise distinct points in the Gromov boundary $\partial_\infty G$. A map $\mathbb{B}: \partial_\infty G^{(4)} \rightarrow \mathbb{R}$ is said to be a 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 G on $\partial_\infty G^{(4)}$. Similar notions appear e.g. in [39,49,50,57].

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 $G = \pi_1 S$, every point of Teichmüller space yields identifications $\partial_\infty G \simeq \partial_\infty \mathbb{H}^2 \simeq \mathbb{RP}^1$ and the projective cross ratio on \mathbb{RP}^1 can be pulled back to an invariant cross ratio¹ on $\partial_\infty G$. The latter uniquely determines the original point of Teichmüller space [11]. More generally, the space of all negatively curved Riemannian metrics on S embeds into the space of invariant cross ratios on $\partial_\infty G$ [56].

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 $\rho: G \rightarrow PSL_n \mathbb{R}$ with a space of Hölder-continuous invariant cross ratios on $\partial_\infty G$ [50].

In this paper, we consider yet another significant geometric structure that groups can be endowed with. More precisely, we study the space of cubulations of a non-elementary hyperbolic group G . Our main result is that the space of cocompact cubulations of G naturally injects² into the space of invariant \mathbb{Z} -valued cross ratios on $\partial_\infty G$ (Theorem A). An important consequence is that most cubulations of G are length-spectrum rigid (Corollary B).

Recall that a cocompact cubulation is a proper cocompact action of G on a CAT(0) cube complex X . A group is said to be cocompactly cubulated if it admits a cocom-

¹ To be precise, one has to take the logarithm of the absolute value of the projective cross ratio if this is to satisfy conditions (i)–(iv).

² Some mild and inevitable assumptions are required; cf. Theorem A below.

compact cubulation. Cocompactly cubulated hyperbolic groups are ubiquitous in geometric group theory: they include surface groups, hyperbolic 3-manifold groups [14], hyperbolic free-by-cyclic groups [45,46], hyperbolic Coxeter groups [54], finitely presented small cancellation groups [64], random groups at low density [58] and many arithmetic lattices in $SO(n, 1)$ [7,44]. Cocompactly cubulated hyperbolic groups are also particularly significant due to recent advances in low-dimensional topology [2,43,65].

Among cubulations of a group G , a subclass is especially relevant for us: that of essential, hyperplane-essential cubulations³. Indeed, due to the extreme flexibility of cube complexes, it is all too easy to perturb any cubulation by adding “insignificant noise” (say, a few loose edges around the space). Essential, hyperplane-essential cube complexes are those from which all “noise” has been removed. A simple procedure for this removal is provided by Sect. 3 in [23] and Theorem A in [42].

More precisely, every cocompact group action on a CAT(0) cube complex X can be collapsed to an action on an essential, hyperplane-essential CAT(0) cube complex X_\bullet . This procedure will preserve most additional properties of the original action. In particular, the collapsing map $X \rightarrow X_\bullet$ has uniformly bounded fibres and it is an equivariant quasi-isometry.

Essential hyperplane-essential cube complexes are the appropriate setting to study cross ratios. Our first result is the following:

Theorem A *Let G be a non-elementary Gromov hyperbolic group.*

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

$$cr_X : \partial_\infty G^{(4)} \rightarrow \mathbb{Z}.$$

There exists a co-meagre⁴ subset $\mathcal{C} \subseteq \partial_\infty G$ such that cr_X is continuous at all points of $\mathcal{C}^{(4)} \subseteq \partial_\infty G^{(4)}$.

- (2) *Let in addition X be hyperplane-essential and consider another action $G \curvearrowright Y$ satisfying the same hypotheses. If there exists a co-meagre subset $\mathcal{D} \subseteq \partial_\infty G$ such that the cross ratios cr_X and cr_Y coincide on $\mathcal{D}^{(4)} \subseteq \partial_\infty G^{(4)}$, then X and Y are G -equivariantly isomorphic.*

In particular, two essential, hyperplane-essential, cocompact cubulations yield the same boundary cross ratio if and only if they are G -equivariantly isomorphic.

Essentiality and hyperplane-essentiality are absolutely crucial to part (2) of Theorem A. Examples 5.4 and 5.5 show that—in a very strong sense—neither of these assumptions can be dropped.

It is not surprising that cr_X takes integer values in Theorem A, after all cube complexes are fundamentally discrete objects. Our cross ratio can be regarded as an exact discretisation of Paulin’s coarse cross ratio on Gromov boundaries of arbitrary Gromov hyperbolic spaces [59].

³ We refer the reader to [23,42] or Sect. 2.1 below for definitions.

⁴ 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 $\partial_\infty G$ are dense.

The main ingredient in the proof of Theorem A is Theorem C below. We will discuss this result at length later in the introduction, but let us first describe one more of its applications.

Let us endow our CAT(0) cube complexes with their ℓ^1 (aka combinatorial) metric d and let us associate to each action $G \curvearrowright X$ the function $\ell_X: G \rightarrow \mathbb{N}$ given by:

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

This is normally known as length function, or marked length spectrum by analogy with the corresponding notion in the setting of Riemannian manifolds. Theorem C below will also have the following consequence.

Corollary B *Let a Gromov hyperbolic group G act properly 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 G -equivariantly isomorphic.*

The same result is conjectured to hold for actions of G on Hadamard manifolds (Problems 3.1 and 3.7 in [8]). This is known as the “marked length-spectrum rigidity conjecture” and it is a notorious open problem. Progress on the conjecture has been remarkably limited, with most results only handling 2-dimensional spaces [16,22,56], or extremely rigid settings such as symmetric spaces [25,40].

In this perspective, Corollary B is particularly interesting as—along with our previous work in [4]—it is the first length-spectrum rigidity result to cover such a broad family of non-positively curved spaces. The proof of Corollary B relies on a reduction—obtained in [4]—to the problem of extending certain boundary maps to isomorphisms of cube complexes. However, we stress that the core argument in the proof of Corollary B lies in the ensuing extension procedure, and this requires completely different techniques from those in [4]. See the statement of Theorem C below and the subsequent discussion for a detailed description.

When X and Y have no free faces (i.e. when their CAT(0) metrics are geodesically complete), Corollary B follows from Theorem A in [4]. Having no free faces, however, is an extremely strong restriction when studying cubulations of hyperbolic groups, as most known cubulating procedures will not yield spaces satisfying this requirement. As an example, consider the case when G is the fundamental group of a closed, oriented surface S of genus ≥ 2 . It is well-known that every finite filling collection of closed curves on S gives rise to an essential, hyperplane-essential, cocompact cubulation of G [14,61,62]. On the other hand, most cube complexes resulting from this construction will have dimension ≥ 3 , which forces the existence of free faces⁵.

Examples 5.4 and 5.5 show that essentiality and hyperplane-essentiality are necessary assumptions on the CAT(0) cube complex X for any form of length-spectrum rigidity to hold. Thus, Corollary B is the optimal result of this type for cubulations of hyperbolic groups. In addition, note that any cocompact cubulation can be made

⁵ More generally, given a CAT(0) cube complex X with no free faces and any group G acting on X cocompactly and with virtually cyclic hyperplane-stabilisers, we necessarily have $\dim X \leq 2$. This can be shown by noticing that \mathbb{R} is the only cube complex with no free faces that admits a cocompact action of the group \mathbb{Z} .

essential and hyperplane-essential by means of the collapsing procedure of [23,42]. In many settings, however, the two assumptions are automatically satisfied, even without resorting to any collapsing: for instance, this is the case for any cubulation of a hyperbolic 3-manifold group arising from Sageev’s construction applied to quasi-Fuchsian immersed surfaces [47].

Remark Although we have preferred to state Theorem A and Corollary B for cube complexes, they more generally hold for cuboid complexes. In such complexes, edges can have arbitrary (positive) real lengths, so the cross ratio cr_X and the length function ℓ_X will take arbitrary real values. The price to pay is that, in both results, we can only conclude that actions with the same cross ratio/length function are G -equivariantly isometric, i.e. said isometries will in general not take vertices to vertices.

All results in this paper equally apply to CAT(0) cuboid complexes, without requiring any significant changes to proofs—although of course all cubical isomorphisms need to be replaced with mere isometries. The reader can consult Sect. 2.5 for a brief discussion of this.

On the proofs of Theorem A and Corollary B. As mentioned, the core result of this paper is an extension procedure for certain partially-defined, cross-ratio preserving boundary maps (Theorem C below). In order to make things precise, let us introduce some terminology.

The horofunction boundary of the cube complex (X, d) is known as the Roller boundary ∂X . In our setting, this space is always compact and totally disconnected—unlike the Gromov/visual⁶ boundary $\partial_\infty X$. As we observed in [4,5], the Roller boundary is naturally endowed with a continuous, \mathbb{Z} -valued cross ratio:

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

Here, the notation $\mathscr{H}(x, z|y, w)$ refers to the collection of hyperplanes of X that separate $x, z \in \partial X$ from $y, w \in \partial X$.

When X is Gromov hyperbolic, the two boundaries ∂X and $\partial_\infty X$ share a “large” subset. More precisely, a co-meagre subset of $\partial_\infty X$ is naturally identified with a subset of ∂X and therefore inherits the cross ratio of ∂X . We will denote this common subset by $\partial_{nt} X$, as it coincides with the collection of non-terminating ultrafilters introduced in [55]. Equivalently, we can describe $\partial_{nt} X \subseteq \partial_\infty X$ as the subset of points that do not lie in the Gromov boundary of any hyperplane of X (Lemma 4.7).

The following is the crucial ingredient in the proofs of part (2) of Theorem A and of Corollary B.

Theorem C *Let a non-elementary Gromov hyperbolic group G act properly and cocompactly on essential, hyperplane-essential CAT(0) cube complexes X and Y . Let $f : \partial_\infty X \rightarrow \partial_\infty Y$ be the unique G -equivariant homeomorphism. Suppose that there exists a nonempty, G -invariant subset $\Omega \subseteq \partial_{nt} X$ such that $f(\Omega) \subseteq \partial_{nt} Y$ and*

⁶ For the visual boundary of a CAT(0) space and Gromov boundary of a Gromov hyperbolic space, we refer the reader, respectively, to Chapters II.8 and III.H.3 in [6]. We denote both boundaries by $\partial_\infty X$, as these coincide whenever both defined.

such that cross ratios of elements of Ω^4 are preserved by f . Then, there exists a unique G -equivariant isomorphism $F: X \rightarrow Y$ extending f .

It is possible that Theorem C will find further application in the proof of rigidity results for certain classes of cubulated hyperbolic groups. Indeed, since ℓ^1 metrics on cube complexes fall in the setting of Section 5.1 of [38], cross-ratio preserving boundary maps should arise naturally from commensurations or quasi-isometries between cubulated hyperbolic groups with suitable properties. At present, however, a major obstruction to pursuing approaches of this kind is that not much is known on conformal dimension and Loewner property for boundaries of cubulated hyperbolic groups. A notable exception are Bourdon groups [12]; also see [9, 10].

It is interesting to compare Theorem C with two old results obtained in different settings. The first is Paulin's classical theorem that homeomorphisms of Gromov boundaries arise from quasi-isometries if and only if they almost preserve the boundary cross ratio (Theorem 1.2 in [59]). Paulin's techniques are of no help in our context, as we want our extensions to be genuine isometries.

A more fitting comparison is with Proposition 2.4.7 in [12], whose statement strikingly resembles that of Theorem C. 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 [35]. 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 special case of part (1) of our Theorem A. Proposition 2.4.7 in [12] and part of Theorem 1.5 in [66] become a special case of Theorem C above.

It is important to remark that, unlike the 2-dimensional setting of Fuchsian buildings, the cube complexes in Theorem C can have arbitrarily high dimension. This will seriously complicate proofs due to a strictly 3-dimensional phenomenon which we now describe.

Given points $x, y, z, w \in \partial X$, the three sets of hyperplanes $\mathscr{W}(x, y|z, w)$, $\mathscr{W}(x, z|y, w)$ and $\mathscr{W}(x, w|y, z)$ are pairwise transverse. If $\dim X \leq 2$, one of these sets must be empty and their three cardinalities can be deduced from their respective differences, i.e. $\text{cr}(x, y, z, w)$, $\text{cr}(y, z, x, w)$ and $\text{cr}(z, x, y, w)$. On the other hand, when $\dim X \geq 3$, it may be impossible to recover all three cardinalities just from cross ratios of 4-tuples involving only the points x, y, z, w (see e.g. Figure 1 in [4] and the related discussion).

In order to resolve part of this issue, we will be led to consider trustworthy 4-tuples $(x, y, z, w) \in (\partial_{\text{nt}} X)^4$, i.e. those 4-tuples for which one of the three sets $\mathscr{W}(x, y|z, w)$, $\mathscr{W}(x, z|y, w)$ and $\mathscr{W}(x, w|y, z)$ is empty. A key point will be that, even in boundaries of high-dimensional cube complexes, it is always possible to find several trustworthy 4-tuples (Lemma 4.22).

We now briefly sketch the proof of Theorem C, denoting 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 $(\mathscr{H}(X), \subseteq, *)$ simply by looking at the Gromov boundary $\partial_\infty X$ and the cross ratio (where defined).

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

- (Ia) For every $h \in \mathcal{H}(X)$, we have $\partial_\infty h \setminus \partial_\infty h^* \neq \emptyset$.
- (Ib) Given $h, k \in \mathcal{H}(X)$, we have $h \subseteq k$ if and only if $\partial_\infty h \subseteq \partial_\infty k$.⁷
- (IIa) For every $w \in \mathcal{W}(X)$, there exists $w' \in \mathcal{W}(Y)$ with $f(\partial_\infty w) = \partial_\infty w'$.
- (IIb) For every $h \in \mathcal{H}(X)$, there exists $h' \in \mathcal{H}(Y)$ with $f(\partial_\infty h) = \partial_\infty h'$.

The boundary homeomorphism f then induces a G -equivariant bijection $f_*: \mathcal{H}(X) \rightarrow \mathcal{H}(Y)$ by Steps (Ia) and (IIb). Step (Ib) shows that f_* preserves inclusion relations and, by general theory of CAT(0) cube complexes, f_* must be induced by a G -equivariant isomorphism $F: X \rightarrow Y$.

Steps (Ia) and (Ib) are the key points where, respectively, essentiality and hyperplane-essentiality come into play. Example 5.4 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 5.5, where *transverse* halfspaces h and k have $\partial_\infty h = \partial_\infty k$.

Regarding Step (IIa), it is not hard to use the cross ratio to characterise which pairs of points $\xi, \eta \in \partial_\infty X$ lie in the Gromov boundary of a common hyperplane (Proposition 4.14). 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 4.24 and Proposition 4.26) based on the fact that $\partial_\infty G$ 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 involved. Given $w \in \mathcal{W}(X)$ bounding $h \in \mathcal{H}(X)$, the set $\partial_\infty h \setminus \partial_\infty w$ is a union of connected components of $\partial_\infty X \setminus \partial_\infty w$. However, $\partial_\infty X \setminus \partial_\infty w$ will in general have many more components than there are halfspaces bounded by w .

The case to keep in mind is when $G = \pi_1 S$, for a closed oriented surface S , and the hyperplane-stabiliser $G_w < G$ is the fundamental group of a subsurface of S with at least 3 boundary components. Not all G_w -invariant partitions of the set of connected components of $\partial_\infty X \setminus \partial_\infty w$ arise from a halfspace of X . Thus, one cannot recover $\partial_\infty h$ from the knowledge of $\partial_\infty w$ purely through topological and dynamical arguments.

We will instead rely again on the cross ratio in order to circumvent these issues. Step (IIb) will finally be completed in Theorem 4.33.

On the relationship between ∂X and $\partial_\infty X$. We still have not discussed the first half of Theorem A, which is mostly based on transferring the cross ratio from ∂X to $\partial_\infty X$. As the required techniques are quite similar, we do not assume hyperbolicity of the CAT(0) cube complex X and we more generally describe the relationship between the Roller boundary ∂X and the contracting boundary $\partial_c X$. The latter was introduced in [24].

Fixing a basepoint $p \in X$, every point of ∂X is represented by a combinatorial ray based at p . We denote by $\partial_{cu} X \subseteq \partial X$ the subset of points that are represented

⁷ This is not true in general, but it is how one should think about things. It only fails when $\partial_\infty h = \partial_\infty k$ and $k \subsetneq h$, in which case k and h are at finite Hausdorff distance anyway.

by contracting combinatorial rays. We endow $\partial_{\text{cu}}X$ with the restriction of the (totally disconnected) topology of ∂X . We moreover denote by $\partial_c^{\text{vis}}X$ 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 [24], we prove:

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

- (1) *There exists a natural continuous surjection $\Phi: \partial_{\text{cu}}X \rightarrow \partial_c^{\text{vis}}X$ with finite fibres. Collapsing its fibres, Φ descends to a homeomorphism.*
- (2) *If X is hyperbolic, we have $\partial_{\text{cu}}X = \partial X$ and $\partial_c^{\text{vis}}X = \partial_{\infty}X$.*

The reader will find additional details on the map Φ in Sect. 3.2, especially in Remark 3.8 and Theorem 3.10. 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, part (1) of Theorem A is obtained by considering a canonical section to the map Φ . The latter is built through a new construction of barycentres for bounded cube complexes, which we describe in Sect. 2.3.

Along with our previous work in [4], Theorem D also allows us to extend Theorem A to the context of non-hyperbolic groups acting on CAT(0) cube complexes with no free faces.

Recall that the Morse boundary of an arbitrary finitely generated group G was introduced in [21]. In accordance with [19], we prefer to refer to it as the *contracting boundary*⁸ of G (denoted $\partial_c G$), as this simplifies notation and terminology (the topology of $\partial_c G$ will not be relevant to us).

Corollary E *Let G be a finitely generated, non-virtually-cyclic group.*

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

$$\text{cr}_X: \partial_c G^{(4)} \rightarrow \mathbb{Z}.$$

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

It is worth pointing out that, under the hypotheses of Corollary E, the contracting boundary $\partial_c G$ is always nonempty. When G is hyperbolic, $\partial_c G$ is naturally identified with the Gromov boundary $\partial_{\infty}G$.

The cross ratio provided by Corollary E is again continuous⁹ at a “large” subset of $\partial_c G^{(4)}$, but it does not make sense to speak of meagre subsets in this context. Indeed, the entire contracting boundary is often itself meagre, even in the Cashen–Mackay

⁸ This is justified by the fact that a quasi-geodesic is Morse if and only if it is sublinearly contracting [1]. However, we acknowledge that the word “contracting” is normally taken to mean “strongly contracting” (as we also do in Sect. 3) and these quasi-geodesics would not provide a satisfactory notion of boundary for a general group G .

⁹ Both endowing $\partial_c G$ with the visual topology and with the topologies of [19,21].

topology. In fact, $\partial_c G$ is likely to be a Baire space if and only if G is hyperbolic. The latter observation follows from Theorem 7.6 in [19] when G is a toral relatively hyperbolic group. We thank Chris Cashen for pointing this out to us.

We conclude the introduction with the following question, which is naturally brought to mind by Theorem A.

Question Let G 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 $\partial_\infty G$ arise from cocompact cubulations of G ?

Theorem 1.1 in [50] 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 $\partial_\infty G$ is endowed with a continuous, invariant, \mathbb{R} -valued cross ratio whenever G 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 [27,53].

We thank the organisers of the following conferences, where part of this work was carried out: *Third GEAR (Junior) Retreat*, Stanford, 2017; *Moduli Spaces*, Ventotene, 2017; *Young Geometric Group Theory VII*, Les Diablerets, 2018; *Topological and Homological Methods in Group Theory*, Bielefeld, 2018; *3-Manifolds and Geometric Group Theory*, Luminy, 2018; *Representation varieties and geometric structures in low dimensions*, Warwick, 2018. EF also thanks M. Incerti-Medici, Viktor Schroeder and Anna Wienhard for his visits at UZH and Universität Heidelberg, respectively.

JB was supported by the Swiss National Science Foundation under Grant 200020\175567. EF was supported by the Clarendon Fund and the Merton Mousouris Scholarship.

2 Preliminaries

2.1 CAT(0) cube complexes

We will assume a certain familiarity with basic properties of cube complexes. The reader can consult for instance [63] and the first sections of [17,20,23,55] for an introduction to the subject. This subsection is mainly meant to fix notation and recall a few facts that we shall rely on.

Let X be a simply connected cube complex satisfying Gromov's no- Δ -condition; see 4.2.C in [37] and Chapter II.5 in [6]. The Euclidean metrics on its cubes fit together to yield a CAT(0) metric on X . We can also endow each cube $[0, 1]^k \subseteq X$ with the restriction of the ℓ^1 metric of \mathbb{R}^k and consider the induced path metric $d(-, -)$. We refer to d as the combinatorial metric (or ℓ^1 metric). In finite dimensional cube complexes, the CAT(0) and combinatorial metrics are bi-Lipschitz equivalent and complete.

The metric space (X, d) is a median space. This means that, given any three points $p_1, p_2, p_3 \in X$, there exists a unique point $m = m(p_1, p_2, p_3) \in X$ such that $d(p_i, p_j) = d(p_i, m) + d(m, p_j)$ for all $1 \leq i < j \leq 3$. We refer to $m(p_1, p_2, p_3)$ as the median of p_1, p_2 and p_3 ; if the three points are vertices of X , so is $m(p_1, p_2, p_3)$. The map $m : X^3 \rightarrow X$ endows X (and its 0-skeleton) with a structure of median algebra. We refer the reader to [15,30,60] for a definition of the latter and more on median geometry.

We will use the more familiar and concise expression ‘‘CAT(0) cube complex’’ with the meaning of ‘‘simply connected cube complex satisfying Gromov’s no- Δ -condition’’. However, unless specified otherwise, all our cube complexes X will be endowed with the combinatorial metric, all points of X will be implicitly assumed to be vertices and all geodesics will be combinatorial geodesics contained in the 1-skeleton. In some situations, especially in Sects. 2.2 and 3.2, 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 cubical subdivision of X . This is the CAT(0) cube complex obtained by adding a vertex $v(c)$ at the centre of each cube $c \subseteq X$; we then join the vertices $v(c)$ and $v(c')$ by an edge 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 write $\mathcal{W}(X)$ and $\mathcal{H}(X)$, respectively, for the sets of hyperplanes and halfspaces of X . Given a halfspace $h \in \mathcal{H}(X)$, we denote its complement by h^* . Endowing $\mathcal{H}(X)$ with the order relation given by inclusions, the involution $*$ is order-reversing. The triple $(\mathcal{H}(X), \subseteq, *)$ is thus a *poset*, in the sense of [63].

We say that two distinct hyperplanes are transverse if they cross. Similarly, we say that two halfspaces—or a halfspace and hyperplane—are transverse if the corresponding hyperplanes are. Halfspaces h and h' are transverse if and only if all four intersections $h \cap h', h^* \cap h', h \cap h'^*$ and $h^* \cap h'^*$ are nonempty. We say that subsets $A, B \subseteq \mathcal{W}(X)$ are transverse if every element of A is transverse to every element of B .

Every hyperplane w can itself be regarded as a CAT(0) cube complex; its cells are precisely the intersections $w \cap c$, where $c \subseteq X$ is a cube. The set of hyperplanes of the cube complex w is naturally identified with the set of hyperplanes of X that are transverse to w . We thus denote by $\mathcal{W}(w)$ this subset of $\mathcal{W}(X)$. We also denote by $C(w)$ the *carrier* of w , i.e. its neighbourhood of radius $\frac{1}{2}$ in X .

Three hyperplanes w_1, w_2 and w_3 form a facing triple if we can choose pairwise disjoint sides h_1, h_2 and h_3 ; the three halfspaces are then also said to form a facing triple. Halfspaces h and h' are nested if either $h \subseteq h'$ or $h' \subseteq h$.

Given a vertex $p \in X$, we denote by $\sigma_p \subseteq \mathcal{H}(X)$ the set of all halfspaces containing p . It satisfies the following properties:

- (1) any two halfspaces in σ_p intersect non-trivially;
- (2) for every hyperplane $w \in \mathcal{W}(X)$, a side of w lies in σ_p ;
- (3) every descending chain of halfspaces in σ_p is finite.

Subsets $\sigma \subseteq \mathcal{H}(X)$ satisfying (1)–(3) are known as DCC ultrafilters. If a set $\sigma \subseteq \mathcal{H}(X)$ only satisfies (1) and (2), we refer to it simply as an ultrafilter.

Let $\iota: X \rightarrow 2^{\mathcal{H}(X)}$ denote the map that takes each vertex p to the set σ_p . Its image $\iota(X)$ coincides with the collection of all DCC ultrafilters. Endowing $2^{\mathcal{H}(X)}$ with the product topology, we can consider the closure $\overline{\iota(X)}$, which happens to coincide 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 [55]; we denote it by \overline{X} .

The Roller boundary ∂X is defined as the difference $\overline{X} \setminus X$. The inclusion $\iota: X \rightarrow \overline{X}$ is always continuous¹⁰. If, moreover, X is locally finite, then ι is a topological embedding, X is open in \overline{X} and ∂X is compact. Even though elements of ∂X are technically just sets of halfspaces, we will rather think of them as points at infinity. In analogy with vertices of X , we will then write $x \in \partial X$ and reserve the notation σ_x for the ultrafilter representing x .

According to an unpublished result of U. Bader and D. Guralnik (and [18,31]), the identity of X extends to a homeomorphism between ∂X and the horofunction boundary of (X, d) (at least when X is proper, see Remark 6.19 in [31]). However, the characterisation of \overline{X} in terms of ultrafilters additionally provides a natural structure of median algebra on \overline{X} , corresponding to the map

$$m(\sigma_x, \sigma_y, \sigma_z) = (\sigma_x \cap \sigma_y) \cup (\sigma_y \cap \sigma_z) \cup (\sigma_z \cap \sigma_x).$$

Under the identification of $p \in X$ and $\sigma_p \subseteq \mathcal{H}(X)$, this map $m: \overline{X}^3 \rightarrow \overline{X}$ restricts to the usual median-algebra structure on X . Given $x, y \in \overline{X}$, the *interval* between x and y is the set $I(x, y) = \{z \in \overline{X} \mid m(x, y, z) = z\}$. If $x, y, z \in \overline{X}$, the median $m(x, y, z)$ is the only point of \overline{X} that lies in all three intervals $I(x, y)$, $I(y, z)$ and $I(z, x)$. Observe that $I(p, q) \subseteq X$ if $p, q \in X$.

In some instances, we will also have to consider the visual boundary of X associated to the CAT(0) metric. To avoid confusion with the Roller boundary ∂X , we will denote the visual boundary by $\partial_\infty X$; note that $\partial_\infty X$ is the horofunction boundary of X with respect to the CAT(0) metric (Theorem II.8.13 in [6]). When X is Gromov hyperbolic, $\partial_\infty X$ is also naturally identified with the Gromov boundary of X , for which we will adopt the same notation.

Given $p \in X$ and $\mathfrak{h} \in \mathcal{H}(X)$, we have $p \in \mathfrak{h}$ if and only if $\mathfrak{h} \in \sigma_p$. By analogy, we say that a point $x \in \overline{X}$ lies in a halfspace $\mathfrak{h} \in \mathcal{H}(X)$ (written $x \in \mathfrak{h}$), if the halfspace \mathfrak{h} is an element of the ultrafilter σ_x . This should be regarded as a way of extending halfspaces into the boundary, yielding a partition $\overline{X} = \mathfrak{h} \sqcup \mathfrak{h}^*$ for every $\mathfrak{h} \in \mathcal{H}(X)$.

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

$$\begin{aligned} \mathcal{H}(A|B) &= \{\mathfrak{h} \in \mathcal{H}(X) \mid B \subseteq \mathfrak{h}, A \subseteq \mathfrak{h}^*\}, \\ \mathcal{W}(A|B) &= \{\mathfrak{w} \in \mathcal{W}(X) \mid \text{one side of } \mathfrak{w} \text{ lies in } \mathcal{H}(A|B)\}. \end{aligned}$$

If $\mathfrak{w} \in \mathcal{W}(A|B)$, we say that \mathfrak{w} *separates* A and B . We denote by $\mathcal{W}(A)$ the set of all hyperplanes separating two points of A . To avoid possible ambiguities, we adopt the convention that hyperplanes \mathfrak{w} are not contained in either of their sides; in particular, $\mathfrak{w} \notin \mathcal{W}(\mathfrak{w}|A)$ for every $A \subseteq \overline{X}$.

¹⁰ Note that ι is only defined on the 0-skeleton, which has the discrete topology.

Lemma 2.1 *Given points $x, y, z, w \in \bar{X}$, the sets $\mathcal{W}(x, y|z, w)$, $\mathcal{W}(x, z|y, w)$ and $\mathcal{W}(x, w|y, z)$ are pairwise transverse.*

Proof Consider $\mathfrak{h} \in \mathcal{H}(x, y|z, w)$ and $\mathfrak{k} \in \mathcal{H}(x, z|y, w)$. 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 $\mathcal{W}(x, y|z, w)$ and $\mathcal{W}(x, z|y, w)$ are transverse; the same argument shows that they are also transverse to $\mathcal{W}(x, w|y, z)$. \square

We will generally conflate all geodesics (and quasi-geodesics) with their images in X . Every geodesic $\gamma \subseteq X$ can be viewed as a collection of edges; distinct edges $e, e' \subseteq \gamma$ must cross distinct hyperplanes. We write $\mathcal{W}(\gamma)$ for the collection of hyperplanes crossed by (the edges of) γ . If two geodesics γ and γ' share an endpoint $p \in X$, their union $\gamma \cup \gamma'$ is again a geodesic if and only if $\mathcal{W}(\gamma) \cap \mathcal{W}(\gamma') = \emptyset$.

Given a ray $r \subseteq X$, we denote by $r(0)$ its origin and, given $n \in \mathbb{N}$, by $r(n)$ the only point of r with $d(r(0), r(n)) = n$. Given a hyperplane $\mathfrak{w} \in \mathcal{W}(X)$, there is exactly one side of \mathfrak{w} that has unbounded intersection with r . The collection of all these halfspaces is an ultrafilter on $\mathcal{H}(X)$ and it therefore determines a point $r(+\infty) \in \partial X$.

Fixing a basepoint $p \in X$, every point of ∂X is of the form $r(+\infty)$ for a ray r with $r(0) = p$. We obtain a bijection between points of ∂X and equivalence classes of rays based at p , where two rays r_1 and r_2 are equivalent if $\mathcal{W}(r_1) = \mathcal{W}(r_2)$. See Proposition A.2 in [34] for details.

Given two vertices $p, q \in X$, we have $d(p, q) = \#\mathcal{W}(p|q)$. By analogy, we can define $d(x, y) = \#\mathcal{W}(x|y)$ for all points $x, y \in \bar{X}$. The resulting function $d: \bar{X} \times \bar{X} \rightarrow \mathbb{N} \cup \{+\infty\}$ 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 \bar{X} ; we refer to its equivalence classes as components.

Given $x \in \bar{X}$, we denote by $Z(x)$ the only component of \bar{X} that contains the point x . When $x \in X$, we have $Z(x) = X$. For every component $Z \subseteq \bar{X}$, 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.


We obtain here a couple of simple results which will be needed later on.

Lemma 2.2 *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 $x \sim y$. There exists an infinite chain $\mathfrak{j}_0 \supseteq \mathfrak{j}_1 \supseteq \dots$ of halfspaces of X transverse to both \mathfrak{h} and \mathfrak{k} .*

Proof Since $\mathfrak{h} \neq \mathfrak{k}^*$, there exists a point $p \in X$ lying in $\mathfrak{h}^* \cap \mathfrak{k}^*$; see Fig. 1. The set $\mathcal{W}(p|x, y)$ is infinite, as $\mathcal{W}(p|x)$ is infinite and $\mathcal{W}(x|y)$ is finite. Since σ_p is a DCC ultrafilter, the sets $\mathcal{W}(p|\mathfrak{h})$ and $\mathcal{W}(p|\mathfrak{k})$ are finite and the set $A = \mathcal{H}(p|x, y) \setminus (\mathcal{H}(p|\mathfrak{h}) \cup \mathcal{H}(p|\mathfrak{k}))$ is infinite. Any halfspace in A is transverse to both \mathfrak{h} and \mathfrak{k} . Any two elements of $\mathcal{H}(p|x, y)$ are either transverse or nested, and any subset of pairwise transverse halfspaces has cardinality at most $\dim X$. The required chain is thus obtained by applying Ramsey’s theorem to A . \square

Lemma 2.3 *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' satisfying $r'(0) = r(0)$, $r'(+\infty) = x'$ and such that the Hausdorff distance $d_{\text{Haus}}(r, r')$ is at most $d(x, x')$.*

Fig. 1 The setup in the proof of Lemma 2.2



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 \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' . □

We say that a subset $C \subseteq X$ is convex if every geodesic with endpoints in C is entirely contained in C ; equivalently, $I(x, y) \subseteq C$ for every $x, y \in C$. Halfspaces are precisely those nonempty convex subsets of X whose complement is convex and nonempty. Given any subset $A \subseteq X$, we denote by $\text{Hull}(A)$ the smallest convex subset of X that contains A . This coincides with the intersection of all halfspaces containing A . The subcomplex $\text{Hull}(A)$ is itself a CAT(0) cube complex and its hyperplane set is identified with $\mathscr{W}(A)$.

Given pairwise-intersecting convex subsets $C_1, \dots, C_k \subseteq X$, we always have $C_1 \cap \dots \cap C_k \neq \emptyset$. This is known as Helly's lemma (Theorem 2.2 in [60]).

Every convex subset $C \subseteq X$ comes equipped with a 1-Lipschitz projection $\pi_C: X \rightarrow C$, with the property that $\pi_C(x) \in I(x, y)$ for every $x \in X$ and every $y \in C$. We refer to π_C as the *gate-projection* to C . For all $x, y \in X$, we have $\mathscr{W}(x|\pi_C(x)) = \mathscr{W}(x|C)$ and $\mathscr{W}(x|y) \cap \mathscr{W}(C) = \mathscr{W}(\pi_C(x)|\pi_C(y))$, so π_C is the nearest-point projection with respect to the combinatorial metric.

Consider now two disjoint halfspaces \mathfrak{h} and \mathfrak{k} . We set

$$M(\mathfrak{h}, \mathfrak{k}) = \{(x, y) \in \mathfrak{h} \times \mathfrak{k} \mid d(x, y) = d(\mathfrak{h}, \mathfrak{k})\}$$

and denote by $B(\mathfrak{h}, \mathfrak{k})$ the union of all intervals $I(x, y)$ with $(x, y) \in M(\mathfrak{h}, \mathfrak{k})$. The set $B = B(\mathfrak{h}, \mathfrak{k})$ is usually known as the *bridge* and it is a convex subcomplex of X . Let S_1 and S_2 denote the projections of $M(\mathfrak{h}, \mathfrak{k}) \subseteq \mathfrak{h} \times \mathfrak{k}$ to the factors \mathfrak{h} and \mathfrak{k} , respectively. We refer to S_1 and S_2 as the *shores*; note that $S_1 = B \cap \mathfrak{h}$ and $S_2 = B \cap \mathfrak{k}$, so shores are also convex subcomplexes.

The restrictions $\pi_{S_1}|_{S_2}$ and $\pi_{S_2}|_{S_1}$ define cubical isomorphisms between S_1 and S_2 . In fact, the intervals $I(x, y)$ associated to pairs $(x, y) \in M(\mathfrak{h}, \mathfrak{k})$ are pairwise disjoint

and all isomorphic to each other, giving rise to isometric splittings $B \simeq I(x, y) \times S$. Here S_1 corresponds to $\{x\} \times S$ and S_2 corresponds to $\{y\} \times S$.

We refer to the cube complex $S = S(\mathfrak{h}, \mathfrak{k})$ simply as the abstract shore when we do not want to identify it with any specific subcomplex of X . It is precisely the restriction quotient of X (in the sense of p. 860 of [23]) associated to the set of hyperplanes transverse to both \mathfrak{h} and \mathfrak{k} . Finally, we remark that, for every $x \in \mathfrak{h}$ and every $y \in \mathfrak{k}$, we have

$$d(x, y) = d(x, S_1) + d(\pi_{S_1}(x), \pi_{S_1}(y)) + d(\mathfrak{h}, \mathfrak{k}) + d(y, S_2). \quad (*)$$

The reader can consult for instance Section 2.G of [17] or Section 2.2 of [29] for a more detailed treatment of bridges and shores.

We say that two disjoint halfspaces \mathfrak{h} and \mathfrak{k} are strongly separated if the corresponding shores are singletons. Equivalently, no hyperplane of X is transverse to both \mathfrak{h} and \mathfrak{k} . Similarly, we say that two hyperplanes are strongly separated if they bound strongly separated halfspaces.

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 [23]); we refer to it as the De Rham decomposition.

Throughout the paper, all groups will be implicitly assumed to be finitely generated. When a group G acts on a CAT(0) cube complex X , we will assume that the action is by cubical automorphisms, i.e. by isometries taking vertices to vertices. We say that $G \curvearrowright X$ is essential if no G -orbit is contained in a metric neighbourhood of a halfspace. Similarly, we say that X is essential if no halfspace is contained in a metric neighbourhood of the corresponding hyperplane. If G acts cocompactly, X is essential if and only if the action $G \curvearrowright X$ is essential.

Remark 2.4 Every essential, Gromov hyperbolic CAT(0) cube complex is irreducible. Indeed, essentiality guarantees that X has no bounded factors in its De Rham decomposition, whereas hyperbolicity implies that there is at most one unbounded factor.

The action $G \curvearrowright X$ is hyperplane-essential if each hyperplane-stabiliser acts essentially on the corresponding hyperplane. Similarly, X is hyperplane-essential if all its hyperplanes are essential cube complexes. Again, if G acts cocompactly, X is hyperplane-essential if and only if the action $G \curvearrowright X$ is hyperplane-essential. This follows from Exercise 1.6 in [63], which we record here for later use:

Lemma 2.5 *Let $\mathfrak{w} \in \mathcal{W}(X)$ be a hyperplane and $G_{\mathfrak{w}} < G$ its stabiliser. If $G \curvearrowright X$ is cocompact, the action $G_{\mathfrak{w}} \curvearrowright \mathfrak{w}$ also is.*

The following notion appeared in Definition 7.3 in [28] and Definition 5.8 in [31]; also see Proposition 7.5 in [28].

Definition 2.6 Let X be irreducible. 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 chain $\mathfrak{h}_0 \supsetneq \mathfrak{h}_1 \supsetneq \dots$ 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.

With reference to the proof sketch of Theorem C in the introduction, the following is the formulation of Steps (Ia) and (Ib) that we will actually use. Note that hyperbolicity is not required here.

Proposition 2.7 *Let X be irreducible, essential and endowed with a proper cocompact action $G \curvearrowright X$ of a non-virtually-cyclic group. Consider two halfspaces h_1, h_2 and a nonempty G -invariant subset $\mathcal{A} \subseteq \partial_{\text{reg}} X$.*

- (1) *The intersections $h_i \cap \mathcal{A}$ and $h_i^* \cap \mathcal{A}$ are always nonempty. If moreover X is hyperplane-essential, the following also hold.*
- (2) *The halfspaces h_1 and h_2 are transverse if and only if the set \mathcal{A} intersects each of the four sectors $h_1 \cap h_2, h_1^* \cap h_2, h_1 \cap h_2^*$ and $h_1^* \cap h_2^*$.*
- (3) *If we have $h_1 \cap \mathcal{A} \subsetneq h_2 \cap \mathcal{A}$, then $h_1 \subsetneq h_2$.*

Proof Part (1) follows from Lemmas 2.9 and 2.19 in [4]. If h_1 and h_2 are not transverse, one of the four intersections $h_1 \cap h_2, h_1^* \cap h_2, h_1 \cap h_2^*, h_1^* \cap h_2^*$ is empty by definition; in particular, it cannot contain any point of \mathcal{A} . This proves one implication of part (2), while the other follows from part (1) and Proposition 2.11 in [4]. We conclude by proving part (3).

If $h_1 \cap \mathcal{A} \subsetneq h_2 \cap \mathcal{A}$, part (2) shows that h_1 and h_2 cannot be transverse. We then have either $h_1 \subsetneq h_2$, or $h_2 \subseteq h_1$, or $h_2^* \subseteq h_1$ or $h_2 \subseteq h_1^*$. In the second case we would have $h_1 \cap \mathcal{A} \subsetneq h_2 \cap \mathcal{A} \subseteq h_1 \cap \mathcal{A}$, and in the third case $h_1 \cap h_2^* \cap \mathcal{A} \supseteq h_2^* \cap \mathcal{A} \neq \emptyset$, which both lead to contradictions. In the fourth case, taking complements we obtain $h_1 \subseteq h_2^*$, hence $h_1 \cap h_2^* \cap \mathcal{A} \supseteq h_1 \cap \mathcal{A} \neq \emptyset$, which is also a contradiction. We conclude that $h_1 \subsetneq h_2$. □

In relation to part (3) of Proposition 2.7, note however that $h_1 \cap \mathcal{A} \subseteq h_2 \cap \mathcal{A}$ does not imply $h_1 \subseteq h_2$, as we might actually have $h_2 \subsetneq h_1$ in this case.

2.2 Combinatorial geodesics vs CAT(0) geodesics

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

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 2.8 *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 $\mathcal{W}(\gamma)$ can be arranged in a sequence $(w_n)_{n \geq 0}$ according to the order in which they are crossed by γ after $\gamma(0)$. We denote this sequence by $\varepsilon(\gamma)$.

Lemma 2.9 *Let $(w_n)_{n \geq 0}$ be a (finite or infinite) sequence of pairwise distinct hyperplanes of X and let $p \in X$ be a vertex. There exists a combinatorial geodesic γ*

based at p such that $\mathfrak{s}(\gamma) = (\mathfrak{w}_n)_{n \geq 0}$ if and only if, for every $n \geq 0$, we have $\mathscr{W}(p|\mathfrak{w}_n) = \{\mathfrak{w}_0, \dots, \mathfrak{w}_{n-1}\} \setminus \mathscr{W}(\mathfrak{w}_n)$.

Proof If there exists a geodesic γ such that $\gamma(0) = p$ and $\mathfrak{s}(\gamma) = (\mathfrak{w}_n)_{n \geq 0}$, 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 $\mathscr{W}(p|\mathfrak{w}_n) = \{\mathfrak{w}_0, \dots, \mathfrak{w}_{n-1}\} \setminus \mathscr{W}(\mathfrak{w}_n)$ for all $n \geq 0$.

Assuming instead that the sequence $(\mathfrak{w}_n)_{n \geq 0}$ satisfies the latter condition, we are going to construct points $p_n \in C(\mathfrak{w}_n)$ with $\mathscr{W}(p|p_n) = \{\mathfrak{w}_0, \dots, \mathfrak{w}_{n-1}\}$. 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 $\mathscr{W}(p|p_n) = \{\mathfrak{w}_0, \dots, \mathfrak{w}_{n-1}\}$, let $p_{n+1} \in X$ be the only point with $\mathscr{W}(p_n|p_{n+1}) = \{\mathfrak{w}_n\}$. As $\mathscr{W}(p|p_{n+1}) = \{\mathfrak{w}_0, \dots, \mathfrak{w}_n\}$, 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 $u \in \mathscr{W}(p_{n+1}|\mathfrak{w}_{n+1})$ and we would have either $u \in \mathscr{W}(p, p_{n+1}|\mathfrak{w}_{n+1})$ or $u \in \mathscr{W}(p_{n+1}|p, \mathfrak{w}_{n+1})$. The former is forbidden by $\mathscr{W}(p|\mathfrak{w}_{n+1}) \subseteq \{\mathfrak{w}_0, \dots, \mathfrak{w}_n\}$, 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 2.8 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 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 2.9, there exists a combinatorial ray r from p to x such that $\mathfrak{s}(\gamma) = (\mathfrak{w}_n)_{n \geq 0}$. It remains to prove that $d_{\text{Haus}}(r, \rho) \leq D$.

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 $\iota: I(p, x) \cap X \rightarrow \mathbb{R}^D$ be an ℓ^1 -isometric cubical embedding; it exists for instance by Theorem 1.14 in [3]. 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 $d(u, v) = d(\iota(u), \iota(v)) \leq D$. □

2.3 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 closed ball containing S ; see e.g. Proposition II.2.7 in [6] or Proposition 3.73 in [26]. 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 Fig. 2.

In Sect. 5 we will need a different notion of barycentre, which we now introduce. It will always be a vertex of the first cubical subdivision S' .




Fig. 2 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 $\frac{1}{2\pi} \sqrt{a^2 + b^2}$ for any $a, b, n \in \mathbb{N}$, so c_S is not a vertex of any iterated cubical subdivision of S

In the discussion below, points of S are assumed to be vertices and we only consider the combinatorial metric on X , as in the rest of the paper.

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

Lemma 2.10 *Given halfspaces $h, k \in \mathcal{H}(S)$ with $h \cap k = \emptyset$ and $h \neq k^*$, at least one of them is light.*

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

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

and similarly $d(y, u) > d(x, w)$. 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 Fig. 2.

We remark that, given bounded CAT(0) cube complexes S_1 and S_2 and an isomorphism $F : S_1 \rightarrow S_2$, we have $F(m_{S_1}) = m_{S_2}$.

2.4 Cross ratios on cube complexes

Let X be a CAT(0) cube complex. Fixing a vertex $p \in X$, the Gromov product of two points $x, y \in \overline{X}$ is:

$$(x \cdot y)_p = \#\mathcal{W}(p|x, y) = d(p, m(p, x, y)) \in \mathbb{N} \cup \{+\infty\}.$$

The following is Lemma 2.3 in [4].

Lemma 2.11 *Consider $x, y, z \in \overline{X}$ and $p \in X$.*

- (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 $(x \cdot y)_p < +\infty$ if and only if $I(x, y)$ intersects X .

Let $\mathcal{A} \subseteq (\overline{X})^4$ be the subset of 4-tuples (x, y, z, w) such that at most one of the values $(x \cdot y)_p + (z \cdot w)_p$, $(x \cdot z)_p + (y \cdot w)_p$ and $(x \cdot w)_p + (y \cdot z)_p$ is infinite; by Lemma 2.11, the set \mathcal{A} is independent of the choice of p .

In our previous work with Incerti-Medici [5], we introduced a cross ratio $\text{cr}: \mathcal{A} \rightarrow \mathbb{Z} \cup \{\pm\infty\}$, which admits the following equivalent characterisations:

- (1) $\text{cr}(x, y, z, w) = \#\mathcal{W}(x, z|y, w) - \#\mathcal{W}(x, w|y, z)$;
- (2) $\text{cr}(x, y, z, w) = (x \cdot z)_p + (y \cdot w)_p - (x \cdot w)_p - (y \cdot z)_p$;
- (3) $\text{cr}(x, y, z, w) = d(x, w) + d(y, z) - d(x, z) - d(y, w)$, if $x, y, z, w \in X$.

In particular, the second characterisation does not depend on the choice of $p \in X$. Note that cr satisfies symmetries (i)–(iv) from the introduction, as long as all involved 4-tuples lie in \mathcal{A} . We will sometimes write cr_X when we wish to specify the cube complex under consideration.

Endowing $\mathcal{A} \subseteq (\overline{X})^4$ with the subspace topology, the following is Proposition 3.3 in [4]. Note that \overline{X} and \mathcal{A} are totally disconnected.

Proposition 2.12 *If X is locally finite, the cross ratio cr is continuous.*

2.5 CAT(0) cuboid complexes

As mentioned in the introduction, all results in this paper equally hold for cube complexes with variable edge lengths: cuboid complexes in our terminology.

For the sake of simplicity and clarity, we will only treat CAT(0) cube complexes in most of the paper. Only very minor changes are required in order to adapt our arguments to general CAT(0) cuboid complexes. We briefly describe them here, along with the relevant definitions.

Consider for a moment a genuine CAT(0) cube complex X . Every function $\mu: \mathcal{W}(X) \rightarrow \mathbb{R}_{>0}$ determines a weighted combinatorial metric d_μ on X . For vertices $v, w \in X$, this is given by:

$$d_\mu(v, w) = \sum_{\mathfrak{w} \in \mathcal{W}(v|w)} \mu(\mathfrak{w}).$$

For instance, the usual combinatorial metric d arises from the function that assigns value +1 to each hyperplane.

Definition 2.13 A CAT(0) cuboid complex \mathbb{X} is any metric cell complex (X, d_μ) arising from this construction.

Two cuboid complexes $\mathbb{X} = (X, d_\mu)$ and $\mathbb{Y} = (Y, d_\nu)$ are isomorphic if there exists an isometric cellular isomorphism $f: \mathbb{X} \rightarrow \mathbb{Y}$. In other words, $f: X \rightarrow Y$ is an isomorphism of CAT(0) cube complexes inducing a map $f_*: \mathcal{W}(X) \rightarrow \mathcal{W}(Y)$ such that $\mu = \nu \circ f_*$.

Note, however, that there can be isometries $\mathbb{X} \rightarrow \mathbb{Y}$ that do not preserve the cellular structures. For instance, consider the cuboid complex \mathbb{X}' arising from the cubical subdivision X' . If we assign each edge of \mathbb{X}' half the length of the corresponding edge of \mathbb{X} , the identity map $\mathbb{X} \rightarrow \mathbb{X}'$ is a (surjective) isometry, but never an isomorphism.

When dealing with CAT(0) cuboid complexes, rather than CAT(0) cube complexes, the following adaptations and conventions are required.

- (1) All group actions on cuboid complexes will be assumed to be by automorphisms (i.e. self-isomorphisms).
- (2) We consider two actions to be the same when they are equivariantly isometric. In Theorems A, C and in Corollaries B, E, equivariant isomorphisms of cube complexes need to be replaced with equivariant isometries of cuboid complexes. It is easy to see that it is not possible to map vertices to vertices in general.
- (3) We define hyperplanes and halfspaces of $\mathbb{X} = (X, d_\mu)$ to coincide with hyperplanes and halfspaces of the underlying cube complex X . This also explains how to interpret notations like $\mathscr{W}(A|B)$ in this context.
- (4) Given a subset $\mathcal{U} \subseteq \mathscr{W}(\mathbb{X}) = \mathscr{W}(X)$, the cardinality $\#\mathcal{U}$ should always be replaced by the weight $\sum_{\mathfrak{w} \in \mathcal{U}} \mu(\mathfrak{w})$. Nonempty subsets are still precisely those that have positive weight.
- (5) The cross ratio cr will no longer take values in $\mathbb{Z} \cup \{\pm\infty\}$, but rather in $M \cup \{\pm\infty\}$, where M is the \mathbb{Z} -module generated by the image of the map μ . A similar observation applies to length functions.

3 The Morse property in cube complexes

Other than the proof of Theorem D (consisting of Theorem 3.10 and Remark 3.8), most of this section will be devoted to collecting more or less well-known facts from the literature. Throughout:

Standing Assumptions Let the CAT(0) cube complex X be finite dimensional and locally finite.

3.1 Contracting geodesics

Recall that we only endow X with its combinatorial metric. All geodesics will be combinatorial in this subsection.

Definition 3.1 Let Y be a proper metric space. Given a closed subset $A \subseteq Y$, we denote by $\pi_A: Y \rightarrow 2^A$ 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 (strongly) contracting if there exists $D > 0$ such that every metric ball B disjoint from A satisfies $\text{diam}(\pi_A(B)) \leq D$.

Given a function $M: [1, +\infty) \rightarrow [0, +\infty)$, 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 [1] for a detailed discussion of contracting subsets and the Morse property in general metric spaces.

Definition 3.2 We say that a geodesic $\gamma \subseteq X$ is C -lean if there do not exist transverse subsets $\mathcal{U} \subseteq \mathcal{W}(\gamma)$ and $\mathcal{V} \subseteq \mathcal{W}(X)$ 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 $C \geq 0$.

The following is due to A. Genevois; see Corollary 3.7 in [34] and Lemma 4.6 in [33].

Theorem 3.3 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 3.4 Let α and γ be rays in X with $\alpha(+\infty) \sim \gamma(+\infty)$. If γ is contracting, then α is at finite Hausdorff distance from γ , hence contracting.

Proof By Lemma 2.3, it suffices to consider the case when $\alpha(+\infty) = \gamma(+\infty)$. By Theorem 3.3, 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 $\mathcal{W}(q|I) = \mathcal{W}(q|x, p)$ is a finite subset of $\mathcal{W}(q|x) = \mathcal{W}(\alpha)$, 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 $\mathcal{W}(m|u)$ and $\mathcal{W}(m|v)$ are contained in $\mathcal{W}(x|p) = \mathcal{W}(\gamma)$. Every halfspace $\mathfrak{h} \in \mathcal{H}(m|u)$ is transverse to every halfspace $\mathfrak{k} \in \mathcal{H}(m|v)$; 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 $\mathcal{W}(m|u)$ and $\mathcal{W}(m|v)$ have the same size and contain no facing triples. Since γ is C -lean, we conclude that $\#\mathcal{W}(m|u) = \#\mathcal{W}(m|v) \leq C$. This shows that $d(u, v) \leq 2C$. \square

3.2 Roller boundaries vs contracting boundaries

Unlike the rest of the paper, this subsection 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. Still, the notation $d(\cdot, \cdot)$ will always refer to the combinatorial metric.

The contracting boundary $\partial_c X$ was introduced in [24]. 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 (see Lemma 3.4 for the equivalence in the definition):

Definition 3.5 We say that a point $x \in \partial X$ is contracting (or a contracting ultrafilter) if one (equivalently, each) combinatorial ray representing x is contracting. We denote the set of contracting ultrafilters by $\partial_{cu} X \subseteq \partial X$.

We stress that our definition of contracting point is not equivalent to the one in Remark 6.7 of [31]; in fact, our notion is weaker.

In general, the inclusion $\partial_{\text{cu}}X \subseteq \partial X$ is strict. If however X is Gromov hyperbolic, every combinatorial ray in X is contracting (see e.g. Theorem 3.3 in [32]) and we have $\partial_{\text{cu}}X = \partial X$.

Lemma 3.4 shows that the set $\partial_{\text{cu}}X$ is a union of \sim -equivalence classes. The following result provides more information.

Lemma 3.6 (1) *Every component of $\partial_{\text{cu}}X$ is bounded.*

(2) *Points $x \in \partial_{\text{cu}}X$ 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 $x \in \partial_{\text{cu}}X$ 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 $d(x, y) \leq C$.

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 $\mathfrak{w} \in \mathscr{W}(x|y)$ 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, all but finitely many hyperplanes in $\mathscr{W}(p|x, y)$ are 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 . \square

Lemma 3.7 *Consider a point $x \in \partial_{\text{cu}}X$. There exists an infinite descending chain of halfspaces $\mathfrak{h}_0 \supseteq \mathfrak{h}_1 \supseteq \dots$ 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 [34] yields an infinite chain of halfspaces $\mathfrak{h}_0 \supseteq \mathfrak{h}_1 \supseteq \dots$ 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 3.6 then shows that $z \in Z(x)$. This proves that $\bigcap \mathfrak{h}_n = Z(x)$ and concludes the proof. \square

Remark 3.8 If X is uniformly locally finite, part (1) of Lemma 3.6 can actually be promoted to say that components of $\partial_{\text{cu}}X$ are finite.

Indeed, let x and $\mathfrak{h}_0 \supseteq \mathfrak{h}_1 \supseteq \dots$ be as in the statement of Lemma 3.7. 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 [24]). Along with Theorem 3.3 and Proposition 2.8, this observation yields the following.

Corollary 3.9 *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 $x \in \partial_{\text{cu}}X$ is represented by combinatorial rays in X . These rays are all contracting by Lemma 3.4 and Corollary 3.9 shows that they are at finite Hausdorff distance from a unique family of pairwise-asymptotic contracting CAT(0) rays. This yields a map

$$\Phi: \partial_{\text{cu}}X \longrightarrow \partial_c X.$$

We endow $\partial_{\text{cu}}X$ with the restriction of the topology of ∂X . We write $\partial_c^{\text{vis}}X$ 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 $\partial_c X$ ([19,24]), it is all that we will need in most of the paper.

The next result describes a few properties of the map Φ . Recall the standing assumption that X be finite dimensional and locally finite.

Theorem 3.10 *The map $\Phi: \partial_{\text{cu}}X \rightarrow \partial_c^{\text{vis}}X$ is a continuous surjection, whose fibres are precisely the \sim -equivalence classes in $\partial_{\text{cu}}X$. Moreover, Φ descends to a homeomorphism $\bar{\Phi}: \partial_{\text{cu}}X/\sim \rightarrow \partial_c^{\text{vis}}X$.*

Proof Surjectivity is immediate from Corollary 3.9. Lemma 3.4 shows that Φ is constant on \sim -equivalence classes. On the other hand, if γ and γ' are combinatorial rays with $\gamma(+\infty) \approx \gamma'(\infty)$, it is clear that the distance $d(\gamma(n), \gamma'(n))$ diverges as n goes to infinity. We conclude that the fibres of Φ are precisely the components of $\partial_{\text{cu}}X$.

We now prove continuity of Φ . Fix a vertex $p \in X$ and let $D = \dim X$. Given points $x, y \in \partial_{\text{cu}}X$, let ρ_x and ρ_y be the CAT(0) rays from p to $\Phi(x)$ and $\Phi(y)$, respectively. Let r_x and r_y be combinatorial rays based at p and satisfying $d_{\text{Haus}}(r_x, \rho_x) \leq D$, $d_{\text{Haus}}(r_y, \rho_y) \leq D$, as provided by Proposition 2.8. The points $x' = r_x(+\infty)$ and $y' = r_y(+\infty)$ lie in the components $Z(x)$ and $Z(y)$, respectively. Theorem 3.3 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 $\partial_{\text{cu}}X$ as follows. Pick halfspaces $\mathfrak{h}_3 \subseteq \mathfrak{h}_2 \subseteq \mathfrak{h}_1$ with $Z(x) \subseteq \mathfrak{h}_3$, $d(p, \mathfrak{h}_1) > n$, $\#\mathcal{W}(\mathfrak{h}_1^*|\mathfrak{h}_2) > C$ and $\#\mathcal{W}(\mathfrak{h}_2^*|\mathfrak{h}_3) > C$. Their existence is ensured by Lemma 3.7. The neighbourhood $U_n(x)$ is then the subset of $\partial_{\text{cu}}X$ 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 $d_{\text{Haus}}(r_x, \rho_x) \leq D$ and $d_{\text{Haus}}(r_y, \rho_y) \leq D$, the claim shows that, given any open neighbourhood V of $\Phi(x)$ in $\partial_c^{\text{vis}}X$, 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 2.2 implies that $x', y' \in \mathfrak{h}_2$. Indeed, as $\#\mathcal{W}(\mathfrak{h}_2^*|\mathfrak{h}_3) > C$ 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 $m = m(p, q_x, q_y)$. Note that the points q_x, q_y and m lie in \mathfrak{h}_1^* as $d(p, \mathfrak{h}_1) > n$. Let $A_x \subseteq \mathcal{H}(m|q_x)$ and $A_y \subseteq \mathcal{H}(m|q_y)$ be the subsets of halfspaces containing \mathfrak{h}_2 ; set $a_x = \#A_x$ and $a_y = \#A_y$.

Each $\mathfrak{k} \in \mathcal{H}(m|q_x) \setminus A_x$ 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 $\mathcal{H}(m|q_x) \setminus A_x$ is transverse to \mathfrak{h}_2 and, similarly, to \mathfrak{h}_1 . Since $\mathcal{W}(\mathfrak{h}_1^*|\mathfrak{h}_2) \subseteq \mathcal{W}(r_x)$ and $\#\mathcal{W}(\mathfrak{h}_1^*|\mathfrak{h}_2) > C$, 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 $A_x \subseteq \mathcal{W}(r_x)$, 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

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

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

Fix a basepoint $p \in X$. Let ρ_n and ρ be CAT(0) rays from p to x_n and x , respectively. Proposition 2.8 yields combinatorial rays r_n and r based at p , with $d_{\text{Haus}}(r_n, \rho_n) \leq D$ and $d_{\text{Haus}}(r, \rho) \leq D$. By Theorem 3.3, there exists $C \geq 0$ such that r is C -lean. Let $\mathfrak{h}_0 \supseteq \mathfrak{h}_1 \supseteq \dots$ be a chain in $\sigma_x \setminus \sigma_p$ with $\bigcap \mathfrak{h}_n = Z(x)$, as provided by Lemma 3.7. Up to passing to a subchain, we can assume that $\#\mathcal{W}(\mathfrak{h}_k^*|\mathfrak{h}_{k+1}) > C$ 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 2.2, 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 (see e.g. Theorem 30.1(b) in [51]). Denote by $p: \partial_{\text{cu}}X \rightarrow \partial_{\text{cu}}X/\sim$ the quotient projection. Suppose for the sake of contradiction that points $x, x_n \in \partial_{\text{cu}}X$ are given so that $\Phi(x_n) \rightarrow \Phi(x)$, but $p(x_n) \not\rightarrow p(x)$. Possibly passing to a subsequence, 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 3.11 Given points $x, x_n \in \partial_{\text{cu}}X$ satisfying $\Phi(x_n) \rightarrow \Phi(x)$ and $Z(x) = \{x\}$, we always have $x_n \rightarrow x$. Otherwise, compactness of ∂X would yield a subsequence x_{n_k} converging to a point $y \in \partial X$ different from x . However, we have shown during the proof of Theorem 3.10 (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 3.12 *If X is Gromov hyperbolic, the map $\Phi: \partial X \rightarrow \partial_\infty X$ descends to a homeomorphism $\bar{\Phi}: \partial X / \sim \rightarrow \partial_\infty X$.*

3.3 Contracting non-terminating ultrafilters

In this subsection:

Standing Assumptions Let, in addition, X be irreducible.

The following definition is originally due to Nevo and Sageev (cf. Section 3.1 in [55]).

Definition 3.13 A point $x \in \partial X$ is non-terminating if $Z(x) = \{x\}$. Equivalently¹¹, the poset (σ_x, \subseteq) does not have minimal elements.

We denote by $\partial_{\text{nt}} X \subseteq \partial X$ the subset of non-terminating ultrafilters and set $\partial_{\text{cnt}} X = \partial_{\text{cu}} X \cap \partial_{\text{nt}} X$. By Theorem 3.10, the restriction to $\partial_{\text{cnt}} X$ of the map $\Phi: \partial_{\text{cu}} X \rightarrow \partial_c X$ is injective. Therefore, we will generally identify the sets $\partial_{\text{cnt}} X \subseteq \partial_{\text{nt}} X$ and $\Phi(\partial_{\text{cnt}} X) \subseteq \partial_c X$, even writing simply $\partial_{\text{cnt}} X \subseteq \partial_c X$.

Remark 3.14 It follows from Theorem 3.10 and Remark 3.11 that $\partial_{\text{cnt}} X$ inherits the same topology from $\partial_{\text{cu}} X \subseteq \partial X$ and $\partial_c^{\text{vis}} X$.

It will be useful to make the following observation.

Lemma 3.15 *We have $\partial_{\text{cnt}} X = \partial_{\text{cu}} X \cap \partial_{\text{reg}} X$.*

Proof We have $\partial_{\text{cnt}} X \supseteq \partial_{\text{cu}} X \cap \partial_{\text{reg}} X$ since it is clear from Definition 2.6 that $\partial_{\text{reg}} X \subseteq \partial_{\text{nt}} X$. For the other inclusion, consider a point $x \in \partial_{\text{cnt}} X$. Lemma 3.7 yields an infinite descending chain of halfspaces $\mathfrak{h}_0 \supseteq \mathfrak{h}_1 \supseteq \dots$ 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 $\bigcap \mathfrak{h}_n = \{x\}$. \square

Contracting non-terminating ultrafilters are plentiful, as they for instance arise in relation to the following notion (see Propositions 3.17 and 3.18).

¹¹ This is because, given an ultrafilter σ and a halfspace $\mathfrak{h} \in \sigma$, the set $(\sigma \setminus \{\mathfrak{h}\}) \cup \{\mathfrak{h}^*\}$ is an ultrafilter (at distance 1 from σ) if and only if \mathfrak{h} is minimal in (σ, \subseteq) .

Definition 3.16 (Definition 4.1 in [4]) We say that $g \in \text{Aut}(X)$ is *neatly contracting* if there exist halfspaces h_1 and h_2 such that $gh_1 \subseteq h_2 \subseteq h_1$ and both pairs (h_2, h_1^*) and (gh_1, h_2^*) are strongly separated.

We collect here various facts on neatly contracting automorphisms that will be needed later on.

Proposition 3.17 *Every neatly contracting automorphism $g \in \text{Aut}(X)$ has exactly two fixed points $g^\pm \in \partial_{\text{cnt}}X$. Given any $x \in \bar{X} \setminus \{g^\pm\}$, we have $g^n x \rightarrow g^+$ and $g^{-n}x \rightarrow g^-$ for $n \rightarrow +\infty$.*

Proof This is essentially Proposition 4.3 in [4]; we only need to prove that the points g^\pm are represented by contracting rays. By Propositions 2.7 and 4.4 in [4], there exists a $\langle g \rangle$ -invariant line $\gamma \subseteq X$ with endpoints g^\pm . As γ must cross all hyperplanes bounding the halfspaces $g^n h_1$, Theorem 3.9 in [34] shows that γ is contracting. This concludes the proof. □

We will write $g_X^\pm \in \partial_{\text{cnt}}X$ when it is necessary to specify the cube complex. The next proposition follows from Lemmas 2.9 and 4.7 in [4], although the main ingredients are actually from [31]. The same result holds for any finite collection of (irreducible, essential, cocompact) cubulations.

Proposition 3.18 *Let a non-virtually-cyclic group G act properly and cocompactly on irreducible, essential CAT(0) cube complexes X and Y . There exists $g \in G$ simultaneously acting as a neatly contracting automorphism on X and Y .*

Given a finitely generated group G , its contracting boundary¹² $\partial_c G$ was introduced in [19,21]; we will not endow $\partial_c G$ with any topology.

Fixing a word metric on G , we say that $g \in G$ is Morse if $n \mapsto g^n$ is a Morse quasi-geodesic in G . This notion is independent of the chosen word metric. Every Morse element $g \in G$ fixes exactly two points $g^{\pm\infty} \in \partial_c G$ (see e.g. Definition 9.1, Theorem 2.2 and Theorem 9.4 in [19]).

Given an action $G \curvearrowright X$ as in the statement of Proposition 3.18, the Milnor–Schwarz lemma shows that orbit maps $G \rightarrow X$ are G -equivariant quasi-isometries. They are at finite distance from each other and all yield the same G -equivariant bijection $o_X : \partial_c G \rightarrow \partial_c X$. This follows from either Proposition 4.2 in [21] or Corollary 6.2 in [19].

Now, if $g \in G$ acts on X as a neatly contracting automorphism, we can consider the points $g^\pm \in \partial_{\text{cnt}}X \subseteq \partial_c X$ introduced in Proposition 3.17. Note that, in this case, g is also a Morse element in G by Theorem 3.3.

Lemma 3.19 *Let $G \curvearrowright X$ be as in Proposition 3.18. For every neatly contracting element $g \in G$, 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 G -fixed points in $\partial_c G$ 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

¹² This is also known as Morse boundary. See the discussion in the introduction for a justification of our terminology.

in [19]). By Theorem 9.4 and Corollary 6.2 in [19], the element g acts on $\partial_c^{\mathcal{FQ}} X$ with north-south dynamics. Note that g also acts on $\partial_c^{\text{vis}} X$ with north-south dynamics, by Proposition 3.17 and Theorem 3.10. Since the identity map $\partial_c^{\mathcal{FQ}} X \rightarrow \partial_c^{\text{vis}} X$ is continuous (Section 7 in [19]), we conclude that the attracting/repelling fixed points of g are the same for the two topologies. \square

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

4 Cubulations of hyperbolic groups

This section is devoted to the proof of Theorem C. Referring to the sketch given in the introduction, we only need to carry out Steps (IIa) and (IIb), as Step (I) is provided by Proposition 2.7. The crucial results are thus Proposition 4.26 in Sect. 4.4 and Theorem 4.33 in Sect. 4.5. Before that, Sects. 4.1, 4.2 and 4.3 develop all necessary ingredients.

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 4.1 *Let G be Gromov hyperbolic and let $H < G$ be an infinite-index quasi-convex subgroup. The limit set $\partial_\infty H \subseteq \partial_\infty G$ is nowhere-dense.*

Proof By Lemma 2.9 in [36], there exists a point $\xi \in \partial_\infty G \setminus \partial_\infty H$. We can assume that H is non-elementary, otherwise 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 $H \cdot \xi \subseteq \partial_\infty G \setminus \partial_\infty H$. Along with the fact that $\partial_\infty H$ is closed in $\partial_\infty G$, this concludes the proof. \square

Lemma 4.2 *Let G be Gromov hyperbolic and let H and K be quasi-convex subgroups. If H and K have the same limit set, they are commensurable.*

Proof Let $L \leq G$ denote the stabiliser of the limit set $\Lambda = \partial_\infty H = \partial_\infty K$. By Lemma 3.8 in [48], the limit set of L coincides with Λ . Since $H \leq L$ and H is quasi-convex in G , Proposition 3.4 in [48] shows that L is quasi-convex in G . Finally, Lemma 2.9 in [36] implies that H and K have finite index in L , hence $H \cap K$ has finite index in both H and K . \square

4.1 Traces at infinity

Throughout this subsection:

Standing Assumptions Let G be a Gromov hyperbolic group (not necessarily non-elementary) with a proper cocompact action on a CAT(0) cube complex X . We denote by $\alpha_X: \partial_\infty G \rightarrow \partial_\infty X$ the only G -equivariant homeomorphism and by $\Phi: \partial X \rightarrow \partial_\infty X$ the map from Corollary 3.12. 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 $\mathfrak{w} \in \mathcal{W}(X)$ 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 4.3 We refer to $\partial_\infty \mathfrak{w} \subseteq \partial_\infty X$ as the *trace at infinity* of \mathfrak{w} .

Denoting by $G_{\mathfrak{w}} \leq G$ the stabiliser of \mathfrak{w} , Lemma 2.5 guarantees that $G_{\mathfrak{w}}$ acts properly and cocompactly on $C(\mathfrak{w})$. In particular, the homeomorphism $o_X^{-1} : \partial_\infty X \rightarrow \partial_\infty G$ takes $\partial_\infty \mathfrak{w}$ to the limit set $\partial_\infty G_{\mathfrak{w}} \subseteq \partial_\infty G$.

We now make a few simple observations on traces at infinity.

Lemma 4.4 Consider a hyperplane $\mathfrak{w} \in \mathcal{W}(X)$ and its two sides \mathfrak{h} and \mathfrak{h}^* .

- (1) The subsets $\partial_\infty \mathfrak{h}$ and $\partial_\infty \mathfrak{h}^*$ are closed in $\partial_\infty X$. Moreover, 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 2.5 ensures that x_n can be chosen so that their gate-projections \bar{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 \bar{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 4.5 In the setting of Lemma 4.4, it is interesting to observe 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 4.6 By Lemma 2.5, 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 3.12 and Remark 3.14, we feel entitled to implicitly identify the sets $\partial_{\text{nt}} X \subseteq \partial X$ and $\Phi(\partial_{\text{nt}} X) \subseteq \partial_\infty X$.

Lemma 4.7 A point $\xi \in \partial_\infty X$ lies in $\partial_{\text{nt}} X$ if and only if there does not exist any $\mathfrak{w} \in \mathcal{W}(X)$ with $\xi \in \partial_\infty \mathfrak{w}$.

Proof By part (1) of Lemma 4.4, 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 $\mathfrak{h} \in \mathcal{H}(X)$ such that $\xi \in \partial_\infty \mathfrak{h} \cap \partial_\infty \mathfrak{h}^*$. 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 3.10 and Definition 3.13). \square

In the rest of the subsection, we collect various facts that will be needed in the proof of Step (IIa). The key ingredient is Proposition 4.11 below.

Given hyperplanes $u, \mathfrak{w} \in \mathcal{W}(X)$, we denote by $B(u, \mathfrak{w}) \subseteq X$ the bridge associated to their carriers. When u and \mathfrak{w} are transverse, $B(u, \mathfrak{w})$ is simply the intersection of the two carriers.

Proposition 4.8 *Given $u, \mathfrak{w} \in \mathcal{W}(X)$, the group $G_u \cap G_{\mathfrak{w}}$ acts properly and cocompactly on $B(u, \mathfrak{w})$. Furthermore, $\partial_\infty B(u, \mathfrak{w}) = \partial_\infty u \cap \partial_\infty \mathfrak{w}$.*

Proof Let S denote the intersection of $B = B(u, \mathfrak{w})$ with the carrier of u and consider its stabiliser $G_S \leq G$. The action $G_S \curvearrowright S$ is cocompact by Proposition 2.7 in [41]. There exists a constant $D > 0$ such that S is contained in the D -neighbourhood of both u and \mathfrak{w} . The finitely many hyperplanes that contain S in their D -neighbourhood are permuted by G_S and it follows that $G_u \cap G_{\mathfrak{w}}$ sits in G_S as a finite-index subgroup. Hence $G_u \cap G_{\mathfrak{w}}$ 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 $G \curvearrowright X$ is proper, the action $G_u \cap G_{\mathfrak{w}} \curvearrowright B$ is also proper.

It is clear that $\partial_\infty B = \partial_\infty S \subseteq \partial_\infty u$. The same argument applied to the other shore of B shows that $\partial_\infty B \subseteq \partial_\infty u \cap \partial_\infty \mathfrak{w}$. Now, consider a point $\xi \in \partial_\infty u \setminus \partial_\infty S$ and a geodesic $\gamma \subseteq C(u)$ representing ξ ; note that the function $t \mapsto d(\gamma(t), S)$ must diverge with t . As soon as $\gamma(t) \notin S$, equation (*) in Sect. 2.1 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 $\partial_\infty u \cap \partial_\infty \mathfrak{w} \subseteq \partial_\infty S = \partial_\infty B$, concluding the proof. \square

We now introduce a preorder on $\mathcal{W}(X)$. We write $u_1 \preceq u_2$ if $\partial_\infty u_1 \subseteq \partial_\infty u_2$ and $u_1 \sim u_2$ if $\partial_\infty u_1 = \partial_\infty u_2$. The latter is an equivalence relation.

Recall that $\mathcal{W}(\mathfrak{w})$ denotes the set of hyperplanes transverse to $\mathfrak{w} \in \mathcal{W}(X)$.

Proposition 4.9 *Let X be hyperplane-essential.*

- (1) *We have $u_1 \preceq u_2$ if and only if $\mathcal{W}(u_1) \subseteq \mathcal{W}(u_2)$. Equivalently, u_1 and u_2 are not transverse and u_1 stays at bounded distance from u_2 .*
- (2) *If there exists \mathfrak{w} with $u_1 \preceq \mathfrak{w}$ and $u_2 \preceq \mathfrak{w}$, the hyperplanes u_1 and u_2 are not transverse.*
- (3) *There exists $N = N(X) \geq 0$ such that, given any hyperplane u with $\partial_\infty u \neq \emptyset$, there exist at most N hyperplanes \mathfrak{w} with $u \preceq \mathfrak{w}$.*

Proof We start by proving part (1). Let B denote the bridge $B(u_1, u_2)$ and let S_1 be the shore $B \cap C(u_1)$. If $\mathcal{W}(u_1) \subseteq \mathcal{W}(u_2)$, we have $u_2 \notin \mathcal{W}(u_1)$ and the carrier $C(u_1)$ is at finite Hausdorff distance from B . Thus, u_1 is contained in a metric neighbourhood of u_2 and this implies that $u_1 \preceq u_2$.

Conversely, suppose that $u_1 \preceq u_2$. Proposition 4.8 implies $\partial_\infty u_1 = \partial_\infty S_1$. Since S_1 is a convex subcomplex of $C(u_1)$, we will also denote by S_1 its projection to a convex subcomplex of the CAT(0) cube complex u_1 . Note that $S_1 = u_1$, or there would exist a halfspace \mathfrak{k} of u_1 with $S_1 \subseteq \mathfrak{k}$. In this case, applying part (2) of Lemma 4.4 to u_1 , we would obtain a point of $\partial_\infty \mathfrak{k}^* \setminus \partial_\infty \mathfrak{k} \subseteq \partial_\infty u_1 \setminus \partial_\infty S_1$, a contradiction. The fact that $S_1 = u_1$ yields $\mathscr{W}(u_1) = \mathscr{W}(S_1) = \mathscr{W}(u_1) \cap \mathscr{W}(u_2) \subseteq \mathscr{W}(u_2)$, completing part (1).

If u_1 and u_2 are transverse and $u_1 \preceq w$, it follows from part (1) that u_2 and w are transverse. In particular $u_2 \not\preceq w$ 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 w_1, w_2 with $\partial_\infty B(w_1, w_2) \neq \emptyset$, we have $d(C(w_1), C(w_2)) \leq D$. If $u \preceq w$, Proposition 4.8 ensures that $\partial_\infty B(u, w) = \partial_\infty u \neq \emptyset$, hence $d(C(u), C(w)) \leq D$. Part (1) now shows that $C(u)$ is contained in the $(D + 1)$ -neighbourhood of $C(w)$. Since X admits a proper cocompact action, it is uniformly locally finite and there exist only uniformly finitely many such hyperplanes w . □

Definition 4.10 Given a hyperplane w , a point $\xi \in \partial_\infty w$ is *generic* if the only hyperplanes u with $\xi \in \partial_\infty u$ are those that satisfy $w \preceq u$.

Proposition 4.11 *For every $w \in \mathscr{W}(X)$, the subset of generic points is dense in $\partial_\infty w$. In particular, generic points exist as soon as $\partial_\infty w \neq \emptyset$.*

Proof Let \mathcal{B} be the family of subsets of $\partial_\infty w$ of the form $\partial_\infty w \cap \partial_\infty u$, with $u \in \mathscr{W}(X)$ and $w \not\preceq u$. By definition, a point $\xi \in \partial_\infty w$ is generic if and only if it does not lie in the union of the elements of \mathcal{B} . Since $\mathscr{W}(X)$ is countable, so is \mathcal{B} . The proposition then follows from Baire’s theorem, if we show that every $\mathfrak{B} \in \mathcal{B}$ is nowhere-dense in $\partial_\infty w$.

If $\mathfrak{B} = \partial_\infty w \cap \partial_\infty u$, set $H = G_w \cap G_u$. Proposition 4.8 shows that the homeomorphism $\sigma_X: \partial_\infty G \rightarrow \partial_\infty X$ takes $\partial_\infty H \subseteq \partial_\infty G_w$ to $\mathfrak{B} \subseteq \partial_\infty w$. Since $w \not\preceq u$, the difference $\partial_\infty w \setminus \mathfrak{B}$ is nonempty and H must have infinite index in G_w . As H is quasi-convex, we conclude via Lemma 4.1. □

Remark 4.12 Assume that X is hyperplane-essential and consider a generic point $\xi \in \partial_\infty w$. Viewing w itself as a CAT(0) cube complex, we have $\xi \in \partial_{nt} w$. This follows from Lemma 4.7 since, if $u \in \mathscr{W}(w)$, part (1) of Proposition 4.9 shows that $w \not\preceq u$ and hence $\xi \notin \partial_\infty u$.

This is a good point to make the following simple observation, which will only be needed in the proof of Proposition 4.28 later on.

Lemma 4.13 *Let X be essential and suppose that $G = \langle g \rangle \simeq \mathbb{Z}$. Then:*

- (1) every hyperplane of X is compact;
- (2) for every $\mathfrak{h} \in \mathcal{H}(X)$, 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 4.4 shows that $\partial_\infty w = \emptyset$ for every $w \in \mathscr{W}(X)$. Lemma 2.5 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. \square

4.2 Towards hyperplane recognition

This subsection is devoted to Propositions 4.14 and 4.18. 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).

Standing Assumptions Let again G be Gromov hyperbolic and let the action $G \curvearrowright X$ be proper and cocompact. Let $\rho_X: \partial_\infty G \rightarrow \partial_\infty X$ denote the only G -equivariant homeomorphism.

Proposition 4.14 *Consider two distinct points $\xi, \eta \in \partial_\infty X$ and four sequences $x_n, y_n, z_n, w_n \in \partial_{nt} X$, 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 hyperplane lying infinitely often in $\mathscr{W}(x_n, y_n | z_n, w_n)$ 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 $\partial_\infty \mathfrak{w}$. \square

Lemma 4.15 *For an infinite-order element $k \in G$ and a hyperplane \mathfrak{w} , the following are equivalent:*

- (1) *a non-trivial power of k preserves \mathfrak{w} ;*
- (2) *the points $\rho_X(k^{\pm\infty})$ 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 $k^n \in G_{\mathfrak{w}}$ for some $n > 0$, we have $k^{\pm\infty} = (k^n)^{\pm\infty} \in \partial_\infty G_{\mathfrak{w}}$. \square

Definition 4.16 If $k \in G$ is infinite-order, we denote by $\mathscr{W}(k)$ the set of hyperplanes satisfying the equivalent conditions in Lemma 4.15. We say that k is *good* if it preserves every halfspace bounded by an element of $\mathscr{W}(k)$.

Remark 4.17 Every infinite-order element has a good power, as the set $\mathscr{W}(k)$ is always finite by Lemma 4.15.

We are interested in good elements because of the following result.

Proposition 4.18 Consider distinct points $x, y \in \partial_{\text{nt}} X$ and a good infinite-order element $k \in G$ 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 $n \rightarrow +\infty$. Proposition 4.14 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 $\mathscr{W}_n \cap \mathscr{W}(k) = \mathscr{W}(x|y) \cap \mathscr{W}(k)$ and this concludes the proof. \square

4.3 Trust issues

Throughout this subsection:

Standing Assumptions Let the Gromov hyperbolic group G act properly and cocompactly on CAT(0) cube complexes X and Y . We fix two subsets $\mathcal{A} \subseteq \partial_{\text{nt}} X$ and $\mathcal{B} \subseteq \partial_{\text{nt}} Y$.

Given $\mathcal{U} \subseteq \mathscr{W}(X)$, we employ the notation $\mathcal{U}(A|B) = \mathscr{W}(A|B) \cap \mathcal{U}$ and

$$\text{cr}_{\mathcal{U}}(x, y, z, w) = \#\mathcal{U}(x, z|y, w) - \#\mathcal{U}(x, w|y, z)$$

for all subsets $A, B \subseteq \overline{X}$ and pairwise distinct points $x, y, z, w \in \partial_{\text{nt}} X$. Given $\mathcal{V} \subseteq \mathscr{W}(Y)$, the same notation applies to subsets of \overline{Y} and points of $\partial_{\text{nt}} Y$.

As mentioned in the introduction, a key complication in the proof of Theorem C is the fact that $\text{cr}_{\mathcal{U}}(x, y, z, w)$ 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 4.19 We say that a 4-tuple $(a, b, c, d) \in (\partial_{\text{nt}} X)^4$ 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. If $\mathcal{U} = \mathcal{W}(X)$, we just say that (a, b, c, d) is trustworthy.

Lemma 4.20 Given subsets $\mathcal{U} \subseteq \mathcal{W}(X)$ and $\mathcal{V} \subseteq \mathcal{W}(Y)$, consider pairwise distinct points $x_1, x_2, x_3, x_4 \in A$ and a bijection $f: A \rightarrow B$ satisfying:

$$\text{cr}_{\mathcal{U}}(x_{\sigma(1)}, x_{\sigma(2)}, x_{\sigma(3)}, x_{\sigma(4)}) = \text{cr}_{\mathcal{V}}(f(x_{\sigma(1)}), f(x_{\sigma(2)}), f(x_{\sigma(3)}), f(x_{\sigma(4)}))$$

for every permutation $\sigma \in \mathfrak{S}_4$. If the 4-tuples $(f(x_1), f(x_2), f(x_3), f(x_4))$ and (x_1, x_2, x_3, x_4) are, respectively, \mathcal{V} -trustworthy and \mathcal{U} -trustworthy, then, for every $\sigma \in \mathfrak{S}_4$, we have:

$$\#\mathcal{U}(x_{\sigma(1)}, x_{\sigma(2)}|x_{\sigma(3)}, x_{\sigma(4)}) = \#\mathcal{V}(f(x_{\sigma(1)}), f(x_{\sigma(2)})|f(x_{\sigma(3)}), f(x_{\sigma(4)})).$$

Proof Since (x_1, x_2, x_3, x_4) is \mathcal{U} -trustworthy, we can permute the four points so that $\mathcal{U}(x_1, x_2|x_3, x_4) = \emptyset$. This implies that $\text{cr}_{\mathcal{U}}(x_1, x_3, x_4, x_2) \geq 0$ and $\text{cr}_{\mathcal{U}}(x_1, x_4, x_3, x_2) \geq 0$. We then have $\text{cr}_{\mathcal{V}}(f(x_1), f(x_3), f(x_4), f(x_2)) \geq 0$ and $\text{cr}_{\mathcal{V}}(f(x_1), f(x_4), f(x_3), f(x_2)) \geq 0$, 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 minimum between $\#\mathcal{V}(f(x_1), f(x_3)|f(x_2), f(x_4))$ 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:

$$\begin{aligned} \#\mathcal{U}(x_1, x_3|x_2, x_4) &= \text{cr}_{\mathcal{U}}(x_1, x_4, x_3, x_2) \\ &= \text{cr}_{\mathcal{V}}(f(x_1), f(x_4), f(x_3), f(x_2)) \\ &= \#\mathcal{V}(f(x_1), f(x_3)|f(x_2), f(x_4)). \end{aligned}$$

The equality between $\#\mathcal{U}(x_1, x_4|x_2, x_3)$ and $\#\mathcal{V}(f(x_1), f(x_4)|f(x_2), f(x_3))$ is obtained similarly. □

The following will be our main source of trustworthy 4-tuples.

Definition 4.21 Let ξ and η be distinct points of $\partial_{\infty} X$. We say that the pair (ξ, η) is *trustworthy* if there do not exist transverse hyperplanes u_1 and u_2 with $\xi, \eta \in \partial_{\infty} u_1 \cap \partial_{\infty} u_2$.

Lemma 4.22 Consider two distinct points $\xi, \eta \in \partial_{\infty} X$ and sequences $a_n, b_n, c_n, d_n \in \partial_{\text{nt}} X$, 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 $\mathcal{W}(a_n, b_n|c_n, d_n)$ and $\mathcal{W}(a_n, d_n|b_n, c_n)$ are all nonempty. Proposition 4.14 provides N such that, for all $n \geq N$, every element of $\mathcal{W}(a_n, b_n|c_n, d_n) \cup \mathcal{W}(a_n, d_n|b_n, c_n)$ contains ξ and η in its trace at infinity. Lemma 2.1, however, shows that the sets $\mathcal{W}(a_n, b_n|c_n, d_n)$ and $\mathcal{W}(a_n, d_n|b_n, c_n)$ are transverse, contradicting the fact that (ξ, η) is trustworthy. □

The next result applies e.g. to the case when no three elements of \mathcal{U} are transverse (Lemma 2.1). It will only be needed in the proof of Theorem 4.33.

Lemma 4.23 *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 $\bar{x}, \bar{y} \in \mathcal{P}$ and $\bar{z}, \bar{w} \in \mathcal{Q}$ such that $\mathcal{U}(\bar{x}, \bar{y}|\bar{z}, \bar{w}) = \mathcal{U}(\mathcal{P}|\mathcal{Q})$.*

Proof Pick points $\bar{x}, \bar{y} \in \mathcal{P}$ and $\bar{z}, \bar{w} \in \mathcal{Q}$ so as to minimise the cardinality of $\mathcal{U}(\bar{x}, \bar{y}|\bar{z}, \bar{w})$; since $\#\mathcal{P}, \#\mathcal{Q} \geq 2$, this set is finite. It is clear that $\mathcal{U}(\bar{x}, \bar{y}|\bar{z}, \bar{w})$ contains $\mathcal{U}(\mathcal{P}|\mathcal{Q})$. Consider any hyperplane $\mathfrak{w} \in \mathcal{U}(\bar{x}, \bar{y}|\bar{z}, \bar{w})$ and let \mathfrak{h} be its side containing \bar{x} and \bar{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 $u_1 \in \mathcal{U}(\bar{x}, \bar{y}|u, \bar{z})$ and $u_2 \in \mathcal{U}(\bar{x}, \bar{y}|u, \bar{w})$. As $\mathcal{U}(\bar{x}, \bar{y}|u, \bar{z})$ and $\mathcal{U}(\bar{x}, \bar{y}|u, \bar{w})$ do not contain \mathfrak{w} and cannot have fewer elements than $\mathcal{U}(\bar{x}, \bar{y}|\bar{z}, \bar{w})$, we must be able to choose u_1 and u_2 outside $\mathcal{U}(\bar{x}, \bar{y}|\bar{z}, \bar{w})$. In conclusion:

$$\begin{aligned} \mathfrak{w} &\in \mathcal{U}(\bar{x}, \bar{y}, u|\bar{z}, \bar{w}) \subseteq \mathcal{U}(\bar{y}, u|\bar{z}, \bar{w}), \\ u_1 &\in \mathcal{U}(\bar{x}, \bar{y}, \bar{w}|u, \bar{z}) \subseteq \mathcal{U}(\bar{y}, \bar{w}|u, \bar{z}), \\ u_2 &\in \mathcal{U}(\bar{x}, \bar{y}, \bar{z}|u, \bar{w}) \subseteq \mathcal{U}(\bar{y}, \bar{z}|u, \bar{w}), \end{aligned}$$

which violates the assumption that every 4-tuple, in particular $(u, \bar{y}, \bar{z}, \bar{w})$, be \mathcal{U} -trustworthy. □

4.4 Traces vs cross ratios

Throughout this subsection:

Standing Assumptions We now assume that the hyperbolic group G is non-elementary. We consider proper cocompact actions of G on essential, hyperplane-essential CAT(0) cube complexes X and Y . These are irreducible by Remark 2.4.

Let $\mathcal{A} \subseteq \partial_{\text{nt}}X$ and $\mathcal{B} \subseteq \partial_{\text{nt}}Y$ be nonempty G -invariant subsets with $f(\mathcal{A}) = \mathcal{B}$, where f is the homeomorphism $o_Y \circ o_X^{-1}: \partial_{\infty}X \rightarrow \partial_{\infty}Y$.

Throughout the subsection, we also fix subsets $\mathcal{U} \subseteq \mathscr{W}(X)$ and $\mathcal{V} \subseteq \mathscr{W}(Y)$ such that:

$$\text{cr}_{\mathcal{U}}(x, y, z, w) = \text{cr}_{\mathcal{V}}(f(x), f(y), f(z), f(w))$$

for all pairwise distinct points $x, y, z, w \in \mathcal{A}$. It will become clear in Sect. 4.5 how this relates to the proof of Theorem C.

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 Sect. 4.5.

Lemma 4.24 *Given a hyperplane $u \in \mathcal{U}$, a generic point $\xi \in \partial_\infty u$ and an arbitrary point $\eta \in \partial_\infty u \setminus \{\xi\}$, there exists a hyperplane $v \in \mathcal{V}$ with $f(\xi), f(\eta) \in \partial_\infty v$.*

Proof Consider a (combinatorial) line $\gamma \subseteq C(u)$ with endpoints ξ and η . Pick $w_0 \in \mathcal{W}(\gamma)$ and label the other elements of $\mathcal{W}(\gamma)$ as $w_n, n \in \mathbb{Z}$, 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 h be any side of the hyperplane u and let h_n be the side of w_n that contains the positive half of γ . For all $n > 0$, part (2) of Proposition 2.7 allows us to pick points of \mathcal{A} as follows: $x_n \in h \cap h_n, y_n \in h \cap h_{-n}^*, z_n \in h^* \cap h_n$ and $w_n \in h^* \cap h_{-n}^*$.

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 3.6 and Corollary 3.12 show that $x_n \in \partial_{nt} X$ converge to ξ . Similarly, z_n converge to ξ , while y_n and w_n converge to η .

By Remark 4.12 and Lemma 3.15, the generic point $\xi \in \partial_\infty u$ is represented by a regular point in the Roller boundary of the CAT(0) cube complex u . For all sufficiently large $n > 0$, it follows that $w_n \cap u$ and $w_{-n} \cap u$ are strongly separated as hyperplanes of the cube complex u . In other words, no hyperplane of X is transverse to $u \in \mathcal{U}(x_n, y_n|z_n, w_n)$ and $w_n, w_{-n} \in \mathcal{W}(x_n, z_n|y_n, w_n)$ at the same time. By Lemma 2.1, the sets $\mathcal{W}(x_n, w_n|y_n, z_n)$ are then all empty for large $n > 0$. It follows that $cr_{\mathcal{U}}(x_n, z_n, y_n, w_n) > 0$, hence $cr_{\mathcal{V}}(f(x_n), f(z_n), f(y_n), f(w_n)) > 0$ and $\mathcal{V}(f(x_n), f(y_n)|f(z_n), f(w_n)) \neq \emptyset$. Since $f(x_n)$ and $f(z_n)$ converge to $f(\xi)$, while $f(y_n)$ and $f(w_n)$ converge to $f(\eta)$, Proposition 4.14 yields the required hyperplane v . □

Definition 4.25 A hyperplane $w \in \mathcal{U}$ is \mathcal{U} -boundary-maximal if $\partial_\infty w \neq \emptyset$ and every hyperplane $u \in \mathcal{U}$ with $w \leq u$ actually satisfies $u \sim w$. When $\mathcal{U} = \mathcal{W}(X)$, we simply speak of *boundary-maximal hyperplanes*.

Part (3) of Proposition 4.9 and Remark 4.6 show that boundary-maximal hyperplanes exist as soon as X is not a tree. In fact, for every hyperplane u there exists a boundary-maximal hyperplane w with $u \leq w$.

The next result can be viewed as Step (IIa) from the introduction.

Proposition 4.26 *If $u \in \mathcal{U}$ is \mathcal{U} -boundary-maximal, there exists a \mathcal{V} -boundary-maximal hyperplane $v \in \mathcal{V}$ with $\partial_\infty v = f(\partial_\infty u)$.*

Proof Proposition 4.11 allows us to pick a generic point $\xi \in \partial_\infty u$. Note that u is acted on properly and cocompactly by its stabiliser (Lemma 2.5), so $\partial_\infty u \setminus \{\xi\} \neq \emptyset$; in particular, Lemma 4.24 shows that there exists a hyperplane $v \in \mathcal{V}$ with $f(\xi) \in \partial_\infty v$. By part (3) of Proposition 4.9, we can take v to be \mathcal{V} -boundary-maximal. Let $\zeta \in \partial_\infty v \setminus \{f(\xi)\}$ be generic.

Since $f(\xi)$ and ζ both lie in $\partial_\infty v$, Lemma 4.24 shows that ξ and $f^{-1}(\zeta)$ lie in the trace at infinity of some $u' \in \mathcal{U}$. Since ξ is generic, we have $u \leq u'$ and, as u is \mathcal{U} -boundary-maximal, we conclude that $\partial_\infty u = \partial_\infty u'$. In particular $f^{-1}(\zeta) \in \partial_\infty u$, showing that the closed subset $f(\partial_\infty u)$ contains every generic point in $\partial_\infty v \setminus \{f(\xi)\}$. By Proposition 4.11, we have $\partial_\infty v \subseteq f(\partial_\infty u)$ 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 4.24 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 $f(\partial_\infty \mathfrak{u}) \subseteq \partial_\infty \mathfrak{v}$. \square

We now prepare for Step (IIb), which will be completed in Theorem 4.33. Given a hyperplane $\mathfrak{w} \in \mathcal{W}(X)$ with $\partial_\infty \mathfrak{w} \neq \emptyset$, we denote by $\mathcal{T}_U(\mathfrak{w})$ the set of hyperplanes $\mathfrak{u} \in \mathcal{U}$ satisfying $\mathfrak{u} \sim \mathfrak{w}$. We also write $\tilde{\mathcal{T}}_U(\mathfrak{w}) \subseteq \mathcal{H}(X)$ for the set of halfspaces bounded by elements of $\mathcal{T}_U(\mathfrak{w})$. By Proposition 4.9, the set $\mathcal{T}_U(\mathfrak{w})$ is finite and no two of its elements are transverse.

Consider for a moment a subset $\mathcal{W} \subseteq \mathcal{W}(X)$ and the collection $\tilde{\mathcal{W}} \subseteq \mathcal{H}(X)$ of halfspaces bounded by the elements of \mathcal{W} . Recall from [23, p. 860] that each such subset \mathcal{W} determines a quotient CAT(0) cube complex whose halfspace-pocset is isomorphic to $\tilde{\mathcal{W}}$. As customary, we will refer to this as the *restriction quotient* of X determined by \mathcal{W} .

Definition 4.27 The dual tree $T(\mathfrak{w})$ is the restriction quotient determined by the subset $\mathcal{T}_U(\mathfrak{w}) \subseteq \mathcal{U} \subseteq \mathcal{W}(X)$.

Note that, since $\mathcal{T}_U(\mathfrak{w})$ is finite and no two of its elements are transverse, the cube complex $T(\mathfrak{w})$ is a finite tree. Being a restriction quotient of X , it comes equipped with a projection $\pi_{\mathfrak{w}} : \bar{X} \rightarrow \overline{T(\mathfrak{w})} = T(\mathfrak{w})$. More precisely, $\pi_{\mathfrak{w}}$ takes the point $x \in \bar{X}$ to the point of $T(\mathfrak{w})$ represented by the ultrafilter $\sigma_x \cap \tilde{\mathcal{T}}_U(\mathfrak{w})$. We will be interested in the pseudo-metric $\delta_{\mathfrak{w}}$ defined on the set $\mathcal{A} \subseteq \partial_{\text{nt}} X$ by the formula

$$\delta_{\mathfrak{w}}(x, y) = d(\pi_{\mathfrak{w}}(x), \pi_{\mathfrak{w}}(y)) = \#(\mathcal{T}_U(\mathfrak{w}) \cap \mathcal{U}(x|y)).$$

The same construction can be applied to hyperplanes $\mathfrak{w} \in \mathcal{W}(Y)$, yielding a subset $\mathcal{T}_V(\mathfrak{w}) \subseteq \mathcal{V}$, a projection to a finite tree $\pi_{\mathfrak{w}} : \bar{Y} \rightarrow T(\mathfrak{w})$ and a pseudo-metric $\delta_{\mathfrak{w}}$ on $\mathcal{B} \subseteq \partial_{\text{nt}} Y$.

The next result is the main ingredient of Step (IIb) (cf. Theorem 4.33). We will employ Definitions 4.16 and 4.21 in its proof.

Proposition 4.28 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 4.2, the subgroups $G_{\mathfrak{u}}$ and $G_{\mathfrak{v}}$ 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 $\mathcal{T}_U(\mathfrak{u}) \subseteq \mathcal{U}(h)$ and $\mathcal{T}_V(\mathfrak{v}) \subseteq \mathcal{V}(h)$ by Lemma 4.15.

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 we have $\mathcal{T}_U(\mathfrak{u}) \cap \mathcal{U}(x|y) = \mathcal{U}(h_0) \cap \mathcal{U}(x|y)$ and $\mathcal{T}_V(\mathfrak{v}) \cap \mathcal{V}(f(x)|f(y)) = \mathcal{V}(h_0) \cap \mathcal{V}(f(x)|f(y))$.*

The claim concludes the proof as follows. Lemma 4.22 shows that the 4-tuples $(h_0^n x, h_0^{-n} x | h_0^n y, h_0^{-n} y)$ are trustworthy for large $n > 0$; the same is true of $(h_0^n f(x), h_0^{-n} f(x) | h_0^n f(y), h_0^{-n} f(y))$. Equivariance of f and Lemma 4.20 then imply that the set $\mathcal{U}(h_0^n x, h_0^{-n} x | h_0^n y, h_0^{-n} y)$ has the same cardinality as $\mathcal{V}(h_0^n f(x), h_0^{-n} f(x) | h_0^n f(y), h_0^{-n} f(y))$ for all large $n > 0$.

Since $x, y \in \partial_{\text{nt}} X$, Lemma 4.7 shows that $x, y \notin \partial_\infty u$. On the other hand h_0 preserves u and its two fixed points in $\partial_\infty X$ lie within $\partial_\infty u$ by Lemma 4.15. We conclude that h_0 does not fix x, y and Proposition 4.18 finally yields $\#(\mathcal{U}(h_0) \cap \mathcal{U}(x|y)) = \#(\mathcal{V}(h_0) \cap \mathcal{V}(f(x)|f(y)))$. By our choice of h_0 , these two cardinalities are precisely $\delta_u(x, y)$ and $\delta_v(f(x), f(y))$.

Proof of claim. The points x, y are regular by Lemma 3.15. It follows that, among hyperplanes of X separating x and y , only finitely many are not strongly separated from u . In light of Proposition 4.8, there are only finitely many hyperplanes u' that separate x and y and satisfy $\partial_\infty u' \cap \partial_\infty u \neq \emptyset$; let us denote by u_1, \dots, u_r those that lie in $\mathcal{U} \setminus \mathcal{T}_U(u)$. Similarly let v_1, \dots, v_s be the elements of $\mathcal{V} \setminus \mathcal{T}_V(v)$ that separate $f(x)$ and $f(y)$ and satisfy $\partial_\infty v_i \cap \partial_\infty v \neq \emptyset$.

By boundary-maximality of u and v , we have $\partial_\infty u_i \cap \partial_\infty u \subsetneq \partial_\infty u$ and $\partial_\infty v_j \cap \partial_\infty v \subsetneq \partial_\infty v$ for all i and j . Lemma 4.1 and Proposition 4.8 then show that each of these subsets is nowhere-dense. Hence, the union $K \subseteq \partial_\infty H$ of all sets $o_X^{-1}(\partial_\infty u_i \cap \partial_\infty u)$ and $o_Y^{-1}(\partial_\infty v_j \cap \partial_\infty v)$ is nowhere-dense.

Observe that there exists $h \in H$ that acts as a neatly contracting automorphism on both u and v . This follows from Proposition 3.18 if H is not virtually cyclic and from Lemma 4.13 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 $o_X(h_0^{+\infty}) \notin \partial_\infty u_i$ and $o_Y(h_0^{+\infty}) \notin \partial_\infty v_j$ for all i and j , which implies that the inclusions $\mathcal{T}_U(u) \cap \mathcal{U}(x|y) \subseteq \mathcal{U}(h_0) \cap \mathcal{U}(x|y)$ and $\mathcal{T}_V(v) \cap \mathcal{V}(f(x)|f(y)) \subseteq \mathcal{V}(h_0) \cap \mathcal{V}(f(x)|f(y))$ are equalities.

Since h_0 acts as a neatly contracting automorphism on u , we have $o_X(h_0^{\pm\infty}) \in \partial_{\text{nt}} u$. Lemma 4.7 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 4.17 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 4.28, we get the commutative diagram:

$$\begin{array}{ccccc}
 \mathcal{A} & \xrightarrow{\pi_u} & \pi_u(\mathcal{A}) & \hookrightarrow & T(u) \\
 \downarrow f & & \downarrow \psi & & \downarrow \Psi \\
 \mathcal{B} & \xrightarrow{\pi_v} & \pi_v(\mathcal{B}) & \hookrightarrow & T(v).
 \end{array}$$

Here $\pi_u(\mathcal{A})$ and $\pi_v(\mathcal{B})$ are precisely the quotient metric spaces associated to the pseudo-metric spaces (\mathcal{A}, δ_u) and (\mathcal{B}, δ_v) . The map ψ is then provided by Proposition 4.28 and it is an isometry.

By part (1) of Proposition 2.7, we see that $\pi_u(\mathcal{A})$ and $\pi_v(\mathcal{B})$ contain all vertices of degree one in $T(u)$ and $T(v)$, respectively. The dashed arrow Ψ is then obtained through the following general fact about trees.

Lemma 4.29 *Let T_1 and T_2 be finite trees with all edges of length one¹³. Let $V_i \subseteq T_i$ be sets of vertices containing all degree-one vertices of T_i . Every distance preserving bijection $\psi: V_1 \rightarrow V_2$ uniquely extends to an isometry $\Psi: T_1 \rightarrow T_2$.*

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. Observe that a point $x \in V_i$ lies outside V'_i precisely when we can find $y, z \in V_i \setminus \{x\}$ with $x = m(x, y, z)$. Since ψ preserves distances, it follows that $\psi(V'_1) = V'_2$.

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 $\phi_i: V'_i \rightarrow \partial T_i$. Observing that the maps ϕ_i preserve cross ratios and ψ is an isometry, we conclude that the bijection $\phi_2 \psi \phi_1^{-1}: \partial T_1 \rightarrow \partial T_2$ preserves cross ratios. It then follows from Theorem 4.3 in [13] that $\phi_2 \psi \phi_1^{-1}$ admits a unique extension to an isometry $\Psi: T_1 \rightarrow T_2$.

Note that $m(x, y, z) = m(\phi_i(x), \phi_i(y), \phi_i(z))$ for all pairwise 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 $m_1 := m(x, y, z)$ 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

$$\begin{aligned} d(m_2, \Psi(x)) &= d(m_1, x) = \frac{1}{2} \cdot [d(x, y) + d(x, z) - d(y, z)] \\ &= \frac{1}{2} \cdot [d(\psi(x), \psi(y)) + d(\psi(x), \psi(z)) - d(\psi(y), \psi(z))] \\ &= d(m_2, \psi(x)), \end{aligned}$$

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 4.30 *Let u and v be as in Proposition 4.28. For all $\mathfrak{h} \in \mathcal{H}(X)$:*

$$\#\{j \in \widetilde{\mathcal{T}}_U(u) \mid j \cap \mathcal{A} = \mathfrak{h} \cap \mathcal{A}\} = \#\{m \in \widetilde{\mathcal{T}}_V(v) \mid m \cap \mathcal{B} = f(\mathfrak{h} \cap \mathcal{A})\}.$$

Proof We begin by observing that it suffices to prove the inequality

$$\#\{j \in \widetilde{\mathcal{T}}_U(u) \mid j \cap \mathcal{A} = \mathfrak{h} \cap \mathcal{A}\} \leq \#\{m \in \widetilde{\mathcal{T}}_V(v) \mid m \cap \mathcal{B} = f(\mathfrak{h} \cap \mathcal{A})\}. \quad (*)$$

This already yields an equality if the right-hand side vanishes. Otherwise, there exists $m' \in \widetilde{\mathcal{T}}_V(v)$ with $m' \cap \mathcal{B} = f(\mathfrak{h} \cap \mathcal{A})$ and we can apply the same argument to f^{-1} and m' to obtain the opposite inequality.

Now, if the left-hand side of (*) vanishes, there is nothing to prove. Otherwise, there exists a halfspace $j' \in \widetilde{\mathcal{T}}_U(u)$ with $j' \cap \mathcal{A} = \mathfrak{h} \cap \mathcal{A}$ and our setup is not affected if we replace \mathfrak{h} with j' . We can thus assume that $\mathfrak{h} \in \widetilde{\mathcal{T}}_U(u)$.

The projection $\pi_u(\mathfrak{h})$ is a halfspace of the tree $T(u)$, with complement $\pi_u(\mathfrak{h}^*)$. Let C and C^* be the convex hulls, respectively, of $\pi_u(\mathfrak{h} \cap \mathcal{A})$ and $\pi_u(\mathfrak{h}^* \cap \mathcal{A})$. These

¹³ When proving Theorem C for cuboid complexes, one should allow edges of arbitrary length in T_1 and T_2 . The proof of the lemma does not change.

are disjoint subtrees of $T(u)$ and the union $C \sqcup C^*$ contains all degree-one vertices of $T(u)$. The complement of $C \sqcup C^*$ must be an open arc $\alpha \subseteq T(u)$ such that every vertex in α has degree two and lies outside $\pi_u(\mathcal{A})$. Thus, the hyperplanes associated to the set $\{j \in \widetilde{\mathcal{T}}_{\mathcal{U}}(u) \mid j \cap \mathcal{A} = \mathfrak{h} \cap \mathcal{A}\}$ are precisely the elements of $\mathcal{T}_{\mathcal{U}}(u)$ that are dual to edges of α .

Observing that $\Psi \circ \pi_u = \pi_v \circ f$, we see that the sets $\Psi(C)$ and $\Psi(C^*)$ are the convex hulls, respectively, of $\pi_v(f(\mathfrak{h} \cap \mathcal{A}))$ and $\pi_v(f(\mathfrak{h}^* \cap \mathcal{A}))$ in $T(v)$. Moreover, the arc $\Psi(\alpha)$ has the same length of α and its edges correspond to pairwise distinct elements of the set $\{m \in \widetilde{\mathcal{T}}_{\mathcal{V}}(v) \mid m \cap \mathcal{B} = f(\mathfrak{h} \cap \mathcal{A})\}$. This yields the desired inequality. □

4.5 Concluding the proof

In this subsection:

Standing Assumptions Let the group G , the cube complexes X and Y and the map f be as in the statement of Theorem C. We set $\mathcal{A} = \Omega$ and $\mathcal{B} = f(\Omega)$, so that we are in the general setup of Section 4.4. Here, however, we do not fix sets \mathcal{U} and \mathcal{V} . Instead, we observe that, by the hypotheses of Theorem C, we have $\text{cr}_X(x, y, z, w) = \text{cr}_Y(f(x), f(y), f(z), f(w))$ for all pairwise distinct points $x, y, z, w \in \mathcal{A}$.

Given subsets $\mathcal{U} \subseteq \mathcal{W}(X)$ and $\mathcal{V} \subseteq \mathcal{W}(Y)$, we denote by $\widetilde{\mathcal{U}} \subseteq \mathcal{H}(X)$ and $\widetilde{\mathcal{V}} \subseteq \mathcal{H}(Y)$ the collections of halfspaces bounded, respectively, by the elements of \mathcal{U} and \mathcal{V} .

Definition 4.31 We say that subsets $\mathcal{U} \subseteq \mathcal{W}(X)$ and $\mathcal{V} \subseteq \mathcal{W}(Y)$ are *well-paired* if, for every $\mathfrak{h} \in \mathcal{H}(X)$ and $\mathfrak{k} \in \mathcal{H}(Y)$, we have:

$$\begin{aligned} \#\{j \in \widetilde{\mathcal{U}} \mid j \cap \mathcal{A} = \mathfrak{h} \cap \mathcal{A}\} &= \#\{m \in \widetilde{\mathcal{V}} \mid m \cap \mathcal{B} = f(\mathfrak{h} \cap \mathcal{A})\}, \\ \#\{m \in \widetilde{\mathcal{V}} \mid m \cap \mathcal{B} = \mathfrak{k} \cap \mathcal{B}\} &= \#\{j \in \widetilde{\mathcal{U}} \mid j \cap \mathcal{A} = f^{-1}(\mathfrak{k} \cap \mathcal{B})\}. \end{aligned}$$

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 4.32 *If the sets $\mathcal{U} \subseteq \mathcal{W}(X)$ and $\mathcal{V} \subseteq \mathcal{W}(Y)$ are well-paired and the points $x, y, z, w \in \mathcal{A}$ are pairwise distinct, we have:*

$$\begin{aligned} \#\mathcal{U}(x, y|z, w) &= \#\mathcal{V}(f(x), f(y)|f(z), f(w)), \\ \text{cr}_{\mathcal{U}}(x, y, z, w) &= \text{cr}_{\mathcal{V}}(f(x), f(y), f(z), f(w)). \end{aligned}$$

The next result provides the correct formulation of Step (IIb) from the introduction, as traces at infinity need to be counted “with multiplicity”.

Theorem 4.33 *The sets $\mathcal{W}(X)$ and $\mathcal{W}(Y)$ 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 4.32 yields

$$cr_{\mathcal{U}_k^c}(x, y, z, w) = cr_{\mathcal{V}_k^c}(f(x), f(y), f(z), f(w)),$$

whenever $x, y, z, w \in \mathcal{A}$ are pairwise distinct. On the other hand, observing that $cr_X = cr_{\mathcal{U}_k} + cr_{\mathcal{U}_k^c}$ and $cr_Y = cr_{\mathcal{V}_k} + cr_{\mathcal{V}_k^c}$, the fact that f takes cr_X to cr_Y implies that

$$cr_{\mathcal{U}_k}(x, y, z, w) = cr_{\mathcal{V}_k}(f(x), f(y), f(z), f(w)).$$

We are now in the setting of Sect. 4.4 and can apply Proposition 4.26. It follows that, for each $u \in \mathcal{U}_k \setminus \mathcal{U}_{k+1}$, there exists $v \in \mathcal{V}_k \setminus \mathcal{V}_{k+1}$ with $\partial_\infty v = f(\partial_\infty u)$. Applying Corollary 4.30 to both f and f^{-1} , we see that the sets $\mathcal{T}_{\mathcal{U}_k}(u)$ and $\mathcal{T}_{\mathcal{V}_k}(v)$ are well-paired. Letting 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 4.9 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

$$cr_{\overline{\mathcal{U}}}(x, y, z, w) = cr_{\overline{\mathcal{V}}}(f(x), f(y), f(z), f(w))$$

for all pairwise distinct points $x, y, z, w \in \mathcal{A}$. By Remark 4.6, no two elements of $\overline{\mathcal{U}}$ or $\overline{\mathcal{V}}$ are transverse; in particular, Lemma 2.1 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 4.20 guarantees that

$$\#\overline{\mathcal{U}}(x, y|z, w) = \#\overline{\mathcal{V}}(f(x), f(y)|f(z), f(w)).$$

Thus, if $\mathfrak{h} \in \mathcal{H}(X)$ 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 4.23 then provides $\bar{x}, \bar{y} \in f(\mathfrak{h} \cap \mathcal{A})$ and $\bar{z}, \bar{w} \in f(\mathfrak{h}^* \cap \mathcal{A})$ with

$$\begin{aligned} \#\overline{\mathcal{V}}(f(\mathfrak{h} \cap \mathcal{A})|f(\mathfrak{h}^* \cap \mathcal{A})) &= \#\overline{\mathcal{V}}(\bar{x}, \bar{y}|\bar{z}, \bar{w}) \\ &= \#\overline{\mathcal{U}}(f^{-1}(\bar{x}), f^{-1}(\bar{y})|f^{-1}(\bar{z}), f^{-1}(\bar{w})) \geq \#\overline{\mathcal{U}}(\mathfrak{h} \cap \mathcal{A}|\mathfrak{h}^* \cap \mathcal{A}) > 0. \end{aligned}$$

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. \square

Proof of Theorem C We are going to show that $f: \partial_\infty X \rightarrow \partial_\infty Y$ induces a G -equivariant pocset isomorphism $f_*: (\mathcal{H}(X), \subseteq, *) \rightarrow (\mathcal{H}(Y), \subseteq, *)$. This will then yield a G -equivariant cubical isomorphism $F: X \rightarrow Y$ by Roller duality; see e.g. [20,52,63]. Uniqueness and the fact that F and f coincide on Ω will be clear from the construction.

We start by observing that the set $\mathcal{H}(\mathfrak{h}) := \{j \in \mathcal{H}(X) \mid j \cap \mathcal{A} = \mathfrak{h} \cap \mathcal{A}\}$ is totally ordered by inclusion for each $\mathfrak{h} \in \mathcal{H}(X)$. Indeed, given $j_1, j_2 \in \mathcal{H}(\mathfrak{h})$, the intersections $j_1 \cap j_2 \supseteq \mathfrak{h} \cap \mathcal{A}$ and $j_1^* \cap j_2^* \supseteq \mathfrak{h}^* \cap \mathcal{A}$ are nonempty by part (1) of Proposition 2.7. Moreover, j_1 and j_2 cannot be transverse, or part (2) of Proposition 2.7 would yield $\mathfrak{h} \cap \mathfrak{h}^* \supseteq j_1 \cap j_2^* \cap \mathcal{A} \neq \emptyset$. Hence $j_1 \subseteq j_2$ or $j_2 \subseteq j_1$.

Note that $\mathcal{H}(\mathfrak{h})$ is finite as, by part (2) of Lemma 3.6, the set $\mathcal{W}(x, y|z, w)$ is finite for all distinct points $x, y \in \mathfrak{h} \cap \mathcal{A}, z, w \in \mathfrak{h}^* \cap \mathcal{A}$. By Theorem 4.33, the set $f_*\mathcal{H}(\mathfrak{h}) := \{m \in \mathcal{H}(Y) \mid m \cap \mathcal{B} = f(\mathfrak{h} \cap \mathcal{A})\}$ 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 $\mathcal{H}(X)$, we have actually defined a map $f_*: \mathcal{H}(X) \rightarrow \mathcal{H}(Y)$. 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 $\mathfrak{h}, \mathfrak{k} \in \mathcal{H}(X)$ 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 2.7 finally yields $f_*(\mathfrak{h}) \subsetneq f_*(\mathfrak{k})$, concluding the proof. \square

Remark 4.34 When dealing with cuboid complexes \mathbb{X} and \mathbb{Y} , the proof of Theorem C needs to be slightly adapted. In general, the sets $\mathcal{H}(\mathfrak{h})$ and $f_*\mathcal{H}(\mathfrak{h})$ will have the same weight, but not the same cardinality. This prevents us from defining an isomorphism f_* between the halfspace pocsets of \mathbb{X}, \mathbb{Y} .

One should instead observe that CAT(0) cuboid complexes are median spaces and, thus, naturally endowed with a structure of space with measured walls [15]. The map $f: \partial_\infty \mathbb{X} \rightarrow \partial_\infty \mathbb{Y}$ then induces a G -equivariant isomorphism of their measured halfspace pocsets (see Sections 2.2 and 3.1 in [30]). In this context, an analogue of Roller duality is provided by Corollaries 3.12 and 3.13 in [30] and we obtain a G -equivariant isometry $F: \mathbb{X} \rightarrow \mathbb{Y}$. In general, F will not take vertices of \mathbb{X} to vertices of \mathbb{Y} .

5 Epilogue

In this section we apply Theorem C to obtain Theorem A and Corollary B. Relying on [4], we also prove Corollary E.

5.1 Cross ratios on contracting boundaries

By Remark 2.4, the following are common hypotheses to Theorem A and Corollary E.

Standing Assumptions Let G be a non-virtually-cyclic group. We consider proper cocompact actions of G on irreducible, essential CAT(0) cube complexes X and Y .

In Sect. 3.3, we defined a G -equivariant bijection $o_X : \partial_c G \rightarrow \partial_c X$ arising from orbit maps. We now want to exploit the map Φ introduced in Section 3.2 to transfer the cross ratio on ∂X to a G -invariant cross ratio on $\partial_c X$ and $\partial_c G$. To this end, we will need the following notion.

Definition 5.1 A subset $\mathcal{A} \subseteq \partial_{cu} X$ is a *section* (of the map Φ) if \mathcal{A} intersects each fibre of Φ at exactly one point. In particular, $\partial_{cnt} X \subseteq \mathcal{A}$.

Remark 5.2 If \mathcal{A} is a section of Φ and $x, y \in \mathcal{A}$ are distinct points, Lemma 3.6 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 $cr(x, y, z, w)$ is well-defined and finite (cf. Lemma 2.11).

Note that it is always possible to find a G -invariant section $\mathcal{A} \subseteq \partial_{cu} X$, up to subdividing X . Indeed, since every component of $\partial_{cu} X$ is finite by Remark 3.8, we can consider the first cubical subdivision X' and pick the median barycentre (cf. Sect. 2.3) of each component of $\partial_{cu} X'$. The resulting subset $\mathcal{A} \subseteq \partial_{cu} X'$ is a G -invariant section of $\Phi' : \partial_{cu} X' \rightarrow \partial_c X' \simeq \partial_c X$.

We stress that 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 $\mathcal{A} \simeq \partial_c G$ via the composition $o_X^{-1} \circ \Phi'$, we obtain an invariant cross ratio

$$cr_X : \partial_c G^{(4)} \rightarrow \mathbb{Z}.$$

By Remark 5.2, the value $cr_X(x, y, z, w)$ is well-defined and finite as soon as $x, y, z, w \in \partial_c G$ are pairwise distinct. We refer to cr_X as the *cubical cross ratio* on $\partial_c G$ associated to the cubulation $G \curvearrowright X$.

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 $cr_X(x, x, y, z) = 0$ and $cr_X(x, y, x, z) = +\infty$, while the other values of cr_X can be recovered using its symmetries.

We can endow $\partial_c G$ with the pull-back of the topology of $\partial_c^{vis} X$, or with one of the topologies of [19,21], but cr_X will rarely be continuous on its entire domain. Indeed, cr_X takes integer values, whereas $\partial_c G$ can be connected (for instance when G is a one-ended hyperbolic group).

Nevertheless, we have the following:

Proposition 5.3 *Endowing $\partial_c G$ with any of the three topologies above, cr_X is continuous at every 4-tuple with coordinates in $\partial_{cnt} X \subseteq \partial_c X$.*

Proof It suffices to consider $\partial_c G$ endowed with the pull-back of the topology of $\partial_c^{vis} X$, as this is the coarsest of the three (see Section 3.1 in [24] and Section 7 in [19]). In this case, the result follows from Proposition 2.12 and Remark 3.14. □

Proof of Theorem A Since X is essential and G acts cocompactly, every hyperplane-stabiliser has infinite index in G . Lemmas 4.1 and 2.5 thus imply that the union of the Gromov boundaries of the hyperplanes of X is meagre in $\partial_\infty X$. By Lemma 4.7, this means that $\mathcal{C} := \partial_{\text{nt}} X$ is co-meagre in $\partial_\infty X$. Part (1) now follows from Proposition 5.3.

Regarding part (2), suppose in addition that X is hyperplane-essential. Let $G \curvearrowright Y$ be another proper cocompact action on an essential, hyperplane-essential CAT(0) cube complex and consider the G -equivariant homeomorphism $f = o_Y \circ o_X^{-1} : \partial_\infty X \rightarrow \partial_\infty Y$. Suppose that cr_X and cr_Y coincide on $\mathcal{D}^{(4)} \subseteq \partial_\infty G^{(4)}$, for a co-meagre subset $\mathcal{D} \subseteq \partial_\infty G$. The set $\mathcal{D} \cap \partial_{\text{nt}} X \cap f^{-1}(\partial_{\text{nt}} Y)$ is co-meagre and so is the intersection of all its G -translates, which we denote by Ω . By Baire's theorem, Ω is nonempty and we conclude by applying Theorem C. \square

Proof of Corollary E Part (1) is immediate from the previous discussion, as cube complexes with no free faces are essential.

Let now $G \curvearrowright X$ and $G \curvearrowright Y$ be two cubulations satisfying the hypotheses of the theorem and inducing the same cubical cross ratio on $\partial_c G$. Proposition 3.18 yields an element $g \in G$ that acts as a neatly contracting automorphism on both X and Y . The G -equivariant bijection $f = o_Y \circ o_X^{-1} : \partial_c X \rightarrow \partial_c Y$ takes $g_X^+ \in \partial_{\text{cnt}} X$ to $g_Y^+ \in \partial_{\text{cnt}} Y$ by Lemma 3.19. Setting $\mathcal{A} = G \cdot g_X^+ \subseteq \partial_{\text{cnt}} X$, we have $f(\mathcal{A}) = G \cdot g_Y^+ \subseteq \partial_{\text{cnt}} Y$ and cross ratios of these points are preserved. The sets $\mathcal{A} \subseteq \partial_{\text{cnt}} X$ and $f(\mathcal{A}) \subseteq \partial_{\text{cnt}} Y$ consist of regular points by Lemma 3.15. We conclude by Theorem E in [4], observing that the actions of G on X and Y are non-elementary (in the sense of [4]) by Lemma 2.9 in [4]. \square

5.2 Marked length-spectrum rigidity

Proof of Corollary B If G is non-elementary, Theorem D in [4] provides a G -equivariant, cross-ratio preserving bijection $f : \mathcal{A} \rightarrow \mathcal{B}$, where $\mathcal{A} \subseteq \partial_{\text{reg}} X$ and $\mathcal{B} \subseteq \partial_{\text{reg}} Y$ are nonempty G -invariant subsets. The reader will not have trouble realising that this map is a restriction of the unique G -equivariant homeomorphism $f : \partial_\infty X \rightarrow \partial_\infty Y$ (see Section 4.2 in [4] for details). Regular points are non-terminating and we conclude by Theorem C.

We are left to consider the case when G is virtually cyclic. If G is finite, X and Y must be single points, by essentiality; so let us assume that G is virtually isomorphic to \mathbb{Z} . By part (1) of Lemma 4.13, every hyperplane of X is compact and, since now X is hyperplane-essential, we must have $X \simeq \mathbb{R}$. The action $G \curvearrowright X$ factors through a faithful action of either \mathbb{Z} or D_∞ . In the former case, the only $g \in G$ with $\ell_X(g) = 0$ are those in the (finite) kernel of the action $G \curvearrowright X$. In the latter case, we have infinitely many $g \in G$ with $\ell_X(g) = 0$, for instance all reflections.

Since $\ell_X = \ell_Y$, the actions $G \curvearrowright X$ and $G \curvearrowright Y$ 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 G -equivariantly isomorphic, since actions $D_\infty \curvearrowright \mathbb{R}$ are determined, up to conjugacy, by the restriction to the maximal \mathbb{Z} subgroup. \square

Fig. 3 Giving a hedgehog back its spines





Fig. 4 Part of a 4-regular tree T and its squarisation $S(T)$

As the next two examples demonstrate, there is no way of removing the essentiality and hyperplane-essentiality requirements from Theorem A and Corollary B.

Example 5.4 Let $G \curvearrowright X$ be any proper cocompact action on a CAT(0) cube complex. Fix a basepoint $x \in X$ and let $G_x \leq G$ denote its stabiliser. The disjoint union $Y = X \sqcup G/G_x$ is endowed with a natural G -action and we give it a structure of CAT(0) cube complex by adding edges connecting gG_x and gx for every $g \in G$. The procedure is depicted in Fig. 3 in a more general context. The CAT(0) cube complex Y is irreducible and the action $G \curvearrowright Y$ 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 G/G_x are vertices of degree 1. It is easy to see that $\ell_X = \ell_Y$ and $cr_X = cr_Y$.

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 Fig. 4.

Let \mathcal{H}' denote the disjoint union of two copies of the pocset $(\mathcal{H}(X), \subseteq, *)$, labelling by h_1 and h_2 the two elements arising from $h \in \mathcal{H}(X)$. We turn \mathcal{H}' into a pocset by declaring that, given $h, k \in \mathcal{H}(X)$ with $h \notin \{k, k^*\}$, we have $h_i \subseteq k_j$ if and only if $h \subseteq k$, no matter what the indices i and j are. On the other hand, the halfspaces h_1 and h_2 are transverse for all $h \in \mathcal{H}(X)$.

Now, $S(X)$ is obtained by applying Sageev’s construction [61,63] to the pocset $(\mathcal{H}', \subseteq, *)$. Note that, if we had instead declared that h_1 and h_2 are nested, we would have obtained the first cubical subdivision X' .

Example 5.5 Let $G \curvearrowright X$ 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 G . It is easy to see that the actions $G \curvearrowright X'$ and $G \curvearrowright S(X)$ determine the same length function, namely the double of the length function associated to $G \curvearrowright X$. Moreover, the (co-meagre) subset of $\partial_\infty G$ arising from non-terminating ultrafilters is the same for X , X' and $S(X)$ and there we have $\text{cr}_{S(X)} = \text{cr}_{X'} = 2 \cdot \text{cr}_X$.

The failure of Theorem A and Corollary B 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 .

Acknowledgements We are indebted to Anthony Genevois for bringing Bourdon's paper [12] to our attention and for his many interesting comments on an earlier version of this preprint. We are grateful to Mark Hagen and Alessandro Sisto for contributing some of the ideas in Sections 2.2 and 2.3, respectively. We also thank Pierre-Emmanuel Caprace, Christopher Cashen, Matt Cordes, Ruth Charney, Cornelia Druţu, Talia Fernós, Ilya Gekhtman, Tobias Hartnick, John Mackay, Gabriel Pallier, Beatrice Pozzetti, Michah Sageev, Viktor Schroeder and Ric Wade for helpful conversations. Finally, we thank Maria Martini for her hedgehog-drawing skills.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Arzhantseva, G.N., Cashen, C.H., Gruber, D., Hume, D.: Characterizations of Morse quasi-geodesics via superlinear divergence and sublinear contraction. *Doc. Math.* **22**, 1193–1224 (2017)
2. Agol, I.: Virtual properties of 3-manifolds. In: *Proceedings of the International Congress of Mathematicians—Seoul 2014*, vol. 1, pp. 141–170. Kyung Moon Sa, Seoul (2014)
3. Brodzki, J., Campbell, S.J., Guentner, E.P., Niblo, G.A., Wright, N.J.: Property A and CAT(0) cube complexes. *J. Funct. Anal.* **256**(5), 1408–1431 (2009)
4. Beyrer, J., Fioravanti, E.: Cross ratios on CAT(0) cube complexes and marked length-spectrum rigidity. *J. Lond. Math. Soc.* ([arXiv:1903.02447v4](https://arxiv.org/abs/1903.02447v4)) (2021)
5. Beyrer, J., Fioravanti, E., Incerti-Medici, M.: CAT(0) cube complexes are determined by their boundary cross ratio. *Groups Geom. Dyn.* **15**(1), 313–333 (2021)
6. Bridson, M.R., Haefliger, A.: *Metric spaces of non-positive curvature. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*, vol. 319. Springer-Verlag, Berlin (1999)
7. Bergeron, N., Haglund, F., Wise, D.T.: Hyperplane sections in arithmetic hyperbolic manifolds. *J. Lond. Math. Soc.* (2) **83**(2), 431–448 (2011)
8. Burns, K., Katok, A.B.: In collaboration with Werner Ballman, Michael Brin, Patrick Eberlein and Robert Osseman. *Manifolds with nonpositive curvature. Ergod. Theory Dyn. Syst.* **5**(2), 307–317 (1985)
9. Bourdon, M., Kleiner, B.: Combinatorial modulus, the combinatorial Loewner property, and Coxeter groups. *Groups Geom. Dyn.* **7**(1), 39–107 (2013)
10. Bourdon, M., Kleiner, B.: Some applications of ℓ_p -cohomology to boundaries of Gromov hyperbolic spaces. *Groups Geom. Dyn.* **9**(2), 435–478 (2015)

11. Bonahon, F.: The geometry of Teichmüller space via geodesic currents. *Invent. Math.* **92**(1), 139–162 (1988)
12. Bourdon, M.: Immeubles hyperboliques, dimension conforme et rigidité de Mostow. *Geom. Funct. Anal.* **7**(2), 245–268 (1997)
13. Beyrer, J., Schroeder, V.: Trees and ultrametric Möbius structures. *p-Adic numbers ultrametric. Anal. Appl.* **9**(4), 247–256 (2017)
14. Bergeron, N., Wise, D.T.: A boundary criterion for cubulation. *Am. J. Math.* **134**(3), 843–859 (2012)
15. Chatterji, I., Druţu, C., Haglund, F.: Kazhdan and Haagerup properties from the median viewpoint. *Adv. Math.* **225**(2), 882–921 (2010)
16. Croke, C.B., Fathi, A., Feldman, J.: The marked length-spectrum of a surface of nonpositive curvature. *Topology* **31**(4), 847–855 (1992)
17. Chatterji, I., Fernós, T., Iozzi, A.: 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)
18. Caprace, P.-E., Lécureux, J.: Combinatorial and group-theoretic compactifications of buildings. *Ann. Inst. Fourier (Grenoble)* **61**(2), 619–672 (2011)
19. Cashen, C.H., Mackay, J.M.: A metrizable topology on the contracting boundary of a group. *Trans. Am. Math. Soc.* **372**(3), 1555–1600 (2019)
20. Chatterji, I., Niblo, G.: From wall spaces to CAT(0) cube complexes. *Int. J. Algebra Comput.* **15**(5–6), 875–885 (2005)
21. Cordes, M.: Morse boundaries of proper geodesic metric spaces. *Groups Geom. Dyn.* **11**(4), 1281–1306 (2017)
22. Croke, C.B.: Rigidity for surfaces of nonpositive curvature. *Comment. Math. Helv.* **65**(1), 150–169 (1990)
23. Caprace, P.-E., Sageev, M.: Rank rigidity for CAT(0) cube complexes. *Geom. Funct. Anal.* **21**(4), 851–891 (2011)
24. Charney, R., Sultan, H.: Contracting boundaries of CAT(0) spaces. *J. Topol.* **8**(1), 93–117 (2015)
25. Dal’Bo, F., Kim, I.: Marked length rigidity for symmetric spaces. *Comment. Math. Helv.* **77**(2), 399–407 (2002)
26. Druţu, C., Kapovich, M.: Geometric group theory. In: AMS series “Colloquium Publications” (2017)
27. Delzant, T., Py, P.: Cubulable Kähler groups. *Geom. Topol.* **23**(4), 2125–2164 (2019)
28. Fernós, T.: The Furstenberg-Poisson boundary and CAT(0) cube complexes. *Ergod. Theory Dyn. Syst.* **38**(6), 2180–2223 (2018)
29. Fioravanti, E.: Superrigidity of actions on finite rank median spaces. *Adv. Math.* **352**, 1206–1252 (2019)
30. Fioravanti, E.: Roller boundaries for median spaces and algebras. *Algebr. Geom. Topol.* **20**(3), 1325–1370 (2020)
31. Fernós, T., Lécureux, J., Mathéus, F.: Random walks and boundaries of CAT(0) cubical complexes. *Comment. Math. Helv.* **93**(2), 291–333 (2018)
32. Genevois, A.: Coning-off CAT(0) cube complexes. *Ann. Institut Fourier* ([arXiv:1603.06513v1](https://arxiv.org/abs/1603.06513v1)), (2016)
33. Genevois, A.: Hyperbolicities in CAT(0) cube complexes. *Enseign. Math.* **65**(1–2), 33–100 (2019)
34. Genevois, A.: Contracting isometries of CAT(0) cube complexes and acylindrical hyperbolicity of diagram groups. *Algebr. Geom. Topol.* **20**(1), 49–134 (2020)
35. Genevois, A., Martin, A.: Automorphisms of graph products of groups from a geometric perspective. *Proc. Lond. Math. Soc.* **119**(6), 1745–1779 (2019)
36. Gitik, R., Mitra, M., Rips, E., Sageev, M.: Widths of subgroups. *Trans. Am. Math. Soc.* **350**(1), 321–329 (1998)
37. Gromov, M.: Hyperbolic groups. In: *Essays in group theory*, vol. 8 of *Math. Sci. Res. Inst. Publ.*, pp. 75–263. Springer, New York (1987)
38. Haïssinsky, P.: Géométrie quasiconforme, analyse au bord des espaces métriques hyperboliques et rigidités [d’après Mostow, Pansu, Bourdon, Pajot, Bonk, Kleiner. . .]. *Astérisque*, (326):Exp. No. 993, ix, 321–362 (2010) 2009
39. Hamenstädt, U.: Cocycles, Hausdorff measures and cross ratios. *Ergod. Theory Dyn. Syst.* **17**(5), 1061–1081 (1997)
40. Hamenstädt, U.: Cocycles, symplectic structures and intersection. *Geom. Funct. Anal.* **9**(1), 90–140 (1999)
41. Hagen, M.F., Susse, T.: On hierarchical hyperbolicity of cubical groups. *Israel J. Math.* **236**(1), 45–89 (2020)

42. Hagen, M.F., Touikan, N.W.M.: Panel collapse and its applications. *Groups Geom. Dyn.* **13**(4), 1285–1334 (2019)
43. Haglund, F., Wise, D.T.: Special cube complexes. *Geom. Funct. Anal.* **17**(5), 1551–1620 (2008)
44. Haglund, F., Wise, D.T.: A combination theorem for special cube complexes. *Ann. Math.* **176**(3), 1427–1482 (2012)
45. Hagen, M.F., Wise, D.T.: Cubulating hyperbolic free-by-cyclic groups: the general case. *Geom. Funct. Anal.* **25**(1), 134–179 (2015)
46. Hagen, M.F., Wise, D.T.: Cubulating hyperbolic free-by-cyclic groups: the irreducible case. *Duke Math. J.* **165**(9), 1753–1813 (2016)
47. Kahn, J., Markovic, V.: Immersing almost geodesic surfaces in a closed hyperbolic three manifold. *Ann. Math.* **175**(3), 1127–1190 (2012)
48. Kapovich, I., Short, H.: Greenberg’s theorem for quasiconvex subgroups of word hyperbolic groups. *Can. J. Math.* **48**(6), 1224–1244 (1996)
49. Labourie, F.: Anosov flows, surface groups and curves in projective space. *Invent. Math.* **165**(1), 51–114 (2006)
50. Labourie, F.: Cross ratios, surface groups, $\mathrm{PSL}(n, \mathbb{R})$ and diffeomorphisms of the circle. *Publ. Math. Inst. Hautes Études Sci.* **106**, 139–213 (2007)
51. James, R.: *Munkres. Topology*. Prentice Hall Inc, Upper Saddle River, NJ (2000)
52. Nica, B.: Cubulating spaces with walls. *Algebr. Geom. Topol.* **4**, 297–309 (2004)
53. Niblo, G.A., Reeves, L.D.: Groups acting on $\mathrm{CAT}(0)$ cube complexes. *Geom. Topol.* **1**, 1–7 (1997)
54. Niblo, G.A., Reeves, L.D.: Coxeter groups act on $\mathrm{CAT}(0)$ cube complexes. *J. Group Theory* **6**(3), 399–413 (2003)
55. Nevo, A., Sageev, M.: The Poisson boundary of $\mathrm{CAT}(0)$ cube complex groups. *Groups Geom. Dyn.* **7**(3), 653–695 (2013)
56. Otal, J.-P.: Le spectre marqué des longueurs des surfaces à courbure négative. *Ann. Math.* **131**(1), 151–162 (1990)
57. Otal, J.-P.: Sur la géométrie symplectique de l’espace des géodésiques d’une variété à courbure négative. *Rev. Mat. Iberoamericana* **8**(3), 441–456 (1992)
58. Ollivier, Y., Wise, D.T.: Cubulating random groups at density less than $1/6$. *Trans. Am. Math. Soc.* **363**(9), 4701–4733 (2011)
59. Paulin, F.: Un groupe hyperbolique est déterminé par son bord. *J. Lond. Math. Soc.* **54**(1), 50–74 (1996)
60. Roller, M.A.: *Poc Sets, Median Algebras and Group Actions*. University of Southampton, An extended study of Dunwoody’s construction and Sageev’s theorem. preprint (1998)
61. Sageev, M.: Ends of group pairs and non-positively curved cube complexes. *Proc. Lond. Math. Soc.* **71**(3), 585–617 (1995)
62. Sageev, M.: Codimension-1 subgroups and splittings of groups. *J. Algebra* **189**(2), 377–389 (1997)
63. Sageev, M.: $\mathrm{CAT}(0)$ cube complexes and groups. In: *Geometric Group Theory*, vol. 21 of IAS/Park City Math. Ser., pp. 7–54. Am. Math. Soc., Providence, RI, (2014)
64. Wise, D.T.: Cubulating small cancellation groups. *Geom. Funct. Anal.* **14**(1), 150–214 (2004)
65. Wise, D. T.: The cubical route to understanding groups. In: *Proceedings of the International Congress of Mathematicians—Seoul 2014*. Vol. II, pp. 1075–1099. Kyung Moon Sa, Seoul (2014)
66. Xie, X.: Quasi-isometric rigidity of Fuchsian buildings. *Topology* **45**(1), 101–169 (2006)

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.