



Homotopy equivalent boundaries of cube complexes

Talia Fernós¹ · David Futer² · Mark Hagen³

Received: 10 May 2023 / Accepted: 15 December 2023 / Published online: 27 January 2024
© The Author(s) 2024

Abstract

A finite-dimensional CAT(0) cube complex X is equipped with several well-studied boundaries. These include the *Tits boundary* $\partial_T X$ (which depends on the CAT(0) metric), the *Roller boundary* $\partial_R X$ (which depends only on the combinatorial structure), and the *simplicial boundary* $\partial_\Delta X$ (which also depends only on the combinatorial structure). We use a partial order on a certain quotient of $\partial_R X$ to define a simplicial Roller boundary $\mathfrak{R}_\Delta X$. Then, we show that $\partial_T X$, $\partial_\Delta X$, and $\mathfrak{R}_\Delta X$ are all homotopy equivalent, $\text{Aut}(X)$ -equivariantly up to homotopy. As an application, we deduce that the perturbations of the CAT(0) metric introduced by Qing do not affect the equivariant homotopy type of the Tits boundary. Along the way, we develop a self-contained exposition providing a dictionary among different perspectives on cube complexes.

Contents

1	Introduction	2
1.1	A plethora of boundaries	3
1.2	Main result	5
1.3	Invariance of CAT(0) boundaries for cubulated groups	6
1.4	Quasiflats	9
1.5	Ingredients of our proof	10
1.6	Expository content	12
1.7	Section Breakdown	13
2	Simplicial complexes and nerve theorems	14
2.1	Equivariant nerve theorems	15
3	Cube complexes and the Roller boundary	19
3.1	CAT(0) cube complexes and metrics	19

✉ Mark Hagen
markfhagen@posteo.net

Talia Fernós
t_fernos@uncg.edu

David Futer
dfuter@temple.edu

¹ Mathematics and Statistics Department, University of North Carolina at Greensboro, Greensboro, NC 27412, USA

² Department of Mathematics, Temple University, Philadelphia, PA 19122, USA

³ School of Mathematics, University of Bristol, Bristol BS8 1UG, UK

3.2	Hyperplanes, half-spaces, crossing, and separation	19
3.3	The Roller boundary $\partial_R X$	21
3.4	Convexity, intervals, and medians	22
3.5	Lifting decompositions and convexity	24
3.6	Gates and bridges	26
3.7	Facing tuples and properness	28
3.8	Combinatorial Geodesic Rays	29
4	Cuboid complexes with modified CAT(0) metrics	30
5	The simplicial Roller boundary $\mathfrak{R}_\Delta X$	33
6	The simplicial boundary	35
6.1	Unidirectional boundary sets	35
6.2	Minimal and dominant UBSeS	37
6.3	Connections between $\partial_\Delta X$ and $\partial_R X$	41
6.4	Order-preserving maps	44
7	The Tits boundary	47
7.1	Deep sets	48
7.2	CAT(0) rays and their UBSeS	50
8	Connections between Tits and Roller boundaries	57
8.1	Tits boundary realizations	58
8.2	Diameter, intrinsic radius, and Tits-convexity of $Q(v)$	59
8.3	Visibility and perturbing the circumcenter	60
8.4	Cuboid generalization	63
9	Open and closed coverings of $\partial_T X$	65
9.1	The closed covering	65
9.2	The open covering	66
9.3	Homotopy equivalence of $\partial_T X$ and \mathcal{N}_T	73
9.4	Cuboid generalization	74
10	Proof of the Main Theorems	75
10.1	Closed cover of $\mathfrak{R}_\Delta^\triangleleft X$	75
10.2	Homotopy equivalence between $\mathfrak{R}_\Delta X$ and $\partial_T X$	76
10.3	Homotopy equivalence between $\partial_\Delta X$ and $\mathfrak{R}_\Delta X$	79
10.4	Cuboid generalization	80
	References	81

Keywords CAT(0) cube complex · Tits boundary · Roller boundary · Nerve theorem

Mathematics Subject Classification 20F65 · 51F99

1 Introduction

CAT(0) cube complexes, which exist in many guises in discrete mathematics (see e.g. [2]), were introduced into group theory by Gromov [36] and have since taken on a central role in that field. As combinatorial objects, CAT(0) cube complexes are ubiquitous due to their flexible, functorial construction from set-theoretic data [24, 52, 58, 60]. This has led to an industry of *cubulating* groups, i.e. constructing group actions on CAT(0) cube complexes in order to transfer information from the highly organized cubical structure to the group. Probably the best known application of this method is the resolution of the virtual Haken and virtual fibering conjectures in 3-manifold theory [1, 63].

The utility of CAT(0) cube complexes comes from the fact that they simultaneously exhibit several types of structures. They have an organized combinatorial structure coming from their *hyperplanes/half-spaces*; this is closely related to the very tractable geometry of their 1-skeleta, which are *median graphs* [20]. On the other hand, endowing the complex with the piecewise-Euclidean metric in which cubes are Euclidean unit cubes, one gets a

CAT(0) space. So, in studying CAT(0) cube complexes, and groups acting on them, one has a wide variety of tools.

This paper is about the interplay between the CAT(0) and combinatorial structures, at the level of *boundaries*. We construe the term “boundary” broadly: we include not just spaces arising as frontiers of injections into compact spaces, but also other spaces encoding some sort of behavior at infinity or large-scale asymptotic structure, such as the Tits boundary.

1.1 A plethora of boundaries

CAT(0) cube complexes have several natural boundaries, each encoding different information, and each defined in terms of either the CAT(0) metric structure or the combinatorial structure coming from the hyperplanes. Fixing a finite-dimensional CAT(0) cube complex X , one has the following list of boundaries:

- (1) The *visual boundary* $\partial_{\triangleleft} X$. The visual boundary can be defined for any CAT(0) space; see [7]. Points are asymptotic equivalence classes of CAT(0) geodesic rays. The visual topology is defined in such a way that, roughly speaking, rays that fellow-travel for a long time are close. When X is locally finite, $X \cup \partial_{\triangleleft} X$ is a compactification of X . While it is a very useful object, we do not study the visual boundary in this paper.
- (2) The *Tits boundary* $\partial_T X$. As with $\partial_{\triangleleft} X$, the points in $\partial_T X$ correspond to asymptotic equivalence classes of CAT(0) geodesic rays, but the topology is finer than the visual topology. Specifically, we equip $\partial_T X$ with the *Tits metric*: by taking a supremum of angles between rays one obtains the *angle metric*, and the induced length metric is the Tits metric, which is CAT(1); see [7, Theorem II.9.13]. The Tits boundary encodes much of the geometry of a CAT(0) space; for example, spherical join decompositions of the Tits boundary correspond to product decompositions of the space [7, Theorem II.9.24]. Although $\partial_T X$ is in general not compact even when X is proper, it is of interest for other reasons, such as encoding “partial flat regions” in X . For example, endpoints of axes of rank-one isometries are isolated points, while flats in X yield spheres in $\partial_T X$. We recall the definition of $\partial_T X$ in Definition 7.1.
- (3) The *simplicial boundary* $\partial_{\Delta} X$, from [39, 40], is an analogue of the Tits boundary depending on the hyperplane structure, rather than on the CAT(0) metric. The idea is that certain sets of hyperplanes—termed *unidirectional boundary sets* (hereafter, *UBSes*)—identify “ways of approaching infinity in X .” Containment of UBSes (modulo finite differences) gives a partial order on UBSes, and this order gives rise to a simplicial complex $\partial_{\Delta} X$ (see Definition 6.9). Here are three helpful examples. First, the simplicial boundary of a tree (or, more generally, of a δ -hyperbolic cube complex) is a discrete set of 0-simplices. Second, the simplicial boundary of the standard square tiling of $[0, \infty)^2$ is a 1-simplex, whose 0-simplices correspond to the sub-UBSes consisting of the vertical hyperplanes and the horizontal hyperplanes. Third, the *staircase* obtained from this square tiling by considering only the cubes below some increasing, unbounded function also has simplicial boundary a 1-simplex. (See Fig. 1.) The maximal simplices of $\partial_{\Delta} X$ encode such “generalized orthants” (namely, convex hulls of ℓ^1 -geodesic rays) in X in roughly the same way that the Tits boundary encodes “partial flats.”

The simplicial boundary has been used to study quasi-isometry invariants like *divergence* [39] and *thickness* [8] for groups acting on cube complexes, and has been generalized in the context of *median spaces* [29] as a tool for proving a Tits alternative for groups acting on such spaces.

- (4) The *Roller boundary* $\partial_R X$ gives another way of compactifying X , using the half-space structure associated to the hyperplanes. Each hyperplane \hat{h} of X has two complementary components, called *half-spaces*, h and h^* , which induce a two-sided partition of $X^{(0)}$. Each vertex $x \in X$ is completely determined by specifying the collection of half-spaces that contain it. This gives an injective map $X^{(0)} \rightarrow 2^{\mathfrak{H}}$, where \mathfrak{H} denotes the set of half-spaces. The closure of the image of $X^{(0)}$ in the Tychonoff topology is the *Roller compactification* \overline{X} , and $\partial_R X = \overline{X} \setminus X^{(0)}$ is the *Roller boundary* of X . (See Definition 3.3.)

Of the boundaries of X that are defined in terms of the cubical structure only, the Roller boundary is perhaps the most well-studied. It has been used to prove a variety of results about CAT(0) cube complexes and groups acting on them.

For example, the Roller boundary plays a key role in the proof that irreducible lattices in $\mathrm{SL}_2\mathbb{C} \times \mathrm{SL}_2\mathbb{C}$ always have global fixed points in any action on a CAT(0) cube complex, while reducible actions always admit proper cocompact actions on CAT(0) cube complexes [19].

Nevo and Sageev identified the Roller boundary as a model for the Furstenberg–Poisson boundary of a cocompact lattice in $\mathrm{Aut}(X)$ [54]. Fernós generalized this result to groups acting properly and non-elementarily on finite dimensional X , and also proved a Tits alternative [28]. The same paper identifies an important subset of the Roller boundary: the *regular points*, which correspond to “hyperbolic” directions in various senses.

In particular, the set of regular points with the induced topology is $\mathrm{Aut}(X)$ -equivariantly homeomorphic with the boundary of the (hyperbolic) *contact graph* [33]. The set of regular points is used in [32] to find rank-one isometries under more general conditions than in previous work [25]. The method in [32] is to use convergence of random walks to regular points to deduce the existence of regular elements without assuming that the ambient group is a lattice.

The set of regular points also features in the proof of marked length spectrum rigidity [5, 6], and was independently identified by Kar and Sageev, who used it to study property P_{naive} and hence C^* -simplicity for cubulated groups [48]. Finally, the Roller boundary has recently been generalized in the context of median spaces [29].

- (5) The *simplicial Roller boundary* $\mathfrak{R}_\Delta X$ is constructed as follows. The Roller boundary $\partial_R X$ carries a natural equivalence relation, first introduced by Guralnik [37], where two points are equivalent if they differ on finitely many half-spaces. The equivalence classes, called *Roller classes*, play a fundamental role in this paper. Part of the reason for this is that the set of Roller classes carries a natural partial order, which admits several equivalent useful characterizations (see Lemma 5.6). The *simplicial Roller boundary* $\mathfrak{R}_\Delta X$ is the simplicial complex realizing this partial order (see Definition 5.7).

In this paper, we build on earlier work relating Roller classes to CAT(0) geodesic rays. Specifically, Guralnik in [37] recognized that each equivalence class of CAT(0) geodesic rays determines a Roller class. Conversely, in [32], it is shown that each Roller class determines a subset of the Tits boundary which admits a canonical circumcenter. These observations are crucial for our arguments: Definition 8.7 depends on the latter and Definition 8.1 is based on the former.

We also note that $\mathfrak{R}_\Delta X$ is isomorphic to the *combinatorial boundary* defined and studied by Genevois [34]. In that work, he explains that the combinatorial boundary is isomorphic to the face-poset of a naturally-defined subcomplex of $\partial_\Delta X$ [34, Proposition A.1] and is isomorphic to $\mathfrak{R}_\Delta X$ [34, Proposition A.5]. (More precisely, Genevois’ work yields the map U_Δ^R from Corollary 6.33 below, although [34] does not explicitly mention a homotopy equivalence.)




Fig. 1 Left: a linear staircase. Right: a sublinear staircase. For both staircases, the simplicial boundary is a 1-simplex, whose vertices correspond to the UBS \mathcal{V} of vertical hyperplanes and the UBS \mathcal{H} of horizontal hyperplanes. In both staircases, \mathcal{H} is not visible (in either the ℓ^1 or the ℓ^2 metric), because any geodesic ray crossing \mathcal{H} must also cross \mathcal{V} . In the sublinear staircase, $\mathcal{V} \sqcup \mathcal{H}$ is also not ℓ^2 -visible

For each of these boundaries, the action of the group $\text{Aut}(X)$ of cubical automorphisms of X extends to an action on the boundary preserving its structure. The actions on $\partial_{\triangleleft} X$ and $\partial_R X$ are by homeomorphisms, the action on $\partial_T X$ is by isometries of the Tits metric, and the actions on $\partial_{\Delta} X$ and $\mathfrak{R}_{\Delta} X$ are by simplicial automorphisms.

The definitions of the simplicial boundary and the simplicial Roller boundary are conceptually similar. Part of the work in this paper is making that similarity precise, by establishing a correspondence between Roller classes and equivalence classes of UBSes. This line of work culminates in Corollary 6.33, which gives explicit maps between $\mathfrak{R}_{\Delta} X$ and a natural subcomplex of $\partial_{\Delta} X$. These maps are $\text{Aut}(X)$ -equivariant homotopy equivalences, although in general they are not simplicial isomorphisms. Then, in Proposition 10.12, we upgrade this to a homotopy equivalence between $\mathfrak{R}_{\Delta} X$ and the whole of $\partial_{\Delta} X$.

One glimpse of a relationship between $\partial_{\Delta} X$ and $\partial_T X$ comes from UBSes associated to geodesic rays. Given a geodesic ray $\alpha \rightarrow X$, in either the CAT(0) metric or the combinatorial metric, the set of hyperplanes crossing α (denoted $\mathcal{W}(\alpha)$) is always a UBS. A UBS \mathcal{U} is called ℓ^1 -visible (resp. ℓ^2 -visible) if it has finite symmetric difference with $\mathcal{W}(\alpha)$ for a combinatorial (resp. CAT(0)) geodesic α . Figure 1 shows examples of both ℓ^1 -invisible and ℓ^2 -invisible UBSes. Despite those caveats, ℓ^2 -visible UBSes provide a way to map points of $\partial_T X$ to classes in $\partial_{\Delta} X$; a similar construction also provides a map to $\mathfrak{R}_{\Delta} X$.

1.2 Main result


Our main theorem relates $\partial_T X$, $\partial_{\Delta} X$, and $\mathfrak{R}_{\Delta} X$ via maps that are $\text{Aut}(X)$ -equivariant up to homotopy. It helps to introduce the following terminology.

Definition 1.1 Suppose that A, B are topological spaces with a group G acting on both spaces by homeomorphisms. A map $f: A \rightarrow B$ is called a G -homotopy equivalence if f is a homotopy equivalence, and furthermore $g \circ f$ and $f \circ g$ are homotopic for every $g \in G$. If such an f exists, we say that A, B are G -homotopy equivalent.

It is immediate to check that the property of being G -homotopy equivalent is an equivalence relation. Indeed, the composition of two G -homotopy equivalences is a G -homotopy

equivalence. Furthermore, any homotopy inverse of a $G \sim$ -homotopy equivalence is itself a $G \sim$ -homotopy equivalence.

Theorem A *Let X be a finite-dimensional CAT(0) cube complex. Then we have the following commutative diagram of $\text{Aut}(X) \sim$ -homotopy equivalences between boundaries of X :*



In particular, the spaces $\partial_{\Delta} X$, $\mathfrak{R}_{\Delta} X$, $\partial_T X$ are all $\text{Aut}(X) \sim$ -homotopy equivalent, where $\partial_T X$ is equipped with the metric topology and $\partial_{\Delta} X$, $\mathfrak{R}_{\Delta} X$ are equipped with either the metric or the weak topology.

In the above diagram, the letter S stands for *Simplicial*, the letter R for *Roller-simplicial*, and the letter T for *Tits*. The map T^S is an $\text{Aut}(X) \sim$ -homotopy equivalence from the simplicial boundary to the Tits boundary, and similarly for the other two. The map T^R is provided by Corollary 10.11, and the map R^S comes from Proposition 10.12. The map T^S is just the composition $T^S = T^R \circ R^S$.

Some intuition behind the theorem can be gleaned from the following simplified situation. A cube complex X is called *fully visible* if every UBS is ℓ^1 -visible. When X is fully visible, UBSes correspond to combinatorial orthants in X (i.e. products of combinatorial geodesic rays) [39], and Theorem A can be understood as relating combinatorial orthants to CAT(0) orthants. See [39, Section 3], where a proof is sketched that the simplicial and Tits boundaries are homotopy equivalent under the restrictive hypothesis of full visibility. However, full visibility is a very strong and somewhat mysterious hypothesis: it is not known to hold even if X admits a proper cocompact group action [45]. Thus Theorem A is more satisfying (and more difficult) because it does not assume this restrictive hypothesis. In order to prove the theorem, we have to understand combinatorial convex hulls of CAT(0) geodesic rays, which is a much more delicate affair when one is not simply assuming that these hulls can be taken to contain Euclidean orthants.

Theorem A does not have an analogue with the Tits boundary replaced by the visual boundary. For example, if X has a proper, cocompact, essential action by a group G , and X is irreducible, then G contains a rank-one isometry [CS11], and therefore the simplicial boundary has two isolated points. Even without a cocompact group action, the boundary has isolated points corresponding to regular points in the Roller boundary, by [28, Proposition 7.4] and [39, Corollary 3.20]. On the other hand, if X is one-ended, then the visual boundary is connected.

Theorem A also motivates the introduction of several new technical tools for relating the geometry of CAT(0) geodesic rays to the combinatorics of hyperplanes. We believe that these tools, explained in Sect. 1.5, are of independent interest.

1.3 Invariance of CAT(0) boundaries for cubulated groups

A group G is called *CAT(0)* if there is a proper CAT(0) space X on which G acts geometrically (properly and cocompactly). Multiple CAT(0) spaces might admit a geometric G -action and witness that G is a CAT(0) group. Since we are often interested in invariants of the group itself, it is natural to look for features of the geometry of X that depend only on G . For instance, in the

context of a Gromov hyperbolic group G , recall that all hyperbolic groups with a geometric G -action have Gromov boundaries that are G -equivariantly homeomorphic. So one might wonder whether a similar result might hold for CAT(0) groups that are not hyperbolic.

A famous result of Croke and Kleiner [21, 22], on which we elaborate more below, shows that this is not the case. Explicitly, they considered the right-angled Artin group

$$G = \langle a, b, c, d : [a, b], [b, c], [c, d] \rangle,$$

whose Cayley 2-complex $X_{\pi/2}$ is a CAT(0) square complex. The subscript $\pi/2$ emphasizes the angles at the corners of the 2-cells. Viewing the 2-cells as Euclidean squares with side length 1, we obtain a CAT(0) space where G acts geometrically. Croke and Kleiner deformed the squares of X into rhombi that have angles α and $\pi - \alpha$, producing a perturbed CAT(0) metric X_α , and showed that this perturbation changes the homeomorphism type of the visual boundary $\partial_{\triangleleft} X_\alpha$. Subsequently, Wilson [61] showed that the visual boundaries of X_α and X_β are homeomorphic only if the angles satisfy $\alpha = \beta$, hence these perturbations produce uncountably many homeomorphism types of boundaries. Further examples are provided by Mooney [50], who exhibited CAT(0) knot groups G admitting uncountably many geometric actions on CAT(0) spaces, all with different visual boundaries. Hosaka [44] has given some conditions on G, X, Y implying that there is an equivariant homeomorphism $\partial_{\triangleleft} X \rightarrow \partial_{\triangleleft} Y$ continuously extending the quasi-isometry $X \rightarrow Y$ coming from orbit maps, but the preceding examples show that any such condition will be hard to satisfy.

So, one has to look for weaker results or less refined invariants, or ask slightly different questions. This has stimulated a great deal of work in various directions.

In the context of 2-dimensional CAT(0) complexes, there are more positive results obtained by replacing the visual boundary $\partial_{\triangleleft} X$ by the Tits boundary $\partial_T X$, and passing to a natural subspace. Specifically, the core of $\partial_T X$ is the union of all embedded circles in $\partial_T X$. Xie [64] showed that if G acts geometrically on CAT(0) 2-complexes X and Y , then the Tits boundaries $\partial_T X$ and $\partial_T Y$ have homeomorphic cores. (This is weaker than Xie’s actual statement; compare [64, Theorem 6.1].) In the Croke–Kleiner example, the core is a G -invariant connected component, while the rest of the Tits boundary consists of uncountably many isolated points and arcs.

There is a closely related result due to Qing, which inspired us to consider *cuboid complexes* in the present paper. Another way to perturb the CAT(0) metric on a CAT(0) cube complex X is to leave the angles alone, but to vary the lengths of the edges in such a way that edges intersecting a common hyperplane are given equal length. In this perturbation, each cube becomes a Euclidean box called a *cuboid*. The resulting path metric is still CAT(0), at least when there are uniform upper and lower bounds on the edge-lengths. In particular, this happens when some group G acts on X with finitely many orbits of hyperplanes, and the edge-lengths are assigned G -equivariantly. We explain the details in Sect. 4, where we define a *G -admissible hyperplane rescaling* in Definition 4.1 and show in Lemma 4.2 that the resulting path-metric space (X, d_X^ρ) is CAT(0) and G continues to act by isometries. Cuboid complexes have been studied by various authors, including Beyrer and Fioravanti [5].

Qing [56] studied the visual boundary of the CAT(0) cuboid metric under perturbation of the edge-lengths. She showed that the visual boundaries of the CAT(0) cuboid complexes obtained from the Croke–Kleiner complex $X_{\pi/2}$ are all homeomorphic. More germane to the present paper, Qing also showed that the Tits boundaries of the rescaled cuboid complexes are homeomorphic, but that no equivariant homeomorphism exists. (Roughly, the cores of any two such Tits boundaries are homeomorphic by the result of Xie mentioned above. Via a cardinality argument, Qing extends the homeomorphism over the whole boundary by sending isolated points/arcs to isolated points/arcs. But this extension cannot be done equivariantly

because equivariance forces some isolated points to be sent to isolated arcs and vice versa.) We will return to cuboid complexes shortly.

There are other results about homeomorphism type as a $\text{CAT}(0)$ group invariant for certain classes of $\text{CAT}(0)$ groups. For example, Bowers and Ruane [17] showed that if G is a product of a hyperbolic $\text{CAT}(0)$ group and \mathbb{Z}^D , then the visual boundaries of the $\text{CAT}(0)$ spaces X and Y with a geometric G -action are equivariantly homeomorphic. However, this equivariant homeomorphism need not come from a continuous extension of a quasi-isometry $X \rightarrow Y$. Later, Ruane [59] extended this result to the case where G is a $\text{CAT}(0)$ direct product of two non-elementary hyperbolic groups.

Around the same time, Buyalo studied groups of the form $G = \pi_1(S) \times \mathbb{Z}$, where S is a hyperbolic surface [18, Section 14]. He constructed two distinct geometric G -actions on $\mathbb{H}^2 \times \mathbb{R}$ and showed that while the two copies of $\mathbb{H}^2 \times \mathbb{R}$ are G -equivariantly quasi-isometric, the Tits boundaries do not admit any G -equivariant quasi-isometry.

Since we are interested in G -equivariant results, it makes sense to seek more general positive results by replacing homeomorphism with a coarser equivalence relation. This has been a successful idea: Bestvina [4] proved that torsion-free $\text{CAT}(0)$ groups have a well-defined visual boundary up to shape-equivalence. This was generalized by Ontaneda [55], whose result removes the “torsion-free” hypothesis.

Homotopy equivalence is a finer equivalence relation than shape equivalence. Might there be some result guaranteeing that the homotopy type of the Tits boundary of X is to some extent independent of the choice of the $\text{CAT}(0)$ space X , even if we allow only homotopy equivalences between boundaries that respect the G -action in the appropriate sense? One precise version of this question appears below as Question 1.3.

Motivated by Qing’s results about perturbing the metric on a $\text{CAT}(0)$ cube complex by changing edge lengths, we consider the situation where G acts on the finite-dimensional $\text{CAT}(0)$ cube complex X , and study the Tits boundaries of the $\text{CAT}(0)$ spaces that result from G -equivariantly replacing cubes by cuboids. Now, the hyperplane combinatorics, and hence the simplicial and Roller boundaries, are unaffected by this change. And, for the unperturbed X , Theorem A tells us that the homotopy type of the Tits boundary is really a feature of the hyperplane combinatorics. So, in order to conclude that the homotopy type of the Tits boundary is unaffected by such perturbations of the $\text{CAT}(0)$ metric, we just need to know that Theorem A holds for cuboid complexes as well as cube complexes. Indeed, we show:

Theorem B *Let X be a finite-dimensional $\text{CAT}(0)$ cube complex and let G be a group acting by automorphisms on X . Let (X, d_X^ρ) be the $\text{CAT}(0)$ cuboid complex obtained from a G -admissible hyperplane rescaling of X . Then the original Tits boundary $\partial_T X$, the perturbed Tits boundary $\partial_T(X, d_X^\rho)$, the simplicial boundary $\partial_\Delta X$, and the simplicial Roller boundary $\mathfrak{R}_\Delta X$ are all G -homotopy equivalent.*

The proof of this result only requires minor modifications to the proof of Theorem A, essentially because convexity of half-spaces and the metric product structure of hyperplane carriers persist in the $\text{CAT}(0)$ cuboid metric d_X^ρ . These small modifications are described in Sects. 8.4, 9.4, and 10.4.

Remark 1.2 The proofs of Theorem A and B yield a slightly stronger conclusion than G -homotopy equivalence. We say that G -spaces A, B are *Borel G -homotopy equivalent* if there is a map $f : A \rightarrow B$ that factors as a composition of finitely many homotopy equivalences, each of which is either G -equivariant or a homotopy inverse of a G -equivariant map. (This definition is somewhat reminiscent of the notion of cell-like equivalence. Com-

pare Guilbault and Mooney [35], specifically the discussion Bestvina’s Cell-like Equivalence Question on page 120.)

In the proof of both theorems, the homotopy equivalence from $\partial_T X$ (or $\partial_T(X, d_X^\rho)$) to $\mathfrak{R}_\Delta X$ is constructed as an explicit composition of maps, each of which is either a G -equivariant homotopy equivalence, a deformation retraction homotopy inverse to a G -equivariant inclusion map (see Proposition 10.5), or a homotopy equivalence coming from one of the nerve theorems; the latter are compositions of G -equivariant maps or homotopy inverses of G -equivariant maps by Remark 2.9. The homotopy equivalence $\mathfrak{R}_\Delta X \rightarrow \partial_\Delta X$ similarly proceeds by composing G -equivariant maps and their homotopy inverses, in view of the same remark about the nerve theorems.

From the definitions, Borel G -homotopy equivalence is a stronger notion than G -homotopy equivalence. Furthermore, letting EG be a classifying space for G , Theorem A and the preceding observations imply that the homotopy quotients $EG \times_G \partial_T X$ and $EG \times_G \mathfrak{R}_\Delta X$ are homotopy equivalent and therefore the G -equivariant (Borel) cohomology of $\partial_T X$ is isomorphic to that of $\mathfrak{R}_\Delta X$ and $\partial_\Delta X$.

Theorem B implies, in particular, that if X admits a geometric action by a group G , then the G -homotopy type of the Tits boundary is unaffected by replacing the standard CAT(0) metric (where all edge lengths are 1) with a cuboid metric obtained by rescaling edges G -equivariantly. This is perhaps evidence in favor of the possibility that the results of Bestvina and Ontaneda about shape equivalence of visual boundaries [4, 55] can be strengthened to results about G -homotopy equivalence if one instead uses the Tits boundary:

Question 1.3 *For which CAT(0) groups G is it the case that any two CAT(0) spaces X, Y on which G acts geometrically have G -homotopy equivalent Tits boundaries?*

If G acts geometrically on CAT(0) cube complexes X, Y , are the simplicial boundaries of X, Y G -homotopy equivalent?

1.4 Quasiflats

A remarkable theorem of Huang [46, Theorem 1.1] says that if X is a d -dimensional CAT(0) cube complex, then any quasi-isometric embedding $\mathbb{R}^d \rightarrow X$ has image that is Hausdorff-close to a finite union of *cubical orthants*, where an orthant is a convex subcomplex that splits as the product of rays. This statement is important for the study of quasi-isometric rigidity in cubical groups, and it is natural to want to strengthen the statement to cover quasi-isometric embeddings $\mathbb{R}^d \rightarrow X$, where d is the largest dimension for which such a map exists, but $\dim X$ is allowed to be larger than d (although still finite). Natural examples where this strengthened form of Huang’s theorem is useful include right-angled Coxeter groups.

We believe that Theorem A could be used as an ingredient in proving such a result. Very roughly, the homotopy equivalence $\partial_T X \rightarrow \partial_\Delta X$ can be used to produce, given a singular d -cycle z in $\partial_T X$, a simplicial d -cycle z' in $\partial_\Delta X$ represented by a finite collection of standard orthants in X . By the latter, we mean there is a finite collection of d -simplices in $\partial_\Delta X$ whose union S carries z' . The proof of Theorem A, specifically Proposition 8.12 and the nerve arguments in Sects. 9 and 10, should provide enough metric control on the map $S \rightarrow \partial_T X$ to invoke a result of Kleiner–Lang [47] to deduce the strengthened quasiflats theorem. Given that the strengthened quasiflats theorem has recently been established by Bowditch [16, Theorem 1.1] and independently Huang–Kleiner–Stadler [43, Theorem 1.7], we have decided not to pursue the matter in this paper.

1.5 Ingredients of our proof

As mentioned above, our primary goal is to elucidate the relationships among three boundaries: the simplicial boundary $\partial_\Delta X$, the simplicial Roller boundary $\mathfrak{R}_\Delta X$, and the Tits boundary $\partial_T X$.

The first two boundaries that we study are combinatorial in nature. The primary difference is that the central objects in $\partial_\Delta X$ are sequences of hyperplanes, whereas the central objects in $\mathfrak{R}_\Delta X$ are sequences of half-spaces. After developing a number of lemmas that translate between the two contexts, we prove the following result:

Corollary 6.33 The barycentric subdivision of $\partial_\Delta X$ contains a canonical, $\text{Aut}(X)$ -invariant subcomplex $\partial_\Delta^{\text{UBS}} X$. There are $\text{Aut}(X)$ -equivariant simplicial maps $\mathbf{R}_\Delta^{\text{U}}: \partial_\Delta^{\text{UBS}} X \rightarrow \mathfrak{R}_\Delta X$ and $\mathbf{U}_\Delta^{\text{R}}: \mathfrak{R}_\Delta X \rightarrow \partial_\Delta^{\text{UBS}} X$, with the following properties:

- (1) $\mathbf{R}_\Delta^{\text{U}}: \partial_\Delta^{\text{UBS}} X \rightarrow \mathfrak{R}_\Delta X$ is surjective.
- (2) $\mathbf{U}_\Delta^{\text{R}}: \mathfrak{R}_\Delta X \rightarrow \partial_\Delta^{\text{UBS}} X$ is an injective section of $\mathbf{R}_\Delta^{\text{U}}$.
- (3) $\mathbf{R}_\Delta^{\text{U}}$ is a homotopy equivalence with homotopy inverse $\mathbf{R}_\Delta^{\text{U}}$.

In fact, all of $\partial_\Delta X$ deformation retracts to $\partial_\Delta^{\text{UBS}} X$ in an $\text{Aut}(X)$ -equivariant way. See Remark 6.34.

In contrast to the the two combinatorial boundaries, the Tits boundary $\partial_T X$ is inherently linked to the geometry of the CAT(0) metric on X . To relate the Tits boundary to the other two boundaries, we need to combinatorialize it: that is, we need to cover $\partial_T X$ by a certain collection of open sets, and then study the nerve of the corresponding covering. Similarly, we cover each of $\partial_\Delta X$ and $\mathfrak{R}_\Delta X$ by simplicial subcomplexes, and study the resulting nerves. We will use two different versions of the nerve theorem to show that a topological space (such as one of our boundaries) is homotopy equivalent to the nerve of a covering. The Open Nerve Theorem 2.7 deals with open coverings and is originally due to Borsuk [13]. The simplicial Nerve Theorem 2.8 deals with coverings by simplicial complexes and is originally due to Björner [11]. In fact, since Theorem A is a G -equivariant statement, we need equivariant versions of both theorems, which have not previously appeared in the literature to our knowledge. Consequently, Sect. 2 contains self-contained proofs of both theorems.

A central object in our construction of nerves is the Tits boundary realization of a Roller class, or point in the Roller boundary (see Definitions 8.4, 8.5, 8.7). In [32], a Roller class v , yields a convex, visually compact subset of the Tits boundary $\partial_T X$. Employing the work of Caprace–Lytchak [23] and Balser–Lytchak [12], this set has radius at most $\pi/2$ and a canonical circumcenter $\chi(v)$. A first approach might be to use these compacta to provide the 0-skeleton of a nerve for $\partial_T X$ that will be homotopy equivalent to $\mathfrak{R}_\Delta X$. However, several issues arise. First, these associated convex closed sets must be made smaller so that their overlaps can be controlled. This is achieved by considering points in the Roller boundary versus their classes.

Secondly, the issue of ℓ^2 -visibility, or rather invisibility must be addressed. A Roller class, (respectively a UBS) is ℓ^2 -visible if it is the intersection (respectively union) of the deep half-spaces (respectively hyperplanes) naturally associated to a CAT(0) geodesic ray α . As Fig. 1 shows, some Roller classes are not ℓ^2 -visible, and invisible classes cause headaches when trying to connect this data to the Tits boundary. In particular, if v is an invisible Roller class, then we will have $Q(v) = Q(w)$ for any Roller class $w < v$. To confront this challenge, we have to find a single class which is maximal among the visible Roller classes represented by rays with endpoints in $Q(v)$. Proving the existence of such a maximal class requires finding

a single geodesic ray that is diagonal, in the sense that its convex hull is the union of the convex hulls of two geodesic rays. In particular, we need the following statement, which is of independent interest:

Proposition 7.16 Let α, β be CAT(0) geodesic rays with $\alpha(0) = \beta(0) \in X^{(0)}$. Suppose that $\mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$ is commensurate with a UBS. Then $a = \alpha(\infty)$ and $b = \beta(\infty)$ are joined by a unique geodesic g in $\partial_T X$. Furthermore, any interior point c of g is represented by a CAT(0) geodesic ray γ such that $\mathcal{W}(\gamma) = \mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$.

Next, we consider maps between $\partial_T X$ and $\mathfrak{R}X$ that relate these boundaries. The first of these maps, called $\psi: \partial_T X \rightarrow \mathfrak{R}X$, is fairly easy to describe. A Tits point $a \in \partial_T X$ is represented by a CAT(0) geodesic ray α and the intersection of half-spaces which are “deep” yields a (principal) Roller class $\psi(a)$. See Definition 8.1 and Lemma 8.2. The reverse map $\varphi: \mathfrak{R}X \rightarrow \partial_T X$ is somewhat more delicate: given a Roller class v , we start with the circumcenter $\chi(v) \in Q(v)$ and then perturb $\chi(v)$ to a nearby point $\varphi(v)$ that has slightly better properties. See Definition 8.17 and Proposition 8.18 for details, and note that the perturbation uses Proposition 7.16 in a crucial way. The upshot is that φ is a section of ψ on exactly the ℓ^2 -visible classes: we have $v = \psi(\varphi(v))$ if and only if v is ℓ^2 -visible (Lemma 8.20).

We can now construct nerves and prove results about them. We begin by defining the set $\text{MaxVis}(X)$ of all visible Roller classes that are maximal among all ℓ^2 -visible classes. Then, we construct a simplicial complex \mathcal{N}_T whose vertex set is $\text{MaxVis}(X)$, with simplices corresponding to collections of Roller classes v_i whose Tits boundary realizations $Q(v_i)$ all intersect. We check that this is indeed the nerve of a cover of $\partial_T X$ (Corollary 9.5). Then, we prove:

Theorem 9.17 There is an $\text{Aut}(X)$ -homotopy equivalence from the simplicial complex \mathcal{N}_T to $\partial_T X$.

The proof of Theorem 9.17 requires an open thickening. The sets $Q(v)$ are closed (in fact, compact), and we do not have a version of the Nerve Theorem for closed covers. Thus we thicken up each compact set $Q(v)$ to an open set $U(v)$, in such a way that the intersection pattern of the open cover $\{U(v) : v \in \text{MaxVis}(X)\}$ is the same as that of the closed cover $\{Q(v) : v \in \text{MaxVis}(X)\}$. The thickening procedure involves some delicate CAT(0) geometry; see Proposition 9.14. As a result, we obtain an open cover of $\partial_T X$ whose nerve is isomorphic to the nerve of the closed cover, namely \mathcal{N}_T . Now, the Equivariant Open Nerve Theorem 2.7 gives an $\text{Aut}(X)$ -homotopy equivalence $\mathcal{N}_T \rightarrow \partial_T X$.

There are two reasons why the $\text{Aut}(X)$ -homotopy equivalence in Theorem 9.17 is not $\text{Aut}(X)$ -equivariant. First, the Equivariant Open Nerve Theorem 2.7 does not provide $\text{Aut}(X)$ -equivariance on the nose. Second, while the collection of Tits boundary realizations $Q(v_i)$ is $\text{Aut}(X)$ -equivariant, the perturbed circumcenter map φ might not be. For these reasons, $\text{Aut}(X)$ -homotopy equivalence is the strongest form of invariance that we can guarantee.

In Sect. 10, we use all of the above results to complete the proof of Theorem A. We already have an open cover of $\partial_T X$ whose nerve is \mathcal{N}_T . Working in $\mathfrak{R}_\Delta X$, we focus attention on a subcomplex $\mathfrak{R}_\Delta^{\triangleleft} X$ whose vertex set corresponds to the ℓ^2 -visible Roller classes, and then construct a cover of $\mathfrak{R}_\Delta^{\triangleleft} X$ by finite simplicial complexes $\{\Sigma_v : v \in \text{MaxVis}X\}$. It is not too hard to check that the nerve \mathcal{N}_Δ of this cover is isomorphic to \mathcal{N}_T , which implies that $\mathfrak{R}_\Delta^{\triangleleft} X$ and $\partial_T X$ are $\text{Aut}(X)$ -homotopy equivalent (Proposition 10.5). To complete the construction of an $\text{Aut}(X)$ -homotopy equivalence $\text{T}^{\mathbb{R}}: \mathfrak{R}_\Delta X \rightarrow \partial_T X$, we build an $\text{Aut}(X)$ -deformation retraction $\mathfrak{R}_\Delta^{\triangleleft} X \rightarrow \mathfrak{R}_\Delta X$; see Proposition 10.6. While the construction of this retraction is

Table 1 Table of notation

Symbol	Meaning	Where
X	Cube complex	Section 3.1
d_1	ℓ^1 , combinatorial metric	Section 3.1
d_X	ℓ^2 , CAT(0) metric	Section 3.1
d_X^ρ	Rescaled (cuboid) CAT(0) metric	Definition 4.1
$\mathfrak{H} = \mathfrak{H}(X)$	Set of half-spaces of X	Section 3.2
\mathfrak{H}_x^+	Set of half-spaces containing x	Section 3.3
$\mathcal{W} = \mathcal{W}(X)$	Set of hyperplanes (walls) of X	Section 3.2
$\mathcal{W}(x, y)$	Hyperplanes separating x from y	Section 3.2
$\mathcal{I}(x, y)$	Vertex interval from x to y	Definition 3.10
$m(x, y, z)$	Median in \bar{X}	Equation (3.11)
$\text{CH}(S)$	Cubical convex hull of set S	Definition 3.7
$\mathcal{J}(x, y)$	Cubical interval from x to y	Definition 3.8, 8.5
γ	Geodesic ray in X , either ℓ^1 or ℓ^2	Definition 3.25
$\gamma(\infty)$	Endpoint of γ at infinity	Lemma 3.26
\bar{X}	Roller compactification of X	Definition 3.3
$\partial_R X$	Roller boundary, $\bar{X} \setminus X$	Definition 3.3
$x \sim y$	Finite distance in \bar{X}	Definition 5.1
$\mathfrak{R}X$	Guralnik quotient, $\partial_R X / \sim$	Definition 5.1
$\mathfrak{R}_\Delta X$	Simplicial Roller boundary	Definition 5.7
$\langle\langle S \rangle\rangle$	Inseparable closure of a set of hyperplanes	Section 6.2
\mathcal{U}, \mathcal{V}	Unidirectional boundary sets (UBS)	Definition 6.3
$\mathcal{U} \sim \mathcal{V}$	Finite symmetric difference	Definition 6.4
$\partial_\Delta X$	Simplicial boundary of X	Definition 6.9

mostly combinatorial, it relies on CAT(0) geometry and the perturbed circumcenter map φ at one crucial step.

Finally, the homotopy equivalence $\mathbb{R}^S: \partial_\Delta X \rightarrow \mathfrak{R}_\Delta X$ is morally very similar to the simplicial map $\mathbb{R}_\Delta^U: \partial_\Delta^{\text{UBS}} X \rightarrow \mathfrak{R}_\Delta X$ constructed in Corollary 6.33. To make the argument precise, we construct isomorphic nerves of simplicial covers of $\partial_\Delta X$ and $\mathfrak{R}_\Delta X$ and apply the Equivariant Simplicial Nerve Theorem 2.8 one final time.

1.6 Expository content

Part of our goal in this paper is to provide exposition of varying aspects of cubical theory. CAT(0) cube complexes are ubiquitous objects that have been well-studied in many different guises. Accordingly, there are several different viewpoints, and various important technical statements are stated and proved in a variety of different ways throughout the literature. Therefore, we have endeavored to give a self-contained discussion of CAT(0) cube complexes combining some of these viewpoints.

Also, throughout the paper, we make heavy use of the *nerve theorem*, for both open and simplicial covers. Results of this sort were originally proved by Borsuk [13], and are now widely used in many slightly different forms. The versions in the literature closest to what we

Table 1 continued

Symbol	Meaning	Where
$\bar{Y}_{\mathcal{U}}$	Umbral of a UBS in $\partial_R X$	Definition 6.16
$Y_{\mathcal{U}}$	Principal class of umbral	Lemma 6.18
\mathcal{U}_Y	UBS representing a Roller class Y	Definition 6.21
$\mathcal{W}(\gamma)$	UBS associated to a geodesic γ	Definition 6.23
$\mathcal{UBS}(X)$	Equivalence classes of UBSes	Theorem 6.27
$\partial_T X$	Tits boundary of X	Definition 7.1
D_a	Deep set of half-spaces for $a \in \partial_T X$	Definition 7.5
$\psi(a)$	Roller class associated to $a \in \partial_T X$	Definition 8.1
$Q(y)$	Tits boundary realization of $y \in \partial_R X$	Definition 8.4, 8.5
$Q(v)$	Tits boundary realization of $v \in \mathfrak{R}X$	Definition 8.7
$\chi(v)$	Circumcenter of $Q(v)$	Definition 8.7
$\varphi(v)$	Pseudocenter, perturbed circumcenter	Definition 8.17
M_v	Maximal Roller class in $\psi(Q(v))$	Lemma 8.14
$\text{Vis}(X)$	ℓ^2 -visible Roller classes, $\psi(\partial_T X)$	Definition 8.19
$\text{MaxVis}(X)$	Maximal visible Roller classes	Definition 9.1
$U(v)$	Open neighborhood of $Q(v)$	Proposition 9.14
\mathcal{N}_T	Nerve of cover of $\partial_T X$ by $Q(\text{MaxVis}(X))$	Corollary 9.5
\mathcal{L}_T	Nerve of cover of $\partial_T X$ by $U(\text{MaxVis}(X))$	Theorem 9.17
$\mathfrak{R}_{\Delta}^{\leq} X$	Visible subcomplex of $\mathfrak{R}_{\Delta} X$	Definition 10.1
Σ_v	Subcomplex of $\mathfrak{R}_{\Delta}^{\leq} X$ for $v \in \text{Vis}(X)$	Definition 10.1
\mathcal{N}_{Δ}	Nerve of simplicial cover of $\mathfrak{R}_{\Delta}^{\leq} X$ by Σ_v 's	Definition 10.3

need here are for open covers of paracompact spaces [42] and for (possibly locally infinite) covers of simplicial complexes by subcomplexes [10, 11]. The recent paper [57] contains generalizations of both statements of the nerve theorem. Since we could not find a written account of these results incorporating an additional conclusion about G -equivalences in the presence of a group action, we have given self-contained proofs (based on arguments in [11, 42]) in Sect. 2.

1.7 Section Breakdown

Section 2 establishes some language about simplicial complexes and proves equivariant versions of two nerve theorems. Section 3 is devoted to background on CAT(0) cube complexes and the Roller boundary. In Sect. 4, we discuss CAT(0) cuboid complexes coming from an admissible rescaling of the hyperplanes. Section 5 introduces the simplicial Roller boundary. In Sect. 6, we introduce UBSes and the simplicial boundary, and relate these to Roller classes and the simplicial Roller boundary. In Sect. 7, we introduce the Tits boundary and prove some technical results relating CAT(0) geodesic rays to Roller classes and UBSes. We apply these in Sect. 8 to analyze the realizations of Roller classes in the Tits boundary. These results are in turn used in Sect. 9 to prove that the Tits boundary is homotopy equivalent to a simplicial complex \mathcal{N}_T arising as the nerve of the covering by open sets associated to certain Roller

classes. In Sect. 10, we realize this nerve as the nerve of a covering of the simplicial Roller boundary by subcomplexes associated to Roller classes, and deduce Theorems A and B.

See Table 1 for a summary of the notation used in this paper.

2 Simplicial complexes and nerve theorems

Throughout the paper, we will make use of simplicial complexes and assume that the reader is familiar with these objects. For clarity, we recall the definition and describe two topologies on a simplicial complex. Then, in Sect. 2.1, we prove two group-equivariant homotopy equivalence theorems about nerves of covers.

A k -simplex is the set

$$\sigma = \left\{ a_0e_0 + \cdots + a_k e_k \mid \sum_{i=0}^k a_i = 1 \text{ and } a_i \geq 0 \text{ for all } i \right\},$$

where e_0, \dots, e_k are the standard basis vectors in \mathbb{R}^{k+1} . A *face* of σ is a j -simplex obtained by restricting all but $j + 1$ of the a_i to 0. Note that σ has a CW complex structure where the 0-cells are the 0-dimensional faces and, more generally, the j -cells are the j -dimensional faces.

A *simplicial complex* is a CW complex \mathcal{N} whose closed cells are simplices, such that

- each (closed) simplex is embedded in \mathcal{N} , and
- if σ, τ are simplices, then $\sigma \cap \tau$ is either empty or a face of both σ and τ . In particular, simplices with the same 0-skeleton are equal.

We often refer to 0-simplices of \mathcal{N} as *vertices* and 1-simplices as *edges*.

As a CW complex, \mathcal{N} is endowed with the weak topology:

Definition 2.1 (*Weak topology*) Let \mathcal{N} be a simplicial complex. The *weak topology* \mathcal{T}_w on \mathcal{N} is characterized by the property that a set $C \subset \mathcal{N}$ is closed if and only if $C \cap \sigma$ is closed for every simplex $\sigma \subset \mathcal{N}$.

In some of the arguments in this section, it will be more convenient to work with the metric topology on the simplicial complex \mathcal{N} .

Definition 2.2 (*Metric topology*) Let \mathcal{N} be a simplicial complex. Let V be the real vector space consisting of functions $f: \mathcal{N}^{(0)} \rightarrow \mathbb{R}$ such that $f(v) \neq 0$ for finitely many $v \in \mathcal{N}^{(0)}$. Then there is a canonical inclusion $\mathcal{N}^{(0)} \hookrightarrow V$, where every vertex $v \in \mathcal{N}^{(0)}$ maps to the corresponding Dirac function $\delta_v \in V$. This map extends affinely over simplices endowed with barycentric coordinates to give an inclusion $\mathcal{N} \hookrightarrow V$.

Equip V with the ℓ^2 norm

$$\|f\|^2 = \sum_{v \in \mathcal{N}^{(0)}} f^2(v),$$

which is well-defined since every $f \in V$ is finitely supported. The restriction of the resulting metric topology on V to the subspace \mathcal{N} is the *metric topology on \mathcal{N}* , denoted \mathcal{T}_m .

The metric topology is coarser than the weak topology, hence the identity map $\text{id}_{\mathcal{N}}: (\mathcal{N}, \mathcal{T}_w) \rightarrow (\mathcal{N}, \mathcal{T}_m)$ is always continuous. When a simplicial complex \mathcal{N} is locally infinite, the inverse map $(\mathcal{N}, \mathcal{T}_m) \rightarrow (\mathcal{N}, \mathcal{T}_w)$ is not continuous. However, Dowker proved

that $\text{id}_{\mathcal{N}}$ is a homotopy equivalence even when \mathcal{N} is locally infinite [27]. Since our interest is in the homotopy type of \mathcal{N} , Dowker’s theorem will be very useful.

Convention 2.3 Unless stated otherwise, a simplicial complex \mathcal{N} is presumed to have the weak topology \mathcal{T}_w . The metric topology will only be needed in Lemma 2.6 and Theorem 2.7, and never afterward.

Some of the simplicial complexes used later will arise from partially ordered sets, as follows.

Definition 2.4 (*Simplicial realization*) Given a partially ordered set (P, \leq) , there is a simplicial complex S in which the k -simplices are the $(k + 1)$ -chains in (P, \leq) , and the face relation is determined by containment of chains. We call S the *simplicial realization* of the partially ordered set (P, \leq) .

In several other places, we will work with simplicial complexes arising as *nerves* of coverings of topological spaces:

Definition 2.5 (*Nerve*) Let Y be a topological space and let $\mathcal{U} = \{Y_i\}_{i \in I}$ be a covering of Y , i.e. a family of subsets with $Y = \bigcup_{i \in I} Y_i$. The *nerve* of \mathcal{U} is the simplicial complex \mathcal{N} with a vertex v_i for each Y_i , and with an n -simplex spanned by v_{i_0}, \dots, v_{i_n} whenever $\bigcap_{j=0}^n Y_{i_j} \neq \emptyset$.

Note that we have defined nerves for arbitrary coverings of arbitrary topological spaces. In practice, we will restrict Y and \mathcal{U} : either Y is a paracompact space and \mathcal{U} is an open covering, or Y is itself a simplicial complex and \mathcal{U} is a covering by subcomplexes. These assumptions provide the settings for the nerve theorems, which relate the homotopy type of Y to that of the nerve of \mathcal{U} .

2.1 Equivariant nerve theorems

In this section, we prove two flavors of nerve theorem that are needed in our proofs of homotopy equivalence. Theorem 2.7 is a group-equivariant version of the classical nerve theorem for open coverings, originally due to Borsuk [13]. Similarly, Theorem 2.8 is an equivariant version of the nerve theorem for simplicial complexes, which is due to Björner [11, Theorem 10.6].

We need the following standard fact:

Lemma 2.6 (Recognizing homotopic maps) *Let Y be a topological space and \mathcal{N} a simplicial complex endowed with the metric topology. Let $f_0, f_1 : Y \rightarrow (\mathcal{N}, \mathcal{T}_m)$ be continuous maps. Suppose that, for all $y \in Y$, there exists a closed simplex σ of \mathcal{N} such that $f_0(y), f_1(y) \in \sigma$. Then f_0 and f_1 are homotopic via a straight line homotopy.*

Proof By abuse of notation, we identify \mathcal{N} with the embedding $\mathcal{N} \hookrightarrow V$ described in Definition 2.2. Then $f_0, f_1 : Y \rightarrow V$ are continuous functions. For $y \in Y$ and $t \in [0, 1]$, define the affine combination $f_t(y) = (1 - t)f_0(y) + tf_1(y)$. Then $(y, t) \mapsto f_t(y)$ is a continuous mapping from $Y \times [0, 1]$ to V . Furthermore, for every $y \in Y$, the image $f_t(y) = (1 - t)f_0(y) + tf_1(y)$ belongs to the same simplex $\sigma \subset \mathcal{N}$ that contains $f_0(y)$ and $f_1(y)$. Thus we get a straight line homotopy $f_t : Y \rightarrow \mathcal{N}$. □

Theorem 2.7 (Equivariant open nerve theorem) *Let Y be a paracompact space, on which a group G acts by homeomorphisms. Let I be a set equipped with a left G -action, and let $\mathcal{U} = \{Y_i\}_{i \in I}$ be an open covering of Y with the following properties:*

- $gY_i = Y_{gi}$ for all $g \in G$ and $i \in I$;
- for any finite $F \subset I$, the intersection $\bigcap_{i \in F} Y_i$ is either empty or contractible.

Let \mathcal{N} be the nerve of the covering \mathcal{U} .

Then G acts on \mathcal{N} by simplicial automorphisms, and there is a G -homotopy equivalence $f: Y \rightarrow (\mathcal{N}, \mathcal{T}_m)$. Consequently, there is also a G -homotopy equivalence $Y \rightarrow (\mathcal{N}, \mathcal{T}_w)$.

Proof If the G -action is trivial, this result is the classical nerve theorem, proved for instance in Hatcher [42, Proposition 4G.2, Corollary 4G.3]. We adapt Hatcher’s line of argument, accounting for the G -action where necessary. Until we say otherwise, at the very end of the proof, the nerve \mathcal{N} will be equipped with the metric topology \mathcal{T}_m .

Simplices comprising the nerve: Let \mathcal{F} be the set of finite subsets $F \subset I$ such that $\bigcap_{i \in F} Y_i \neq \emptyset$. Then, by Definition 2.5, every finite set $F \in \mathcal{F}$ can be canonically identified with the set of vertices of a simplex $\sigma(F) \subset \mathcal{N}$. The partial order on \mathcal{F} given by set inclusion corresponds to the face relation on \mathcal{N} . The action of G on \mathcal{U} induces an action on \mathcal{F} , hence a simplicial action on \mathcal{N} .

The space ΔY and its quotient maps: Let $\sigma = \sigma(F)$ be a simplex of \mathcal{N} , corresponding to a finite set $F \in \mathcal{F}$. Let $U_F = \bigcap_{i \in F} Y_i$, a contractible open set, and consider the product $U_F \times \sigma(F) \subset Y \times \mathcal{N}$. Then, define

$$\Delta Y = \bigcup_{F \in \mathcal{F}} (U_F \times \sigma(F)) \subset Y \times \mathcal{N}.$$

The diagonal G -action on $Y \times \mathcal{N}$ induces a G -action on ΔY . The projections from $Y \times \mathcal{N}$ to its factors restrict to G -equivariant projection maps

$$p: \Delta Y \rightarrow Y, \quad q: \Delta Y \rightarrow \mathcal{N}.$$

The homotopy equivalence $\Delta Y \rightarrow \mathcal{N}$: Given a point $x \in \mathcal{N}$ contained in the interior of a simplex $\sigma(F)$, the fiber $q^{-1}(x) = U_F \times \{x\}$ is contractible by hypothesis. Thus, by [42, Proposition 4G.1], the projection $q: \Delta Y \rightarrow \mathcal{N}$ is a homotopy equivalence.

The fiber $p^{-1}(y)$: Fix $y \in Y$, and consider the fiber $p^{-1}(y)$. Let $(y, x) \in p^{-1}(y) \subset Y \times \mathcal{N}$. Let $\sigma(F)$ be a simplex of \mathcal{N} containing $x = q(y, x)$. The vertices of σ correspond to the elements of F , hence to subsets Y_i for $i \in F$. We endow $\sigma(F)$ with barycentric coordinates, so that $x = \sum_{i \in F} \beta_i Y_i$ where $\sum \beta_i = 1$. Thus

$$p^{-1}(y) = \left\{ \left(y, \sum_{i \in F} \beta_i Y_i \right) \mid y \in U_F, F \in \mathcal{F} \right\} = \{y\} \times \Sigma_y$$

for a simplicial subcomplex $\Sigma_y \subset \mathcal{N}$. Since the projection $p: \Delta Y \rightarrow Y$ is G -equivariant, we have

$$\{gy\} \times g\Sigma_y = g(\{y\} \times \Sigma_y) = gp^{-1}(y) = p^{-1}(gy) = \{gy\} \times \Sigma_{gy},$$

and in particular $g\Sigma_y = \Sigma_{gy}$.

For any pair of points $x, x' \in \Sigma_y$, belonging to simplices $\sigma(F), \sigma(F')$ of \mathcal{N} , we have $y \in U_F \cap U_{F'} = U_{F \cup F'}$. Thus there is also a simplex $\sigma(F \cup F')$ containing both $\sigma(F)$ and $\sigma(F')$ as faces. It follows that x, x' are connected by an affine line segment in the common simplex $\sigma(F \cup F')$. In other words, Σ_y is convex.

The section $s : Y \rightarrow \Delta Y$: Since Y is paracompact, there is a partition of unity $\{\phi_\alpha\}_{\alpha \in A}$ subordinate to $\mathcal{U} = \{Y_i\}_{i \in I}$. That is, for each $\alpha \in A$, we have $\text{supp}(\phi_\alpha) \subset Y_{i(\alpha)}$ for some $i(\alpha) \in I$. Given $y \in Y$, let $A(y) = \{\alpha \in A : y \in \text{supp}(\phi_\alpha)\}$, a finite set. Then, define

$$f(y) = \sum_{\alpha \in A(y)} \phi_\alpha(y) Y_{i(\alpha)} \in \Sigma_y \quad \text{and} \quad s(y) = (y, f(y)) \in p^{-1}(y).$$

We will eventually show that $f : Y \rightarrow \mathcal{N}$ is the homotopy equivalence claimed in the theorem statement. For now, we focus on s .

Checking that s is a homotopy inverse of p : Since $s(y) \in p^{-1}(y)$, we have $p \circ s = id_Y$ by construction. We claim that the other composition $s \circ p$ is homotopic to $id_{\Delta Y}$. To prove this, consider the following pair of functions from $\Delta Y \rightarrow \mathcal{N}$:

$$q_0(y, x) = q(y, x) = x, \quad q_1(y, x) = f(y).$$

For every pair $(y, x) \in \Delta Y$, the image points $x = q_0(y, x)$ and $x' = q_1(y, x) = f(y)$ are both contained in Σ_y . Above, we have checked both image points x, x' are contained in a common simplex $\sigma(F \cup F') \subset \Sigma_y \subset \mathcal{N}$. Recall that we are working with the metric topology on \mathcal{N} . Thus, by Lemma 2.6, q_0 is homotopic to q_1 via a straight line homotopy q_t . For every y , this straight-line homotopy runs through Σ_y .

We can now define a homotopy $h_t : \Delta Y \rightarrow Y \times \mathcal{N}$ as follows:

$$h_t(y, x) = (y, q_t(y, x)).$$

Observe that $p \circ h_t$ is continuous because $p \circ h_t(y, x) = y$, and $q \circ h_t$ is continuous because $q \circ h_t(y, x) = q_t(y, x)$. Thus, by the universal property of the product topology, h_t is continuous. Restricting t to be 0 or 1 gives

$$h_0(y, x) = (y, q_0(y, x)) = (y, x), \quad h_1(y, x) = (y, q_1(y, x)) = (y, f(y)) = s \circ p(y, x).$$

In fact, for every (y, x) , the path $t \mapsto q_t(y, x)$ has image in Σ_y , hence $h_t(y, x)$ has image in $\{y\} \times \Sigma_y \subset \Delta Y$. Thus h_t is a homotopy from $s \circ p$ to $id_{\Delta Y}$, and s is a homotopy inverse of p .

A $G \sim$ homotopy equivalence, for both topologies on \mathcal{N} : Thus far, we have homotopy equivalences $s : Y \rightarrow \Delta Y$ and $q : \Delta Y \rightarrow \mathcal{N}$. Hence the composition $f = q \circ s : Y \rightarrow \mathcal{N}$ is a homotopy equivalence as well.

It remains to check that f is a $G \sim$ homotopy equivalence. We saw above that s is a homotopy inverse for p ; in particular p is an isomorphism in the homotopy category of spaces, with inverse s . Moreover, each $g \in G$ acts as a homeomorphism, and in particular an isomorphism in the homotopy category, on both Y and ΔY . By construction, we have $g \circ p = p \circ g$. So, letting \simeq denote homotopy of maps, and using that $s \circ p \simeq Id_{\Delta Y}$ and $p \circ s \simeq Id_Y$, we have

$$s \circ g \simeq s \circ g \circ (p \circ s) \simeq s \circ (p \circ g) \circ s \simeq g \circ s,$$

which is to say that s is a $G \sim$ homotopy equivalence. Since q is a G -equivariant homotopy equivalence, it is a $G \sim$ homotopy equivalence, and hence $f = q \circ s : Y \rightarrow \mathcal{N}$ is as well. Since $id_{\mathcal{N}} : (\mathcal{N}, \mathcal{T}_w) \rightarrow (\mathcal{N}, \mathcal{T}_m)$ is a G -equivariant homotopy equivalence by [27], f is thus a $G \sim$ homotopy equivalence for either topology on \mathcal{N} . □

We are now finished with having the metric topology on our nerves, and work only with the weak topology in the remainder of the paper.

The following result is stated (without the group action) as [11, Theorem 10.6]. The proof given there assumes that $\mathcal{V} = \{\Sigma_i\}$ is a locally finite cover, and a proof in full generality is

given in [10, Lemma 1.1]. (See [57] for a more general result.) Since we need to adapt the proof slightly to account for the group action, we write down a proof using Theorem 2.7, without any assumptions on local finiteness of the cover $\{\Sigma_i\}$.

Theorem 2.8 (Equivariant simplicial nerve theorem) *Let Y be a simplicial complex, and let G be a group acting on Y by simplicial automorphisms. Let I be a set equipped with a left G -action, and let $\mathcal{V} = \{\Sigma_i\}_{i \in I}$ be a covering of Y by subcomplexes, with the following properties:*

- $g\Sigma_i = \Sigma_{gi}$ for all $g \in G$ and $i \in I$;
- for any finite $F \subset I$, the intersection $\bigcap_{i \in F} \Sigma_i$ is either empty or contractible.

Let \mathcal{N} be the nerve of the covering \mathcal{V} . Then there is a G -homotopy equivalence $f : Y \rightarrow (\mathcal{N}, \mathcal{T}_w)$.

Proof Let Y' be the barycentric subdivision of Y . We think of Y and Y' as the same underlying topological space, with different simplicial structures.

Every subcomplex $\Sigma \subset Y$ can be viewed as a subcomplex of Y' . Indeed, $\Sigma \subset Y$ is a full subcomplex of Y' , in the following sense: for every simplex $\tau \subset Y'$ whose vertices belong to Σ , it follows that $\tau \subset \Sigma$. This holds because the vertices of τ correspond to simplices of Y ; if these vertices belong to Σ , then so do the corresponding cells of Y , hence $\tau \subset \Sigma$.

Complementary complexes and open neighborhoods: For a subcomplex $\Sigma \subset Y$, define the complementary complex T_Σ to be the union of all closed simplices of Y' disjoint from Σ . Then T_Σ is also a full subcomplex of Y' . Indeed, consider a simplex $\tau \subset Y'$ whose vertices are disjoint from Σ . Then $\tau \cap \Sigma$ cannot contain any faces of τ , hence $\tau \cap \Sigma = \emptyset$. Since T_Σ is a subcomplex of Y' , it is closed (in the weak topology).

Given a subcomplex $\Sigma \subset Y$, we define an open neighborhood $U_\Sigma = Y \setminus T_\Sigma$. Thus U_Σ is the union of all the open simplices of Y' whose closures intersect Σ .

For every subcomplex Σ_i , where $i \in I$, we write $U_i = U_{\Sigma_i}$. Observe that $U_{gi} = gU_i$ for all $g \in G$ and $i \in I$. We will prove the theorem by replacing the subcomplex cover $\mathcal{V} = \{\Sigma_i\}_{i \in I}$ by the open cover $\mathcal{U} = \{U_i\}_{i \in I}$.

Same nerve: Let Σ_α and Σ_β be a pair of subcomplexes of Y , not necessarily belonging to \mathcal{V} . We claim that $\Sigma_\alpha \cap \Sigma_\beta = \emptyset$ if and only if $U_{\Sigma_\alpha} \cap U_{\Sigma_\beta} = \emptyset$. The “if” direction is obvious. For the “only if” direction, suppose that $\Sigma_\alpha \cap \Sigma_\beta = \emptyset$. Then $T_{\Sigma_\alpha} \cup T_{\Sigma_\beta} = T_{\Sigma_\alpha \cap \Sigma_\beta} = T_\emptyset = Y'$, hence $U_{\Sigma_\alpha} \cap U_{\Sigma_\beta} = \emptyset$, as desired.

Now, let $F \subset I$ be a finite set. By induction on $|F|$, combined with the argument of the above paragraph, we see that $\bigcap_{i \in F} \Sigma_i = \emptyset$ if and only if $\bigcap_{i \in F} U_{\Sigma_i} = \emptyset$. Hence the open cover $\mathcal{U} = \{U_i\}_{i \in I}$ and the subcomplex cover $\mathcal{V} = \{\Sigma_i\}_{i \in I}$ have the same nerve \mathcal{N} .

Same homotopy type: For every subcomplex $\Sigma \subset Y$, we claim that the open neighborhood U_Σ deformation retracts to Σ . This is a standard fact about simplicial complexes, proved e.g. in Munkres [51, Lemma 70.1]. The proof constructs a straight-line homotopy in every simplex $\sigma \subset Y'$ that belongs to neither Σ nor T_Σ . The proof uses the fullness of Σ and T_Σ in Y' , but does not depend on any finiteness properties of these complexes.

Now, let $F \subset I$ be a finite set such that $\bigcap_{i \in F} \Sigma_i \neq \emptyset$. Define $\Sigma_F = \bigcap_{i \in F} \Sigma_i$, and recall that by hypothesis, Σ_F is contractible. Since $U_{\Sigma_F} = \bigcap_{i \in F} U_{\Sigma_i}$, and U_{Σ_F} deformation retracts to Σ_F , it follows that $\bigcap_{i \in F} U_i = \bigcap_{i \in F} U_{\Sigma_i}$ is also contractible.

Conclusion: By a theorem of Bourgin [14], the simplicial complex Y is paracompact (with the weak topology \mathcal{T}_w). We have shown that the open cover $\mathcal{U} = \{U_i\}_{i \in I}$ and the subcomplex cover $\mathcal{V} = \{\Sigma_i\}_{i \in I}$ have the same nerve \mathcal{N} . We have checked that $U_{gi} = gU_i$ for all $g \in G$ and $i \in I$. Furthermore, for every finite set $F \subset I$, the intersection $\bigcap_{i \in F} U_i$

is either empty or contractible. Thus, by Theorem 2.7, we have a $G\sim$ homotopy equivalence $f: Y \rightarrow (\mathcal{N}, \mathcal{T}_w)$. □

Remark 2.9 The proof of Theorem 2.7 yields a slightly stronger conclusion than $G\sim$ homotopy equivalence. Indeed, we produced G -equivariant homotopy equivalences $\Delta Y \rightarrow Y$ and $\Delta Y \rightarrow \mathcal{N}$, and $G\sim$ homotopy equivalences between Y and \mathcal{N} were then obtained by composing one with a homotopy inverse of the other. The same is true for Theorem 2.8, since it is proved by applying Theorem 2.7 to a G -invariant open cover. Hence, under the hypotheses of either theorem, we have actually shown that Y and \mathcal{N} are Borel $G\sim$ homotopy equivalent (see Remark 1.2).

3 Cube complexes and the Roller boundary

This section recalls some background about CAT(0) cube complexes and their Roller boundaries. Many facts about CAT(0) cube complexes, median graphs, wallspaces, and related structures occur in various places in the literature, in many equivalent guises. We have endeavored to connect some of the perspectives, in part because we will need to use these connections in the sequel. We refer the reader to [60] and [62] for more background.

3.1 CAT(0) cube complexes and metrics

A *cube* is a copy of a Euclidean unit cube $[-\frac{1}{2}, \frac{1}{2}]^n$ for $0 \leq n < \infty$. A *face* of $[-\frac{1}{2}, \frac{1}{2}]^n$ is a subspace obtained by restricting some of the coordinates to $\pm\frac{1}{2}$. A *cube complex* is a CW complex whose cells are cubes and whose attaching maps restrict to isometries on faces. A cube complex with embedded cubes is *nonpositively curved* if, for all vertices v and all edges e_1, \dots, e_k incident to v such that e_i, e_j span a 2-cube for all $i \neq j$, we have that e_1, \dots, e_k span a unique k -cube. If X is nonpositively-curved and simply connected, then X is a CAT(0) cube complex. The *dimension* of X , denoted $\dim X$, is the supremum of the dimensions of cubes of X .

Throughout this paper, X will denote a finite-dimensional CAT(0) cube complex. We let $\text{Aut}(X)$ denote the group of cellular automorphisms of X . We do not assume that X is locally finite.

The cube complex X carries a metric d_X such that (X, d_X) is a CAT(0) space, the restriction of d_X to each cube is the Euclidean metric on a unit cube, and each cube is geodesically convex (see [7, 36]). We refer to d_X as the ℓ^2 metric or the CAT(0) metric on X .

One can view $X^{(1)}$ as a graph whose vertices are the 0-cubes and whose edges are the 1-cubes. (We often refer to 0-cubes as vertices of X , and 1-cubes as edges.) We let d_1 denote the ℓ^1 metric on the vertex set $X^{(0)}$, which is the restriction of the usual graph metric on $X^{(1)}$. We refer to d_1 as the *combinatorial metric* on $X^{(0)}$. A *combinatorial geodesic* between $x, y \in X^{(0)}$ is an edge path in $X^{(1)}$ that realizes $d_1(x, y)$.

3.2 Hyperplanes, half-spaces, crossing, and separation

Let $[-\frac{1}{2}, \frac{1}{2}]^n$ be an n -dimensional cube. A *midcube* of $[-\frac{1}{2}, \frac{1}{2}]^n$ is the subset obtained by restricting one coordinate to be 0.

A *hyperplane* of a CAT(0) cube complex X is a connected subset whose intersection with each cube is either a midcube of that cube, or empty. The open $1/2$ -neighborhood of

a hyperplane \hat{h} in the metric d_X is denoted $N(\hat{h})$ and is called the *open hyperplane carrier* of \hat{h} . It is known that every open carrier $N(\hat{h})$ is geodesically convex in d_X . Furthermore, $X \setminus N(\hat{h})$ consists of two connected components, each of which is also convex [60, Theorem 4.10].

A component of $X \setminus N(\hat{h})$ is called a *CAT(0) half-space*. The intersection of one of these components with the vertex set $X^{(0)}$ is called a *vertex half-space*.

The two vertex half-spaces associated to a hyperplane \hat{h} are denoted h, h^* . Note that $\hat{h}^* = \hat{h}$. Given a vertex half-space h , the corresponding CAT(0) half-space is the union of all cubes of X whose vertices lie in h . We will use the unmodified term *half-space* to mean either a CAT(0) half-space or a vertex half-space when the meaning is clear from context.

The collection of all vertex half-spaces is denoted by \mathfrak{H} , or $\mathfrak{H}(X)$ if we wish to specify the space X . If $h \in \mathfrak{H}$, then $h^* = X^{(0)} \setminus h$ is exactly the complementary half-space associated to the same hyperplane \hat{h} .

The collection of all hyperplanes of X is denoted \mathcal{W} (for “walls”). Generally speaking, calligraphic letters denote collections of hyperplanes, while gothic letters denote collections of half-spaces. There is a two-to-one map $\mathfrak{H} \rightarrow \mathcal{W}$, namely $h \mapsto \hat{h}$. Given a subset $\mathcal{S} \subset \mathcal{W}$, an *orientation* on \mathcal{S} is a choice of a lift $\mathcal{S} \rightarrow \mathfrak{H}$.

A pair of half-spaces $h, k \in \mathfrak{H}$ are called *transverse* (denoted $h \pitchfork k$) if the four intersections $h \cap k, h \cap k^*, h^* \cap k$ and $h^* \cap k^*$ are nonempty. In terms of the underlying hyperplanes \hat{h} and \hat{k} , transversality is equivalent to the condition that $\hat{h} \neq \hat{k}$ and $\hat{h} \cap \hat{k} \neq \emptyset$. In this case, we also write $\hat{h} \pitchfork \hat{k}$. We sometimes say that transverse hyperplanes *cross*.

Given subsets $A, B \subset X$, we say that \hat{h} *separates* A, B if A is contained in one CAT(0) half-space associated to \hat{h} , and B is contained in the other. We will usually be interested in situations where A, B are subcomplexes, sets of vertices, or hyperplanes, in which case this notion is equivalent to another notion of separation: namely that A, B lie in distinct components of $X \setminus \hat{h}$. Later in the paper, we will occasionally use the latter notion when we need to talk about points in $N(\hat{h})$ being separated by \hat{h} . (Elsewhere in the literature, e.g. [60], the half-spaces associated to \hat{h} are defined to be the components of $X \setminus \hat{h}$, or sometimes their closures. This small difference in viewpoint is usually irrelevant.)

If e is an edge, i.e. a 1-cube, of X , then there is a unique hyperplane \hat{h} separating the two vertices of e . We say that \hat{h} is *dual to* e , and vice versa. Note that \hat{h} is the unique hyperplane intersecting e . More generally, if $k \geq 0$ and c is a k -cube, then there are exactly k hyperplanes intersecting c . These hyperplanes pairwise cross, and their intersection contains the barycenter of c . We say that c is *dual to* this family of hyperplanes. Any set of k pairwise-crossing hyperplanes is dual to at least one k -cube. In particular, $\dim X$ bounds the cardinality of any set of pairwise crossing hyperplanes in X .

For any points $x, y \in X$, define $\mathcal{W}(x, y)$ to be the set of hyperplanes separating x from y . If $x, y \in X^{(0)}$, then an edge path γ from x to y is a geodesic in $X^{(1)}$ if and only if γ never crosses the same hyperplane twice. Consequently,

$$d_1(x, y) = |\mathcal{W}(x, y)|.$$

In view of this, the following standard lemma relates the metrics d_X and d_1 . See Caprace and Sageev [25, Lemma 2.2] or Hagen [41, Lemma 3.6] for proofs.

Lemma 3.1 *There are constants $\lambda_0 \geq 1, \lambda_1 \geq 0$, depending on $\dim X$, such that the following holds. For any pair of points $x, y \in X$,*

$$\frac{1}{\lambda_0} d_X(x, y) - \lambda_1 \leq |\mathcal{W}(x, y)| \leq \lambda_0 d_X(x, y) + \lambda_1.$$

For an arbitrary set $B \subset X$, let $\mathcal{W}(B)$ be the union of all $\mathcal{W}(x, y)$ for all $x, y \in B$. In practice, we will often be interested in the following special cases. First, if B is a hyperplane, then $\mathcal{W}(B)$ coincides with the set of hyperplanes \hat{h} that cross B . If B is a convex subcomplex (see Definition 3.7), then $\mathcal{W}(B)$ coincides with the set of \hat{h} with $\hat{h} \cap B \neq \emptyset$. If B is a combinatorial geodesic in $X^{(1)}$, then, because of the above characterization of d_1 , the set $\mathcal{W}(B)$ is the set of hyperplanes dual to edges of B , or equivalently the set of hyperplanes intersecting B , or equivalently the set of hyperplanes separating the endpoints of B .

Remark 3.2 One can extend the ℓ^1 metric on $X^{(0)}$ to all of X as follows. On a single cube c , let $d_1|_{c \times c} \rightarrow \mathbb{R}$ be the standard ℓ^1 metric. This can be extended to a path-metric on all of X by concatenating paths in individual cubes. Miesch [49] showed that this procedure gives a geodesic path metric on X that extends the graph metric on $X^{(1)}$. With this extended definition, d_1 becomes bilipschitz to d_X , with the Lipschitz constant λ_0 depending only on $\dim(X)$.

3.3 The Roller boundary $\partial_R X$

Next, we define the *Roller boundary* and *Roller compactification* of X .

Every vertex $v \in X^{(0)}$ defines the collection $\mathfrak{H}_v^+ = \{h \in \mathfrak{H} : v \in h\}$ of half-spaces containing v . Going in the opposite direction, the collection \mathfrak{H}_v^+ uniquely determines v :

$$\bigcap_{h \in \mathfrak{H}_v^+} h = \{v\}.$$

We endow the set $2^{\mathfrak{H}}$ with the product topology, which is compact by Tychonoff’s theorem. Recall that a basis for this topology consists of *cylinder sets* defined by the property that finitely many coordinates (i.e. finitely many half-spaces) take a specified value (0 or 1). Cylinder sets are both open and closed.

There is a continuous one-to-one map $X^{(0)} \hookrightarrow 2^{\mathfrak{H}}$ defined by $v \mapsto \mathfrak{H}_v^+$. (In fact, $X^{(0)}$ is homeomorphic to its image if and only if X is locally compact. We will not need this fact.)

Definition 3.3 (*Roller compactification, Roller boundary*) The *Roller Compactification* of $X^{(0)}$, denoted by \bar{X} or $\bar{X}(\mathfrak{H})$, is the closure of the image of $X^{(0)}$ in $2^{\mathfrak{H}}$. The *Roller Boundary* of X is $\partial_R X = \bar{X} \setminus X^{(0)}$.

Recall that $\text{Aut}(X)$ is the group of cubical automorphisms of X . We observe that the action of $\text{Aut}(X)$ extends to a continuous action on \bar{X} and therefore on $\partial_R X$.

Definition 3.4 (*Extended half-spaces*) Let $h \in \mathfrak{H}(X)$. The *extension* of h to \bar{X} is defined to be the intersection in $2^{\mathfrak{H}}$ between \bar{X} and the cylinder set corresponding to h . It is straightforward to check that the extensions of h and h^* form a partition of \bar{X} . We think of the extensions of h and h^* as complementary (vertex) half-spaces in \bar{X} . By a slight abuse of notation, we use the same symbol h to refer to both a half-space in X and its extension in \bar{X} .

By the above discussion of the topology on $2^{\mathfrak{H}}$, the basic open sets in \bar{X} are intersections between \bar{X} and cylinder sets. Therefore, every basis set is a finite intersections of (extended) half-spaces.

The duality between points and half-spaces in X , described above, extends to all of $\bar{X} = X^{(0)} \cup \partial_R X$. Let $y \in \bar{X}$, and let \mathfrak{H}_y^+ be the set of extended half-spaces that contain y . Then

y and \mathfrak{H}_y^+ determine one another:

$$\bigcap_{h \in \mathfrak{H}_y^+} h = \{y\}. \tag{3.5}$$

Chatterji and Niblo [24], and independently Nica [52], extended Sageev’s construction [60] to prove the following.

Theorem 3.6 *Let $\mathfrak{H}' \subset \mathfrak{H}(X)$ be an involution-invariant collection of half-spaces. Then \mathfrak{H}' determines a CAT(0) cube complex $X(\mathfrak{H}')$ and the Roller compactification $\overline{X}(\mathfrak{H}')$. In particular, $X(\mathfrak{H}(X)) = X$.*

If $x \in X^{(0)}$, the corresponding set $\mathfrak{H}_x^+ \in 2^{\mathfrak{H}}$ satisfies the *descending chain condition*: any decreasing sequence of elements of \mathfrak{H}_x^+ is eventually constant. On the other hand, the collection of half-spaces \mathfrak{H}_y^+ corresponding to a boundary point $y \in \partial_R X$ will fail the descending chain condition. Namely, for $y \in \overline{X}$ we have that $y \in \partial_R X$ if and only if there exists a sequence $\{h_n\}_{n \in \mathbb{N}} \subset \mathfrak{H}_y^+$ such that $h_{n+1} \subsetneq h_n$ for each $n \in \mathbb{N}$.

3.4 Convexity, intervals, and medians

Here, we discuss the related notions of convexity and intervals in X and \overline{X} . Since we are interested in two different metrics on X , there are several distinct definitions of convexity (ℓ^1 geodesic convexity, cubical convexity, interval convexity) that turn out to be equivalent for cubical subcomplexes of X .

We then introduce a median operation on \overline{X} and its restriction to X .

Definition 3.7 (*Convexity in $X^{(0)}$, convex subcomplexes*) A set $A \subset X^{(0)}$ is called *vertex-convex* if it is the intersection of vertex half-spaces in X . For a set $S \subset X^{(0)}$, the *vertex convex hull* of S is the intersection of all vertex-convex sets containing S , or equivalently the intersection of all vertex half-spaces containing S .

A subcomplex $C \subset X$ is called *cubically convex* if it is the intersection of CAT(0) half-spaces. For a set $S \subset X$, the *cubical convex hull*, denoted $\text{CH}(S)$, is the intersection of all cubically convex sets containing S , or equivalently the intersection of all CAT(0) half-spaces containing S . We remark that for subcomplexes of X , this definition coincides with geodesic convexity in the CAT(0) metric d_X ; see [38, Remark 2.10] and [60, Theorem 4.10].

Observe that if C is a full subcomplex of X (meaning, C contains a cube c whenever it contains the 0-skeleton of c), then C is cubically convex whenever $C^{(0)}$ is vertex-convex. For this reason, we can use the two notions of convexity interchangeably when referring to full subcomplexes.

Definition 3.8 (*Intervals in X*) Given $x, y \in X^{(0)}$, define the (*vertex*) *interval* $\mathcal{I}(x, y)$ to be vertex convex hull of $\{x, y\}$. Define the *cubical interval* $\mathcal{J}(x, y)$ to be the cubical convex hull of $\{x, y\}$.

Observe that $\mathcal{J}(x, y)$ is a full, convex subcomplex, and is the union of all the cubes whose vertex sets are contained in $\mathcal{I}(x, y)$. In addition, $\mathcal{J}^{(1)}(x, y)$ is the union of all of the combinatorial geodesics from x to y . In fact, intervals can be used to characterize convexity as follows:

Lemma 3.9 *A set $A \subset X^{(0)}$ is vertex-convex if and only if the following holds: for all $x, y \in A$, the interval $\mathcal{I}(x, y) \subset A$.*

A subcomplex $C \subset X$ is cubically convex if and only if the following holds: for all $x, y \in C^{(0)}$, the cubical interval $\mathcal{J}(x, y) \subset C$.

Proof Suppose C is a cubically convex subcomplex and $x, y \in C$. Since $\mathcal{J}(x, y)$ is the intersection of all convex subcomplexes containing x, y , and C is one such set, we have $\mathcal{J}(x, y) \subset C$.

Toward the converse, suppose that C is interval-convex, meaning $\mathcal{J}(x, y) \subset C$ for every pair x, y . We wish to prove that C is convex.

We claim the following: for every hyperplane \hat{h} that intersects C and is dual to an edge e with endpoints x, y such that $x \in C$, the other endpoint y belongs to C as well. This can be shown as follows. Since \hat{h} intersects C , there is an edge $e' \subset C$ that is dual to \hat{h} . Let x', y' be the endpoints of e' , such that x and x' are on the same side of \hat{h} . Then $\mathcal{W}(x, y') = \mathcal{W}(x, x') \sqcup \hat{h} = \hat{h} \sqcup \mathcal{W}(y, y')$. In particular, there exists a combinatorial geodesic from x to y' whose initial edge is e . Since C is interval-convex, this proves the claim.

We can now show that C is cubically convex. Suppose $z \in X^{(0)} \setminus C$. Let $\gamma \subset X^{(0)}$ be a shortest combinatorial geodesic from z to C , and let x be the terminus of γ , and let e be the last edge of γ . Since the next-to-last vertex of γ is not in C , the above claim implies that the hyperplane \hat{h} dual to e is disjoint from C , and separates C from z . Thus C is contained in a CAT(0) half-space disjoint from z . Since z was arbitrary, it follows that C is cubically convex.

The proof for vertex-convex subsets is identical, up to replacing all the appropriate sets by their 0-skeleta. □

Our next goal is to extend the notion of convexity to \overline{X} . It turns out that in \overline{X} , the correct definition of convexity begins with intervals. After developing some definitions and tools, we will eventually prove a generalization of Lemma 3.9 to \overline{X} , using a very different argument.

Definition 3.10 (*Intervals and convexity in \overline{X}*) Given $x, y \in \overline{X}$, define the *interval* $\mathcal{I}(x, y)$ to be the intersection of all (extended, vertex) half-spaces that contain both x and y . Equivalently,

$$\mathcal{I}(x, y) = \{z \in \overline{X} : \mathfrak{H}_z^+ \supset (\mathfrak{H}_x^+ \cap \mathfrak{H}_y^+)\}.$$

A nonempty set $A \subset \overline{X}$ is called *convex* if it contains $\mathcal{I}(x, y)$ for all $x, y \in A$.

As a very particular case, half-spaces and hence also their intersections are convex in \overline{X} .

Observe that for $x, y \in X^{(0)}$, the interval $\mathcal{I}(x, y)$ will in fact be contained in $X^{(0)}$, and coincides with the previous definition of $\mathcal{I}(x, y)$. Consequently, $X^{(0)}$ is convex in \overline{X} . By Lemma 3.9, any vertex-convex subset of $X^{(0)}$ is also convex in \overline{X} .

Since $\overline{X} \subset 2^{\mathfrak{H}(X)}$ is by definition a set of vertices, we do not define convex subcomplexes of \overline{X} .

As in [58], the fact that \overline{X} is a *median space* is captured by the following property: for every $x, y, z \in \overline{X}$ there is a unique point $m = m(x, y, z) \in \overline{X}$ such that

$$\{m\} = \mathcal{I}(x, y) \cap \mathcal{I}(y, z) \cap \mathcal{I}(x, z). \tag{3.11}$$

This unique point $m = m(x, y, z)$ is called the *median* of x, y , and z . We will not need the general definition of an (extended) median metric space; the only property of medians that we will need is that $m : \overline{X} \times \overline{X} \times \overline{X} \rightarrow \overline{X}$ satisfies (3.11). In terms of half-spaces, we have:

$$\mathfrak{H}_m^+ = (\mathfrak{H}_x^+ \cap \mathfrak{H}_y^+) \cup (\mathfrak{H}_y^+ \cap \mathfrak{H}_z^+) \cup (\mathfrak{H}_x^+ \cap \mathfrak{H}_z^+).$$

Remark 3.12 (Median in $X^{(0)}$) Since intervals between points in $X^{(0)}$ are contained in $X^{(0)}$, we see that $m(x, y, z) \in X^{(0)}$ if $x, y, z \in X^{(0)}$. So, the median m on \overline{X} restricts to a median m on $X^{(0)}$. The point $m(x, y, z)$ is the unique vertex m for which any two of x, y, z can be joined by a geodesic in $X^{(1)}$ passing through m . In other words, we also get an analogue of (3.11) with $\mathcal{I}(\cdot, \cdot)$ replaced by $\mathcal{J}(\cdot, \cdot)$:

$$\{m\} = \mathcal{J}(x, y) \cap \mathcal{J}(y, z) \cap \mathcal{J}(x, z) \quad \text{for } x, y, z \in X^{(0)}.$$

In fact, a graph with this property—called a *median graph*—is always the 1-skeleton of a uniquely determined CAT(0) cube complex [20].

3.5 Lifting decompositions and convexity

A set of half-spaces $\mathfrak{s} \subset \mathfrak{H}$ is called *consistent* if the following two conditions hold: if $h \in \mathfrak{s}$ then $h^* \notin \mathfrak{s}$, and if $k \supset h \in \mathfrak{s}$ then $k \in \mathfrak{s}$. Given a subset $\mathfrak{H}' \subset \mathfrak{H}$, a *lifting decomposition* for \mathfrak{H}' is a consistent set of half-spaces $\mathfrak{s} \subset \mathfrak{H} \setminus \mathfrak{H}'$ such that $\mathfrak{H} = \mathfrak{H}' \sqcup \mathfrak{s} \sqcup \mathfrak{s}^*$. Lifting decompositions do not necessarily exist, and are not necessarily unique.

Lifting decompositions naturally occur in the following way. Consider a set $A \subset \overline{X}$. Analogous to \mathfrak{H}_A^+ , we define a set $\mathfrak{H}_A^+ = \{h \in \mathfrak{H} : A \subset h\}$ of half-spaces that contain A , and observe that \mathfrak{H}_A^+ is consistent. Define $\mathfrak{H}_A^- = (\mathfrak{H}_A^+)^*$ to be the set of half-spaces disjoint from A , and finally a set

$$\mathfrak{H}_A = \{h \in \mathfrak{H} : A \cap h \neq \emptyset, A \cap h^* \neq \emptyset\}$$

of half-spaces that cut A . Then we get a lifting decomposition $\mathfrak{H} = \mathfrak{H}_A \sqcup \mathfrak{H}_A^+ \sqcup \mathfrak{H}_A^-$, where $\mathfrak{s} = \mathfrak{H}_A^+$.

Remark 3.13 Observe that if $A \subset B$, then $\mathfrak{H}_B^+ \subset \mathfrak{H}_A^+$.

Given points $x, y \in \overline{X}$, let $\mathfrak{H}(x, y) = \mathfrak{H}_x^+ \Delta \mathfrak{H}_y^+$ denote the set of half-spaces that separate x from y . More generally, given two disjoint convex sets $A, B \subset \overline{X}$, let $\mathfrak{H}(A, B) = \mathfrak{H}_A^+ \Delta \mathfrak{H}_B^+$ denote the set of half-spaces that separate A from B .

In an analogous fashion, we generalize the definition of $\mathcal{W}(x, y)$ from Sect. 3.2 to convex subsets $A, B \subset \overline{X}$. For subsets of this form, $\mathcal{W}(A, B) = \widehat{\mathfrak{H}(A, B)}$ consists of all the hyperplanes that separate A from B . We also extend the ℓ^1 metric on $X^{(0)}$ to \overline{X} , where $d_1(x, y) = |\mathcal{W}(x, y)| = \frac{1}{2}|\mathfrak{H}_x^+ \Delta \mathfrak{H}_y^+|$ is allowed to take the value ∞ .

The following result says that lifting decompositions are in one-to-one correspondence with Roller-closed subcomplexes of \overline{X} .

Proposition 3.14 *The following are true:*

- (1) *Suppose that $\mathfrak{H}' \subset \mathfrak{H}(X)$. If there exists a lifting decomposition \mathfrak{s} for \mathfrak{H}' , then there is a d_1 -isometric embedding $\overline{X}(\mathfrak{H}') \hookrightarrow \overline{X}$ induced from the map $2^{\mathfrak{H}'} \hookrightarrow 2^{\mathfrak{H}(X)}$ where $U \mapsto U \sqcup \mathfrak{s}$. The image of this embedding is*

$$\bigcap_{h \in \mathfrak{s}} h \subset \overline{X}.$$

- (2) *Similarly, if $\mathfrak{s} \subset \mathfrak{H}(X)$ is a consistent set of half-spaces, then, setting $\mathfrak{H}_{\mathfrak{s}} = \mathfrak{H}(X) \setminus (\mathfrak{s} \sqcup \mathfrak{s}^*)$ we get an isometric embedding $\overline{X}(\mathfrak{H}_{\mathfrak{s}}) \hookrightarrow \overline{X}$, obtained as above, onto*

$$\bigcap_{h \in \mathfrak{s}} h \subset \overline{X}.$$

(3) The set \mathfrak{s} satisfies the descending chain condition if and only if the image of $X(\mathfrak{H}_{\mathfrak{s}})$ is in X .

Furthermore, if the set \mathfrak{s} is G -invariant, for some group $G \leq \text{Aut}(X)$, then, with the restricted action on the image, the above natural embeddings are G -equivariant.

This result is a slightly strengthened version of [19, Lemma 2.6]. See also [28, Proposition 2.11] for a very similar statement.

We remark that Proposition 3.14 includes the possibility that $\mathfrak{s} = \emptyset$ and hence $\mathfrak{H} \sqcup \emptyset \sqcup \emptyset$ is a legitimate lifting decomposition of \mathfrak{H} . In this case, recall that an intersection over an empty collection of sets is everything, i.e. $\bigcap_{h \in \emptyset} h = \bar{X}$.

Proof Suppose that there is a lifting decomposition \mathfrak{s} for \mathfrak{H}' . Then, since $\mathfrak{s} \cap \mathfrak{s}^* = \emptyset$ and $\mathfrak{H} = \mathfrak{H}' \sqcup \mathfrak{s} \sqcup \mathfrak{s}^*$, we deduce that $\mathfrak{H}' = \mathfrak{H} \setminus (\mathfrak{s} \sqcup \mathfrak{s}^*)$ is involution invariant.

Next, we claim that $\mathfrak{H}' \subset \mathfrak{H}$ has a property called *tightly nested*, meaning that for every pair $h, k \in \mathfrak{H}'$, and for every $\ell \in \mathfrak{H}$ with $h \subset \ell \subset k$, we have $\ell \in \mathfrak{H}'$. Assume, for a contradiction, that $h, k \in \mathfrak{H}'$, and $\ell \in \mathfrak{H}$ with $h \subset \ell \subset k$, but $\ell \notin \mathfrak{H}'$. Then (up to replacing the three half-spaces with their complements) we may assume that $\ell \in \mathfrak{s}$. But then $k \in \mathfrak{s}$, a contradiction.

With the verification that \mathfrak{H}' is involution invariant and tightly nested, our hypotheses imply those of [19, Lemma 2.6]. Now, conclusions (1)–(3) follow from that lemma.

Finally, the conclusion about G -invariance follows because the embeddings are canonically determined by the associated half-spaces. □

Remark 3.15 In Proposition 3.14, the isometric embeddings of $\bar{X}(\mathfrak{H}')$ and $\bar{X}(\mathfrak{H}_{\mathfrak{s}})$ restrict to isometric embeddings of $X(\mathfrak{H}')$ and $X(\mathfrak{H}_{\mathfrak{s}})$, respectively. Although $\partial_R X$ is typically not closed, the images under the isometric embedding of $\partial_R X(\mathfrak{H}')$ and $\partial_R X(\mathfrak{H}_{\mathfrak{s}})$ always lie in $\partial_R X$. Indeed, recall from the discussion following Theorem 3.6 that a point belongs to the Roller boundary if and only if there is an infinite descending chain of half-spaces containing the point in question. Hence, if U has an infinite descending chain, then so does $U \sqcup \mathfrak{s}$.

As a first application of Proposition 3.14, we will prove a generalization of Lemma 3.9 to \bar{X} . Observe that a naive generalization of Lemma 3.9 does not hold, because X is (interval) convex in \bar{X} by Definition 3.10, but is not an intersection of extended half-spaces. However, the following lemma says that a *closed*, convex set in \bar{X} is always an intersection of extended half-spaces.

Lemma 3.16 *Let $A \subset \bar{X}$ be a convex set. Then the Roller closure \bar{A} in \bar{X} is the image of $\bar{X}(\mathfrak{H}_A)$ under the embedding of Proposition 3.14, and*

$$\bar{A} = \bigcap_{h \in \mathfrak{H}_A^+} h = \bigcap_{A \subset h} h.$$

Proof Observe that \mathfrak{H}_A^+ is a consistent set of half-spaces. Thus Proposition 3.14.(2) gives an embedding $\bar{X}(\mathfrak{H}_A) \hookrightarrow \bar{X}$ whose image is

$$W = \bigcap_{A \subset h} h = \bigcap_{h \in \mathfrak{H}_A^+} h = \bigcap_{h \in \mathfrak{H}_A^+} h.$$

Here, the first equality is the definition of W , the second equality is the definition of \mathfrak{H}_A^+ , and the third equality holds because half-spaces are closed, meaning $\mathfrak{H}_A^+ = \mathfrak{H}_A^{\pm}$.

It remains to show that $\bar{A} = W$. The inclusion $\bar{A} \subset W$ is obvious, because $\bar{A} \subset h$ for every $h \in \mathfrak{H}_A^+$. To prove the reverse inclusion, suppose for a contradiction that there is a point $y \in W \setminus \bar{A}$. Since $\bar{X} \setminus \bar{A}$ is open, there is a basic open neighborhood of y of the form $\bigcap_{i=1}^n h_i$, such that

$$y \in \left(\bigcap_{i=1}^n h_i \right) \subset \bar{X} \setminus \bar{A}.$$

Therefore, $A \subset h_1^* \cup \dots \cup h_n^*$. Without loss of generality, assume that n is minimal in the sense that for each $j = 1, \dots, n$ there is a point

$$a_j \in A \setminus \left(\bigcup_{j \neq i=1}^n h_i^* \right) = A \cap \left(\bigcap_{j \neq i=1}^n h_i \right).$$

One crucial observation is that $n \geq 2$: otherwise, if $A \subset h_1^*$, then $h_1 \cap W = \emptyset$, but we know that $y \in h_1 \cap W$.

Now, define $p = m(y, a_1, a_2)$. Recall that $\mathcal{I}(a_1, a_2) \subset A$, because A is convex. By the definition of medians, we have

$$p \in \mathcal{I}(a_1, a_2) \subset A, \quad p \in \mathcal{I}(y, a_1) \subset \left(\bigcap_{j=2}^n h_j \right), \quad p \in \mathcal{I}(y, a_2) \subset h_1.$$

Consequently, we have $p \in A \cap \left(\bigcap_{i=1}^n h_i \right) = \emptyset$, a contradiction. Thus $\bar{A} = W$. □

3.6 Gates and bridges

A very useful property of a convex subset $A \subset X^{(0)}$ is the existence of a *gate map* $\pi_A: X \rightarrow A$; this is a well-known notion, see e.g. [41, Section 2.2] and the citations therein. Metrically, $\pi_A: X^{(0)} \rightarrow A^{(0)}$ is just closest-point projection, but gates can be characterized entirely in terms of the median (see e.g. [30, Section 2.1] and citations therein). Since we saw that \bar{X} has a median satisfying (3.11), convex subsets of \bar{X} do similarly admit gates. This is the content of Proposition 3.17, which has the advantage of showing an important relationship between gates and walls.

Recall that $\mathcal{W}(x, Y)$ is the set of hyperplanes separating x from Y , where $Y \subset \bar{X}$ is convex.

Proposition 3.17 (Gate projection in \bar{X}) *Let $C \subset \bar{X}$ be closed and convex. There is a unique projection $\pi_C: \bar{X} \rightarrow C$ such that for any $x \in \bar{X}$ we have*

$$\mathcal{W}(x, \pi_C(x)) = \mathcal{W}(x, C).$$

Proof Let $\mathfrak{H} = \mathfrak{H}_C \sqcup \mathfrak{H}_C^+ \sqcup \mathfrak{H}_C^-$ be the lifting decomposition associated to C . The map $2^{\mathfrak{H}} \rightarrow 2^{\mathfrak{H}_C}$ defined by $U \mapsto U \setminus (\mathfrak{H}_C^+ \sqcup \mathfrak{H}_C^-)$ restricts to a surjection $\bar{X} \rightarrow \bar{X}(\mathfrak{H}_C)$. Next, Proposition 3.14 gives an embedding $\bar{X}(\mathfrak{H}_C) \rightarrow \bar{X}$, namely $V \mapsto V \sqcup \mathfrak{H}_C^+$. Composing the surjection $\bar{X} \rightarrow \bar{X}(\mathfrak{H}_C)$ with the embedding $\bar{X}(\mathfrak{H}_C) \rightarrow \bar{X}$ yields a map $\pi_C: \bar{X} \rightarrow \bar{X}$. The image $\pi_C(\bar{X})$ is C since, by Lemma 3.16, the image of the embedding $\bar{X}(\mathfrak{H}_C) \rightarrow \bar{X}$ is exactly $C = \bar{C}$.

To prove the equality in the statement, first observe that $\mathcal{W}(x, C) \subset \mathcal{W}(x, \pi_C(x))$ because $\pi_C(x) \in C$. For the opposite inclusion, consider a hyperplane $\hat{h} \in \mathcal{W}(x, \pi_C(x))$, and choose an orientation $h \in \mathfrak{H}_x^+$. By the definition of π_C , we have

$$\mathfrak{H}_{\pi_C(x)}^+ = (\mathfrak{H}_x^+ \setminus \mathfrak{H}_C^-) \cup \mathfrak{H}_C^+.$$

Since $h \in \mathfrak{H}_x^+ \setminus \mathfrak{H}_{\pi_C(x)}^+$, we must have $h \in \mathfrak{H}_C^-$, hence $\hat{h} \in \mathcal{W}(x, C)$. Thus $\mathcal{W}(x, C) = \mathcal{W}(x, \pi_C(x))$.

Finally, for $x \in C$, we must have $x = \pi_C(x)$ because $\mathcal{W}(x, \pi_C(x)) = \mathcal{W}(x, C) = \emptyset$. Thus π_C is indeed a projection to C . □

Note that π_C restricts to the identity on C . We call π_C the *gate map to C* , which is consistent with the terminology used in the theory of median spaces because of the following lemma:

Lemma 3.18 *Let π_C be the map from Proposition 3.17. Then for all $x \in \bar{X}$ and $y \in C$, we have $m(x, y, \pi_C(x)) = \pi_C(x)$. In particular, $\pi_C(x) \in \bigcap_{y \in C} \mathcal{I}(x, y)$.*

Proof By Proposition 3.17 and the fact that $y \in C$, we have $\mathcal{W}(x, \pi_C(x)) \subset \mathcal{W}(x, y)$, so $\mathcal{I}(x, \pi_C(x)) \subset \mathcal{I}(x, y)$. So, $\pi_C(x) \in \mathcal{I}(x, y) \cap \mathcal{I}(x, \pi_C(x)) \cap \mathcal{I}(y, \pi_C(x))$, hence $m = \pi_C(x)$. □

The above lemma has two corollaries. First, we can characterize medians in terms of projections:

Corollary 3.19 *Suppose that $x, y, z \in X^{(0)}$. Then $\pi_{\mathcal{I}(x,y)}(z) = m(x, y, z)$.*

Proof By Lemma 3.18, we have $\pi_{\mathcal{I}(x,y)}(z) \in \mathcal{I}(x, z)$ and $\pi_{\mathcal{I}(x,y)}(z) \in \mathcal{I}(y, z)$. On the other hand, $\pi_{\mathcal{I}(x,y)}(z) \in \mathcal{I}(x, y)$ by definition. So, by (3.11), we get $\pi_{\mathcal{I}(x,y)}(z) = m(x, y, z)$. □

Second, Lemma 3.18 allows us to project from X to convex subcomplexes of X .

Corollary 3.20 *Let $A \subset X$ be a convex subcomplex. Then, for all vertices $x \in X^{(0)}$, we have $\pi_{\bar{A}}(x) \in A$. Consequently, we get a projection $\pi_A = \pi_{\bar{A}}: X^{(0)} \rightarrow A^{(0)}$.*

Proof Let $x \in X^{(0)}$ and $y \in A^{(0)}$. By Lemma 3.18, we have $\pi_{\bar{A}}(x) \in \mathcal{I}(x, y) \subset X^{(0)}$. But $\bar{A} \cap X^{(0)} = A^{(0)}$. □

Remark 3.21 In fact, it is possible to extend π_A over higher-dimensional cubes to get a map $\pi_A: X \rightarrow A$. This is done as follows (see also [9, Section 2.1]). First let e be a 1-cube of X joining 0-cubes x, y and dual to a hyperplane \hat{h} . From Proposition 3.17, one has the following: if \hat{h} does not cross A , then $\pi_A(x) = \pi_A(y)$, and, otherwise, $\pi_A(x)$ and $\pi_A(y)$ are joined by an edge $\bar{e} \subset A$ dual to \hat{h} . In the former case, π_A extends over e by sending every point to $\pi_A(x) = \pi_A(y)$; in the latter case, we extend by declaring $\pi_A: e \rightarrow \bar{e}$ to be the obvious isometry.

Now, if c is a d -cube for $d \geq 2$, then for some $s \in \{0, \dots, d\}$, there are 1-cubes $e_1, \dots, e_s \subset c$ that span an s -cube c' and have the property that the hyperplanes $\hat{h}_1, \dots, \hat{h}_s$ respectively dual to e_1, \dots, e_s are precisely the hyperplanes intersecting both c and A . (The case $s = 0$ corresponds to the situation where no hyperplane intersecting c intersects A , and c' is an arbitrary 0-cube of c .) So for $1 \leq i \leq s$, we have that $\bar{e}_i = \pi_A(e_i)$ is a 1-cube of A dual to \hat{h}_i , and the 1-cubes $\bar{e}_1, \dots, \bar{e}_s$ span an s -cube c' . We extend π_A over c' using the obvious cubical isometry $c' \rightarrow \bar{c}'$ extending the map π_A on 1-cubes. By construction, for each $v \in c^{(0)}$, there is a unique 0-cube $v' \in c'$ such that $\pi_A(v) = \pi_A(v')$, the assignment $v \mapsto v'$ extends to a cubical map $c \rightarrow c'$ (collapsing the hyperplanes not in $\{\hat{h}_1, \dots, \hat{h}_s\}$), and composing this with $\pi_A: c' \rightarrow A$ gives the (extended) gate map $\pi_A: c \rightarrow A$. Note that $c' \subset c$ is only unique up to parallelism, but any two allowable choices of c' give the same map $c \rightarrow A$.

We conclude that gates enjoy the following properties:

Lemma 3.22 *Let $A \subset X$ be a convex subcomplex. Then:*

- (1) *For all $x, y \in X$, we have $\mathcal{W}(\pi_A(x), \pi_A(y)) = \mathcal{W}(x, y) \cap \mathcal{W}(A)$, and $\mathcal{W}(x, \pi_A(x)) = \mathcal{W}(x, A)$.*
- (2) *The map π_A is 1-lipschitz on $X^{(0)}$ with the d_1 metric and the CAT(0) metric d_X .*

Proof For 0-cubes, the first part of conclusion (1) appears in various places in the literature; see e.g. [41, Lemma 2.5], or one can deduce it easily from Proposition 3.17. The second part of the first assertion (again for 0-cubes) restates Proposition 3.17. The generalization of (1) to arbitrary points in X follows immediately from the corresponding assertion for 0-cubes, together with the construction in Remark 3.21.

Finally, (1) implies the first part of (2), since $d_1(x, y) = |\mathcal{W}(x, y)|$ for 0-cubes x, y . For the second part, first note that the restriction of π_A to each cube is 1-lipschitz for the CAT(0) metric, since it factors as the natural projection of the cube onto one of its faces (which is lipschitz for the Euclidean metric) composed with an isometric embedding. Now, fix $x, y \in X$ and let γ be the CAT(0) geodesic joining them, which decomposes as a finite concatenation of geodesics, each lying in a single cube. Each such geodesic is mapped to a path whose length has not increased, so $\pi_A \circ \gamma$ is a path from $\pi_A(x)$ to $\pi_A(y)$ of length at most $d_X(x, y)$, as required. □

The following standard application of gates is often called the *bridge lemma*.

Lemma 3.23 (Bridge Lemma) *Let $I, J \subset X$ be convex subcomplexes. Let $\pi_I : X \rightarrow I$ and $\pi_J : X \rightarrow J$ be the gate maps. Then the following hold:*

- (1) *$\pi_I(J)$ and $\pi_J(I)$ are convex subcomplexes.*
- (2) *$\mathcal{W}(\pi_I(J)) = \mathcal{W}(\pi_J(I)) = \mathcal{W}(I) \cap \mathcal{W}(J)$.*
- (3) *The map $\pi_I : \pi_J(I) \rightarrow \pi_I(J)$ is an isomorphism of CAT(0) cube complexes.*
- (4) *The cubical convex hull $\text{CH}(\pi_I(J) \cup \pi_J(I))$ is a CAT(0) cube complex isomorphic to $\pi_J(I) \times \mathcal{I}(x, \pi_J(x))$ for any vertex $x \in \pi_I(J)$.*
- (5) *If $I \cap J \neq \emptyset$, then $\pi_I(J) = \pi_J(I) = I \cap J$, and $\mathcal{W}(I \cap J) = \mathcal{W}(I) \cap \mathcal{W}(J)$.*

Proof This can be assembled from results in the literature in various ways. For example, the first statement is part of [30, Lemma 2.2]; the second and last statements follow from [9, Lemmas 2.1 and 2.6]. The third follows from [9, Lemma 2.4]. For the fourth, see [19, Lemma 2.18] or [9, Lemmas 2.4 and 2.6]. □

3.7 Facing tuples and properness

A *facing k -tuple* of hyperplanes is a set $\{\hat{h}_1, \dots, \hat{h}_k\}$ of hyperplanes such that, for each $i \leq k$, we can choose a half-space h_i such that $h_i \cap h_j = \emptyset$ for $i \neq j$. We are often particularly interested in facing triples, since many of the subcomplexes of X considered later in the paper will be CAT(0) cube complexes that do not contain facing triples.

One useful application of the notion of a facing tuple is the following lemma, which in practice will be used to guarantee properness of certain subcomplexes of X . See [41, Section 3] for a more detailed discussion. See also [19, Lemma 2.33] for a related result, where X is required to isometrically embed into \mathbb{Z}^D for some D .

Lemma 3.24 *Let X be a CAT(0) cube complex of dimension $D < \infty$. Suppose that there exists T such that X does not contain a facing T -tuple. Then (X, d_X) is a proper CAT(0) space.*

Proof Let $R \geq 0$ and let B be a ball of radius R in X (with respect to the CAT(0) metric d_X). Let $\mathcal{W}(B)$ be the set of hyperplanes intersecting B .

We say that hyperplanes $\hat{h}_1, \dots, \hat{h}_n$ in $\mathcal{W}(B)$ form a *chain* if (up to relabelling), each \hat{h}_i separates \hat{h}_{i-1} from \hat{h}_{i+1} for $2 \leq i \leq n - 1$.

We claim that there exists N , depending only on D and R , such that any chain in $\mathcal{W}(B)$ has cardinality at most N . Indeed, if $\hat{h}_1, \dots, \hat{h}_n$ is a chain, then any edge dual to \hat{h}_1 lies at ℓ^1 distance at least $n - 2$ from any edge dual to \hat{h}_n . Now, Lemma 3.1 implies there exists N (depending only on D and R) such that $n > N$ implies that $d_X(\hat{h}_1, \hat{h}_n) > 2R$. This is impossible, since \hat{h}_1, \hat{h}_n intersect a common R -ball.

On the other hand, [41, Proposition 3.3] provides a constant $K = K(D, T)$ such that any finite set $\mathcal{F} \subset \mathcal{W}(B)$ must contain a chain of length at least $|\mathcal{F}|/K$. So, $|\mathcal{W}(B)| \leq KN$.

Let $\text{CH}(B)$ be the cubical convex hull of B . The set of hyperplanes of the CAT(0) cube complex $\text{CH}(B)$ is $\mathcal{W}(B)$, which we have just shown is finite. Hence $\text{CH}(B)$ is a compact CAT(0) cube complex, and hence proper. Since $\widehat{B} \hookrightarrow X$ is an isometric embedding (with respect to CAT(0) metrics), B is a ball in $\text{CH}(B)$, and is therefore compact. This completes the proof. □

The use of the above lemma will be the following. Given $x \in X^{(0)}$ and $y \in \partial_R X$, we can consider the interval $\mathcal{I}(x, y)$ in \overline{X} . The intersection $\mathcal{I}(x, y) \cap X$ is a vertex-convex set, which is the 0-skeleton of a uniquely determined convex subcomplex A . Note that $\mathcal{W}(A) = \mathcal{W}(x, y)$. This infinite set of hyperplanes cannot contain a facing triple. Hence, Lemma 3.24 shows that A is proper. We will need this in the proof of Lemma 9.9, to arrange for a sequence of CAT(0) geodesic segments in A to converge uniformly on compact sets to a CAT(0) geodesic ray. In this way, we avoid a blanket hypothesis that X is proper.

Note that we therefore only use a special case of the lemma, namely properness of cubical intervals. Properness of intervals follows from a much stronger statement: intervals endowed with the ℓ_1 metric isometrically embed into $\mathbb{R}^{\dim(X)}$ (i.e. are *Euclidean*). It follows that intervals are proper in the CAT(0) metric, since the two metrics are bilipschitz. This embedding into $\mathbb{R}^{\dim(X)}$ goes back at least to [3, Theorem 1.16]. This embedding is closely related to [19, Lemma 2.33], which bounds the cardinality of facing tuples in Euclidean CAT(0) cube complexes, which include, but are more general than, cubical intervals.

3.8 Combinatorial Geodesic Rays

Next, we develop some basic facts about combinatorial geodesics in CAT(0) cube complexes that will be needed when we relate different types of boundaries.

Definition 3.25 (*Combinatorial geodesic rays*) A map $\gamma : [0, \infty) \rightarrow X^{(1)}$ is said to be a *combinatorial geodesic ray* if $\gamma(\mathbb{N}) \subset X^{(0)}$ and for each $n, m \in \mathbb{N}$ we have $d_1(\gamma(n), \gamma(m)) = |n - m|$ and γ is an isometry on each interval $[n, n + 1]$.

Lemma 3.26 Let $\gamma : [0, \infty) \rightarrow X$ be a combinatorial geodesic ray. There exists a unique point $y \in \partial_R X$ such that $\gamma(n) \rightarrow y$ in \overline{X} as $n \rightarrow \infty$.

In the sequel, we will write $y = \gamma(\infty)$ to mean $\gamma(n) \rightarrow y$.

Proof Consider the set

$$D(\gamma) = \{h \in \mathfrak{H} : \gamma([N, \infty) \cap \mathbb{N}) \subset h \text{ for some } N \in \mathbb{N}\}.$$

This is clearly a consistent choice of half-spaces. We claim that for each $h \in \mathfrak{H}$, either h or h^* belongs to $D(\gamma)$. Indeed, up to replacing h by h^* , there must be a monotonic sequence $n_k \rightarrow \infty$ such that $\gamma(n_k) \in h$. Since h is convex, it follows that $\gamma([n_1, \infty) \cap \mathbb{N}) \subset h$.

Now, define $C = \bigcap_{h \in D(\gamma)} h$. We claim that C consists of exactly one point. Note that $C \neq \emptyset$ because it is a nested intersection of compact sets in \bar{X} . On the other hand, if $x, x' \in C$ are distinct points, then consider a half-space h such that $x \in h$ and $x' \in h^*$. This half-space must satisfy both $h \in D(\gamma)$ and $h^* \in D(\gamma)$, a contradiction. Thus $C = \{y\}$, as desired. \square

Lemma 3.27 *For every $x \in X^{(0)}$ and every $y \in \partial_R X$, there is a combinatorial geodesic ray $\gamma: [0, \infty) \rightarrow X$ such that $\gamma(0) = x$ and $\gamma(\infty) = y$.*

The proof of Lemma 3.27 is modeled on the construction of normal cubed paths, introduced by Niblo and Reeves in [53].

Proof Without loss of generality, let us assume that $\bar{X} = \mathcal{I}(x, y)$, that is, $\mathfrak{H}(X) = \mathfrak{H}_x^+ \Delta \mathfrak{H}_y^+$. The collection of half-spaces $\mathfrak{H}_x^+ \setminus \mathfrak{H}_y^+$ is partially ordered by set inclusion. Consider $h, k \in \mathfrak{H}_x^+ \setminus \mathfrak{H}_y^+$. Since $x \in h \cap k$ and $y \in h^* \cap k^*$, the possible relations between h and k are: $h \subset k$ or $k \subset h$ or $h \pitchfork k$. In other words, incomparable elements are transverse. In particular, all minimal elements of $\mathfrak{H}_x^+ \setminus \mathfrak{H}_y^+$ must be pairwise transverse. By finite dimensionality, the set M_1 of minimal elements is finite. Let x_1 be the vertex obtained by replacing every $h \in M_1$ by h^* , that is, $\mathfrak{H}_{x_1}^+ = (\mathfrak{H}_x^+ \setminus M_1) \cup M_1^*$, a clearly consistent and total choice of half-spaces. In other words, x_1 is diagonally across from x in the unique maximal cube containing x in the vertex interval $\mathcal{I}(x, y)$.

Proceeding by induction, suppose that x_{n-1} has been defined. Let M_n be the set of minimal elements of $\mathfrak{H}_{x_{n-1}}^+ \setminus \mathfrak{H}_y^+$. Let x_n be the vertex obtained by replacing every $h \in M_n$ by h^* , corresponding to the half-space interval $\mathfrak{H}_{x_n}^+ = (\mathfrak{H}_{x_{n-1}}^+ \setminus M_n) \cup M_n^*$.

Observing that $M_k^* \cap M_j = \emptyset$ for $j < k$, we conclude that

$$\mathfrak{H}_{x_n}^+ = \left(\mathfrak{H}_x^+ \setminus \bigcup_{j=1}^{n-1} M_j \right) \cup \left(\bigcup_{j=1}^{n-1} M_j^* \right).$$

Fix $k < m \in \mathbb{N}$. We claim that for any integer $\ell \in [k, m]$, x_ℓ lies on every geodesic path from x_k to x_m . Equivalently, we claim that the median of x_k, x_ℓ , and x_m is x_ℓ . Recall that this will be the case if every half-space containing x_k and x_m also contains x_ℓ . This containment can be established as follows

$$\begin{aligned} \mathfrak{H}_{x_k}^+ \cap \mathfrak{H}_{x_m}^+ &= \left(\mathfrak{H}_x^+ \setminus \bigcup_{j=1}^{m-1} M_j \right) \cup \left(\bigcup_{j=1}^{k-1} M_j^* \right) \\ &\subset \left(\mathfrak{H}_x^+ \setminus \bigcup_{j=1}^{\ell-1} M_j \right) \cup \left(\bigcup_{j=1}^{\ell-1} M_j^* \right) = \mathfrak{H}_{x_\ell}^+. \end{aligned}$$

This proves the claim.

Set $x_0 = x$ and choose an edge-geodesic $\gamma_k: [0, d(x_{k-1}, x_k)] \rightarrow \mathcal{I}(x_{k-1}, x_k)$ from x_{k-1} to x_k . Finally, define γ as the infinite concatenation $\gamma_1 \gamma_2 \dots$. Then, by construction, $\gamma(\infty) = y$. \square

We will discuss further properties of combinatorial geodesic rays in Sect. 6.

4 Cuboid complexes with modified CAT(0) metrics

A *cuboid* is a box of the form $[0, a_1] \times \dots \times [0, a_n] \subset \mathbb{R}^n$. In this section, we describe a way to modify the metric on a finite-dimensional CAT(0) cube complex X to produce a

new CAT(0) space whose cells are cuboids rather than cubes. Our results relating the Tits boundary and simplicial boundary of a CAT(0) cube complex will carry over to this context with small changes to some proofs, which we will indicate in the relevant places.

The reader only interested in CAT(0) cube complexes with the standard CAT(0) metric can safely skip this section, apart from the statements of Lemmas 4.6 and 4.7. Those lemmas are cuboid generalizations of standard results in the literature, namely [25, Lemma 2.5] and [5, Proposition 2.8]. Thus, in the standard metric, one can substitute those results from the literature for the lemmas of this section.

Let X be a finite-dimensional CAT(0) cube complex and let $G \rightarrow \text{Aut}(X)$ be a group acting on X by cubical automorphisms. Let \mathcal{W} be the set of hyperplanes in X .

Definition 4.1 (*G-admissible hyperplane rescaling*) A hyperplane rescaling of X is a map $\rho: \mathcal{W} \rightarrow (0, \infty)$. The rescaling is called *G-admissible* if it is G -invariant, and in addition we have $m_\rho = \inf_{\hat{h} \in \mathcal{W}} \rho(\hat{h}) > 0$ and $M_\rho = \sup_{\hat{h} \in \mathcal{W}} \rho(\hat{h}) < \infty$.

Let ρ be a G -admissible hyperplane rescaling of X . Fix a cube c of X . For each hyperplane \hat{h}_i intersecting c , let $\rho_i = \rho(\hat{h}_i)$. Regarding c as an isometric copy of $\prod_i [0, \rho_i]$, with the Euclidean metric, let d_X^ρ be the resulting piecewise-Euclidean path metric on X .

The following result is a special case of a theorem of Bowditch [15], but we are able to give a direct proof.

Lemma 4.2 *Suppose that G acts on X by cubical automorphisms, and let ρ be a G -admissible hyperplane rescaling of X . Then (X, d_X^ρ) is a CAT(0) space, and the action of G on (X, d_X^ρ) is by isometries.*

Proof Let $x \in X$ be an arbitrary point. Then the link of x in (X, d_X^ρ) is isometric to the link of x in (X, d_X) , and is hence a CAT(1) space. Since the amount of rescaling is bounded, there is a small ϵ depending on x and ρ such that the ϵ -balls about x in the two metrics are (abstractly) isometric. Thus (X, d_X^ρ) is locally CAT(0). Recalling that X is contractible (the homeomorphism type has not changed), we conclude that (X, d_X^ρ) is globally CAT(0) by Cartan–Hadamard theorem [7, Theorem II.4.1]. Since ρ is G -equivariant, G takes cuboids to d_X^ρ -isometric cuboids, hence G acts by isometries. □

One can also prove Lemma 4.2 using a result of Bridson and Haefliger [7, Theorem II.5.2]. Our proof is essentially the same argument, but the cubical structure allows us to assume bounded rescaling rather than finitely many shapes.

Definition 4.3 (*CAT(0) cuboid metric*) Consider a finite-dimensional CAT(0) cube complex X , a group G acting on X by cubical automorphisms, and a G -admissible hyperplane rescaling ρ of X . We say that the CAT(0) metric d_X^ρ from Lemma 4.2 is a *CAT(0) cuboid metric* on X , and (X, d_X^ρ) is a *CAT(0) cuboid complex*.

Definition 4.4 (*Automorphism group of cuboid complex*) Suppose that G acts on X by cubical automorphisms, and let ρ be a G -admissible hyperplane rescaling of X . Define $\text{Aut}(X^\rho) \subset \text{Aut}(X)$ to be the automorphism group of the rescaled complex (X, d_X^ρ) . By construction, the representation $G \rightarrow \text{Aut}(X)$ has image in $\text{Aut}(X^\rho)$.

Corollary 4.5 *Every cubically convex subcomplex $Y \subset X$ is also geodesically convex in the rescaled metric d_X^ρ . In particular, every CAT(0) half-space $\text{CH}(k)$ is convex in d_X^ρ .*

Proof By Definition 3.7, a cubically convex subcomplex is an intersection of CAT(0) half-spaces. Since any intersection of d_X^ρ -convex sets will be d_X^ρ -convex, it suffices to prove that CAT(0) half-spaces are d_X^ρ -convex.

Now, let $k \in \mathfrak{H}(X)$, and let $\text{CH}(k)$ be the the corresponding CAT(0) half-space. Observe that $\text{CH}(k)$ is a CAT(0) cube complex in its own right with half-space structure $\mathfrak{H}(k) = \{h \in \mathfrak{H} : h \uparrow k\}$. Let $G_k = \text{stab}_G(k)$ which then acts by isometries on $\text{Aut}(\text{CH}(k))$. Clearly the restriction $\rho|_{\mathfrak{H}(k)}$ is G_k -admissible. By Lemma 4.2 we have $(\text{CH}(k), d_{\text{CH}(k)}^{\rho|_{\mathfrak{H}(k)}})$ is a CAT(0) space, so that

$$d_X^\rho|_{\text{CH}(k)} = d_{\text{CH}(k)}^{\rho|_{\mathfrak{H}(k)}}.$$

Therefore, the identity inclusion $\text{CH}(h) \hookrightarrow X$ is isometric with respect to this new CAT(0) metric, and in particular $\text{CH}(k)$ is d_X^ρ -convex. \square

The following result is a generalization of [25, Lemma 2.5] for CAT(0) cuboid complexes.

Lemma 4.6 *Let X be a CAT(0) cube complex and let ρ be a G -admissible hyperplane rescaling. Suppose that $\mathcal{W}(X) = \mathcal{W}_1 \sqcup \mathcal{W}_2$, where \hat{h}, \hat{v} cross whenever $\hat{h} \in \mathcal{W}_1, \hat{v} \in \mathcal{W}_2$.*

Then $X = X_1 \times X_2$ (as a product cube complex), where X_1 and X_2 are CAT(0) cube complexes dual to \mathcal{W}_1 and \mathcal{W}_2 .

Letting $p_i : X \rightarrow X_i$ denote the natural projection, for $i \in \{1, 2\}$, the map p_i induces a bijection on the sets of hyperplanes such that the preimages of the hyperplanes in X_i are the hyperplanes in \mathcal{W}_i .

Moreover, restricting ρ to \mathcal{W}_i , we obtain rescaled CAT(0) metrics $d_{X_i}^\rho$ such that $(X_i, d_{X_i}^\rho)$ is a CAT(0) cuboid complex and we have

$$d_X^\rho(x, y) = \sqrt{d_{X_1}^\rho(p_1(x), p_1(y))^2 + d_{X_2}^\rho(p_2(x), p_2(y))^2}$$

for all $x, y \in X$.

Proof (Sketch.) The statement about X decomposing as a product cube complex follows from [25, Lemma 2.5] and the discussion in [25] preceding it. So, it remains to prove the statement about rescaled CAT(0) metrics.

Admissibility of the rescalings on X_1, X_2 follows from admissibility of ρ and the fact that p_i is a restriction quotient in the sense of [25, Section 2.3], meaning it sends hyperplanes to hyperplanes. So, Lemma 4.2 makes each $(X_i, d_{X_i}^\rho)$ a CAT(0) cuboid complex.

We now verify the metric statement. Let $x, y \in X$. Let α be the d_X^ρ -geodesic from $x = (p_1(x), p_2(x))$ to $y = (p_1(y), p_2(y))$. Let β be the d_X^ρ -geodesic from $(p_1(x), p_2(y))$ to $(p_1(x), p_2(x))$, and let γ be the d_X^ρ -geodesic from $(p_1(y), p_2(y))$ to $(p_1(x), p_2(y))$. So, $\alpha\gamma\beta$ is a geodesic triangle.

Note that every point on β has the form $(p_1(x), z)$ and every point on γ has the form $(z, p_2(y))$, because the p_i are restriction quotients and therefore have convex fibers by Lemma 4.5. In particular, $\|\beta\| = d_{X_2}^\rho(p_2(x), p_2(y))$ and $\|\gamma\| = d_{X_1}^\rho(p_1(x), p_1(y))$ and p_1, p_2 are respectively isometries on γ and β . Moreover, the Alexandrov angle (see Definition 7.1) formed by β and γ at their common point is $\pi/2$. Perform the same construction with $(p_1(x), p_2(y))$ replaced by $(p_1(y), p_2(x))$. This yields a geodesic quadrilateral in (X, d_X^ρ) , with all angles $\pi/2$, two adjacent sides having length $d_{X_2}^\rho(p_2(x), p_2(y))$ and $d_{X_1}^\rho(p_1(x), p_1(y))$, and x and y as opposite corners. The Flat Quadrilateral Theorem ([7, Theorem II.2.11]) now implies that

$$d_X^\rho(x, y) = \sqrt{d_{X_1}^\rho(p_1(x), p_1(y))^2 + d_{X_2}^\rho(p_2(x), p_2(y))^2},$$

as required. □

We also need the following standard lemma. In the context of CAT(0) cube complexes, this is [5, Proposition 2.8]. We state it here for cuboid complexes. The proof from [5] goes through with tiny changes that we indicate below.

Lemma 4.7 *Let X be a D -dimensional CAT(0) cube complex and let ρ be a G -admissible hyperplane rescaling. Let (X, d_X^ρ) be the resulting CAT(0) cuboid complex. Let $\gamma : [0, \infty) \rightarrow X$ be a geodesic ray for the metric d_X^ρ , with $\gamma(0) \in X^{(0)}$. Then X contains a combinatorial geodesic ray α in $X^{(1)}$ such that $\alpha(0) = \gamma(0)$ and such that α fellow-travels with γ at distance depending only on D and ρ .*

Proof The proof from [5, Section 2.2] works with almost no change. First, [5, Lemma 2.9] is about the combinatorial structure of X only, which does not change when we rescale edges to pass from the CAT(0) metric d_X to the CAT(0) metric d_X^ρ . The proof of [5, Proposition 2.8] needs this lemma, plus CAT(0) convexity of half-spaces, which continues to hold in view of Corollary 4.5. The proof of [5, Proposition 2.8] produces a combinatorial ray α such that each point of α lies in a common cube as a point of γ , and vice versa. Hence the two rays fellow travel at CAT(0) distance bounded by the diameter of a cube. In the metric d_X , this is bounded in terms of D ; in the metric d_X^ρ , this is bounded in terms of D and the constant M_ρ from Definition 4.1, as required by the statement. □

Since d_X and d_X^ρ are path-metrics, and for each cube c of X , the standard CAT(0) metric d_X on c is bilipschitz to the restriction of d_X^ρ (with constant depending only on the rescaling constant), the identity map $(X, d_X) \rightarrow (X, d_X^\rho)$ is bilipschitz. Combining this fact with Lemma 3.1 yields a version of Lemma 3.1 for the rescaled metric d_X^ρ :

Lemma 4.8 *There are constants $\lambda_0^\rho \geq 1, \lambda_1^\rho \geq 0$, depending on $\dim X$ and ρ , such that the following holds. For any pair of points $x, y \in X$,*

$$\frac{1}{\lambda_0^\rho} d_X^\rho(x, y) - \lambda_1^\rho \leq |\mathcal{W}(x, y)| \leq \lambda_0^\rho d_X^\rho(x, y) + \lambda_1^\rho.$$

In a similar fashion, by Remark 3.2, there is a bilipschitz relationship between (X, d_X^ρ) and (X, d_1) .

5 The simplicial Roller boundary $\mathfrak{R}_\Delta X$

In this section, we discuss one of the main objects in the paper: the *simplicial Roller boundary*.

Definition 5.1 (*Roller class*) Let $x, y \in \overline{X}$. Recall that $\mathfrak{H}_x^+ \triangle \mathfrak{H}_y^+$ is the set of half-spaces that separate x from y . We say that x is *commensurate* to y , and write $x \sim y$, if $|\mathfrak{H}(x, y)| < \infty$. Observe that $X^{(0)}$ forms a single commensurability class in \overline{X} . Commensurability classes contained in $\partial_R X = \overline{X} \setminus X^{(0)}$ are also called *Roller classes*.

Definition 5.2 (*Guralnik quotient*) The *Guralnik quotient* of the Roller Boundary is $\mathfrak{R}X = \partial_R X / \sim$, the set of commensurability classes of points in the Roller boundary [37]. We shall often implicitly consider elements $[x] \in \mathfrak{R}X$ as subsets $[x] \subset \overline{X}$.

In most cases, the quotient topology on $\mathfrak{R}X = \partial_R X / \sim$ is not Hausdorff (in fact, it does not even satisfy the T_1 axiom). We think of $\mathfrak{R}X$ as only a set, without a topology. In

Definition 5.7, we will use a partial order on $\mathfrak{R}X$ to define a simplicial complex $\mathfrak{R}_\Delta X$ with far nicer topological properties.

Following [54], we say a set $S \subset \mathfrak{H}$ is *non-terminating* if S contains no minimal elements: given any half-space $h \in S$, there is a half-space $k \subsetneq h$ with $k \in S$.

Lemma 5.3 *Let $C \subset \partial_R X$ be a closed, convex set. The following are equivalent:*

- (1) *There is a point $y \in \partial_R X$ such that $C = \overline{[y]}$.*
- (2) *There exists $k \leq \dim(X)$ and a family $(h_m^1)_{m \geq 0}, (h_m^2)_{m \geq 0}, \dots, (h_m^k)_{m \geq 0}$ of descending chains of half-spaces such that*

$$C = \bigcap_{i=1}^k \bigcap_{m \geq 0} h_m^i.$$

- (3) *C is a union of Roller classes.*
- (4) *The set of half-spaces \mathfrak{H}_C^+ is non-terminating.*

Proof We will prove (1) \Rightarrow (2) \Rightarrow (3) \Rightarrow (1) and \neg (3) \Leftrightarrow \neg (4).

The claim that (1) implies (2) is a restatement of [32, Lemma 6.17].

To prove that (2) implies (3), assume that $y \in C = \bigcap_{i=1}^k \bigcap_{m \geq 0} h_m^i$ and $y' \sim y$ meaning that $|\mathfrak{H}(y, y')| < \infty$. Since $\mathfrak{H}_y^+ \setminus \mathfrak{H}_{y'}^+$ is finite, it follows that there exists an N for which $h_m^i \in \mathfrak{H}_{y'}^+$ (i.e. $y' \in h_m^i$) for each $m \geq N$ and $i = 1, \dots, k$. That is to say,

$$y' \in \bigcap_{i=1}^k \bigcap_{m \geq N} h_m^i = \bigcap_{i=1}^k \bigcap_{m \geq 0} h_m^i = C.$$

To prove (3) implies (1), suppose that C is a union of Roller classes. Let \mathfrak{H}_C be the collection of half-spaces that cut C . By Theorem 3.6, there is an associated cube complex $X(\mathfrak{H}_C)$. Since C is closed and convex, Proposition 3.14 and Lemma 3.16 give an embedding $i_C : \overline{X}(\mathfrak{H}_C) \hookrightarrow \overline{X}$ whose image is C , and furthermore $C = \bigcap_{h \in \mathfrak{H}_C^+} h$.

Let $y_0 \in X(\mathfrak{H}_C)^{(0)}$ and note that $\overline{[y_0]} = \overline{X}(\mathfrak{H}_C)$. Let $y = i_C(y_0) \in C$ and note that $i_C(\overline{[y_0]}) \subset [y]$, hence $C = i_C(\overline{[y_0]}) \subset [y]$. Since C is a union of equivalence classes, we have $[y] \subset C$. Since C is closed, we have that $\overline{[y]} \subset C$ and hence $C = \overline{[y]}$.

To prove \neg (4) implies \neg (3), suppose there exists a minimal half-space $h_0 \in \mathfrak{H}_C^+$. We claim that $\mathfrak{H}_C^+ \setminus \{h_0\}$ is consistent. Indeed, if $h \in \mathfrak{H}_C^+ \setminus \{h_0\}$ then $h^* \notin \mathfrak{H}_C^+ \setminus \{h_0\}$. If $h \subset k$ then $k \in \mathfrak{H}_C^+$, and since h_0 is minimal, $h_0 \neq k$ so $k \in \mathfrak{H}_C^+ \setminus \{h_0\}$.

We may therefore consider the CAT(0) cube complex associated to $\mathfrak{H}_C \sqcup \{h_0, h_0^*\}$. Fix $y, y' \in X(\mathfrak{H}_C \sqcup \{h_0, h_0^*\})$ such that $y \in h_0$ and $y' \in h_0^*$. Then, by Proposition 3.14 there is a unique point $\tilde{y} \in C$ associated to the total and consistent choice of half-spaces

$$\{h \in \mathfrak{H}_C : y \in h\} \sqcup \mathfrak{H}_C^+.$$

Similarly, there is a unique point $\tilde{y}' \in \partial_R X \setminus C$ associated to

$$\{h \in \mathfrak{H}_C : y' \in h\} \sqcup \{h_0^*\} \sqcup (\mathfrak{H}_C^+ \setminus \{h_0\}).$$

By construction, $|\mathfrak{H}(\tilde{y}, \tilde{y}')| = 2$, hence C is not a union of Roller classes.

Finally, to prove \neg (3) implies \neg (4), suppose that C is not a union of Roller classes. Then there are commensurate elements $y \in C$ and $y' \in \partial_R X \setminus C$. Convexity means

$$C = \bigcap_{h \in \mathfrak{H}_C^+} h,$$

hence there is a half-space $h \in \mathfrak{H}_C^+$ with $y \in C \subset h$ and $y' \notin h$. Under the assumption that $|\mathfrak{H}(y, y')| < \infty$, we may choose a minimal such h and conclude that \mathfrak{H}_C^+ fails to be non-terminating. □

Definition 5.4 (*Principal class*) Given a closed, convex set C satisfying one of the equivalent conditions of Lemma 5.3, the class $[y]$ whose closure is C is called the *principal class* of C . Compare Guralnik [37, Section 3.1].

Definition 5.5 (*Partial order on $\mathfrak{R}X$*) Let $[x], [y] \subset \partial_R X$ be Roller classes. We define a partial order on $\mathfrak{R}X$ by setting $[x] \leq [y]$ whenever $\mathfrak{H}_{[x]}^+ \subset \mathfrak{H}_{[y]}^+$, or equivalently $\mathfrak{H}_{[x]}^+ \subset \mathfrak{H}_{[y]}^+$.

The following result is essentially a corollary of Lemma 5.3.

Lemma 5.6 *For Roller classes $[x], [y] \subset \partial_R X$, the following are equivalent:*

- (1) $[x] \leq [y]$
- (2) $\mathfrak{H}_{[x]}^+ \subset \mathfrak{H}_{[y]}^+$
- (3) $[y] \subset \overline{[x]}$
- (4) $[y] \cap \overline{[x]} \neq \emptyset$

Proof Conditions (1) and (2) are equivalent by Definition 5.5. Since $\mathfrak{H}_{[x]}^+ = \mathfrak{H}_{\overline{[x]}}^+$, by Lemma 3.16, Remark 3.13 implies that (2) and (3) are equivalent. Next, (3) trivially implies (4). Finally, (4) implies (3) by Lemma 5.3. □

Definition 5.7 (*Simplicial Roller boundary*) The *simplicial Roller boundary* $\mathfrak{R}_\Delta X$ is the simplicial realization of the partial order of Definition 5.5. That is: vertices of $\mathfrak{R}_\Delta X$ correspond to points of $\mathfrak{R}X$, or equivalently Roller classes in $\partial_R X$. Simplices in $\mathfrak{R}_\Delta X$ correspond to totally ordered chains of Roller classes.

Observe that the action of $\text{Aut}(X)$ on $\partial_R X$ preserves commensurability classes and the order \leq , yielding a bijective action on $\mathfrak{R}X$ and a simplicial action on $\mathfrak{R}_\Delta X$.

Remark 5.8 [Dimension of $\mathfrak{R}_\Delta X$] The characterization of Lemma 5.3.(2) implies that any simplex in $\mathfrak{R}_\Delta X$ has $k \leq \dim X$ vertices, hence $\dim \mathfrak{R}_\Delta X \leq \dim X - 1$.

6 The simplicial boundary

In this section, we recall the definition of the simplicial boundary $\partial_\Delta X$ [39]; see Definition 6.9. Then, we develop some connections between $\partial_\Delta X$, the Roller boundary $\partial_R X$, and the simplicial Roller boundary $\mathfrak{R}_\Delta X$. See Theorem 6.27 and Corollary 6.33 for the most top-level statement. The lemmas used to prove those top-level results will be extensively used in later sections.

6.1 Unidirectional boundary sets

Recall that \mathfrak{H} is the set of half-spaces of X and \mathcal{W} is the set of hyperplanes of X . In this section, we mainly focus on collections of hyperplanes.

Definition 6.1 (*Facing triple*) A triple $h_1, h_2, h_3 \in \mathfrak{H}$ is called a *facing triple* if their complements h_1^*, h_2^*, h_3^* are pairwise disjoint. Similarly, a triple of hyperplanes $\hat{h}_1, \hat{h}_2, \hat{h}_3 \in \mathcal{W}$ is called a *facing triple* if there exists a choice of orientation such that h_1, h_2, h_3 form a facing triple.

Caprace and Sageev [25] showed that in a non-elementary essential CAT(0) cube complex, given any half-space $h \in \mathfrak{H}$ there exist half-spaces $k, \ell \in \mathfrak{H}$ such that $\{h, k, \ell\}$ is a facing triple.

Definition 6.2 (*Unidirectional, inseparable*) A collection of hyperplanes $\mathcal{S} \subset \mathcal{W}$ is called *unidirectional* if for every $\hat{h} \in \mathcal{S}$, at most one side of \hat{h} contains infinitely many other elements of \mathcal{S} . A collection $\mathcal{S} \subset \mathcal{W}$ is called *inseparable* if whenever $\hat{h}, \hat{k} \in \mathcal{S}$ and $\hat{\ell}$ is a hyperplane that separates \hat{h} and \hat{k} , then $\hat{\ell} \in \mathcal{S}$. This is closely related to the “tightly nested” condition in Proposition 3.14.

Definition 6.3 (*UBS*) A set $\mathcal{U} \subset \mathcal{W}$ is a *unidirectional boundary set* (abbreviated *UBS*) if it is infinite, unidirectional, inseparable, and does not contain a facing triple.

Definition 6.4 (*Partial order on UBSes*) Let \mathcal{U} and \mathcal{V} be UBSes. Define a relation \preceq where $\mathcal{U} \preceq \mathcal{V}$ if all but finitely many elements of \mathcal{U} lie in \mathcal{V} . Say that \mathcal{U} and \mathcal{V} are *commensurate* (denoted $\mathcal{U} \sim \mathcal{V}$) if $\mathcal{U} \preceq \mathcal{V}$ and $\mathcal{V} \preceq \mathcal{U}$.

The relation \preceq on UBSes descends to a partial order on commensurability classes of UBSes. We say that \mathcal{V} is *minimal* if its commensurability class is minimal—that is, if $\mathcal{U} \preceq \mathcal{V}$ implies $\mathcal{U} \sim \mathcal{V}$. Similarly, we say that \mathcal{V} is *maximal* if $\mathcal{V} \preceq \mathcal{V}'$ implies $\mathcal{V} \sim \mathcal{V}'$.

Example 6.5 For disjoint subsets $A, B \subset \bar{X}$, let $\mathcal{W}(A, B)$ be the set of hyperplanes separating A from B . It is straightforward to check that $\mathcal{W}(A, B)$ is inseparable and does not contain a facing triple. Furthermore, if $A \subset X$ then $\mathcal{W}(A, B)$ is unidirectional. If, in addition, $B \subset \partial_R X$, then $\mathcal{W}(A, B)$ is infinite. Under these last two conditions, we conclude that $\mathcal{W}(A, B) = \mathcal{W}(A, \bar{B})$ is a UBS. Compare Lemma 6.20.

Definition 6.6 (*Pruned UBS, canonical half-space*) Let \mathcal{U} be a UBS. We say that \mathcal{U} is *pruned* if every hyperplane $\hat{h} \in \mathcal{U}$ has exactly one associated half-space h containing infinitely many elements of \mathcal{U} . In this situation, we call h the *canonical orientation* of \hat{h} or the *canonical half-space associated to \hat{h}* .

Lemma 6.7 Let \mathcal{U} be a UBS. Then \mathcal{U} contains a canonical pruned sub-UBS $\mathcal{U}' \subset \mathcal{U}$ that is commensurate to \mathcal{U} .

Proof Let $F_{\mathcal{U}}$ be the set the hyperplanes of \mathcal{U} with the property that *both* associated half-spaces contain finitely many elements of \mathcal{U} . We claim that $F_{\mathcal{U}}$ is finite. To see this, let $\hat{h}_1, \dots, \hat{h}_D$ be a maximal family of pairwise transverse hyperplanes in $F_{\mathcal{U}}$. This family is finite because $D \leq \dim X$. For any $\hat{h} \in F_{\mathcal{U}} \setminus \{\hat{h}_1, \dots, \hat{h}_D\}$ there is at least one i for which \hat{h} is parallel (i.e. not transverse) to \hat{h}_i . By the definition of $F_{\mathcal{U}}$, there are finitely many hyperplanes in \mathcal{U} parallel to each \hat{h}_i , hence $F_{\mathcal{U}}$ is finite.

Now, define $\mathcal{U}' = \mathcal{U} \setminus F_{\mathcal{U}}$. By construction, every $\hat{h} \in \mathcal{U}'$ has exactly one associated half-space containing infinitely many elements of \mathcal{U} , hence infinitely many elements of \mathcal{U}' as well.

We claim that \mathcal{U}' is a UBS. Since $F_{\mathcal{U}}$ is finite, it follows that \mathcal{U}' is infinite. Unidirectionality and the lack of facing triples are properties that pass to subsets of \mathcal{U} , hence to \mathcal{U}' . Next, suppose that $\hat{h}, \hat{k} \in \mathcal{U}'$ with $\hat{\ell}$ in between. Since $\hat{h}, \hat{k} \in \mathcal{U}$ and \mathcal{U} is inseparable, observe that $\hat{\ell} \in \mathcal{U}$. But $\hat{\ell} \notin F_{\mathcal{U}}$, because \hat{h}, \hat{k} both bound half-spaces containing infinitely many elements of \mathcal{U} . We conclude that $\hat{\ell} \in \mathcal{U}'$, verifying that \mathcal{U}' is a UBS. Since \mathcal{U} and \mathcal{U}' differ by finitely many hyperplanes, they are commensurate. □

Hagen proved that every (commensurability class of) UBS \mathcal{U} decomposes in a well-defined fashion into a disjoint union of minimal UBSes. The following is a restatement of [39, Theorem 3.10] and [40, Theorem A].

Lemma 6.8 *Let \mathcal{U} be a UBS. Then there is a decomposition $\mathcal{U} \sim \bigsqcup_{i=1}^k \mathcal{U}_i$, where:*

- (1) *each \mathcal{U}_i is a minimal UBS;*
- (2) *any UBS $\mathcal{V} \subset \mathcal{U}$ is commensurate with the disjoint union of some of the \mathcal{U}_i ;*
- (3) *if $1 \leq j < i \leq k$, then every hyperplane in \mathcal{U}_i crosses all but finitely many of the hyperplanes in \mathcal{U}_j .*

Consequently, there are finitely many commensurability classes of UBSes $\mathcal{V} \preceq \mathcal{U}$.

Lemma 6.8 enables the following definition.

Definition 6.9 (*Simplicial boundary*) For a UBS \mathcal{U} , the *dimension* of \mathcal{U} is the number k appearing in the decomposition of Lemma 6.8. This is an invariant of the commensurability class. An immediate consequence of Lemma 6.8.(3) is that $\dim(\mathcal{U}) \leq \dim(X)$.

The *simplicial boundary* $\partial_\Delta X$ is the following simplicial complex. For each commensurability class of a UBS \mathcal{U} , there is a $(k - 1)$ -simplex $\sigma_\mathcal{U}$, where k is the dimension of the class of \mathcal{U} . The face relation is defined as follows: $\sigma_\mathcal{U}$ is a face of $\sigma_\mathcal{V}$ if and only if $\mathcal{U} \preceq \mathcal{V}$.

Note that $\partial_\Delta X$ has a vertex for each commensurability class of a minimal UBS. Furthermore, a finite collection of vertices spans a maximal simplex if and only if the union of the corresponding minimal UBSes is commensurate with a maximal UBS. However, there can exist non-maximal simplices that do not correspond to a UBS; see Example 6.32. For this reason, $\partial_\Delta X$ is *not* the simplicial realization of the partial order \preceq .

Observe that $\text{Aut}(X)$ acts on UBSes in a way that preserves the \preceq relation, and hence preserves commensurability. Thus $\text{Aut}(X)$ acts simplicially on $\partial_\Delta X$.

6.2 Minimal and dominant UBSes

Next, we establish some structural results about minimal UBSes. Given a set S consisting of hyperplanes, the *inseparable closure* of S , denoted $\langle\langle S \rangle\rangle$, is the intersection of all inseparable sets containing S .

Lemma 6.10 *Let $\mathcal{U} \subset \mathcal{W}(X)$ be a minimal UBS. Then there exists an infinite descending chain $\{h_n : n \in \mathbb{N}\} \subset \mathfrak{H}$ with $\hat{h}_n \in \mathcal{U}$, such that $\mathcal{U} \sim \langle\langle \{\hat{h}_n : n \in \mathbb{N}\} \rangle\rangle$.*

Conversely, for every descending chain $\{h_n : n \in \mathbb{N}\} \subset \mathfrak{H}$, there exists $N \geq 0$ such that $\langle\langle \{\hat{h}_n : n \geq N\} \rangle\rangle$ is a minimal UBS.

Proof Let \mathcal{U} be a minimal UBS. Let $\{h_n : n \in \mathbb{N}\} \subset \mathfrak{H}$ be a descending chain such that $\{\hat{h}_n\} \subset \mathcal{U}$, which exists because \mathcal{U} is infinite, unidirectional, contains no facing triple, and contains no infinite collection of pairwise-crossing hyperplanes. Let $\mathcal{U}' = \langle\langle \{\hat{h}_n\} \rangle\rangle$ be the inseparable closure of $\{\hat{h}_n\}$. Since \mathcal{U} is inseparable, we have $\mathcal{U}' \subset \mathcal{U}$ by the definition of the inseparable closure. Note that \mathcal{U}' is a UBS (compare the proof of [39, Lemma 3.7] or [19, Caprace Lemma B.6]). Since $\mathcal{U}' \subset \mathcal{U}$ and \mathcal{U} is minimal, we have $\mathcal{U}' \sim \mathcal{U}$, which proves the first assertion of the lemma.

Now we prove the second assertion. For each $m \geq 0$, let $C_m = \{h_n : n \geq m\}$ and let $\langle\langle C_m \rangle\rangle$ be its inseparable closure. As before, each $\langle\langle C_m \rangle\rangle$ is a UBS. Note that $\langle\langle C_m \rangle\rangle \subset \langle\langle C_{m'} \rangle\rangle$ whenever $m \geq m'$. In particular, $\langle\langle C_m \rangle\rangle \preceq \langle\langle C_{m'} \rangle\rangle$ for $m' \geq m$. So, by Lemma 6.8, there exists $M \geq 0$ such that $\langle\langle C_m \rangle\rangle \sim \langle\langle C_M \rangle\rangle$ for all $m \geq M$. Hence, if $m \geq M$ and $\{k_i\}$ is any descending chain of half-spaces corresponding to hyperplanes in $\langle\langle C_M \rangle\rangle$, then there exists $I \geq 0$ such that $\{\hat{k}_i : i \geq I\} \subset \langle\langle C_m \rangle\rangle$. This implies that $\langle\langle C_M \rangle\rangle$ is commensurate with $\langle\langle \{\hat{k}_i : i \geq 0\} \rangle\rangle$, by, for example [29, Lemma 4.4]. See also the proof of [29, Proposition 4.7.(2)].

Now let $\mathcal{V} \subset \langle\langle C_M \rangle\rangle$ be a minimal UBS, which exists by [39, Lemma 3.7]. By the first part of the lemma, \mathcal{V} is commensurate with the inseparable closure of some chain contained in $\langle\langle C_M \rangle\rangle$. By the preceding paragraph, $\langle\langle C_M \rangle\rangle \sim \mathcal{V}$, so $\langle\langle C_M \rangle\rangle$ is minimal, as required. \square

Remark 6.11 Lemma 6.10 has been generalized to the setting of median spaces by Fioravanti [29, Proposition 4.7]. In fact, the same result of Fioravanti also generalizes many other basic facts about UBSes, such as [39, Lemma 3.7 and Theorem 3.10].

The second assertion of Lemma 6.10 corrects a slight misstatement in [19, Caprace Lemma B.6]: it is true that minimal UBSes are commensurate with inseparable closures of chains, but inseparable closures of chains need not be minimal; one might first need to pass to a (cofinite) sub-chain (see e.g. [29, Figure 1] and the paragraph following). We emphasize that this does not affect the use of [19, Caprace Lemma B.6] in [19], because what’s being used is the (correct) statement that, up to commensurability, minimal UBSes arise as inseparable closures of chains.

In this paper, we primarily use the first assertion of Lemma 6.10, in conjunction with Lemma 6.8, to decompose a UBS \mathcal{U} as a disjoint union of minimal UBSes (up to commensurability), each of which has the additional property that it is the inseparable closure of a chain.

The following consequence of Lemma 6.10 neatly sums up the special properties of chains whose inseparable closures are minimal UBSes:

Corollary 6.12 *Let $\{h_n : n \geq 0\}$ be an infinite descending chain of half-spaces. For each $m \geq 0$, let \mathcal{U}_m be the inseparable closure of $\{\hat{h}_n : n \geq m\}$. Then the following are equivalent:*

- (1) *The UBS \mathcal{U}_m is minimal for all $m \geq 0$.*
- (2) *$\mathcal{U}_m \sim \mathcal{U}_0$ for all $m \geq 0$.*

In particular, if \mathcal{U}_0 is minimal, then $\mathcal{U}_0 \sim \mathcal{U}_m$ for all $m \geq 0$.

Proof Since $\mathcal{U}_m \leq \mathcal{U}_0$ for all m , the first statement immediately implies the second. For the reverse direction, Lemma 6.10 provides $M \geq 0$ such that \mathcal{U}_m is minimal for $m \geq M$. If the second assertion holds, it then follows that the first must also. \square

Definition 6.13 (*Chain of hyperplanes*) Following Corollary 6.12, for the rest of this subsection we shall refer to a sequence of disjoint hyperplanes $\{\hat{h}_n\}$ whose inseparable closure is a minimal UBS as a *chain of hyperplanes*.

In particular, a chain of hyperplanes $\{\hat{h}_n\}_{n \geq 0}$ defines a descending chain of half-spaces $\{h_n\}_{n \geq 0}$. A descending chain of half-spaces $\{h_n\}_{n \geq 0}$ has the property that for all sufficiently large N , the sequence $\{\hat{h}_n\}_{n \geq N}$ is a chain of hyperplanes, in view of Lemma 6.10.

Definition 6.14 (*Dominant*) Let \mathcal{U} be a UBS. A hyperplane $\hat{h} \in \mathcal{U}$ is called *dominant for \mathcal{U}* if $\hat{h} = \hat{h}_0$ is the base of a chain of hyperplanes $\{h_n\} \subset \mathcal{U}$ such that \hat{h} crosses all but finitely many hyperplanes in $\mathcal{U} \setminus \langle\langle \{h_n\} \rangle\rangle$.

Recall that by definition of a *chain of hyperplanes*, $\langle\langle \{h_n\} \rangle\rangle$ is a minimal UBS. See Fig. 2 for an example. By Lemma 6.8, the inseparable closure $\langle\langle \{h_n\} \rangle\rangle$ is commensurate to some minimal UBS \mathcal{U}_i in the decomposition $\mathcal{U} \sim \bigsqcup_{i=1}^k \mathcal{U}_i$ of that lemma. Furthermore, \mathcal{U}_i is unique: since a dominant hyperplane $\hat{h} \in \mathcal{U}$ is the base of a chain in \mathcal{U}_i , then \hat{h} must intersect all but finitely many hyperplanes of \mathcal{U}_j (for $j \neq i$), hence \hat{h} cannot be disjoint from a chain in \mathcal{U}_j .




Fig. 2 Part of a CAT(0) cube complex for which the set of hyperplanes is a UBS. Each horizontal hyperplane is the base of a chain of horizontal hyperplanes whose inseparable closure consists of all but finitely many of the horizontal hyperplanes. Each horizontal hyperplane crosses all but finitely many vertical ones. So each horizontal hyperplane is dominant. The vertical hyperplanes are not dominant. Indeed, if v is vertical, it is the base of a chain of vertical hyperplanes whose inseparable closure consists of vertical hyperplanes, while v fails to cross infinitely many horizontal hyperplanes. On the other hand, v is also the base of a sequence of hyperplanes consisting of v and infinitely many horizontal hyperplanes, whose canonical half-spaces form a descending chain. The inseparable closure of this set of hyperplanes contains all but finitely many of the hyperplanes, and therefore is not a minimal UBS, so this set is not a chain of hyperplanes

We also observe that if \mathcal{U} is minimal and pruned, then every hyperplane of \mathcal{U} is vacuously dominant.

The following lemma is only stated in the special case where $\mathcal{W}(X)$ is already a UBS. In practice, we will apply Lemma 6.15 to convex subcomplexes (such as the convex hull of a geodesic ray) where this hypothesis holds.

Lemma 6.15 (Dominant hyperplane combinatorial properties) *Suppose that $\mathcal{U} = \mathcal{W}(X)$ is a pruned UBS of dimension k . Then \mathcal{U} contains minimal UBSes $\mathcal{D}_d, \dots, \mathcal{D}_k$ such that the following hold:*

- (1) *Each \mathcal{D}_i is a set of dominant hyperplanes for \mathcal{U} , and all but finitely many dominant hyperplanes belong to $\bigcup_{i=d}^k \mathcal{D}_i$.*
- (2) *For each $i \in \{d, \dots, k\}$ and each $\hat{h}_0 \in \mathcal{D}_i$, there is a chain $\{\hat{h}_n\}_{n \geq 0} \subset \mathcal{D}_i$ such that \mathcal{D}_i is commensurate with $\langle\langle \{\hat{h}_n\} \rangle\rangle$, the inseparable closure of the chain.*
- (3) *For any such chain $\{\hat{h}_n\}$, there is $N > 0$, such that for all $n \geq N$, we have*

$$\mathcal{W}(\hat{h}_{n+1}) \setminus \mathcal{D}_i \subset \mathcal{W}(\hat{h}_n) \setminus \mathcal{D}_i.$$

Furthermore, for all $n \geq 0$, we have


$$|\mathcal{W}(\hat{h}_n) - \mathcal{W}(\hat{h}_{n+1})| < \infty.$$

- (4) *For every hyperplane $\hat{u} \in \mathcal{U}$, there exists $i \in \{d, \dots, k\}$ and a dominant hyperplane $\hat{h} \in \mathcal{D}_i$ such that \hat{h} lies in the canonical half-space associated to \hat{u} .*

We will not use item (3) later in the paper, but we record it as a potentially useful fact about UBSes.

Proof By Lemma 6.8, there is a decomposition $\mathcal{U} \sim \bigsqcup_{i=1}^k \mathcal{U}_i$, with an ordering such that for $j < i$, every hyperplane in \mathcal{U}_i crosses all but finitely many of the hyperplanes in \mathcal{U}_j . The ordering is only partially determined: if every hyperplane in \mathcal{U}_i crosses all but finitely many

Fig. 3 The sets $\mathcal{U}'_i, \mathcal{U}_j$ in the proof of Lemma 6.15.(4)



hyperplanes in \mathcal{U}_{i+1} , then we can reverse the order of $\mathcal{U}_i, \mathcal{U}_{i+1}$ while keeping the conclusion of Lemma 6.8.(3). Let d be the smallest index i such that \mathcal{U}_i can be placed at the top of the order in such a decomposition.

Now, fix an index $i \geq d$. By Lemma 6.10, \mathcal{U}_i is commensurate with $\mathcal{D}_i = \langle\langle \{\hat{h}_n\} \rangle\rangle$, where $\{\hat{h}_n\} \subset \mathcal{U}_i$ is a chain. Thus $\mathcal{D}_i \sim \mathcal{U}_i$ is a minimal UBS.

We claim that every $\hat{h} \in \mathcal{D}_i$ is dominant. Given $\hat{h} \in \mathcal{D}_i$ and the canonical associated half-space h (recall Definition 6.6), there exists $N \geq 0$ such that $\hat{h}_N \subset h$. Thus \hat{h} is the base of the infinite chain $\{\hat{h}, \hat{h}_N, \hat{h}_{N+1}, \dots\}$. Furthermore, $\hat{h} \in \mathcal{D}_i \subset \mathcal{U}_i$ intersects all but finitely many hyperplanes in \mathcal{U}_j for $j \neq i$, because we have chosen an index $i \geq d$. Thus \hat{h} is dominant by Definition 6.14. Thus we have found dominant minimal UBSes $\mathcal{D}_d, \dots, \mathcal{D}_k$ satisfying conclusions (1) and (2).

Now we verify assertion (3). Fix $i \geq d$ and let $\hat{h}, \hat{h}' \in \mathcal{D}_i$. Since the minimal UBS \mathcal{D}_i is the inseparable closure of a chain, \hat{h}, \hat{h}' both cross finitely many hyperplanes in \mathcal{D}_i . Indeed, if $\mathcal{W}(\hat{h}) \cap \mathcal{D}_i$ is infinite, it must contain a chain of hyperplanes whose inseparable closure \mathcal{D}' belongs to \mathcal{D}_i , and is contained in $\mathcal{W}(\hat{h})$ since the latter is inseparable. Since \mathcal{D}_i is minimal, $\mathcal{D}' \sim \mathcal{D}_i$. Hence $\mathcal{W}(\hat{h})$ contains all but finitely many elements of \mathcal{D}_i ; since $\hat{h} \in \mathcal{D}_i$, this contradicts that \mathcal{D}_i is the inseparable closure of a chain. Thus $\mathcal{W}(\hat{h}) \cap \mathcal{D}_i$ is finite. Moreover, since \mathcal{D}_i is dominant, \hat{h}, \hat{h}' both cross all but finitely many hyperplanes of $\mathcal{U} \setminus \mathcal{D}_i$, and we can therefore conclude that

$$|\mathcal{W}(\hat{h}) \Delta \mathcal{W}(\hat{h}')| < \infty.$$

So, it suffices to show that, for all sufficiently large n , we have $\mathcal{W}(\hat{h}_{n+1}) \setminus \mathcal{D}_i \subset \mathcal{W}(\hat{h}_n) \setminus \mathcal{D}_i$. If this doesn't hold, then for infinitely many n we can find a hyperplane $\hat{u}_n \notin \mathcal{D}_i$ that crosses \hat{h}_{n+1} but not \hat{h}_n ; note that \hat{u}_n does not separate any two elements of our chain since $\hat{u}_n \notin \mathcal{D}_i$. Since there are no facing triples, \hat{u}_n crosses \hat{h}_m for $m \geq n + 1$. Since the collection of violating hyperplanes $\{\hat{u}_n\}$ is infinite and contains no hyperplanes crossing \hat{h}_0 , dominance of \hat{h}_0 is therefore contradicted. Thus there exists N such that $\mathcal{W}(\hat{h}_{n+1}) \setminus \mathcal{D}_i \subset \mathcal{W}(\hat{h}_n) \setminus \mathcal{D}_i$ for $n \geq N$, establishing (3).

We now prove assertion (4). Let $\hat{u} \in \mathcal{U}$ be a hyperplane, and let u be the canonical half-space associated to \hat{u} . Given the fixed decomposition $\mathcal{U} \sim \bigsqcup_{i=1}^k \mathcal{U}_i$, let i be the largest index such that u contains a chain of hyperplanes in \mathcal{U} , whose inseparable closure \mathcal{U}'_i is commensurate to some \mathcal{U}_i . Such an i exists because \mathcal{U} is pruned. We claim that $i \geq d$. This claim implies that u contains infinitely many hyperplanes of \mathcal{D}_i , which will suffice to prove the lemma.

Suppose for a contradiction that $i < d$. Fix $\hat{v} \in \mathcal{U}'_i$, and consider a minimal UBS \mathcal{U}_j for $j > i$. If any hyperplane in \mathcal{U}_j fails to cross \hat{v} , then it fails to cross \hat{u} , hence u contains a chain in \mathcal{U}_j , contradicting the maximality of i . See Fig. 3.

Furthermore, since $\hat{v} \in \mathcal{U}_i$, it must cross all but finitely many hyperplanes of \mathcal{U}_ℓ for $\ell < i$ by Lemma 6.8.(3). Thus \hat{v} crosses all but finitely many hyperplanes of $\mathcal{U} \setminus \mathcal{U}'_i$, hence \hat{v} is dominant. Since $\hat{v} \in \mathcal{U}'_i$ was arbitrary, this shows that $i \geq d$ by the definition of d . \square

6.3 Connections between $\partial_\Delta X$ and $\partial_R X$

The next several lemmas develop tools for comparing the Roller boundary and the simplicial boundary.

Definition 6.16 (*Umbra of a UBS*) Let $\mathcal{U} \subset \mathcal{W}(X)$ be a UBS, and let $\mathcal{U}' \sim \mathcal{U}$ be a pruned UBS commensurate to \mathcal{U} . For every $\hat{h} \in \mathcal{U}'$, let $h \in \mathfrak{H}(\bar{X})$ be the canonical orientation of \hat{h} , that is, the (extended, vertex) half-space containing infinitely many hyperplanes of \mathcal{U}' . The umbra $\bar{Y}_\mathcal{U}$ associated to \mathcal{U} is

$$\bar{Y}_\mathcal{U} = \bigcap_{\hat{h} \in \mathcal{U}'} h \subset \bar{X}.$$

By Definition 3.10, the umbra $\bar{Y}_\mathcal{U}$ is a convex subset of $\partial_R X$, one may think of it as the Roller realization of the UBS \mathcal{U} . In Lemma 6.18, we will show that $\bar{Y}_\mathcal{U}$ satisfies the equivalent conditions of Lemma 5.3. Then, $Y_\mathcal{U}$ will denote the principal class of $\bar{Y}_\mathcal{U}$: that is, the unique Roller class whose closure is $\bar{Y}_\mathcal{U}$.

Lemma 6.17 Let $\mathcal{U} \subset \mathcal{W}(X)$ be a UBS. Then the umbra $\bar{Y}_\mathcal{U}$ associated to \mathcal{U} is a nonempty subset of $\partial_R X$. Furthermore,

- If $\mathcal{U} \sim \mathcal{V}$ then $\bar{Y}_\mathcal{U} = \bar{Y}_\mathcal{V}$.
- If $\mathcal{U} \sim \mathcal{U}_1 \cup \dots \cup \mathcal{U}_N$ for some UBSes $\mathcal{U}_1, \dots, \mathcal{U}_N$ then

$$\bar{Y}_\mathcal{U} = \bigcap_{i=1}^N \bar{Y}_{\mathcal{U}_i}.$$

Proof Let \mathcal{U}' be the pruned UBS commensurate to \mathcal{U} . As in Definition 6.16, let h be the canonical orientation of every $\hat{h} \in \mathcal{U}'$, and define $\mathcal{U}' = \{h : \hat{h} \in \mathcal{U}'\} \subset \mathfrak{H}$. Then $\bar{Y}_\mathcal{U} := \bigcap_{h \in \mathcal{U}'} h$.

Since \bar{X} is compact, to show that $\bar{Y}_\mathcal{U}$ is not empty, it is sufficient to show that \mathcal{U}' has the finite intersection property. By the Helly property for half-spaces [58, Section 2.2] this is equivalent to showing that every pair of half-spaces $h, k \in \mathcal{U}'$ has nonempty intersection. But a disjoint pair $h, k \in \mathcal{U}'$ contradicts the unidirectionality of \mathcal{U}' , as both sides of \hat{h} would contain infinitely many elements of \mathcal{U}' .

Using the same argument, we observe that for every $h \in \mathcal{U}'$ there is an infinite descending chain $\{h_n\}$ with $h_0 = h$ and $h_n \in \mathcal{U}'$. Indeed, fix $h \in \mathcal{U}'$. Since \mathcal{U}' is pruned, there are infinitely many hyperplanes in \mathcal{U}' contained in h . Since \mathcal{U}' is unidirectional, $\hat{k} \subset h$ and $\hat{k} \in \mathcal{U}'$ implies $k \subset h$. Applying induction, we find the infinite descending chain. Thus $\bar{Y}_\mathcal{U} \subset \partial_R X$.

We now turn to the commensurability class of \mathcal{U} . Suppose that \mathcal{V} is a UBS commensurate to \mathcal{U} . Since commensurability is an equivalence relation, we conclude that \mathcal{U}' and \mathcal{V}' are commensurate as well. Let $\mathcal{V}' = \{h : \hat{h} \in \mathcal{V}'\} \subset \mathfrak{H}$ be the set of canonical half-spaces associated to \mathcal{V}' .

We claim that \mathcal{U}' and \mathcal{V}' are also commensurate. Since the map taking half-spaces to hyperplanes is 2-to-1, it suffices to show that under this map, the set $\mathcal{U}' \setminus \mathcal{V}'$ has image in the finite set $\mathcal{U}' \setminus \mathcal{V}'$, and vice versa. Let $h \in \mathcal{U}' \setminus \mathcal{V}'$. Then $\hat{h} \in \mathcal{U}'$. By definition of \mathcal{U}' , we

know that the set $\{\hat{k} \in \mathcal{U} : \hat{k} \subset h\}$ is infinite. The commensurability of \mathcal{U} and \mathcal{V} implies that $\{\hat{k} \in \mathcal{V} : \hat{k} \subset h\}$ is infinite as well. But $h \notin \mathcal{V}'$, hence it must be the case that $\hat{h} \notin \mathcal{V}'$.

To see that $\bar{Y}_{\mathcal{U}} = \bar{Y}_{\mathcal{V}}$ it is sufficient to show that if $h \in \mathcal{U}'$ then there is $k \in \mathcal{V}' \cap \mathcal{U}'$ with $k \subset h$. But this is immediate. As was proved above, for any $h \in \mathcal{U}'$ there is an infinite descending chain $\{h_n\} \subset \mathcal{U}'$ with $h_0 = h$. Since \mathcal{U}' and \mathcal{V}' are commensurate, for N sufficiently large, we must have that $h_N \in \mathcal{V}'$.

Finally, the construction above makes it clear that if \mathcal{U} is commensurate to $\mathcal{U}_1 \cup \dots \cup \mathcal{U}_N$ for some UBSes $\mathcal{U}_1, \dots, \mathcal{U}_N$ then $\bar{Y}_{\mathcal{U}} = \bigcap_{i=1}^N \bar{Y}_{\mathcal{U}_i}$. □

Lemma 6.18 *For every UBS \mathcal{U} , there is a Roller class $Z \subset \partial_R X$ such that $\bar{Y}_{\mathcal{U}} = \bar{Z}$. Furthermore, the assignment $\mathcal{U} \mapsto Z = Y_{\mathcal{U}}$ is $\text{Aut}(X)$ -equivariant.*

Proof By Lemmas 6.8 and 6.10, \mathcal{U} is commensurate to a UBS $\mathcal{V} = \bigsqcup_{i=1}^k \mathcal{U}_i$, where $\mathcal{U}_i = \langle\langle \hat{h}_n^i \rangle\rangle$ for a descending chain of half-spaces $\{h_n^i\}$. Note that every \mathcal{U}_i is pruned. Thus

$$\bar{Y}_{\mathcal{U}} = \bigcap_{1 \leq i \leq k} \bar{Y}_{\mathcal{U}_i} = \bigcap_{1 \leq i \leq k} \bigcap_{\hat{h} \in \mathcal{U}_i} h = \bigcap_{1 \leq i \leq k} \bigcap_{n \in \mathbb{N}} h_n^i. \tag{6.19}$$

Here, the first equality holds by Lemma 6.17, the second equality holds by Definition 6.16, and the third equality holds because the hyperplanes added in the inseparable closure do not affect the intersection. By Lemma 5.3, a set of this form is the closure of a unique (principal) Roller class $Y_{\mathcal{U}}$.

We observe that the above construction is independent of the choice of descending chains with the desired properties, and hence the map $\mathcal{U} \mapsto Y_{\mathcal{U}}$ is $\text{Aut}(X)$ -equivariant. □

Lemma 6.20 *Let $x \in X^{(0)}$ and let $Y \subset \partial_R X$ be a convex set. Then the set of separating hyperplanes $\mathcal{W}(x, Y) = \mathcal{W}(x, \pi_Y(x))$ is a pruned UBS. If $x' \in X^{(0)}$, then $\mathcal{W}(x', Y) \sim \mathcal{W}(x, Y)$. Finally, $\mathcal{W}(x, [y]) \sim \mathcal{W}(x, y)$ for every $y \in \partial_R X$.*

Proof We begin by showing that $\mathcal{V} := \mathcal{W}(x, Y)$ is a UBS. Recall that $\mathcal{W}(x, Y)$ is the set of hyperplanes associated to $\mathfrak{H}_x^+ \Delta \mathfrak{H}_Y^+$. Since $x \in X^{(0)}$, it follows that \mathfrak{H}_x^+ satisfies the descending chain condition and since $Y \subset \partial_R X$, it follows that $\mathfrak{H}_Y^+ = \mathfrak{H}_Y^+$ has an infinite descending chain. Hence, \mathcal{V} is unidirectional and infinite. Since $\mathcal{V} = \mathcal{W}(x, Y)$ is an interval, it follows that it is inseparable and does not contain a facing triple, hence \mathcal{V} is a UBS. The fact that $\mathcal{W}(x, Y) = \mathcal{W}(x, \pi_Y(x))$ follows from Proposition 3.17 and shows that \mathcal{V} is pruned.

Let $x, x' \in X^{(0)}$. We claim that $\mathcal{W}(x, Y) \Delta \mathcal{W}(x', Y) \subset \mathcal{W}(x, x')$. By symmetry, it suffices to observe that a hyperplane \hat{h} separating x from Y but not x' from Y necessarily separates x from x' . Since the cardinality of $\mathcal{W}(x, x')$ is the distance $d(x, x') < \infty$, we conclude that the two UBSes are commensurate.

Finally, fix $y \in \partial X$. We show that $\mathcal{W}(x, [y]) \sim \mathcal{W}(x, y)$. Proposition 3.17 says that $\mathcal{W}(x, [y]) = \mathcal{W}(x, \pi_{[y]}(x))$, which implies that $\mathcal{W}(x, [y]) \cup \mathcal{W}(\pi_{[y]}(x), y) = \mathcal{W}(x, y)$. But $\pi_{[y]}(x) \in [y]$ by definition, hence $|\mathcal{W}(\pi_{[y]}(x), y)| < \infty$, completing the proof. □

Definition 6.21 (UBS of a Roller class) Let $Y \in \mathfrak{R}X$ be a Roller class, and fix a basepoint $x \in X^{(0)}$. Following Lemma 6.20, we call $\mathcal{W}(x, Y)$ a UBS representing the class Y , and denote it \mathcal{U}_Y . Since the commensurability class $[\mathcal{U}_Y]$ is independent of x , by Lemma 6.20, we do not include the basepoint x in the notation \mathcal{U}_Y .

It follows immediately that the assignment $Y \mapsto [\mathcal{U}_Y]$ is $\text{Aut}(X)$ -equivariant.

With this notation in place, we conclude:

Corollary 6.22 Fix $x \in X^{(0)}$ and $Y \in \mathfrak{R}X$. Then $Y = Y_{\mathcal{U}_Y}$. In particular, every Roller class arises as the principal class of the umbra of some UBS.

Proof Since $\mathcal{U}_Y = \mathcal{W}(x, Y)$ is a UBS, by Lemma 6.18 it is commensurate to $\bigsqcup_{i=1}^k \langle \langle \hat{h}_n^i \rangle \rangle$ for descending chains $\{h_n^i\}$, hence Eq. (6.19) gives

$$\bar{Y}_{\mathcal{U}_Y} = \bigcap_{1 \leq i \leq k} \bigcap_{n \in \mathbb{N}} h_n^i.$$

By construction, $\bar{Y} \subset h_n^i$ for every $\hat{h}_n^i \in \mathcal{U}_i$, hence $\bar{Y} \subset \bar{Y}_{\mathcal{U}_Y}$.

As in Lemma 6.18, let $Y_{\mathcal{U}_Y}$ denote the principal Roller class of $\bar{Y}_{\mathcal{U}_Y}$. We claim that $Y = Y_{\mathcal{U}_Y}$. So far, we have shown that $Y_{\mathcal{U}_Y} \leq Y$ in the partial order of Lemma 5.6 and so we are left to show that $Y \leq Y_{\mathcal{U}_Y}$. To this end, we fix a half-space $h \in \mathfrak{H}$ with $Y \subset h$ and now show that $Y_{\mathcal{U}_Y} \subset h$.

Let $x' \in h^* \cap X^{(0)}$ and set $\mathcal{U}'_Y = \mathcal{W}(x', Y)$. Observe that $\hat{h} \in \mathcal{U}'_Y$ by construction. Then by Lemma 6.20 we have that \mathcal{U}'_Y is a pruned UBS commensurate to \mathcal{U}_Y and so Lemma 6.17 implies $\bar{Y}_{\mathcal{U}_Y} = \bar{Y}_{\mathcal{U}'_Y}$. Finally, by Definition 6.16 we have that $Y_{\mathcal{U}_Y} \subset \bar{Y}_{\mathcal{U}_Y} \subset h$. □

Definition 6.23 (ℓ^1 -visible) Let $\gamma \subset X$ be a geodesic ray, in either the CAT(0) or combinatorial metric. Let $\mathcal{W}(\gamma)$ be the collection of hyperplanes intersecting γ , which is a UBS by [39, Section 3.5.1]. If \mathcal{U} is a UBS and $\mathcal{U} \sim \mathcal{W}(\gamma)$ for some combinatorial geodesic ray γ , we say that the class of \mathcal{U} is ℓ^1 -visible.

Remark 6.24 Let $Y \in \mathfrak{R}X$. Then, for any $x \in X^{(0)}$ and any $y \in Y$, Lemma 3.27 provides a combinatorial geodesic ray γ such that $\gamma(0) = x$ and $\gamma(\infty) = y \in Y$. By Definition 6.23 and Lemma 3.16, we have

$$\mathcal{W}(\gamma) = \mathcal{W}(x, y) \sim \mathcal{W}(x, [y]) = \mathcal{W}(x, \bar{[y]}).$$

Furthermore, by Lemma 6.20, the commensurability class $[\mathcal{W}(\gamma)] = [\mathcal{U}_Y]$ does not depend on the choice of $x \in X^{(0)}$ and $y \in Y$. In particular, every UBS representing a class Y is ℓ^1 -visible.

Lemma 6.25 Let \mathcal{U} be a UBS. Let $y \in \bar{Y}_{\mathcal{U}}$ and let γ be a combinatorial geodesic in X such that $\gamma(\infty) = y$. Then $\mathcal{U} \preceq \mathcal{W}(\gamma)$. In particular, if \mathcal{U} is maximal, then $\mathcal{U} \sim \mathcal{W}(\gamma)$.

Proof As in Lemma 6.18, \mathcal{U} is commensurate to $\bigsqcup_{i=1}^k \mathcal{U}_i$, where the \mathcal{U}_i are minimal and $\mathcal{U}_i = \langle \langle \hat{h}_n^i \rangle \rangle$ for a descending chain of half-spaces $\{h_n^i\}$. Thus Eq. (6.19) gives

$$\bar{Y}_{\mathcal{U}} = \bigcap_{1 \leq i \leq k} \bar{Y}_{\mathcal{U}_i} = \bigcap_{1 \leq i \leq k} \bigcap_{\hat{h} \in \mathcal{U}_i} h = \bigcap_{1 \leq i \leq k} \bigcap_{n \in \mathbb{N}} h_n^i.$$

Since $\gamma(m) \rightarrow y \in \bar{Y}_{\mathcal{U}}$, we must have $y \in h_n^i$ for every i and every n . At the same time, since the sequence $\{h_n^i\}$ is a descending chain for each i , we have that $x = \gamma(0) \notin h_n^i$ for sufficiently large n . Thus all but finitely many hyperplanes \hat{h}_n^i belong to $\mathcal{W}(x, y) = \mathcal{W}(\gamma)$. The inseparable closure of this cofinite subset of $\{\hat{h}_n^i\}$ must belong to $\mathcal{W}(\gamma)$, and by minimality of \mathcal{U}_i and Corollary 6.12, it follows that all but finitely many elements of each \mathcal{U}_i belong to $\mathcal{W}(\gamma)$. □

Corollary 6.26 Let \mathcal{U} be a maximal UBS. Then, for every $y \in Y_{\mathcal{U}}$, there is a combinatorial geodesic γ limiting to y , such that $\mathcal{U} \sim \mathcal{W}(\gamma)$. In particular, the commensurability class $[\mathcal{U}]$ is ℓ^1 -visible.

Proof Let $y \in Y_{\mathcal{U}}$. By Lemma 3.27, there is a combinatorial geodesic ray γ limiting to y . By Lemma 6.25, we have $\mathcal{U} \sim \mathcal{W}(\gamma)$. Thus, by Definition 6.23, \mathcal{U} is ℓ^1 -visible. \square

In Corollary 6.26, the conclusion that $[\mathcal{U}]$ is visible was previously shown by Hagen [39, Theorem 3.19]. The above proof, using the Roller boundary viewpoint, is considerably simpler.

6.4 Order-preserving maps

Recall from Definitions 5.5 and 6.4 that the partial order \leq on Roller classes is set containment of the associated sets of half-spaces, while the partial order \preceq on commensurability classes of UBSes is set containment up to commensurability. The following theorem says that these partial orders are closely related.

Theorem 6.27 *Let $\mathcal{UBS}(X)$ be the collection of commensurability classes of UBSes in X . Then there are $\text{Aut}(X)$ -equivariant functions $R^U: \mathcal{UBS}(X) \rightarrow \mathfrak{R}X$ and $U^R: \mathfrak{R}X \rightarrow \mathcal{UBS}(X)$, with the following properties.*

- (1) *The assignment $R^U: [\mathcal{U}] \mapsto Y_{\mathcal{U}}$ is order-preserving and onto.*
- (2) *The assignment $U^R: Y \mapsto [\mathcal{U}_Y]$ is an order-preserving section of R^U .*
- (3) *R^U restricts to a bijection between the set of \preceq -maximal classes of UBSes and the set of \leq -maximal Roller classes, with inverse U^R .*

The notation U^R stands for ‘‘Roller to unidirectional,’’ while R^U stands for ‘‘unidirectional to Roller.’’ This is inspired by the classical notation denoting the collection of maps from a set A to a set B as B^A .

Proof To prove (1), let \mathcal{U} be a UBS. By Lemma 6.18, there is a well-defined Roller class $Y_{\mathcal{U}}$ such that the umbra $\overline{Y}_{\mathcal{U}}$ is the Roller closure of $Y_{\mathcal{U}}$. By Lemma 6.17, $Y_{\mathcal{U}}$ only depends on the commensurability class $[\mathcal{U}]$. Hence the assignment $[\mathcal{U}] \mapsto Y_{\mathcal{U}}$ gives a well-defined function $R^U: \mathcal{UBS}(X) \rightarrow \mathfrak{R}X$. This function R^U is $\text{Aut}(X)$ -equivariant by Lemma 6.18, and onto by Corollary 6.22.

Next, we check that R^U respects the ordering. Lemma 6.17 implies that if $\mathcal{U} \preceq \mathcal{V}$, then $\overline{Y}_{\mathcal{V}} \subset \overline{Y}_{\mathcal{U}}$, hence the Roller classes satisfy $Y_{\mathcal{U}} \leq Y_{\mathcal{V}}$ by Lemma 5.6. Thus R^U is order-preserving, proving (1).

To prove (2), let $Y \in \mathfrak{R}X$. By Definition 6.21, there is a well-defined commensurability class of UBS representing Y , namely $[\mathcal{U}_Y] = R^U(Y) \in \mathcal{UBS}(X)$. Thus U^R is well-defined. Equivariance also follows from Definition 6.21. Next, observe that

$$R^U \circ U^R(Y) = R^U([\mathcal{U}_Y]) = Y_{\mathcal{U}_Y} = Y,$$

where the last equality follows from Corollary 6.22. Thus U^R is a section of R^U .

To check that U^R is order-preserving, let Y, Z be Roller classes with $Y \leq Z$. By Lemma 5.6, this means $\overline{Z} \subset \overline{Y}$. Fix a base vertex $x \in X^{(0)}$. By Lemma 3.27, there exist combinatorial geodesic rays γ_y, γ_z , originating at x , such that $\gamma_y(\infty) = y \in Y$ and $\gamma_z(\infty) = z \in Z$. Now, Remark 6.24 says that

$$\mathcal{U}_Y \sim \mathcal{W}(x, Y) \subset \mathcal{W}(x, Z) \sim \mathcal{U}_Z,$$

hence $[\mathcal{U}_Y] \preceq [\mathcal{U}_Z]$, as desired.

To prove (3), let \mathcal{U} be a maximal UBS. Choose a point $y \in Y_{\mathcal{U}} = \mathbb{R}^{\mathcal{U}}([\mathcal{U}])$. By Corollary 6.26, there is a combinatorial geodesic ray γ limiting to y , such that $\mathcal{U} \sim \mathcal{W}(\gamma)$. Thus, by Remark 6.24, we have

$$[\mathcal{U}] = [\mathcal{W}(\gamma)] = [\mathcal{U}_{Y_{\mathcal{U}}}] = \mathbb{U}^{\mathbb{R}}(Y_{\mathcal{U}}) = \mathbb{U}^{\mathbb{R}} \circ \mathbb{R}^{\mathcal{U}}([\mathcal{U}]). \tag{6.28}$$

Thus the restriction of $\mathbb{R}^{\mathcal{U}}$ to maximal classes in $\mathcal{UBS}(X)$ is invertible, and in particular bijective.

Equation (6.28) also shows that $\mathbb{R}^{\mathcal{U}}$ sends maximal classes in $\mathcal{UBS}(X)$ to maximal classes in $\mathfrak{R}X$. Indeed, suppose that $[\mathcal{U}]$ is maximal but $\mathbb{R}^{\mathcal{U}}([\mathcal{U}]) = Y_{\mathcal{U}} < Z$ for some Roller class Z . Then we would have $[\mathcal{U}] = \mathbb{U}^{\mathbb{R}}(Y_{\mathcal{U}}) < \mathbb{U}^{\mathbb{R}}(Z)$, contradicting the maximality of \mathcal{U} . That $\mathbb{U}^{\mathbb{R}}$ sends maximal classes to maximal classes is checked in the same way. \square

The following example shows that $\mathbb{R}^{\mathcal{U}}$ can fail to be injective.

Example 6.29 Consider the standard 1/8th-plane staircase, corresponding to squares whose vertices (x, y) satisfy $x \geq y \geq 0$. In this case, there are exactly 2 equivalence classes of minimal UBS, corresponding to vertical hyperplanes \mathcal{U}_V and horizontal hyperplanes \mathcal{U}_H . These are almost transverse, hence $\mathcal{U}_H \approx \mathcal{U}_V \cup \mathcal{U}_H$. However, $\mathbb{R}^{\mathcal{U}}[\mathcal{U}_H] = \mathbb{R}^{\mathcal{U}}([\mathcal{U}_V \cup \mathcal{U}_H])$. Thus $\mathbb{R}^{\mathcal{U}}$ is not injective.

On the other hand, the following corollary shows that $\mathbb{R}^{\mathcal{U}}$ is finite-to-one.

Corollary 6.30 *Let $Y \in \mathfrak{R}X$ be a Roller class. Then there are finitely many other Roller classes Z such that $Z \leq Y$.*

Proof For each $Z \leq Y$, Theorem 6.27.(2) says that $\mathbb{U}^{\mathbb{R}}(Z) \leq \mathbb{U}^{\mathbb{R}}(Y)$. Furthermore, $\mathbb{U}^{\mathbb{R}}$ is injective. By Lemma 6.8, there are only finitely many commensurability classes $[\mathcal{W}] \leq \mathbb{U}^{\mathbb{R}}(Y)$. Thus there are finitely many $Z \leq Y$. \square

Definition 6.31 Let $(\partial_{\Delta} X)'$ be the barycentric subdivision of $\partial_{\Delta} X$. Recall from Definition 6.9 that $\partial_{\Delta} X$ is the union of simplices corresponding to commensurability classes of UBSes. However, a simplex of $\partial_{\Delta} X$ may have proper faces that do not come from commensurability classes, as illustrated in Example 6.32 below. Hence, a vertex of $(\partial_{\Delta} X)'$ might fail to be associated to a UBS commensurability class.

Let $\partial_{\Delta}^{\text{UBS}} X$ be the simplicial realization of the partial order \leq on $\mathcal{UBS}(X)$. Then $\partial_{\Delta}^{\text{UBS}} X$ naturally embeds into $(\partial_{\Delta} X)'$. By the above paragraph, this embedding is not in general onto.

The following example, constructed by Dan Guralnik and Alessandra Iozzi, illustrates a subtlety in the definition of $\partial_{\Delta} X$ which shows the difference between $\partial_{\Delta} X$ and $\partial_{\Delta}^{\text{UBS}} X$.

Example 6.32 (Weird faces of simplices of $\partial_{\Delta} X$) Let X be the 3-dimensional CAT(0) cube complex given by the following data:

- The set of hyperplanes has the form $\mathcal{W}(X) = \{B_i\}_{i \geq 0} \sqcup \{S_j\}_{j \geq 0} \sqcup \{D_k\}_{k \geq 0}$.
- For each $A \in \{B, S, D\}$ and each $n \geq 1$, the hyperplanes $A_{n \pm 1}$ are separated by A_n .
- B_i crosses S_j if and only if $i \geq j$.
- S_j crosses D_k if and only if $j \geq k$.
- B_i crosses D_k if and only if $i \geq k$.
- If $i < j < k$, then S_j separates B_i from D_k .




Fig. 4 The cube complex X in Example 6.32, whose simplicial boundary has a “weird” 1-simplex. The hyperplanes B_i are shown in broken lines, the hyperplanes S_j are solid, and the hyperplanes D_k are dotted. This non-perspective drawing represents an embedding into \mathbb{R}^3 where the solid hyperplanes have also been rescaled by a factor of 2

Part of X is shown in Fig. 4.

Then $\partial_\Delta X$ is a single 2-simplex represented by the UBS $\{B_i\}_{i \geq 0} \cup \{S_j\}_{j \geq 0} \cup \{D_k\}_{k \geq 0}$. The 0-simplices are represented by the minimal UBSes $\{B_i\}, \{S_j\}, \{D_k\}$. The sets $\{B_i\} \cup \{S_j\}$ and $\{S_j\} \cup \{D_k\}$ are UBSes representing 1-simplices, but $\{B_i\} \cup \{D_k\}$ is not a UBS, because it is not inseparable (consider the last bullet). The complex $\partial_\Delta^{\text{UBS}} X$ is therefore obtained from the barycentric subdivision of the 2-simplex by deleting all of the cells that contain a vertex at the barycenter of $\{B_i\} \cup \{D_k\}$.

We remark that by [39, Theorem 3.23], a face of a simplex of $\partial_\Delta X$ can only fail to correspond to a commensurability class of UBSes when one of the 0-simplices of that face is ℓ^1 -invisible.

Having passed from $\partial_\Delta X$ to $\partial_\Delta^{\text{UBS}} X$, Theorem 6.27 has the following corollary.

Corollary 6.33 *The order-preserving maps R^U and U^R of Theorem 6.27 induce $\text{Aut}(X)$ -equivariant simplicial maps $R_\Delta^U: \partial_\Delta^{\text{UBS}} X \rightarrow \mathfrak{R}_\Delta X$ and $U_\Delta^R: \mathfrak{R}_\Delta X \rightarrow \partial_\Delta^{\text{UBS}} X$, as follows:*

- (1) $R_\Delta^U: \partial_\Delta^{\text{UBS}} X \rightarrow \mathfrak{R}_\Delta X$ is surjective.
- (2) $U_\Delta^R: \mathfrak{R}_\Delta X \rightarrow \partial_\Delta^{\text{UBS}} X$ is an injective section of R_Δ^U .
- (3) R_Δ^U is a homotopy equivalence with homotopy inverse U_Δ^R .

Proof By definition, $\mathfrak{R}_\Delta X$ is the simplicial realization of the partial order \leq on $\mathfrak{R}X$. Similarly, $\partial_\Delta^{\text{UBS}} X$ is the simplicial realization of the partial order \preceq on $\mathcal{UBS}(X)$. So, the first two enumerated assertions follow from Theorem 6.27.

To prove (3), let Y be a Roller class. Let $[\mathcal{V}] \in (\mathbb{R}^U)^{-1}(Y)$. Then $Y_{\mathcal{V}} = Y_{\mathcal{U}_Y}$. By Remark 6.24 and Lemma 6.25, we have $\mathcal{V} \preceq \mathcal{U}_Y$.

Using the notation Y for the 0-simplex of $\mathfrak{R}_\Delta X$ corresponding to the Roller class Y , we have just shown that $\mathbb{R}_\Delta^{U^{-1}}(Y)$ is a subcomplex of $\partial_\Delta^{\text{UBS}} X$ spanned by 0-simplices adjacent or equal to the 0-simplex $[Y]$, so $\mathbb{R}_\Delta^{U^{-1}}(Y)$ is topologically a cone. In particular, $\mathbb{R}_\Delta^{U^{-1}}(Y)$ is contractible. Quillen’s fiber theorem (see e.g. [11, Theorem 10.5]) implies that \mathbb{R}_Δ^U is a homotopy equivalence, such that any section is a homotopy inverse. In particular, \mathbb{U}_Δ^R is a homotopy inverse of \mathbb{R}_Δ^U . □

Remark 6.34 [Alternate strategy] The preceding corollary provides an explicit $\text{Aut}(X)$ -equivariant homotopy equivalence $\partial_\Delta^{\text{UBS}} X \rightarrow \mathfrak{R}_\Delta X$. Furthermore, it is possible, although somewhat tedious, to construct an explicit, $\text{Aut}(X)$ -equivariant homotopy equivalence $\partial_\Delta X \rightarrow \partial_\Delta^{\text{UBS}} X$. Composing these maps gives a homotopy equivalence $\partial_\Delta X \rightarrow \mathfrak{R}_\Delta X$, which is $\text{Aut}(X)$ -equivariant on the nose. This is slightly stronger than Proposition 10.12, which asserts that $\partial_\Delta X$ and $\mathfrak{R}_\Delta X$ are $\text{Aut}(X) \sim$ homotopy equivalent. We have chosen a different proof of Proposition 10.12 that is somewhat shorter, relying on the nerve theorem, and only yielding $\text{Aut}(X)$ -equivariance up to homotopy. Our ultimate goal is to relate the homotopy type of $\partial_\Delta X$ and $\mathfrak{R}_\Delta X$ to that of the Tits boundary $\partial_T X$, and in any case our homotopy equivalence $\mathfrak{R}_\Delta X \rightarrow \partial_T X$, provided by Theorem 9.17, is only equivariant up to homotopy.

7 The Tits boundary

In this section, we recall the definition of the Tits boundary $\partial_T X$, and establish some connections between points on the Tits boundary and UBSes. The main result of this section, Proposition 7.16, is a technical result about combining UBSes that will be crucial in the next section.

In this section, we will work extensively with the CAT(0) metric d_X on the cube complex X , hence most half-spaces considered here are CAT(0) half-spaces. Recall from Sect. 3.2 that if \hat{h} is a hyperplane with open carrier $N(\hat{h})$, the two vertex half-spaces associated to \hat{h} are denoted h, h^* . The two CAT(0) half-spaces associated to \hat{h} are the cubical convex hulls $\text{CH}(h)$ and $\text{CH}(h^*)$, respectively. Observe that $\text{CH}(h)$ and $\text{CH}(h^*)$ are exactly the components of $X \setminus N(\hat{h})$.

Definition 7.1 (The metric d_T) The Tits boundary of X , denoted $\partial_T X$, is the set of equivalence classes of CAT(0) geodesic rays in X , where rays γ, γ' are equivalent if the Hausdorff distance between them is finite. As a set, $\partial_T X$ coincides with the visual boundary of X .

We endow $\partial_T X$ with a metric, as follows. Given CAT(0) geodesic rays $\alpha, \beta: [0, \infty) \rightarrow X$ with $\alpha(0) = \beta(0) = x$, the Alexandrov angle formed by α, β at x is

$$\angle_x(\alpha, \beta) = \cos^{-1} \left(\lim_{t, t' \rightarrow 0} \frac{d_X(\alpha(t), \beta(t'))^2 - t^2 - (t')^2}{2tt'} \right) \tag{7.2}$$

In other words, $\angle_x(\alpha, \beta)$ is computed by applying the law of cosines to a triangle with sides along α and β , and then taking a limit as the triangle shrinks. Let $a = [\alpha] \in \partial_T X$ and

$b = [\beta] \in \partial_T X$ be the the equivalence classes of α, β respectively. Following [7, Definition II.9.4], we define

$$\angle_T(a, b) = \angle_T([\alpha], [\beta]) := \sup_{x \in X} \angle_x(\alpha_x, \beta_x), \tag{7.3}$$

where α_x, β_x are the rays emanating from x and representing a, b respectively. Then \angle_T induces a length metric d_T on $\partial_T X$ in the standard way, making it a CAT(1) space. See [7, Theorem II.9.13 and Definition II.9.18]. We always equip $\partial_T X$ with the metric d_T . By [7, Proposition II.9.5], the automorphism group $\text{Aut}(X)$ acts on $\partial_T X$ by isometries.

We remark that the original definition of the Alexandrov angle [7, Definition I.1.12] is somewhat more involved. In our setting, since X is a CAT(0) space, [7, Proposition II.3.1] says the limit in (7.2) exists and is equal to the Alexandrov angle in the original definition. Compare [7, Remark I.1.13.(4)].

Definition 7.1 has the following generalization to the cuboid setting.

Definition 7.4 (*Cuboid Tits boundary*) Fix a group G acting on X by cubical automorphisms, as well as a G -admissible hyperplane rescaling ρ , as in Definition 4.3. Then the cuboid metric d_X^ρ is still a CAT(0) metric by Lemma 4.2. Applying Definition 7.1 to the CAT(0) space (X, d_X^ρ) , we obtain the *cuboid Tits boundary* $\partial_T^\rho X = \partial_T(X, d_X^\rho)$, with an associated CAT(1) metric d_T^ρ . The restricted automorphism group $\text{Aut}(X^\rho)$ acts on $\partial_T^\rho X$ by isometries. Since ρ is a G -admissible rescaling, we have $G \subset \text{Aut}(X^\rho)$ by Definition 4.4, hence $\partial_T^\rho X$ still admits a G -action.

All of the results of this section extend to a cuboid metric d_X^ρ . See Remarks 7.8, 7.12, and 7.15 for details.

Throughout this section, we use Greek letters α, β, γ to denote geodesic rays, and the corresponding Roman letters a, b, c to denote equivalence classes. We write $a = [\alpha] = \alpha(\infty)$ to denote the endpoint of α in $\partial_T X$.

7.1 Deep sets

Every point of the Tits boundary has an associated *deep set* of half-spaces.

Definition 7.5 (*Deep set of a Tits point*) Let $a \in \partial_T X$, and let $\alpha : [0, \infty) \rightarrow X$ be a CAT(0) geodesic representing a . We define the *deep set*

$$D_a = \{h \in \mathcal{H} : \text{for all } R > 0 \text{ there is } t_R > 0 \text{ s.t. } \mathcal{N}_R(\alpha[t_R, \infty)) \subset \text{CH}(h)\}.$$

It is immediate to check that the definition does not depend on the representative geodesic α . If $h \in D_a$, we say that h is *deep* for a .

Lemma 7.6 *Let $\alpha : [0, \infty) \rightarrow X$ be a CAT(0) geodesic ray with $\alpha(0) \in X^{(0)}$, and let $\hat{h} \in \mathcal{W}(\alpha)$. Then the following hold:*

(1) *Define $f(t) = d_X(\alpha(t), \hat{h})$, and let $t_0 = \alpha^{-1}(\hat{h})$. Then, for $t_0 \leq s \leq t$, we have*

$$f(t) - f(s) \geq A_0 \cdot (t - s),$$

where $A_0 > 0$, and A_0 may depend on \hat{h} and α but not on s, t .

(2) *The hyperplane \hat{h} has a deep half-space $h \in D_a$.*

Consequently, $\mathcal{W}(\alpha)$ is a pruned UBS.

Proof By [7, Corollary II.2.5], f is a convex, 1-Lipschitz function. By convexity, f has left and right derivatives at every point in $(0, \infty)$, and these are non-decreasing. Since f is Lipschitz and convex, the left and right derivatives coincide almost everywhere, hence f' is well-defined almost everywhere, and non-decreasing at all points where it is defined.

Let $t_0 = \alpha^{-1}(\hat{h})$. Since the carrier $N(\hat{h}) \cong \hat{h} \times (-\frac{1}{2}, \frac{1}{2})$ is an ℓ^2 product and $\alpha(t_0) \in \hat{h}$ lies on the midcube of a cube, there is an initial segment of α after crossing \hat{h} on which $f(t)$ is an increasing linear function. Therefore, $f'(t) = f'(t_0) = A_0 > 0$ on this segment. (In fact, $f'(t_0) = \sin \theta_0$, where θ_0 is the angle between α and \hat{h} in the ℓ^2 metric on a cube.) Hence $f'(t) \geq A_0$ for all $t \geq t_0$ at which $f'(t)$ is defined. Since f is Lipschitz, it follows that for all $t_0 \leq s \leq t$, we have

$$f(t) - f(s) = \int_s^t f'(z) dz \geq A_0 \cdot (t - s),$$

proving (1). Now, (1) immediately implies (2). The conclusion that $\mathcal{W}(\alpha)$ is pruned is then immediate from Definition 6.6. □

The following lemma is a strengthening of [19, Lemma 2.27]. While [19, Lemma 2.27] proves that D_a contains an infinite descending chain, we prove the stronger statement that D_a is non-terminating.

Lemma 7.7 *For a point $a \in \partial_T X$, define $\bar{Y}_a = \bigcap_{h \in D_a} h$. Then D_a and \bar{Y}_a have the following properties:*

- (1) D_a is consistent, and therefore $D_a = \mathfrak{H}_{\bar{Y}_a}^+$.
- (2) Every half-space $h_0 \in D_a$ belongs to a descending chain $\{h_0, h_1, \dots\} \subset D_a$. Furthermore, if $\hat{h}_0 \in \mathcal{W}(\alpha)$ for a ray $\alpha \in a$, then $\hat{h}_n \in \mathcal{W}(\alpha)$ for all n .
- (3) $\bar{Y}_a \subset \partial_R X$ is the closure of a (unique, principal) Roller class Y_a .

Proof The consistency of D_a is immediate from the definition. Thus, by Proposition 3.14.(2), we obtain $D_a = \mathfrak{H}_{\bar{Y}_a}^+$. This proves (1).

To prove (2), let $h_0 \in D_a$. By [7, Proposition II.8.2], there exists a geodesic ray α representing a such that $\alpha(0) \in h_0^*$. Since $h_0 \in D_a$, Definition 7.5 implies that α must deviate arbitrarily far from \hat{h}_0 . Thus, by Lemma 3.1, a point $\alpha(t)$ for large t has the property that $|\mathcal{W}(\alpha(t), y)| \geq 1$ for every $y \in N(\hat{h}_0)$. Now, Proposition 3.17 says that there is at least one hyperplane \hat{h}_1 separating $\alpha(t)$ from $N(\hat{h}_0)$. Observe that $\hat{h}_1 \in \mathcal{W}(\alpha)$ by construction, and orient \hat{h}_1 by choosing the half-space $h_1 \subset h_0$. Now, Lemma 7.6 says that $h_1 \in D_a$. Continuing inductively, we obtain a descending chain $\{h_0, h_1, \dots\} \subset D_a$ where $\hat{h}_n \in \mathcal{W}(\alpha)$ for all n , proving (2).

By (1), $D_a = \mathfrak{H}_{\bar{Y}_a}^+$ is consistent, and by (2), D_a contains an infinite descending chain. Thus $\bar{Y}_a = \bigcap_{h \in D_a} h$ is a convex, nonempty subset of $\partial_R X$. Since D_a is non-terminating by (2), Lemma 5.3 implies that \bar{Y}_{D_a} is the closure of a principal Roller class $Y_a \in \mathfrak{R}X$. (Compare Definition 5.4.) This proves (3). □

Remark 7.8 The previous lemmas generalize readily to a cuboid metric d_X^ρ . Indeed, Lemma 7.6 generalizes because hyperplane carriers are still ℓ^2 products, and because [7, Corollary II.2.5] works in any CAT(0) space. Similarly, Lemma 7.7 generalizes because [7, Proposition II.8.2] works in every CAT(0) space and, because d_X^ρ is quasi-isometric to the ℓ^1 -metric on $X^{(0)}$.

7.2 CAT(0) rays and their UBSes

Given a CAT(0) geodesic ray γ in X , starting at a point of $X^{(0)}$, recall that $\mathcal{W}(\gamma)$ denotes the set of hyperplanes h that intersect γ . It is easy to see (cf. [39, Section 3]) that $\mathcal{W}(\gamma)$ is a UBS.

Throughout this subsection, fix $a, b \in \partial_T X$ and let α, β be CAT(0) geodesic rays representing a and b , respectively. We can assume that $\alpha(0) = \beta(0) = x \in X^{(0)}$.

The following lemma says that when $\mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$ has finite symmetric difference with a UBS, it actually is a UBS. Several subsequent lemmas build toward Proposition 7.16, which will show that $\mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$ is in fact a UBS associated to a ray γ .

Lemma 7.9 *Suppose that $\mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$ is commensurate with a UBS. Define sets $\mathcal{A} = \mathcal{W}(\alpha) \setminus \mathcal{W}(\beta)$ and $\mathcal{B} = \mathcal{W}(\beta) \setminus \mathcal{W}(\alpha)$ and $\mathcal{C} = \mathcal{W}(\alpha) \cap \mathcal{W}(\beta)$. Then*

- (1) *Each of \mathcal{A}, \mathcal{B} is infinite or empty.*
- (2) *Every element of \mathcal{A} crosses every element of \mathcal{B} .*
- (3) *$\mathcal{W}(\alpha) \cup \mathcal{W}(\beta) = \mathcal{A} \sqcup \mathcal{B} \sqcup \mathcal{C}$ is a UBS.*

Proof Let $\mathcal{U} = \mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$. Since \mathcal{U} is commensurate with a UBS by hypothesis, it must be infinite and unidirectional. We will check the other properties of a UBS at the end of the proof.

By Lemma 7.6, every hyperplane $\hat{h} = \hat{h}_0 \in \mathcal{W}(\alpha)$ corresponds to a deep half-space $h_0 \in D_a$, and by Lemma 7.7, h_0 is the start of a descending chain $\{h_0, h_1, \dots\} \subset D_a$, where $h_n \in \mathcal{W}(\alpha)$ for every n . Similarly, every hyperplane $k_0 \in \mathcal{W}(\beta)$ defines a descending chain $\{k_0, k_1, \dots\} \subset D_b$, where $k_m \in \mathcal{W}(\beta)$ for every m .

If $\mathcal{A} = \mathcal{W}(\alpha) \setminus \mathcal{W}(\beta) \neq \emptyset$, then any hyperplane $\hat{h}_0 \in \mathcal{A}$ defines a chain $\{\hat{h}_0, \hat{h}_1, \dots\} \subset \mathcal{W}(\alpha)$, as above. Furthermore, \hat{h}_0 separates every \hat{h}_n from β . Thus $\hat{h}_n \in \mathcal{A}$, and \mathcal{A} is infinite. Similarly, \mathcal{B} must be empty or infinite, proving (1).

Next, suppose for a contradiction that $\hat{h}_0 \in \mathcal{A}$ is disjoint from $\hat{k}_0 \in \mathcal{B}$. Then the deep half-spaces $h_0 \in D_a$ and $k_0 \in D_b$ are also disjoint. But then the disjoint chains $\{\hat{h}_n\} \subset \mathcal{W}(\alpha)$ and $\{\hat{k}_m\} \subset \mathcal{W}(\beta)$ contradict the fact that $\mathcal{U} = \mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$ is unidirectional. This proves (2).

Finally, we show that \mathcal{U} is a UBS. We have already checked that \mathcal{U} is infinite and unidirectional. To check inseparability, suppose that $\hat{\ell}$ is a hyperplane of X that separates $\hat{h}, \hat{k} \in \mathcal{U}$. Since we already know that $\mathcal{W}(\alpha)$ and $\mathcal{W}(\beta)$ are inseparable, it suffices to check the case that $\hat{h} \in \mathcal{A}$ and $\hat{k} \in \mathcal{B}$. But then \hat{h} and \hat{k} must cross by (2), a contradiction. Similarly, since $\mathcal{W}(\alpha)$ and $\mathcal{W}(\beta)$ contain no facing triples, any potential facing triple in \mathcal{U} must contain at least one hyperplane $\hat{h} \in \mathcal{A}$ and at least one hyperplane $\hat{k} \in \mathcal{B}$. But then \hat{h} and \hat{k} must cross, hence \mathcal{U} cannot contain any facing triples. Thus \mathcal{U} is a UBS, proving (3). \square

Recall from Definition 6.16 that given a UBS \mathcal{U} , we have the umbra $\overline{Y}_{\mathcal{U}} \subset \overline{X}$. Following Proposition 3.17, let $\pi_{Y_{\mathcal{U}}} = \pi_{\overline{Y}_{\mathcal{U}}} : \overline{X} \rightarrow Y_{\mathcal{U}}$ be the gate projection to $\overline{Y}_{\mathcal{U}}$.

Lemma 7.10 *Suppose α and β are geodesic rays starting at $x \in X^{(0)}$, and that $\mathcal{U} = \mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$ is commensurate with a UBS. Then $\mathcal{U} = \mathcal{W}(x, \pi_{Y_{\mathcal{U}}}(x))$.*

Proof Since the geodesic rays start at the vertex x , Lemma 7.7 says that all the hyperplanes they cross are deep, i.e. $\mathcal{W}(x, Y_a) = \mathcal{W}(\alpha)$ and $\mathcal{W}(x, Y_b) = \mathcal{W}(\beta)$, hence $\mathcal{U} = \mathcal{W}(x, Y_a) \cup \mathcal{W}(x, Y_b)$. In addition, Lemma 7.9 says that \mathcal{U} is a UBS.

Now, Lemma 6.17 gives $\overline{Y}_{\mathcal{U}} = \overline{Y}_a \cap \overline{Y}_b$, and in particular $\mathfrak{H}_{\overline{Y}_{\mathcal{U}}}^+ = D_a \cup D_b$. By Lemma 6.20 and Remark 6.24, we have $\mathcal{W}(x, \overline{Y}_{\mathcal{U}}) = \mathcal{W}(x, Y_{\mathcal{U}}) = \mathcal{W}(x, \pi_{Y_{\mathcal{U}}}(x))$. We must show

that this UBS equals $\mathcal{U} = \mathcal{W}(x, Y_a) \cup \mathcal{W}(x, Y_b)$. To this end we have by Proposition 3.14:

$$\begin{aligned} \mathfrak{H}_{\pi_{Y_a}(x)}^+ &= [\mathfrak{H}_x^+ \setminus D_a^*] \cup D_a, \\ \mathfrak{H}_{\pi_{Y_b}(x)}^+ &= [\mathfrak{H}_x^+ \setminus D_b^*] \cup D_b, \\ \mathfrak{H}_{\pi_{Y_{\mathcal{U}}}(x)}^+ &= [\mathfrak{H}_x^+ \setminus (D_a^* \cup D_b^*)] \cup (D_a \cup D_b). \end{aligned}$$

Then a standard Venn diagram argument shows that

$$\begin{aligned} \mathfrak{H}_x^+ \setminus \mathfrak{H}_{\pi_{Y_{\mathcal{U}}}(x)}^+ &= \mathfrak{H}_x^+ \setminus ([\mathfrak{H}_x^+ \setminus (D_a^* \cup D_b^*)] \cup (D_a \cup D_b)) \\ &= (\mathfrak{H}_x^+ \cap (D_a^* \cup D_b^*)) \setminus (D_a \cup D_b). \end{aligned}$$

The set of hyperplanes on the left-hand side is $\mathcal{W}(x, \pi_{Y_{\mathcal{U}}}(x))$, while the set of hyperplanes on the right-hand side is $\mathcal{W}(x, Y_a) \cup \mathcal{W}(x, Y_b)$. □

Lemma 7.11 (Tits distance bound) *Suppose that $\mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$ is commensurate with a UBS. Then $d_T(a, b) \leq \pi/2$.*

Proof We will show that $\pi/2$ is an upper bound on the Alexandrov angle between rays representing a, b based at an arbitrary point, and then use (7.3). The bound will be produced by finding a convex subspace, isometric to a Euclidean cube, which has one corner at the basepoint and which contains nontrivial initial segments of the rays.

Angle bound at a vertex: Let $y \in X$ be an arbitrary vertex. Let α', β' be CAT(0) geodesic rays emanating from y and representing a, b respectively (so, α is asymptotic to α' and β is asymptotic to β'). Let $\mathcal{U} = \mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$, and let $\mathcal{U}' = \mathcal{W}(\alpha') \cup \mathcal{W}(\beta')$. Then \mathcal{U} is commensurate to \mathcal{U}' , and hence commensurate to a UBS. Lemma 7.9 then implies that \mathcal{U}' is actually a UBS.

Now let $X_y \subset X$ be the cubical convex hull of $\alpha' \cup \beta'$. Then the set of hyperplanes crossing X_y is exactly \mathcal{U}' . Moreover, by Lemma 7.10, there exists a point \bar{y} in the Roller boundary such that \mathcal{U}' is precisely the set of hyperplanes separating y from \bar{y} .

Let \triangleleft be the partial order on \mathcal{U}' defined by declaring that $\hat{h} \triangleleft \hat{h}'$ if the hyperplane \hat{h} separates y from \hat{h}' . Since every hyperplane of X_y separates y from \bar{y} , any hyperplanes $\hat{h}, \hat{h}' \in \mathcal{U}'$ are either \triangleleft -comparable, or they cross.

Note that if $\hat{h} \in \mathcal{U}'$ is dual to a 1-cube of X_y incident to y , then \hat{h} is \triangleleft -minimal. Hence the set \mathcal{U}'_{\min} of \triangleleft -minimal hyperplanes is nonempty. Since \triangleleft -incomparable hyperplanes cross, \mathcal{U}'_{\min} is a set of pairwise-crossing hyperplanes of X . Since X_y is convex and every hyperplane in \mathcal{U}' crosses X_y , we have a cube $C \subset X_y$ such that the hyperplanes crossing C are exactly those in \mathcal{U}'_{\min} .

Note $y \in C$. Also note that any cube C' of X_y with $y \in C'$ satisfies $C' \subset C$. Thus any nontrivial CAT(0) geodesic segment in X_y emanating from y has a nontrivial initial subsegment lying in C . In particular, α' and β' have nontrivial initial segments lying in C . Since the ambient CAT(0) metric restricts to the Euclidean metric on the cube C , the Alexandrov angle made by α', β' at y satisfies

$$\angle_y(\alpha', \beta') = \lim_{t, t' \rightarrow 0} \angle_y(\alpha'(t), \beta'(t)) \leq \pi/2,$$

because any two segments in a cube emanating from a common corner make an angle at most $\pi/2$.

Angle bound at arbitrary points: We now bound the angle made by rays representing a, b and emanating from non-vertex points, by an identical argument taking place in an appropriate cubical subdivision of X . Let $y \in X$ be an arbitrary point. Perform a cubical

subdivision of X to obtain a CAT(0) cuboid complex X' in which y is a vertex, so that the CAT(0) cuboid metric coincides with the original metric on X .

More precisely, for each hyperplane \hat{h} , identify its carrier with $\hat{h} \times [-\frac{1}{2}, \frac{1}{2}]$. For each \hat{h} such that $y \in \hat{h} \times (-\frac{1}{2}, \frac{1}{2})$, let $\epsilon_{\hat{h}}$ be such that $y \in \hat{h} \times \{\epsilon_{\hat{h}}\}$. The (geodesically convex) subspace $\hat{h} \times \{\epsilon_{\hat{h}}\}$ has a natural cubical structure with an n -cube for each $(n + 1)$ -cube intersecting \hat{h} . We subdivide X so that the cubes of $\hat{h} \times \{\epsilon_{\hat{h}}\}$ are subdivided cubes whose edges are segments in X whose lengths are inherited from X . The result is a cuboid complex X' such that the identity map $X \rightarrow X'$ is an isometry in the CAT(0) metric, preserves the median, and sends y to a vertex. Since the hyperplanes of X' are parallel copies of hyperplanes of X , the set of hyperplanes of X' that cross $\alpha \cup \beta$ continues to be a UBS, so Lemma 7.10 still applies. (Compare Remark 7.8.) We can thus argue exactly as before to see that $\angle_y(\alpha', \beta') \leq \pi/2$.

Conclusion: We have shown that, for all $y \in X$, letting α', β' be the rays based at y and representing a, b respectively, we have $\angle_y(\alpha', \beta') \leq \pi/2$. Taking the supremum over all $y \in X$, as in (7.3), we see that the angular metric satisfies $\angle_T(a, b) \leq \pi/2$. Now, Definition 7.1 gives $d_T(a, b) \leq \pi/2$, as required. \square

Remark 7.12 Lemmas 7.9 and 7.10 extend immediately to the cuboid setting, because they use only hyperplane combinatorics and prior lemmas. The above proof of Lemma 7.11 also extends to cuboids, because it uses prior lemmas combined with CAT(0) geometry. Indeed, the above argument uses cuboids in combination with Remark 7.8 to prove the desired angle bound for arbitrary basepoints in X . The key conceptual reason why the argument works for cuboids is that every pair of geodesic segments in a cuboid, emanating from a corner, meets at angle at most $\pi/2$.

The reliance on cuboids to prove the desired result for cubes can be avoided, as follows. First, prove the angle bound at a vertex of X , exactly as above. Then, let X' be the cubical subdivision of X , and observe that the CAT(0) geodesics in X' are exactly the same as those in X . (See e.g. [31, Section 2.1.6].) Thus the same argument proves the angle bound at every vertex of X' .

Continuing to subdivide by induction, we obtain a set $V_\infty \subset X$ of points that are vertices of the subdivision at some (and hence all subsequent) stages. Note that V_∞ intersects each of the original cubes of X in a dense subset, and in particular contains all the original vertices. The desired angle bound holds when the rays in question are based at any point in V_∞ . Now, fixing $a, b \in \partial_T X$, we can consider the function $x \mapsto \angle_x(\alpha, \beta)$, where α, β are the rays representing a, b and starting at x . This function need not be continuous on X (it is upper semicontinuous [7, Proposition II.9.(2)]), but it can be shown to be continuous on each open cube of X (of any dimension). Since it is bounded above by $\pi/2$ on a dense subset of each such cube, we conclude that $\sup_x \angle_x(\alpha, \beta) = \angle_T(a, b) \leq \pi/2$.

Finally, observe that the statement and proof of Lemma 7.11 fail completely if we modify the angles of the cubes, precisely because two segments emanating from a vertex of a parallelogram might meet at a large angle.

Our main goal in this section is to show that if α, β are geodesic rays with common initial point, and $\mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$ is a UBS, then any geodesic ray γ representing an interior point on the Tits geodesic from $\alpha(\infty)$ to $\beta(\infty)$ must cross all the hyperplanes in $\mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$. Before proving this in Proposition 7.16, we will need a few auxiliary results, starting with a corollary of Lemma 6.15.

Corollary 7.13 (Crossing all dominants implies crossing everyone) *Suppose $U = \mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$ is a UBS. Let Y be the cubical convex hull of $\alpha \cup \beta$. Let $\gamma : [0, \infty) \rightarrow X$ be a CAT(0)*

geodesic ray with $\gamma(0) = \alpha(0) = \beta(0)$. Suppose that γ is contained in Y and crosses every dominant hyperplane in \mathcal{U} . Then γ crosses every hyperplane in \mathcal{U} .

Proof This follows immediately from Lemma 6.15.(4) applied to $\mathcal{U} = \mathcal{W}(Y)$. □

Next, we study angles at which rays cross dominant hyperplanes.

Lemma 7.14 (Lower angle bound for single rays) *Let α be a CAT(0) geodesic ray with $\alpha(0) \in X^{(0)}$. Then there exists $\kappa > 0$ such that the following holds. Let $\mathcal{D}_d, \dots, \mathcal{D}_k \subset \mathcal{W}(\alpha)$ be the dominant minimal UBSes provided by applying Lemma 6.15 to $\mathcal{W}(\alpha)$, and let $i \in \{d, \dots, k\}$. Then \mathcal{D}_i contains a chain $\{\hat{u}_n\}_{n \geq 0}$ of hyperplanes such that \mathcal{D}_i is commensurate with the inseparable closure of $\{\hat{u}_n\}_{n \geq 0}$, and*

$$\angle_{y_n}(\alpha, \hat{u}_n) \geq \kappa$$

for all $n \geq 0$, where y_n is the point $\alpha \cap \hat{u}_n$.

Proof Recall from Lemma 7.6 that $\mathcal{W}(\alpha)$ is a pruned UBS. Thus Lemma 6.15 applies to $\mathcal{W}(\alpha)$.

Now, fix $i \in \{d, \dots, k\}$. Lemma 6.15 says that $\mathcal{D}_i = \langle\langle \{\hat{h}_m\}_{m \geq 0} \rangle\rangle$, where each \hat{h}_m is dominant. This means each \hat{h}_m crosses all but finitely many hyperplanes in $\mathcal{W}(\alpha) \setminus \mathcal{D}_i$.

The divergence $f(t)$ of α from \hat{h}_0 : Define $f(t) = d_X(\alpha(t), \hat{h}_0)$, and let $T_0 = \alpha^{-1}(\hat{h}_0)$. By Lemma 7.6.(1), there is a constant $A_0 > 0$ (depending on \hat{h}_0 but not s, t) such that

$$f(t) - f(s) \geq A_0 \cdot (t - s)$$

for all $T_0 \leq s \leq t$.

Hyperplane count: We wish to produce a constant C such that any length- C subpath of α crosses an element of \mathcal{D}_i .

Let $s_0 \geq T_0$ be sufficiently large that $\alpha([0, s_0])$ crosses all of the finitely many hyperplanes in $\mathcal{W}(\alpha) \setminus \mathcal{D}_i$ not crossing \hat{h}_0 . Suppose that $s_0 \leq s \leq t$. Then all hyperplanes crossing $\alpha([s, t])$ either cross \hat{h}_0 or belong to \mathcal{D}_i . Assume that $\alpha([s, t])$ does not cross any element of \mathcal{D}_i .

Let $H_{st} = \text{CH}(\alpha([s, t]))$ denote the cubical convex hull of $\alpha([s, t])$. There may be hyperplanes crossing H_{st} that do not cross $\alpha([s, t])$, because $\alpha(s), \alpha(t)$ need not be vertices. However, we can bound the number of such hyperplanes as follows. Let $\mathcal{W}_{\text{bad}} \subset \mathcal{W}(H_{st}) \subset \mathcal{W}(\alpha)$ be the set of hyperplanes that cross H_{st} but do not cross $\alpha([s, t])$. Note that $\mathcal{W}(H_{st})$ inherits a partial order from $\mathcal{W}(\alpha)$, by restricting the partial order \triangleleft from the proof of Lemma 7.11. Suppose that $\hat{h} \in \mathcal{W}(H_{st})$ is neither \triangleleft -maximal nor \triangleleft -minimal in $\mathcal{W}(H_{st})$. Then there exist $\hat{u}, \hat{v} \in \mathcal{W}(H_{st})$ such that \hat{h} separates \hat{u} from \hat{v} . By the definition of the convex hull, \hat{h} cannot separate \hat{u} or \hat{v} from $\alpha([s, t])$, so \hat{h} must cross $\alpha([s, t])$. Thus every hyperplane in \mathcal{W}_{bad} is \triangleleft -minimal or \triangleleft -maximal. Since incomparable hyperplanes have to cross, we conclude that $|\mathcal{W}_{\text{bad}}| \leq 2 \dim(X)$, hence there are at most $D = 2 \dim X$ hyperplanes that cross H_{st} but not $\alpha([s, t])$. The rest of the hyperplanes crossing H_{st} must cross \hat{h}_0 , since they cross $\alpha([s, t])$, and we have chosen $s \geq s_0$ and assumed $\mathcal{W}(\alpha([s, t])) \cap \mathcal{D}_i = \emptyset$.

Consider the CAT(0) closest-point projections $p: X \rightarrow \hat{h}_0$ and $q: X \rightarrow H_{st}$. By [46, Lemma 2.10], the d_X -convex hull of $p(H_{st}) \cup q(\hat{h}_0)$ is isometric to $p(H_{st}) \times [0, d_X(\hat{h}_0, H_{st})]$, where $p(H_{st})$ is identified with $p(H_{st}) \times \{0\}$, and $q(\hat{h}_0)$ is identified with $p(H_{st}) \times \{d_X(\hat{h}_0, H_{st})\}$. By [46, Lemma 2.14], the hyperplanes that cross $q(\hat{h}_0)$ are precisely those that cross \hat{h}_0 and H_{st} . So, there are at most D hyperplanes crossing H_{st} but not $q(\hat{h}_0)$. By

Lemma 3.1, the points $\alpha(s), \alpha(t)$ are thus both within distance $\lambda_0 D + \lambda_1$ of points in $q(\hat{h}_0)$, where $\lambda_0 D + \lambda_1$ depends only on D . Thus

$$d_X(\alpha(t), \hat{h}_0) \leq \lambda_0 D + \lambda_1 + d_X(H_{st}, \hat{h}_0) \leq \lambda_0 D + \lambda_1 + d_X(\alpha(s), \hat{h}_0).$$

In other words, since $s, t \geq s_0 \geq T_0$, we have shown

$$A_0(t - s) \leq f(t) - f(s) \leq \lambda_0 D + \lambda_1.$$

Thus, for any $0 \leq s \leq t$ (without assuming $s \geq s_0$), we get

$$t - s \leq \frac{\lambda_0 D + \lambda_1}{A_0} + s_0.$$

Letting $C = \frac{\lambda_0 D + \lambda_1}{A_0} + s_0 + 1$, we have shown that any subsegment of α of length at least C crosses an element of \mathcal{D}_i . Hence, for any $0 \leq s \leq t$, we have that $\alpha([s, t])$ crosses at least $\lfloor (t - s)/C \rfloor$ elements of \mathcal{D}_i .

Big angle hyperplanes: Let σ be a subsegment of α . Consider the hyperplanes in \mathcal{D}_i crossing σ . Any two such hyperplanes are either disjoint or not, and any collection of pairwise crossing hyperplanes has size at most $\dim X$. So, if there are more than $\text{Ram}(3, \dim X + 1)$ hyperplanes in \mathcal{D}_i crossing σ , then there are three disjoint such hyperplanes, where $\text{Ram}(\cdot, \cdot)$ denotes the Ramsey number. Hence, if σ has length $|\sigma| = C \cdot \text{Ram}(3, \dim X + 1) + 1$, then σ crosses three disjoint hyperplanes $\hat{u}, \hat{v}, \hat{w} \in \mathcal{D}_i$. Without loss of generality, say \hat{v} separates \hat{u}, \hat{w} . Hence $\alpha \cap N(\hat{v})$ is a subpath of α lying between the points $\alpha \cap \hat{u}$ and $\alpha \cap \hat{w} \in \sigma$, so

$$|\alpha \cap N(\hat{v})| \leq C \cdot \text{Ram}(3, \dim X + 1) + 1 = L.$$

Note that L is independent of the hyperplane \hat{v} . Indeed, L depends on $A_0, \lambda_0, \lambda_1, s_0$, and hence is determined by X , the hyperplane \hat{h}_0 and the geodesic α . Now, using the fact that $N(\hat{v})$ is isometric to a product of the form $\hat{v} \times (-\frac{1}{2}, \frac{1}{2})$, we see that

$$\angle_{\alpha \cap \hat{v}}(\hat{v}, \alpha) \geq \sin^{-1}(\frac{1}{L}),$$

which we denote by κ .

Since we can do the above procedure for infinitely many disjoint length- L segments in α , we find infinitely many hyperplanes in \mathcal{D}_i making an angle at least κ with α . Since any infinite subset of a UBS contains a chain, we thus have a chain $\{\hat{u}_n\}_{n \geq 0} \subset \mathcal{D}_i$ with this property. Finally, the inseparable closure of $\{\hat{u}_n\}_{n \geq 0}$ is a UBS contained in \mathcal{D}_i , and hence commensurate with \mathcal{D}_i since \mathcal{D}_i is minimal. This verifies the statement for the given \mathcal{D}_i , and we conclude by replacing κ with the minimal κ for the various $i \in \{d, \dots, k\}$. \square

Remark 7.15 Extending Lemma 7.14 to a cuboid metric d_X^ρ requires a bit of care. The above proof relies on some results of of Huang [46], which are written in the context of a cube complex with finitely many isometry types of cells. This hypothesis may fail in (X, d_X^ρ) . Fortunately, we only need to use Huang’s results in a *finite* cuboid complex, namely the convex hull of $H_{st} \cup p(H_{st})$, where H_{st} is itself finite and p is a certain CAT(0) projection. Thus Huang’s results [46, Lemmas 2.10 and 2.14] apply to the subcomplex we need.

Some constants in the proof would need to be adjusted for the cuboid metric. The constants λ_0, λ_1 of Lemma 3.1, used in the definition of C , would have to be replaced by the constants $\lambda_0^\rho, \lambda_1^\rho$ of Lemma 4.8. In the constant $\sin^{-1}(\frac{1}{L})$, the numerator 1 is the thickness of a hyperplane carrier, and would have to be replaced by the minimal thickness of a hyperplane carrier in the metric d_X^ρ .

Now we can prove the main result of this section.

Proposition 7.16 (Combining UBSes for interior points of Tits geodesic) *Let α, β be CAT(0) geodesic rays with $\alpha(0) = \beta(0) \in X^{(0)}$. Suppose that $\mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$ is commensurate with a UBS. Then $a = \alpha(\infty)$ and $b = \beta(\infty)$ are joined by a unique geodesic g in $\partial_T X$. Furthermore, any interior point c of g is represented by a CAT(0) geodesic ray γ such that $\mathcal{W}(\gamma) = \mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$.*

Proof By Lemma 7.11, we have $d_T(a, b) \leq \pi/2$, so the CAT(1) space $\partial_T X$ contains a unique geodesic g from a to b . This proves the first assertion of the lemma.

Let $c \in \partial_T X$ be an interior point of g . Let γ be the unique CAT(0) geodesic ray starting at $\alpha(0)$ and representing c . We need to show that $\mathcal{W}(\gamma) = \mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$.

Working in the convex hull of $\alpha \cup \beta$: Let Y be the cubical convex hull of $\alpha \cup \beta$. We claim that g lies in $\partial_T Y \subset \partial_T X$. Indeed, applying Lemma 7.11 inside of Y shows that a, b can be joined by a unique geodesic g' in $\partial_T Y$. Since the Tits distance in $\partial_T Y$ from a to b depends only on $\{d_Y(\alpha(t), \beta(t)) : t \geq 0\}$ and $d_Y(\alpha(t), \beta(t)) = d_X(\alpha(t), \beta(t))$ by convexity of Y , we have that g' is a geodesic of $\partial_T X$. Since g' has length less than π , it is the unique geodesic in $\partial_T X$ from a to b , i.e. $g = g'$ and g lies in $\partial_T Y$.

For the rest of the proof, we work entirely in Y . For any $y \in Y$, we can choose rays α_y, β_y with $\alpha_y(0) = \beta_y(0) = y$ and $\alpha_y(\infty) = a, \beta_y(\infty) = b$, and note that the cubical convex hull Y_y of $\alpha_y \cup \beta_y$ is contained in Y , and crosses all but finitely many of the hyperplanes crossing Y . By convexity, $\partial_T Y_y = \partial_T Y$. So, when convenient, we can move the basepoint.

What we will actually verify: The plan is as follows. Let \hat{h} be a hyperplane of Y that is dominant in the UBS $\mathcal{U} = \mathcal{W}(Y) = \mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$. We will show that $c \notin \partial_T \hat{h}$. Since the shallow side of any hyperplane in Y is contained in a neighborhood of that hyperplane, it follows that γ must cross \hat{h} . Once we show that γ crosses every dominant hyperplane, Corollary 7.13 will then imply that γ crosses every hyperplane, i.e. $\mathcal{W}(\gamma) = \mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$, as required. So, it remains to argue that $c \notin \partial_T \hat{h}$ when \hat{h} is a dominant hyperplane.

Initial segments of α and β inside a cube: For each $r \geq 0$, let $y_r = \alpha(r)$. Let α_{y_r} and β_{y_r} be geodesics asymptotic to α and β , defined as above. Note that $Y_{y_r} \subset Y$ contains a ray based at y_r and representing each element of $\partial_T Y$. As in the proof of Lemma 7.11, there is a single cube $C_{y_r} \subset Y_{y_r}$, with $y_r \in C_{y_r}$, such that any ray in Y_{y_r} emanating from y_r has a nontrivial initial segment in C_{y_r} .

Let $\alpha'_{y_r}, \beta'_{y_r}, \gamma'_{y_r}$ be the maximal (nontrivial) segments of $\alpha_{y_r}, \beta_{y_r}, \gamma_{y_r}$ that lie in C_{y_r} . Then the Alexandrov angle between any two of the rays $\alpha_{y_r}, \beta_{y_r}, \gamma_{y_r}$ at y_r is just the Euclidean angle in C_{y_r} between the corresponding segments.

By [7, Proposition II.9.8.2], $\angle_{y_r}(\alpha_{y_r}, \beta_{y_r})$ converges to $\angle_T(a, b) = d_T(a, b)$, where the equality follows from Lemma 7.11 because the distance is less than π .

Angle computation: Let \hat{h} be a dominant hyperplane in \mathcal{U} . By Lemma 6.15, we can assume that $\hat{h} \in \mathcal{D}$, where $\mathcal{D} = \mathcal{D}_j$ is a minimal UBS consisting of dominant hyperplanes. Since $\mathcal{D} \subset \mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$ is a minimal UBS, we must have either $\mathcal{D} \leq \mathcal{W}(\alpha)$ or $\mathcal{D} \leq \mathcal{W}(\beta)$. We assume without loss of generality that $\mathcal{D} \leq \mathcal{W}(\alpha)$.

Note that every hyperplane in \mathcal{D} is dominant in $\mathcal{W}(\alpha)$, because any hyperplane crossing all but finitely many hyperplanes in $\mathcal{U} \setminus \mathcal{D}$ crosses all but finitely many hyperplanes in $\mathcal{W}(\alpha) \setminus \mathcal{D}$. We now apply Lemma 7.14 to $\mathcal{W}(\alpha)$ and \mathcal{D} to produce a constant κ and a chain $\{\hat{h}_n\} \subset \mathcal{D}$ whose inseparable closure is \mathcal{D} and whose (necessarily dominant) hyperplanes all cross α at an angle at least κ :

$$\angle_{y_n}(\alpha_{y_n}, \hat{h}_n) \geq \kappa,$$




Fig. 5 The figure shows a portion of the cube C_{y_n} , where the back wall is $M_n = C_{y_n} \cap \hat{h}_n$. The angle between α'_{y_n} and M_n is at least κ , so the angle between η_n and M_n is bounded away from 0 by $\kappa_1 = \kappa(1 - s)$. But the angle between γ'_{y_n} and η_n can be made much smaller than $\kappa(1 - s)$ by taking n large. So γ'_{y_n} makes a positive angle with \hat{h}_n

where $y_n = \hat{h}_n \cap \alpha$ and α_{y_n} is the sub-ray of α emanating from y_n . Since all but finitely many of the \hat{h}_n lie on the deep side of \hat{h} , it suffices to show that γ crosses \hat{h}_n for sufficiently large n . As explained above, it suffices to show that $c \notin \partial_T \hat{h}_n$ for sufficiently large n .

Let β_{y_n} be the ray emanating from y_n and asymptotic to β . Working in the cube C_{y_n} constructed above, we have that the segment α'_{y_n} makes a Euclidean angle at least κ with the midcube $M_n = C_{y_n} \cap \hat{h}_n$. Meanwhile, the segment β'_{y_n} makes some angle $\theta \geq 0$ with M_n . See Fig. 5.

Let $s \in (0, 1)$ be such that $d_T(a, c) = s \cdot d_T(a, b)$ and $d_T(c, b) = (1 - s) \cdot d_T(a, b)$. Such an s exists since c is an interior point of the Tits geodesic g from a to b .

Let $\eta_n \subset C_{y_n}$ be the maximal segment that emanates from y_n and makes an angle $s \cdot \angle_{y_n}(\alpha'_{y_n}, \beta'_{y_n})$ with α'_{y_n} and an angle $(1 - s) \cdot \angle_{y_n}(\alpha'_{y_n}, \beta'_{y_n})$ with β'_{y_n} . Then there exists $\kappa_1 = (1 - s)\kappa > 0$ such that $\angle_{y_n}(\eta_n, M_n) \geq \kappa_1$. Crucially, κ_1 is independent of n . On the other hand, as $n \rightarrow \infty$, we have convergent sequences

$$\begin{aligned} \angle_{y_n}(\alpha'_{y_n}, \beta'_{y_n}) &\rightarrow d_T(a, b), \\ \angle_{y_n}(\alpha'_{y_n}, \gamma'_{y_n}) &\rightarrow d_T(a, c) = s \cdot d_T(a, b), \\ \angle_{y_n}(\beta'_{y_n}, \gamma'_{y_n}) &\rightarrow d_T(b, c) = (1 - s) \cdot d_T(a, b), \end{aligned}$$

Thus, for all sufficiently large n ,

$$\angle_{y_n}(\eta_n, \gamma'_{y_n}) < \frac{\kappa_1}{2},$$

so by the triangle inequality for Alexandrov angles, $\angle_{y_n}(\gamma'_n, M_n) > \kappa_1/2 > 0$. Hence the ray γ_{y_n} is not contained in \hat{h}_n and, since its initial point is in \hat{h}_n , we thus have $c \notin \partial_T \hat{h}_n$. This completes the proof. \square

We observe that the proof of Proposition 7.16 extends to the cuboid setting with minimal effort. The proof combines prior results and the local Euclidean geometry of a cube; all of the local arguments work equally well in a rescaled cuboid.

8 Connections between Tits and Roller boundaries

In this section, we establish some important connections between $\partial_T X$ and $\mathfrak{R}X$. Every Roller class $v \in \mathfrak{R}X$ is assigned a canonical CAT(1)-convex Tits boundary realization $Q(v) \subset \partial_T X$. We define a pair of maps $\psi: \partial_T X \rightarrow \mathfrak{R}X$ and $\varphi: \mathfrak{R}X \rightarrow \partial_T X$ that will play a major role in the proof of our main theorem. See Proposition 8.12 for the properties of $Q(v)$ and Proposition 8.18 for the relationship between φ and ψ .

Toward the end of this section, we focus on ℓ^2 -visible Roller classes, namely all $v \in \mathfrak{R}X$ such that $\psi \circ \varphi(v) = v$. In Sect. 9, we will use the visible Roller classes to construct a closed covering of $\partial_T X$ that is compatible with a covering of a large part of $\mathfrak{R}_\Delta X$, denoted $\mathfrak{R}_\Delta^< X$.

All of the results of this section extend with minimal effort to a cuboid metric d_X^ρ obtained via a G -admissible hyperplane rescaling (recall Definitions 4.3 and 7.4). This extension is described in Sect. 8.4.

Definition 8.1 (Map $\psi: \partial_T X \rightarrow \mathfrak{R}X$) Given $a \in \partial_T X$, define $\psi(a) = Y_a \in \mathfrak{R}X$ to be the Roller class Y_a constructed in Lemma 7.7. That is, $\psi(a)$ is the principal Roller class of the intersection of half-spaces in the deep set D_a .

Our definition of ψ generalizes Guralnik’s boundary decomposition map [37, Definition 4.8], because $\partial_T X$ agrees as a set with the visual boundary $\partial_{\triangleleft} X$.

Lemma 8.2 Let $a = [\alpha] \in \partial_T X$. Let $\mathcal{W}(\alpha)$ be the UBS consisting of hyperplanes crossing the geodesic α . Then $\psi(a) = Y_a = Y_{\mathcal{W}(\alpha)}$.

Proof By Lemma 7.7 we have that $D_a = \mathfrak{H}_{Y_a}^+$ and $Y_a \in \mathfrak{R}X$. Fix $x \in X^{(0)}$. By definition, $\widehat{D_a} \setminus \mathfrak{H}_x^+ = \mathcal{W}(x, Y_a) = \mathcal{W}(x, [y])$ for any $y \in Y_a$. Finally by Lemma 6.20 $\mathcal{W}(x, [y]) \sim \mathcal{W}(x, y)$.

By Lemma 3.27, there is a combinatorial geodesic $\alpha': [0, \infty) \rightarrow X$ with $\alpha'(0) = x$ and $\alpha'(\infty) = y$. By Remark 6.24, we have $\mathcal{W}(\alpha') = \mathcal{W}(x, y)$.

Now, Lemma 4.7 implies that α and α' can be chosen to lie at finite Hausdorff distance. Thus α and α' cross the same hyperplanes, except possibly for finitely many. In symbols, $\mathcal{W}(\alpha') \sim \mathcal{W}(\alpha)$. By Lemma 6.17, we conclude that $Y_a = Y_{\mathcal{W}(\alpha)}$. \square

8.1 Tits boundary realizations

Defining a map $\varphi: \partial_R X \rightarrow \partial_T X$ that serves as a partial inverse to ψ takes considerably more effort. As a first step, we will define a *Tits boundary realization* $Q(v)$ associated to a Roller class v .

A family \mathcal{F} of subsets of X is called *filtering* if for every $E, F \in \mathcal{F}$ there is a $D \in \mathcal{F}$ such that $D \subset E \cap F$. For example, for $y \in \partial_R X$, the family

$$\mathcal{F} := \{\text{CH}(h_1) \cap \dots \cap \text{CH}(h_n) : h_1, \dots, h_n \in \mathfrak{H}_y^+, n \in \mathbb{N}\}$$

is a filtering family of closed convex subspaces.

The following theorem combines results of Caprace–Lytchak [23, Theorem 1.1] and Baleser–Lytchak [12, Proposition 1.4].

Theorem 8.3 *Let $\{X_i\}_{i \in I}$ be a filtering family of closed convex subsets of a finite-dimensional CAT(0) space X . If the intersection $\bigcap_{i \in I} X_i$ is empty, then the intersection $\bigcap_{i \in I} \partial_T X_i$ of their boundaries is nonempty, and furthermore $\bigcap_{i \in I} \partial_T X_i$ has intrinsic radius at most $\pi/2$ and therefore a canonical circumcenter.*

Next, we present two definitions that will turn out to be equivalent (compare Lemma 8.6). The first of these definitions appears in [32, Corollary 6.2].

Definition 8.4 Let $y \in \partial_R X$. Define $Q(y) = \bigcap_{h \in \mathfrak{H}_y^+} \partial_T \text{CH}(h)$. Note that $Q(y) \neq \emptyset$ by Theorem 8.3. We call $Q(y)$ the *Tits boundary realization* of y .

Definition 8.5 Let $x \in X^{(0)}$ and let $y \in \partial_R X$. Recall that $\mathcal{I}(x, y) \subset \bar{X}$ is the vertex interval between x and y , so that $\mathcal{I}(x, y) \cap X$ is a vertex–convex subset of $X^{(0)}$. Generalizing Definition 3.8, let $\mathcal{J}(x, y)$ be the union of all of the cubes in X whose 0-skeleta lie in $\mathcal{I}(x, y) \cap X$. Then $\mathcal{J}(x, y)$ is a convex subcomplex of X , and hence has a well-defined Tits boundary naturally embedded in $\partial_T X$. We define $Q'(y) = \partial_T \mathcal{J}(x, y)$.

In fact, the two definitions are equivalent:

Lemma 8.6 *Let $y \in \partial_R X$ and let $x, x' \in X^{(0)}$. Then:*

- (1) $\partial_T \mathcal{J}(x, y) = \partial_T \mathcal{J}(x', y)$, hence $Q'(y)$ is well-defined.
- (2) $Q(y) = Q'(y)$.
- (3) If $y \sim y'$, then $Q'(y) = Q'(y')$ and $Q(y) = Q(y')$.

Proof. To prove conclusion (1), observe that $\mathcal{J}(x, y)$ and $\mathcal{J}(x', y)$ lie at Hausdorff distance bounded by $d_1(x, x')$. Hence $\mathcal{J}(x, y)$ and $\mathcal{J}(x', y)$ have the same Tits boundary.

Next, we consider conclusion (2). For any given vertex half-space $h \in \mathfrak{H}_y^+$, let $\text{CH}(h)$ be the associated CAT(0) half-space containing y . By (1), we can assume that $x \in \text{CH}(h)$. Hence $\mathcal{J}(x, y) \subset \text{CH}(h)$, so $Q'(y) = \partial_T \mathcal{J}(x, y) \subset \partial_T \text{CH}(h)$. Therefore $Q'(y) \subset Q(y)$.

For the reverse inclusion, suppose $q \notin Q'(y)$. Then there is a CAT(0) half-space $\text{CH}(h)$ such that $\mathcal{J}(x, y) \subset \text{CH}(h)$ and $q \notin \partial_T \text{CH}(h)$. (To find such a half-space, let γ be any CAT(0) geodesic ray representing q . Since γ leaves every finite neighborhood of $\mathcal{J}(x, y)$, it must cross a hyperplane \hat{h} disjoint from $\mathcal{J}(x, y)$.) The associated vertex half-space h satisfies $h \in \mathfrak{H}_y^+$ but $q \notin \partial_T \text{CH}(h)$. By Definition 8.4, it follows that $q \notin Q(y)$. Hence $Q(y) \subset Q'(y)$.

Finally, consider conclusion (3). If $y \sim y'$, the Hausdorff distance from $\mathcal{J}(x, y)$ to $\mathcal{J}(x, y')$ is bounded by $d_1(y, y') < \infty$. Using (2), we obtain

$$Q(y) = Q'(y) = \partial_T \mathcal{J}(x, y) = \partial_T \mathcal{J}(x, y') = Q'(y') = Q(y'). \quad \square$$

Following Theorem 8.3 and Lemma 8.6, we can make the following definition.

Definition 8.7 (*Tits boundary realization*) Let $v \in \mathfrak{R}X$ be a Roller class. Define $Q(v) = Q(y)$ for any representative element $y \in v$. The reader can think of $Q(y)$ according to either Definition 8.4 or Definition 8.5. We call $Q(v)$ the *Tits boundary realization* of the Roller class v .

Observe that the collection of half-spaces \mathfrak{H}_y^+ containing y has empty intersection in X . Thus, by Theorem 8.3, $Q(v)$ has a canonical circumcenter. We define $\chi(v) \in \partial_T X$ to be the circumcenter of $Q(v)$.

Corollary 8.8 *For every $v \in \mathfrak{R}X$ and every $g \in \text{Aut}(X)$, we have $Q(gv) = gQ(v)$ and $\chi(gv) = g\chi(v)$. Furthermore, $Q(v)$ and $Q(gv)$ have the same intrinsic radius.*

Proof Let $y \in \partial_R X$, and let $v = [y]$. Then, for every $h \in \mathfrak{H}_y^+$, the map g gives an isometry from $\partial_T \text{CH}(h)$ to $\partial_T \text{CH}(gh)$. Thus $g: Q(v) \rightarrow Q(gv)$ is an isometry, hence $Q(v)$ and $Q(gv)$ have the same intrinsic radius. By Theorem 8.3, $\chi(v)$ and $\chi(gv)$ are uniquely determined by the geometry of $Q(v)$ and $Q(gv)$, respectively, hence $\chi(gv) = g\chi(v)$. \square

Consider a Roller class $v \in \mathfrak{R}X$. We say that a combinatorial geodesic ray γ in X represents v if $\gamma(\infty) = y \in v$, or equivalently if the UBS $\mathcal{W}(\gamma)$ represents the class v . Setting $x = \gamma(0)$, observe that $\mathcal{J}(x, y)$ is the cubical convex hull of γ , and $\mathcal{W}(\gamma) = \mathcal{W}(x, y)$ is exactly the collection of hyperplanes crossing $\mathcal{J}(x, y)$. This leads to

Lemma 8.9 *Let $v \in \mathfrak{R}X$ be a Roller class and let $a \in Q(v)$. Then $\psi(a) \leq v$.*

Proof Let γ be a combinatorial geodesic ray representing v , with $\gamma(0) = x$ and $\gamma(\infty) = y \in v$. Then the cubical convex hull of γ , namely $\mathcal{J}(x, y)$, is also CAT(0) geodesically convex. Since $a \in Q(v) = \partial_T \mathcal{J}(x, y)$ by Lemma 8.6, we may choose a CAT(0) geodesic α representing a such that $\alpha \subset \mathcal{J}(x, y)$. Then $\mathcal{W}(\alpha) \subset \mathcal{W}(\gamma) = \mathcal{W}(x, y)$. Thus every hyperplane crossing α crosses $\mathcal{J}(x, y)$ and hence γ . Hence $\psi(a) \leq v$. \square

Remark 8.10 In Definition 8.17, we will define a map $\varphi: \mathfrak{R}X \rightarrow \partial_T X$ with the property that $\varphi(v) \in Q(v)$ for each Roller class v . This map will be defined by slightly perturbing the circumcenter $\chi(v)$. In view of Lemma 8.9, we will have $\psi(\varphi(v)) \leq v$.

Lemma 8.11 *Let $v, w \in \mathfrak{R}X$ satisfy $w \leq v$. Then $Q(w) \subset Q(v)$.*

Proof Fix a basepoint $x_0 \in X^{(0)}$ and let γ_v, γ_w be combinatorial geodesic rays emanating from x_0 and representing v, w respectively. Since $w \leq v$, Theorem 6.27.(2) says that $\mathcal{W}(\gamma_w) \preceq \mathcal{W}(\gamma_v)$. Without moving x_0 , we can replace γ_w by its image under the gate map to the cubical convex hull of γ_v , ensuring that $\mathcal{W}(\gamma_w) \subset \mathcal{W}(\gamma_v)$. Hence, setting $y_w = \gamma_w(\infty) \in w$ and $y_v = \gamma_v(\infty) \in v$, the cubical convex hulls of these geodesics satisfy $\mathcal{J}(x_0, y_w) \subset \mathcal{J}(x_0, y_v)$. Taking boundaries and applying Lemma 8.6 gives $Q(w) \subset Q(v)$. \square

8.2 Diameter, intrinsic radius, and Tits-convexity of $Q(v)$

Let $v \in \mathfrak{R}X$ be a Roller class. Recall that $Q(v)$ is the intersection of the Tits boundaries of the half-spaces corresponding to a representative element of v . We will need the following properties of $Q(v)$:

Proposition 8.12 (Features of $Q(v)$) *For any Roller class v , the Tits boundary realization $Q(v)$ has the following properties:*

- (1) $Q(v)$ has diameter at most $\pi/2$.
- (2) $Q(v)$ is Tits-convex.
- (3) $Q(v)$ is contractible in the Tits metric topology and compact in the visual topology.
- (4) $Q(v)$ has intrinsic radius $r_v < \pi/2$.
- (5) $Q(v) \subset \overline{B}_T(\chi(v), r_v)$.

Proof Let v be an arbitrary Roller class and $x_0 \in X^{(0)}$ an arbitrary basepoint. Let γ be a combinatorial geodesic ray based at x_0 and representing v . Let $H = H(x_0, v, \gamma)$ denote the cubical convex hull of γ . Then, for an arbitrary point $y \in [y] = v$, the cubical hull H and the interval $\mathcal{J}(x_0, y)$ lie within bounded Hausdorff distance of each other. Thus, by Lemma 8.6, $Q(v)$ is equal to the closed subset $\partial_T H = \partial_T \mathcal{J}(x_0, y) \subset \partial_T X$.

Diameter: To bound $\text{diam}_T(Q(v))$, let $a, b \in Q(v)$. Represent a, b by CAT(0) rays α, β with initial point x_0 . Note that $\alpha, \beta \subset H$ because H is a convex subcomplex. Then $\mathcal{W}(\alpha), \mathcal{W}(\beta)$ are both subsets of $\mathcal{W}(\gamma)$, which is a UBS representing v . So, $\mathcal{W}(\alpha) \cup \mathcal{W}(\beta) \subset \mathcal{W}(\gamma)$. Now, $\mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$ is infinite. Moreover, $\mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$ is unidirectional and contains no facing triple, since $\mathcal{W}(\gamma)$ has those properties. Thus $\mathcal{W}(\alpha) \cup \mathcal{W}(\beta)$ is a UBS, hence Lemma 7.11 implies $d_T(a, b) \leq \pi/2$.

Convexity: Since $H = H(x_0, v, \gamma) \hookrightarrow X$ is an isometric embedding in the CAT(0) metric, the inclusion $\partial_T H = Q(v) \hookrightarrow \partial_T X$ is a Tits-metric isometric embedding. Indeed, let $a, b \in Q(v)$ be represented by CAT(0) rays $\alpha, \beta \subset H$ with initial point x_0 . Since $Q(v)$ has diameter less than π , we have $d_T(a, b) = \angle(a, b)$ by [7, Remark II.9.19.(2)]. Since $\alpha, \beta \subset H$, the angle $\angle_T(a, b)$ is determined entirely by the set $\{d_X(\alpha(t), \beta(t))\}_{t \geq 0}$, by [7, Proposition II.9.8.(4)]. By convexity of H , this set coincides with $\{d_H(\alpha(t), \beta(t))\}_{t \geq 0}$. Hence the Tits distance from a to b measured in $\partial_T X$ is the same as the Tits distance measured in $\partial_T H = Q(v)$.

Now, $Q(v)$ contains a unique geodesic joining a, b . Since this geodesic realizes the distance from a to b in $\partial_T X$, this shows that $Q(v)$ is Tits-convex.

Intrinsic radius: Since $Q(v)$ is a CAT(1) space of diameter less than π , the intrinsic radius r_v satisfies $r_v < \text{diam } Q(v)$, by [12, Proposition 1.2]. Since $\text{diam } Q(v) \leq \pi/2$, we have $r_v < \pi/2$, proving assertion (4). Now, observe that $Q(v) \subset \overline{B}_T(\chi(v), r_v)$ by Definition 8.7 and Theorem 8.3, proving (5).

Topological properties: Since $Q(v)$ is a uniquely geodesic CAT(1) space, a standard straight-line homotopy allows one to deformation retract $Q(v)$ to a single point. Thus $Q(v)$ is contractible.

Finally, recall from Lemma 8.6 (and from earlier in this proof) that $Q(v) = \partial_T \mathcal{J}(x_0, y)$ for an arbitrary point $y \in [y] = v$. By Lemma 3.24, the subcomplex $\mathcal{J}(x_0, y)$ is proper, hence $Q(v)$ is compact in the visual topology. (The visual topology is defined in [7, Section II.8], where it is called the *cone topology*.) \square

8.3 Visibility and perturbing the circumcenter

Recall the $\text{Aut}(X)$ -equivariant circumcenter map $\chi : \mathfrak{R}X \rightarrow \partial_T X$, mentioned in Corollary 8.8. In this subsection, we perturb χ to obtain a map $\varphi : \mathfrak{R}X \rightarrow \partial_T X$ with slightly nicer properties. First, like χ , the map φ will have the property that $\varphi(v) \in Q(v)$ for each Roller class v , and hence $\psi(\varphi(v)) \leq v$. Compare Remark 8.10.

We will define φ so that Roller classes v for which $\psi(\varphi(v)) = v$ are exactly those for which there is some CAT(0) geodesic ray representing v . (Note that a combinatorial geodesic




Fig. 6 For the cube complex X in Example 8.13, the Tits boundary $\partial_T X$ is the spherical triangle shown here. We have $\partial_T X = Q(v)$ for a single Roller class v , and $\chi(v) = \chi(w)$ for a Roller class $w < v$

exists for every v , but a CAT(0) geodesic is not guaranteed.) Later, we will work with only such Roller classes, which we term ℓ^2 -visible; see Definition 8.19. See also Lemma 8.20 for several equivalent characterizations of visibility. While the map φ is not guaranteed to be equivariant, the set of ℓ^2 -visible Roller classes will still be invariant.

The following example illustrates the point that $\chi(v)$ can lie in $Q(w) \subsetneq Q(v)$ for some $w < v$.

Example 8.13 Consider the cone of \mathbb{R}^3 cut out by the planes $z = 0$, $z = y$, and $z = x$, and let X be the union of all cubes (in the standard cubical tiling of \mathbb{R}^3) contained in this cone. Then $\mathfrak{R}X$ has four classes, described in coordinates as $u = [(\infty, 0, 0)]$, $u' = [(0, \infty, 0)]$, $w = [(\infty, \infty, 0)]$, and $v = [(\infty, \infty, \infty)]$, all of which are represented by ℓ^2 geodesic rays. Furthermore, $Q(v) = Q(\infty, \infty, \infty) = \partial_T X$ is isometric to an isosceles triangle of S^2 whose base has length $\pi/2$ and whose height is $\arctan(1/\sqrt{2})$. The base of the triangle is $Q(w) = Q(\infty, \infty, 0)$. See Fig. 6.

First, observe that $\chi(w)$ is the midpoint of the geodesic segment $Q(w)$. Next, we claim that $\chi(v) = \chi(w)$. To see this, let β be the altitude from $\chi(w)$ to the apex of $Q(v)$. Because $Q(v)$ has a reflective symmetry in β , its circumcenter must be contained in β . In addition, since $\text{len}(\beta) = \arctan(1/\sqrt{2}) < \pi/4$, any point $b \in \beta$ must have distance less than $\pi/4$ to the apex but distance at least $\pi/4$ to the two endpoints of $Q(w)$. Thus $\chi(w) = \chi(v)$ even though $w < v$ and $Q(w) \subsetneq Q(v)$.

The phenomenon of Example 8.13 will be inconvenient later. To remedy this problem, we will define $\varphi(v)$ by perturbing $\chi(v)$ slightly, to get a point in the interior of $Q(v)$ that retains the property that all of $Q(v)$ is contained in the r_v -neighborhood of $\varphi(v)$ for some $r_v < \pi/2$. (This latter property will also be necessary later.)

To achieve this, we need some preliminary discussion and lemmas.

Lemma 8.14 For every Roller class v , there is a unique Roller class $M_v \in \psi(Q(v))$ that is maximal among all Roller classes in $\psi(Q(v))$. Furthermore $Q(M_v) = Q(v)$.

Proof If $w \in \psi(Q(v))$, then $w = \psi(a)$ for some $a \in Q(v)$, so $w \leq v$ by Lemma 8.9. Since there is a bound on the length of \leq -chains, it follows that $\psi(Q(v))$ contains \leq -maximal elements.

Suppose that $m, m' \in \psi(Q(v))$ are \leq -maximal. Let $a, a' \in Q(v)$ be such that $\psi(a) = m$ and $\psi(a') = m'$. Let α, α' be CAT(0) geodesic rays in X representing a, a' respectively,

chosen so that $\alpha(0) = \alpha'(0)$. Then $\mathcal{U}_m \sim \mathcal{W}(\alpha)$ and $\mathcal{U}_{m'} \sim \mathcal{W}(\alpha')$. Since $m, m' \leq v$, Theorem 6.27 gives $\mathcal{W}(\alpha) \preceq \mathcal{U}_v$ and $\mathcal{W}(\alpha') \preceq \mathcal{U}_v$, hence $\mathcal{W}(\alpha) \cup \mathcal{W}(\alpha')$ is commensurate with a UBS by Lemma 6.8. Thus, by Proposition 7.16, there exists a CAT(0) geodesic ray β with $\mathcal{W}(\beta) = \mathcal{W}(\alpha) \cup \mathcal{W}(\alpha')$. By construction,

$$m \leq \psi(\beta(\infty)) \leq v, \quad m' \leq \psi(\beta(\infty)) \leq v$$

where the first inequality for m or m' is strict if $m \neq m'$. So, $\beta(\infty) \in Q(v)$ by Lemma 8.11 and the second inequality. Hence the first inequality contradicts the maximality of m, m' in $\psi(Q(v))$, unless $m = m'$. This proves the uniqueness of a maximal element M_v .

Next, $M_v \leq v$, so $Q(M_v) \subset Q(v)$ by Lemma 8.11. On the other hand, if $a \in Q(v)$, then $\psi(a) \in \psi(Q(v))$, so $\psi(a) \leq M_v$. Hence $a \in Q(\psi(a)) \subset Q(M_v)$ by Lemma 8.11. Thus $Q(v) \subset Q(M_v)$, and we conclude that $Q(v) = Q(M_v)$. \square

For a Roller class v , define $Q_0(v) = Q(v) \setminus \bigcup_{u < M_v} Q(u)$.

Lemma 8.15 *Let v, w be Roller classes, and let $a \in Q(v)$. Then $a \in Q_0(v)$ if and only if $\psi(a) = M_v$. Furthermore, if $Q(v) = Q(w)$, then $M_v = M_w$ and $Q_0(v) = Q_0(w)$.*

Proof First, suppose that $a \in Q_0(v)$. Then $\psi(a) \leq M_v$ by the definition of M_v . If $\psi(a) = u < M_v$, then $a \in Q(u)$, contradicting the definition of $Q_0(v)$. Thus $\psi(a) = M_v$, as desired.

Conversely, suppose that $\psi(a) = M_v$. If $a \in Q(u)$ for some $u < M_v$, then $\psi(a) \leq u$ by Lemma 8.9, so $\psi(a) < M_v$, a contradiction. Hence $a \in Q_0(v)$.

Finally, suppose $Q(v) = Q(w)$. Then Lemma 8.14 says that $M_v = M_w$ is the unique maximal element of $\psi(Q(v)) = \psi(Q(w))$. Thus, by the above equivalence, $Q_0(v) = Q_0(w)$. \square

Lemma 8.16 *Let v be a Roller class. Then $Q_0(v)$ contains points arbitrarily close to $\chi(v)$ in the Tits metric.*

Proof By Lemma 8.14, there is a point $a \in Q(v)$ such that $\psi(a) = M_v$. By Lemma 8.15, we have $a \in Q_0(v)$. Let g be the geodesic in $\partial_T X$ from a to $b = \chi(v)$. By Proposition 8.12.(2), g lies in $Q(v)$. For any $\epsilon > 0$, choose $c \in g$ so that $0 < d_T(c, \chi(v)) < \epsilon$. We will show $c \in Q_0(v)$.

By Proposition 7.16, the Tits points a, b, c are represented by CAT(0) geodesic rays α, β, γ such that $\mathcal{W}(\gamma) = \mathcal{W}(\alpha) \cup \mathcal{W}(\beta) \supset \mathcal{W}(\alpha)$. Thus, by Lemma 8.2, we have $\psi(c) = Y_{\mathcal{W}(\gamma)} \subset Y_{\mathcal{W}(\alpha)} = \psi(a)$. By Lemma 5.6, we have $\psi(c) \geq \psi(a) = M_v$, hence $\psi(c) = M_v$ by the maximality of M_v . Therefore, $c \in Q_0(v)$ by Lemma 8.15. \square

Definition 8.17 *(The pseudocenter φ)* Let v be a Roller class, and fix $Q = Q(v)$. Let $r < \pi/2$ be the intrinsic radius of Q . Using Lemma 8.16, choose $p_Q \in Q_0(v)$ such that $d_T(p_Q, \chi(v)) \leq \pi/4 - r/2$. We call p_Q the *pseudocenter* of Q , and denote it $\varphi(v)$.

Using Lemma 8.15, we ensure that $\varphi(v) = \varphi(w) = p_Q$ for every Roller class w such that $Q(w) = Q(v) = Q$.

In the following proposition we collect the facts about the pseudocenter map φ that will be needed later, in conjunction with Lemma 8.20 and the definition of φ .

Proposition 8.18 *(φ and ψ facts) The pseudocenter map φ has the following properties. Let $v \in \mathfrak{R}X$. Then:*

- (1) $\varphi(v) \in Q(v)$.

- (2) $\psi(\varphi(v)) \leq v$.
- (3) $Q(v) = Q(\psi(\varphi(v)))$.
- (4) If $Q(w) = Q(v)$, then $\varphi(w) = \varphi(v)$.

Proof Assertion (1) holds by construction, and, together with Lemma 8.9, implies assertion (2). Since $\varphi(v) \in Q_0(v)$, we have $\psi(\varphi(v)) = M_v$ by Lemma 8.15, and $Q(v) = Q(M_v)$ by Lemma 8.14, yielding assertion (3). Assertion (4) follows from the definition of φ and Lemma 8.15. □

We can now define (ℓ^2)-visible Roller classes:

Definition 8.19 (ℓ_2 -visible) Let $v \in \mathfrak{R}X$. Then v is ℓ^2 -visible if $v \in \psi(\partial_T X)$. Let $\text{Vis}(X) = \psi(\partial_T X)$ denote the set of visible Roller classes. Observe that $\text{Aut}(X)$ stabilizes $\text{Vis}(X)$, because the map ψ is $\text{Aut}(X)$ -equivariant.

Lemma 8.20 (Characterizing visibility) Let $v \in \mathfrak{R}X$. Then the following are equivalent:

- (1) v is ℓ^2 -visible.
- (2) There exists a CAT(0) geodesic ray α such that $\mathcal{W}(\alpha)$ represents v .
- (3) $v = M_v$.
- (4) $v = \psi(\varphi(v))$.

Proof To begin with, we have (1) \Leftrightarrow (2) by Lemma 8.2.

For (2) \Rightarrow (3), let α be a CAT(0) geodesic ray such that $\mathcal{W}(\alpha)$ represents v . Let $a = \alpha(\infty) \in \partial_T X$.

Then, by Lemma 8.2, we have $\psi(a) = v$. Now, recall from Lemma 8.14 that M_v is maximal among all Roller classes in $\psi(Q(v))$. By Lemma 8.9, every class $w \in \psi(Q(v))$ satisfies $w \leq v$, hence $M_w \leq v$. On the other hand, $v = \psi(a) \in \psi(Q(v))$, hence $v \leq M_v$ by the maximality of M_v . Thus $v = M_v$.

For (3) \Rightarrow (4), suppose $v = M_v$. By Definition 8.17, we have $\varphi(v) \in Q_0(v)$. Thus, by Lemma 8.15, we have $\psi(\varphi(v)) = M_v = v$.

For (4) \Rightarrow (2), suppose $v = \psi(\varphi(v))$. Then $\varphi(v) \in Q_0(v) \subset Q(v)$. Choose a CAT(0) geodesic ray α representing $\varphi(v)$. The associated UBS $\mathcal{W}(\alpha)$ represents $\psi(\varphi(v)) = v$. □

Lemma 8.21 Let $v \in \mathfrak{R}X$ be a \leq -minimal Roller class. Then v is ℓ^2 -visible.

Proof Lemma 8.9 gives $v \geq \psi(\varphi(v))$. By minimality of v , we thus have $v = \psi(\varphi(v))$. □

8.4 Cuboid generalization

Here, we explain how to generalize the results of this section to a cuboid metric d_X^ρ . A reader who is only interested in cube complexes with the standard ℓ^1 and ℓ^2 metrics is invited to skip ahead to Sect. 9.

For the duration of this subsection, fix a G -admissible hyperplane rescaling ρ (Definition 4.3) and the resulting cuboid metric d_X^ρ . Recall that the cuboid Tits boundary $\partial_T^\rho X = \partial_T(X, d_X^\rho)$ was defined in Definition 7.4. All of the constructions and results involving $\partial_T^\rho X$ will be invariant under the restricted automorphism group $\text{Aut}(X^\rho)$ of Definition 4.4.

To start, we can define a map $\psi^\rho : \partial_T^\rho X \rightarrow \mathfrak{R}X$ exactly as in Definition 8.1. Generalizing Lemma 8.2, we can characterize $\psi^\rho(a) = \psi^\rho([\alpha])$ as the principal class of the umbra of a UBS representing a ray α :

$$\psi^\rho(a) = \psi^\rho([\alpha]) = Y_{\mathcal{W}(\alpha)}.$$

This characterization holds because the proof of Lemma 8.2 works for ψ^ρ : its proof combines previous lemmas and combinatorial properties of hyperplane sets.

Now, consider a point $y \in \partial_R X$. Since (X, d_X^ρ) is a CAT(0) space and the collection of half-spaces \mathfrak{H}_y^+ is a filtering family, we can define the Tits boundary realization $Q^\rho(y)$ exactly as in Definition 8.4. Following Lemma 8.6, whose proof extends to cuboids because cubical half-spaces are convex in d_X^ρ , we learn that the boundary realization $Q^\rho(y)$ is also the boundary of a cubical interval, and depends only on the Roller class $v = [y]$. Thus we may define a Tits-convex set $Q^\rho(v)$ and its circumcenter $\chi^\rho(v)$, as in Definition 8.7. Following Corollary 8.8, both $Q^\rho(v)$ and $\chi^\rho(v)$ are invariant under the restricted automorphism group $\text{Aut}(X^\rho)$.

The partial order properties of the Tits boundary realization $Q(v)$ that are proved in Lemmas 8.9 and 8.11 still hold for $Q^\rho(v)$, because the proofs of those results are essentially combinatorial. The features of $Q(v)$ described in Proposition 8.12 also hold for $Q^\rho(v)$, because the proof of the Proposition uses prior results and the convexity of the cubical convex hull of a geodesic ray.

Next, Sect. 8.3 contains several results about perturbing the circumcenter $\chi(v) \in Q(v)$ to a pseudocenter $\varphi(v)$. The definition of the pseudocenter $\varphi(v)$ is enabled by Lemmas 8.14, 8.15, and 8.16. All of these lemmas generalize immediately to the cuboid setting, because their proofs are essentially a top-level assembly of prior results. Thus we may generalize Definition 8.17 to define a pseudocenter $\varphi^\rho(v) \in Q^\rho(v)$. This pseudocenter $\varphi^\rho(v)$ has all of the properties described in Proposition 8.18, because the proof of that proposition merely assembles previously established results.

Finally, Definition 8.19 generalizes immediately to define $\text{Vis}^\rho(X) = \psi^\rho(\partial_T^\rho X)$, the set of Roller classes that are visible after the rescaling ρ . Since ψ^ρ is $\text{Aut}^\rho(X)$ -equivariant, the set $\text{Vis}^\rho(X)$ is $\text{Aut}^\rho(X)$ -equivariant as well. Lemmas 8.20 and 8.21 generalize immediately to the cuboid setting, because their proofs are top-level assemblies of previous lemmas.

Remark 8.22 [*The set $\text{Vis}(X)$ and cuboids*] In Definition 8.19, the set $\text{Vis}(X) = \psi(\partial_T X)$ is defined in terms of the CAT(0) metric on X . When we change the metric on X using a G -admissible rescaling ρ , there is no *a priori* reason to expect $\text{Vis}(X) = \psi(\partial_T X)$ to coincide with $\text{Vis}^\rho(X) = \psi^\rho(\partial_T^\rho X)$. Instead, our whole argument simply goes through for whichever of $\text{Vis}(X)$ or $\text{Vis}^\rho(X)$ we are considering.

In fact, it turns out that $\text{Vis}(X) = \text{Vis}^\rho(X)$. Since we do not use this fact in any proofs, we only sketch a proof. Let $v \in \text{Vis}(X)$ be a visible Roller class, and let $\alpha : [0, \infty) \rightarrow X$ be a CAT(0) geodesic ray for the metric d_X with $\mathcal{W}(\alpha)$ representing v . We will produce a geodesic ray β for the CAT(0) metric d_X^ρ , satisfying $\mathcal{W}(\alpha) = \mathcal{W}(\beta)$, from which it follows that $v \in \text{Vis}^\rho(X)$. This shows $\text{Vis}(X) \subseteq \text{Vis}^\rho(X)$, and a symmetric argument gives the reverse inclusion.

Since cubical convexity of subcomplexes is independent of which CAT(0) metric we consider, we assume for convenience that X is the cubical convex hull of α , i.e. $\mathcal{W}(\alpha) = \mathcal{W}(X)$.

First, use Lemma 7.14 to produce chains $\{\hat{h}_n^1\}_n, \dots, \{\hat{h}_n^k\}_n$ of dominant hyperplanes such that there exists a constant $C \geq 0$ with the property that, for all $j \leq k$, any subpath of α of length at least C crosses \hat{h}_n^j for some n . The lemma allows us to choose these to be dominant hyperplanes in the UBS $\mathcal{W}(\alpha)$; therefore, if we construct a d_X^ρ -geodesic β that crosses each h_n^j , then by Lemma 6.15, we get $\mathcal{W}(\alpha) = \mathcal{W}(\beta)$.

By admissibility of the rescaling, there exists $m \geq 1$ such that the identity map $(X, d_X) \rightarrow (X, d_X^\rho)$ is m -bilipschitz. Since X is the convex hull of α and contains no facing triples, Lemma 3.24 guarantees properness of (X, d_X) and implies properness of (X, d_X^ρ) .

Thus, letting β_t be the d_X^ρ -geodesic from $\alpha(0)$ to $\alpha(t)$, and letting t tend to infinity, the β_t subconverge uniformly on compact sets to a d_X^ρ -geodesic β with $\beta(0) = \alpha(0)$.

Fix $j \leq k$. Now, for any t , the hyperplanes crossing β_t are exactly those crossing $\alpha([0, t])$. Indeed, any hyperplane separating $\alpha(0)$, $\alpha(t)$ crosses β_t , and any hyperplane crossing β_t does so in a single point (and hence separates $\alpha(0)$, $\alpha(t)$), since β_t is a d_X^ρ -geodesic. In particular, $\alpha([0, t])$ crosses $h_1^j, \dots, h_{N_t}^j$, where $N_t = \lfloor t/C \rfloor$. Hence β_t crosses the same hyperplanes while having d_X^ρ -length at most mt . From this one deduces the following: there exists L such that for all n , and all sufficiently large t , the path β_t crosses h_1^j, \dots, h_n^j and intersects each of those carriers in a subpath of length at most L . In other words, there is a uniform bound on how long each β_t can follow-travel any of the \hat{h}_n^j , whence β cannot be parallel to a ray in any \hat{h}_n^j . This implies that β must cross \hat{h}_n^j for all j, n , and hence, as explained above, $\mathcal{W}(\beta) = \mathcal{W}(\alpha)$ and $v \in \text{Vis}^\rho(X)$.

Looking ahead to Definition 9.1, the above argument also shows that $\text{MaxVis}(X) = \text{MaxVis}^\rho(X)$ for any admissible rescaling ρ . In the next section, we will heavily use $\text{MaxVis}(X)$ to construct open and closed coverings of $\partial_T X$, so that the nerves of those coverings can be used to prove homotopy equivalence. In an analogous fashion, we will use $\text{MaxVis}^\rho(X)$ to construct open and closed coverings of $\partial_T^\rho X$. One can then apply Lemma 8.9 and Lemma 8.11 to show that the nerves of the coverings of $\partial_T X$ coincide with the nerves of the corresponding coverings of $\partial_T^\rho X$, and one can use Lemma 10.4 to see that the covering of $\mathfrak{R}_\Delta^< X$ coming from visible Roller classes of X coincides with the corresponding covering constructed using visible classes of X^ρ . However, as we mentioned, the cuboid version of our argument does not rely on this; one instead just substitutes $\text{MaxVis}^\rho(X)$ for $\text{MaxVis}(X)$ everywhere, and runs all the arguments.

9 Open and closed coverings of $\partial_T X$

In this section, we study the closed covering of $\partial_T X$ by the sets $Q(v)$ corresponding to maximal ℓ^2 -visible Roller classes (see Definition 9.1). Our goal, achieved in Theorem 9.17, is to show that $\partial_T X$ is homotopy equivalent to the nerve \mathcal{N}_T of this cover.

Definition 9.1 (*Maximal visible Roller classes*) Let $\text{MaxVis}(X)$ denote the set of Roller classes v such that:

- v is ℓ^2 -visible;
- if w is visible and $v \leq w$, then $w = v$.

The plan for this section is as follows. We will first show that $\{Q(v) : v \in \text{MaxVis}(X)\}$ is a closed covering of $\partial_T X$. Then, we will thicken each closed set $Q(v)$ to an open set $U(v)$, in such a way that the intersection pattern of the open cover $\{U(v) : v \in \text{MaxVis}(X)\}$ is the same as that of the closed cover $\{Q(v) : v \in \text{MaxVis}(X)\}$. Then we will apply the Nerve Theorem to conclude that $\partial_T X$ is homotopy equivalent to the nerves of these covers.

The main result of this section is Theorem 9.17. Section 9.1 is about the initial closed covering. Section 9.2 describes the thickening procedure. Section 9.4 describes how to generalize these results to a cuboid metric d_X^ρ .

9.1 The closed covering

The maximal visible Tits boundary realizations $Q(v)$ provide a closed covering of the Tits boundary. We use this covering to define a nerve, as follows.

Definition 9.2 (*Simplicial complex \mathcal{N}_T*) Let \mathcal{N}_T be the simplicial complex with vertex set $\text{MaxVis}(X)$, where vertices $v_0, \dots, v_n \in \text{MaxVis}(X)$ span an n -simplex if and only if $\bigcap_{i=0}^n Q(v_i) \neq \emptyset$.

Lemma 9.3 (*Covering $\partial_T X$*) Let $v \in \mathfrak{R}X$. Then there exists $w \in \text{MaxVis}(X)$ with $\varphi(v) \in Q(w)$. Hence $\{Q(v) : v \in \text{MaxVis}(X)\}$ covers $\partial_T X$.

Proof Note that $\psi(\varphi(v)) \in \text{Vis}(X)$, by Definition 8.19. Hence there exists $w \in \text{MaxVis}(X)$ such that $\psi\varphi(v) \leq w$. By Lemma 8.11, $\varphi(v) \in Q(\psi\varphi(v)) \subset Q(w)$.

Now, given a point $a \in \partial_T X$, we can apply the above argument to $v = \psi(a)$. Then, for some $w \in \text{MaxVis}(X)$, Lemma 8.20 implies

$$a \in Q(\psi(a)) = Q(\psi\varphi\psi(a)) \subset Q(w).$$

Hence $\{Q(v) : v \in \text{MaxVis}(X)\}$ covers $\partial_T X$. □

Lemma 9.4 Let $v, w \in \text{Vis}(X)$. Then $Q(v) = Q(w)$ if and only if $v = w$. In particular, the assignment $v \mapsto Q(v)$ gives a bijection from $\text{MaxVis}(X)$ to its image.

Proof Suppose $Q(v) = Q(w)$. Then Proposition 8.18.(4) implies $\varphi(v) = \varphi(w)$. By visibility and Lemma 8.20, we have $\psi\varphi(v) = v$ and $\psi\varphi(w) = w$, hence $v = w$. □

Now, Lemmas 9.3 and 9.4 combine to yield:

Corollary 9.5 \mathcal{N}_T is the nerve of the covering of $\partial_T X$ by the collection of closed sets $\{Q(v) : v \in \text{MaxVis}(X)\}$.

9.2 The open covering

The goal of this subsection is to thicken up the closed sets $Q(v)$ for $v \in \text{MaxVis}(X)$ to be open sets $U(v)$, such that the intersection pattern of the $U(v)$ is the same as that of the $Q(v)$. This is needed since the Nerve Theorem 2.7 works for open covers only.

Lemma 9.6 Let $v, w \in \text{Vis}(X)$ be points such that $Q(v) \cap Q(w) = \emptyset$. Then $d_T(Q(v), Q(w)) > \epsilon_0$, where $\epsilon_0 > 0$ is a constant depending only on $\dim X$.

Proof Let $a \in Q(v)$ and $b \in Q(w)$. By definition, $a \in Q(\psi(a))$. Since $v \geq \psi(a)$, we have $Q(\psi(a)) \subset Q(v)$ and similarly $Q(\psi(b)) \subset Q(w)$, by Lemma 8.11. Since we have assumed $Q(v) \cap Q(w) = \emptyset$, it follows that $Q(\psi(a)) \cap Q(\psi(b)) = \emptyset$.

Choose CAT(0) geodesic rays α, β such that $\alpha(0) = \beta(0) \in X^{(0)}$ and $\alpha(\infty) = a$, $\beta(\infty) = b$. Since X is finite-dimensional, Lemma 4.7 provides ℓ^1 geodesic rays $\bar{\alpha}, \bar{\beta}$ with common basepoint $\alpha(0) = \beta(0)$, that lie at uniformly bounded Hausdorff distance from α, β respectively. Note that $\mathcal{W}(\alpha) \sim \mathcal{W}(\bar{\alpha})$ and $\mathcal{W}(\beta) \sim \mathcal{W}(\bar{\beta})$.

We claim that $\mathcal{W}(\bar{\alpha}) \cap \mathcal{W}(\bar{\beta})$ is finite. Suppose not. By Lemma 8.6, $Q(\psi(a)) = \partial_T \mathcal{J}_{\bar{\alpha}}$, where $\mathcal{J}_{\bar{\alpha}} = \mathcal{J}(\bar{\alpha}(0), \bar{\alpha}(\infty))$ is the cubical convex hull of $\bar{\alpha}$. Similarly, $Q(\psi(b)) = \partial_T \mathcal{J}_{\bar{\beta}}$. Since $C = \mathcal{J}_{\bar{\alpha}} \cap \mathcal{J}_{\bar{\beta}} \neq \emptyset$, we have $\mathcal{W}(C) = \mathcal{W}(\bar{\alpha}) \cap \mathcal{W}(\bar{\beta})$ by Lemma 3.23.(5). Now $\mathcal{J}_{\bar{\alpha}}$ is a proper CAT(0) space, by Lemma 3.24. Thus C is also a proper CAT(0) space, hence $|\mathcal{W}(C)| = \infty$ implies C is unbounded and $\partial_T C$ is a nonempty subspace of $Q(\psi(a))$. Similarly, $\partial_T C$ is a nonempty subspace of $Q(\psi(b))$. Thus $\partial_T C \subset Q(\psi(a)) \cap Q(\psi(b)) = \emptyset$, a contradiction. This proves the claim.

Since $\mathcal{W}(\alpha) \sim \mathcal{W}(\bar{\alpha})$ and $\mathcal{W}(\beta) \sim \mathcal{W}(\bar{\beta})$, we may define $N = |\mathcal{W}(\alpha) \cap \mathcal{W}(\beta)| < \infty$. Recall that $\alpha(0) = \beta(0)$. Then, for all $t \geq 0$, we have

$$|\mathcal{W}(\alpha(t), \beta(t))| \geq |\mathcal{W}(\alpha(t), \alpha(0))| + |\mathcal{W}(\beta(0), \beta(t))| - 2N.$$

Applying Lemma 3.1 to both sides gives

$$\lambda_0 \cdot d_X(\alpha(t), \beta(t)) + \lambda_1 \geq \left(\frac{t}{\lambda_0} - \lambda_1\right) + \left(\frac{t}{\lambda_0} - \lambda_1\right) - 2N = \frac{2t}{\lambda_0} - 2(N + \lambda_1).$$

Here, $\lambda_0 \geq 1$ and λ_1 are constants depending only on $\dim X$. Taking limits gives

$$\lim_{t \rightarrow \infty} \frac{d_X(\alpha(t), \beta(t))}{2t} \geq \frac{1}{\lambda_0^2}.$$

By [7, Proposition II.9.8.(4)], we thus have $\sin(\angle_T(a, b)/2) \geq 1/\lambda_0^2$. Now,

$$d_T(a, b) \geq \angle_T(a, b) \geq \sin^{-1}(1/\lambda_0^2) > 1/\lambda_0^2.$$

Setting $\epsilon_0 = 1/\lambda_0^2$ completes the proof. □

In the next lemma, we show that, given a collection of Roller classes, the intersection of the associated Tits boundary realizations coincides with the Tits boundary of the intersection of the convex hulls of the corresponding CAT(0) geodesic rays.

Lemma 9.7 *Let $v_1, \dots, v_k \in \text{Vis}(X)$ be Roller classes. Then there exist convex subcomplexes $\mathcal{J}_1, \dots, \mathcal{J}_k \subset X$ such that the following holds. Each \mathcal{J}_i is the cubical convex hull of a CAT(0) geodesic ray γ_i , with a common basepoint. Furthermore, for every subset $J \subset \{1, \dots, k\}$, we have $\bigcap_{i \in J} Q(v_i) = \partial_T(\bigcap_{i \in J} \mathcal{J}_i)$.*

Proof Fix a basepoint $x \in X^{(0)}$. For $1 \leq i \leq k$, let $\pi_{v_i}: X \rightarrow v_i$ be the gate map of Proposition 3.17, and let $y_i = \pi_{v_i}(x)$. Let $\mathcal{J}_i = \mathcal{J}(x, y_i)$ be the convex subcomplex of X determined by the vertex interval between x and y_i as in Definition 8.5. Then $\mathcal{W}(\mathcal{J}_i) = \mathcal{W}(x, y_i) = \mathcal{W}(x, v_i)$, where the last equality holds by Proposition 3.17. By Lemma 8.6, $Q(v_i) = \partial_T \mathcal{J}_i$.

For each i , let γ_i be a CAT(0) geodesic ray that emanates from x and represents $\varphi(v_i)$. We want to show that $\mathcal{J}_i = \text{CH}(\gamma_i)$, the cubical convex hull of γ_i .

First, we claim that $\text{CH}(\gamma_i) \subset \mathcal{J}_i$. Since \mathcal{J}_i is convex in the CAT(0) metric d_X , it contains a CAT(0) geodesic from every point in the interior to every point on the boundary. In particular, it contains the unique CAT(0) geodesic γ_i from $\gamma_i(0) = x \in \mathcal{J}_i$ to $\gamma_i(\infty) \in Q(v_i) = \partial_T \mathcal{J}_i$. Thus $\gamma_i \subset \mathcal{J}_i$. Since \mathcal{J}_i is cubically convex, it follows that $\text{CH}(\gamma_i) \subset \mathcal{J}_i$.

Next, we claim that $\mathcal{J}_i \subset \text{CH}(\gamma_i)$, or equivalently $\mathcal{W}(\mathcal{J}_i) \subset \mathcal{W}(\gamma_i)$. Since $\gamma_i(0) = x \in \mathcal{J}_i$, any hyperplane $\hat{h} \in \mathcal{W}(\mathcal{J}_i) \setminus \mathcal{W}(\gamma_i)$ would separate the entire ray γ_i from the entire class v_i , contradicting that $\psi(\gamma_i(\infty)) = \psi(\varphi(v_i)) = v_i$, where the second equality holds by visibility and Lemma 8.20. Hence $\mathcal{W}(\mathcal{J}_i) \subset \mathcal{W}(\gamma_i)$, and we conclude that $\mathcal{J}_i = \text{CH}(\gamma_i)$.

Finally, consider an arbitrary subset $J \subset \{1, \dots, k\}$. The inclusion $\partial_T(\bigcap_{i \in J} \mathcal{J}_i) \subset \bigcap_{i \in J} Q(v_i)$ is immediate. For the other inclusion, let $b \in \bigcap_{i \in J} Q(v_i)$. Then each convex set \mathcal{J}_i contains the unique CAT(0) ray β starting at x and ending at b . Thus $\beta \subset \bigcap_{i \in J} \mathcal{J}_i$, hence $b \in \partial_T(\bigcap_{i \in J} \mathcal{J}_i)$. □

Remark 9.8 The following related statement about ℓ^1 geodesics can be proved by an easier analogue of the proof of Lemma 9.7. Let v_1, \dots, v_k be arbitrary Roller classes. Then there exist convex subcomplexes $\mathcal{J}_1, \dots, \mathcal{J}_k$ (defined in exactly the same way, namely $\mathcal{J}_i =$

$\mathcal{J}(x, \pi_{v_i}(x))$ such that each $\mathcal{W}(\mathcal{J}_i)$ is a UBS representing v_i . Furthermore, each \mathcal{J}_i is the cubical convex hull of a combinatorial geodesic, with a common basepoint. The intersection $\bigcap_i \mathcal{J}_i$ itself has the property that $\mathcal{W}(\bigcap_i \mathcal{J}_i) = \bigcap_i \mathcal{W}(\mathcal{J}_i)$. In particular, if $\mathcal{W}(\bigcap_i \mathcal{J}_i)$ is infinite, then it is an ℓ^1 -visible UBS.

The next lemma is crucial. We will later replace the closed covering of $\partial_T X$ by Tits boundary realizations with an open covering, by slightly thickening each Tits realization. The lemma says, roughly, that for a collection of Tits boundary realizations corresponding to a simplex in the nerve of the closed covering, any point in the intersection of small ϵ -neighborhoods of the Tits boundary realizations is still in a $f(\epsilon)$ -neighborhood of the intersection of the Tits boundary realizations.

Lemma 9.9 *There exists a constant $\epsilon_1 > 0$, depending on $\dim(X)$, and a function $f : (0, \epsilon_1) \rightarrow (0, \pi/2)$ such that $\lim_{\epsilon \rightarrow 0} f(\epsilon) = 0$ and such that the following holds. Fix $\epsilon \in (0, \epsilon_1)$ and let $v_0, \dots, v_k \in \text{Vis}(X)$ satisfy $\bigcap_{i=0}^k Q(v_i) \neq \emptyset$. For every $a \in \partial_T X$ such that $d_T(a, Q(v_i)) \leq \epsilon$ for all i , there is a point $b \in \bigcap_{i=0}^k Q(v_i)$ such that $d_T(a, b) \leq f(\epsilon)$.*

Proof We begin by defining the constant ϵ_1 and the function $f : (0, \epsilon_1) \rightarrow (0, \pi/2)$. Let λ_0, λ_1 be the constants of Lemma 3.1, which depend only on $\dim(X)$. Define $K = 4\lambda_0^2 \dim(X)$. Then, we set

$$\epsilon_1 = \frac{4}{5K}, \quad f(\epsilon) = \cos^{-1}(1 - K\epsilon) \leq \cos^{-1}\left(\frac{1}{5}\right) < \frac{\pi}{2}.$$

The bound $f(\epsilon) \leq \cos^{-1}(\frac{1}{5})$ holds because $K\epsilon < K\epsilon_1 = \frac{4}{5}$. The property $\lim_{\epsilon \rightarrow 0} f(\epsilon) = 0$ is now immediate.

Plan of the proof. Let $v_0, \dots, v_k \in \text{Vis}(X)$ be Roller classes such that $\bigcap_{i=0}^k Q(v_i) \neq \emptyset$. By Lemma 9.7, there exist convex subcomplexes $\mathcal{J}_0, \dots, \mathcal{J}_k$ such that $\bigcap_{i=0}^k Q(v_i) = \partial_T(\bigcap_{i=0}^k \mathcal{J}_i)$. Choose a basepoint $x \in \bigcap_{i=0}^k \mathcal{J}_i$.

Fix $\epsilon < \epsilon_1$, and let $a \in \partial_T X$ be a point such that $d_T(a, Q(v_i)) \leq \epsilon$ for all i . Let $\alpha : [0, \infty) \rightarrow X$ be a CAT(0) geodesic ray, such that $\alpha(0) = x$ and $\alpha(\infty) = a$. To prove the lemma, we will find a point $b \in \bigcap_{i=0}^k Q(v_i)$ such that $d_T(a, b) \leq f(\epsilon)$.

We locate the point b using gate projections. Let $\pi : X \rightarrow \bigcap_{i=1}^k \mathcal{J}_i$ be the gate map. For $t \geq 0$, let $x_t = \pi(\alpha(t))$. Since $\bigcap_{i=1}^k \mathcal{J}_i$ is a finite-dimensional CAT(0) cube complex with no facing triple of hyperplanes, it is proper by Lemma 3.24. We will check that the sequence $\{x_t\}_{t \in \mathbb{N}}$ is unbounded. Thus the segments joining x_0 to the elements of $\{x_t\}_{t \in \mathbb{N}}$ subconverge uniformly on compact sets to a CAT(0) geodesic ray representing some point $b \in \bigcap_{i=0}^k Q(v_i)$.

To prove that $d_T(a, b) \leq f(\epsilon)$, we will estimate the Alexandrov angle at $x = \alpha(0)$ between $\alpha(t)$ and x_t . In order to do this, we will control the number of hyperplanes that separate $\alpha(t)$ from $\bigcap_{i=0}^k \mathcal{J}_i$.

Main hyperplane estimate. For each $t \geq 0$, let $\mathcal{U}(t)$ be the set of hyperplanes that separate $\alpha(t)$ from $\bigcap_{i=0}^k \mathcal{J}_i$. We will prove the following double-sided estimate for all large t :

$$\frac{1}{\lambda_0} d_X(\alpha(t), x_t) - \lambda_1 \leq |\mathcal{U}(t)| \leq \dim(X)(\lambda_0 \cdot 2t\epsilon + \lambda_1) + D, \tag{9.10}$$

where D is a constant independent of t .

The lower bound of (9.10) is straightforward. By Lemma 3.22, the gate map $\pi : X \rightarrow \bigcap_{i=0}^k \mathcal{J}_i$ is characterized by the property that a hyperplane \hat{h} separates $\alpha(t)$ from $x_t = \pi(\alpha(t))$

if and only if \hat{h} separates $\alpha(t)$ from $\bigcap_{i=0}^k \mathcal{J}_i$. In other words, $\mathcal{U}(t) = \mathcal{W}(\alpha(t), x_t)$. Hence Lemma 3.1 gives

$$\frac{1}{\lambda_0} d_X(\alpha(t), x_t) - \lambda_1 \leq |\mathcal{U}(t)|,$$

as desired. By contrast, the upper bound requires some hyperplane combinatorics.

Hyperplane sets. For each i , let $\mathcal{B}_i(t)$ be the set of hyperplanes that separate $\alpha(t)$ from \mathcal{J}_i . Note that $\mathcal{U}(t) \supset \bigcup_{i=0}^k \mathcal{B}_i(t)$. We claim that $\mathcal{U}(t) = \bigcup_{i=0}^k \mathcal{B}_i(t)$.

Suppose that $\hat{h} \in \mathcal{U}(t) \setminus \bigcup_{i=0}^k \mathcal{B}_i(t)$. Then \hat{h} crosses each \mathcal{J}_i , so by a standard cubical convexity argument, \hat{h} crosses $\bigcap_{i=0}^k \mathcal{J}_i$ and hence cannot separate any point from $\bigcap_{i=0}^k \mathcal{J}_i$. This is a contradiction, and thus $\mathcal{U}(t) = \bigcup_{i=0}^k \mathcal{B}_i(t)$.

Let $\mathcal{A}_0(t) = \mathcal{B}_0(t)$. For $i \geq 1$, let $\mathcal{A}_i(t) = \mathcal{B}_i(t) \setminus \bigcup_{j < i} \mathcal{A}_j(t)$. Then the sets $\mathcal{A}_0(t), \dots, \mathcal{A}_k(t)$ are pairwise disjoint, and their union is $\mathcal{U}(t)$. For each i , let $\mathcal{A}_i = \bigcup_{t \geq 0} \mathcal{A}_i(t)$. Observe that every \mathcal{A}_i is inseparable, by definition. Furthermore, \mathcal{A}_i is unidirectional. Indeed, $\mathcal{A}_i \subset \mathcal{B}_i = \bigcup_{t \geq 0} \mathcal{B}_i(t)$, and \mathcal{B}_i separates \mathcal{J}_i from the tail of the ray α . Unidirectionality passes to subsets, hence \mathcal{A}_i is unidirectional. Consequently \mathcal{A}_i is either finite or a UBS.

Let $J \subset \{0, \dots, k\}$ be the set of i such that \mathcal{A}_i is a UBS. Observe that $\mathcal{A}_i \cap \mathcal{A}_j = \emptyset$ for $i \neq j$ and $\bigsqcup_{i \in J} \mathcal{A}_i \subset \mathcal{W}(\alpha)$, which is itself a UBS. Thus, by Lemma 6.8, we have $|J| \leq \dim(X)$. We define $D = \sum_{i \notin J} |\mathcal{A}_i|$.

The upper bound of (9.10). Next, we will use the decomposition $\mathcal{U}(t) = \bigsqcup_{i=1}^k \mathcal{A}_i(t)$ to prove the upper bound of (9.10).

For each $i \in \{0, \dots, k\}$, let ξ_i be a ray in \mathcal{J}_i issuing from $x = \alpha(0)$, such that $\xi_i(\infty)$ is a closest point in $Q(v_i) = \partial_T \mathcal{J}_i$ to a . Now, for all i and all t , any hyperplane $\hat{h} \in \mathcal{A}_i(t)$ separates $\alpha(t)$ from \mathcal{J}_i and hence separates $\alpha(t)$ from $\xi_i(t)$. Thus Lemma 3.1 gives

$$|\mathcal{A}_i(t)| \leq |\mathcal{W}(\alpha(t), \xi_i(t))| \leq \lambda_0 \cdot d_X(\alpha(t), \xi_i(t)) + \lambda_1.$$

Now, by [7, Proposition II.9.8.(4)], for any $\delta > 0$ there exists t_0 such that for all $t > t_0$, we have $d_X(\alpha(t), \xi_i(t)) \leq t(2 \sin \frac{\epsilon}{2} + \delta)$. Since ϵ is independent of t , it follows that for all sufficiently large t , we have

$$d_X(\alpha(t), \xi_i(t)) \leq 4t \sin \frac{\epsilon}{2} < 2t\epsilon.$$

Summing $|\mathcal{A}_i(t)|$ over all i gives

$$\begin{aligned} |\mathcal{U}(t)| &= \sum_{i \in J} |\mathcal{A}_i(t)| + \sum_{i \notin J} |\mathcal{A}_i(t)| \\ &\leq \sum_{i \in J} \left(\lambda_0 \cdot d_X(\alpha(t), \xi_i(t)) + \lambda_1 \right) + \sum_{i \notin J} |\mathcal{A}_i| \\ &\leq |J|(\lambda_0 \cdot 2t\epsilon + \lambda_1) + D, \end{aligned}$$

which proves (9.10) because $|J| \leq \dim(X)$.

Angle estimate. Let $x = \alpha(0)$. Set $A = d_X(\alpha(0), \alpha(t)) = t$ (latter equality since α is a geodesic), and $B = d_X(\alpha(0), x_t)$, and $C = d_X(\alpha(t), x_t)$. Let $\bar{\theta}_t$ be the Euclidean angle at the vertex corresponding to x in the Euclidean comparison triangle with sides of length A, B, C , i.e. the comparison angle between $\alpha(t)$ and x_t at x . Let θ_t be the Alexandrov angle at $x = \alpha(0)$ between x_t and $\alpha(t)$, so that $\theta_t \leq \bar{\theta}_t$.

We can estimate $\bar{\theta}_t$ using the law of cosines, starting with an estimate of C . Multiplying every term of (9.10) by λ_0 and choosing t sufficiently large yields

$$C = d_X(\alpha(t), x_t) \leq 4\lambda_0^2 \dim(X) \cdot \epsilon t = K\epsilon t. \tag{9.11}$$

Meanwhile, B can be estimated as follows:

$$t(1 - K\epsilon) \leq d_X(\alpha(0), \alpha(t)) - d_X(\alpha(t), x_t) \leq d_X(\alpha(0), x_t) = B \leq t. \tag{9.12}$$

Here, the first inequality is by (9.11), the second inequality is the triangle inequality, and the final inequality holds because the projection π is 1-Lipschitz for the CAT(0) metric by Lemma 3.22.

By the cosine law and the CAT(0) inequality, we have

$$2AB \cos(\theta_t) \geq 2AB \cos(\bar{\theta}_t) = A^2 + B^2 - C^2.$$

The equality is the cosine law applied to the Euclidean comparison triangle, and the inequality holds for the corresponding CAT(0) triangle. We are interested in the former, since we will need a bound on $\bar{\theta}_t$ (which will incidentally bound θ_t). Substituting the bounds (9.11) and (9.12) into the cosine law shows that for all large t , we have

$$2t^2 \cos \bar{\theta}_t \geq t^2 + t^2(1 - K\epsilon)^2 - t^2(K\epsilon)^2,$$

which simplifies to $\cos \bar{\theta}_t \geq 1 - K\epsilon$. Since we have chosen $\epsilon \in (0, \frac{4}{5K})$, this ensures that $\bar{\theta}_t$ is a small angle; more precisely $\bar{\theta}_t \leq \cos^{-1}(1 - K\epsilon)$.

Conclusion. Since $K\epsilon < \frac{4}{5}$, equation (9.12) implies $d_X(\alpha(0), x_t) \geq (1 - K\epsilon)t > \frac{1}{5}t$. Thus the sequence $\{x_t\}_{t \in \mathbb{N}}$ is unbounded, and some subsequence converges to a point $b \in \bigcap_{i=0}^k Q(v_i)$, as mentioned above. By [7, Lemma II.9.16], which relates the Tits angle $\angle(a, b)$ to the Euclidean comparison angles $\bar{\theta}_t$, we have

$$\angle(a, b) \leq \liminf_{t \rightarrow \infty} \bar{\theta}_t \leq \cos^{-1}(1 - K\epsilon) = f(\epsilon).$$

Since we have already checked that $f(\epsilon) < \frac{\pi}{2}$, this ensures $\angle(a, b) \leq f(\epsilon) < \pi/2$. In particular, $d_T(a, b) = \angle(a, b) \leq f(\epsilon)$. □

Recall from Proposition 8.12.(4),(5) that for each $v \in \text{Vis}(X)$, there is a radius $r_v < \pi/2$ such that $Q(v)$ is contained in the closed ball $\bar{B}_T(\chi(v), r_v)$ in the Tits metric.

Definition 9.13 (*Thickening constant $\epsilon(w)$*) Let ϵ_0 be the constant of Lemma 9.6. Let $f: (0, \epsilon_1) \rightarrow \mathbb{R}$ be the function from Lemma 9.9. Since $\lim_{\epsilon \rightarrow 0} f(\epsilon) = 0$, that lemma allows us to choose a constant $\epsilon_2 < \epsilon_1$, depending only on $\dim(X)$, so that $f(\epsilon) + \epsilon \leq \epsilon_0$ whenever $\epsilon \leq \epsilon_2$. For each Roller class w , fix a *thickening constant* $\epsilon(w) = \min\{\epsilon_0/2, \epsilon_2, r_w/4\} > 0$.

Since the constants ϵ_0, ϵ_2 depend only on $\dim(X)$, and the intrinsic radius r_w depends only on the $\text{Aut}(X)$ -orbit of w by Corollary 8.8, it follows that $\epsilon(w)$ itself depends only on the $\text{Aut}(X)$ -orbit of w .

Proposition 9.14 (Open neighborhood $U(v)$) *For every $v \in \text{Vis}(X)$, there exists a subset $U(v) \subset \partial_T X$ such that the following hold:*

- $Q(v) \subset U(v) \subset \mathcal{N}_{\epsilon(w)}(Q(v))$;
- $U(v)$ has diameter less than $3\pi/4$;
- $U(v)$ is open;
- $U(v)$ is convex and contractible;

- $U(gv) = gU(v)$ for all $g \in \text{Aut}(X)$.

Proof Let Y be the Euclidean cone on $\partial_T X$, with cone-point denoted 0. By [7, Theorem II.3.14], the usual cone metric d_Y on Y has the property that (Y, d_Y) is a CAT(0) geodesic space. We use the notation $v_\delta(A)$ to denote the open δ -neighborhood of a set $A \subset Y$ in this CAT(0) metric.

We identify $\partial_T X$ with the unit sphere about 0 in Y . We also set $Q = Q(v)$ for simplicity of notation. Recall from Proposition 8.12.(4) that the intrinsic radius of Q is $r_v \in [0, \pi/2)$.

The convex sub-cone $Z = Z(v)$: Let $Z \subset Y$ be the subspace arising as the union of all rays in Y emanating from 0 and passing through $Q \subset \partial_T X$. We claim that Z is convex in (Y, d_Y) . Indeed, let $x, y \in Z$. Then x, y respectively lie on rays α_x, α_y emanating from 0 and intersecting Q in points x', y' . By [7, Proposition I.5.10(1)], the rays α_x, α_y determine a convex sector $S \subset Y$ intersecting $\partial_T X$ in the unique Tits geodesic from x' to y' , which lies in Q by Tits convexity of Q . Hence S lies in Z , and thus so does the d_Y -geodesic from x to y , since $x, y \in S$.

The horoball $B = B(v)$: Let $c = \chi(v)$ be the circumcenter of $Q = Q(v)$. Let $\alpha = \alpha_c$ be the parametrized geodesic ray in Y such that $\alpha(0) = 0$ and $\alpha(1) = c$. Let $p: Y \rightarrow \mathbb{R}$ be the associated horofunction, defined by

$$p(y) = \lim_{s \rightarrow \infty} [d_Y(\alpha(s), y) - s].$$

Let $t_0 = \cos r_v$, which is positive since $r_v \in [0, \pi/2)$.

Let $B = B(v) = p^{-1}((-\infty, -t_0])$. Recall from e.g. [26, Lemma 3.88] that B is a convex subset of Y . Let $K(v) = B \cap Z$. See Fig. 7.

Q lies in $K(v)$: We claim that $Q \subset K(v)$. Since Q is the unit sphere in Z by the definition of Y and Z , it suffices to show that $Q \subset B$. To that end, fix $q \in Q$, and observe that $d_Y(q, 0) = 1$. Let $s \gg 0$ and consider the convex Euclidean triangle in Y with vertices at $0, q, \alpha(s)$. The angle in this triangle at 0 is denoted θ , and we note that $\theta \leq r_v$. From the cosine law, we obtain $d_Y(q, \alpha(s))^2 = s^2 + 1 - 2s \cos \theta$. Hence

$$p(q) = \lim_{s \rightarrow \infty} \left[\sqrt{s^2 + 1 - 2s \cos \theta} - s \right] = -\cos \theta \leq -\cos r_v = -t_0.$$

Thus, by the definition of B , we have $q \in B$, as required. Since $Z \cap \partial_T X = Q$, we have in fact shown that $K(v) \cap \partial_T X = Q$.

$K(v)$ satisfies $d(0, K(v)) = t_0$: We claim that $K(v)$ is disjoint from the open t_0 -ball in Y about 0. Indeed, let $b \in K(v)$. Then $p(b) \leq -t_0$. Hence, for any fixed $\delta > 0$, and any sufficiently large $s > 0$, we have $d_Y(b, \alpha(s)) < s - t_0 + \delta$. By the triangle inequality, $s \leq d_Y(0, b) + s - t_0 + \delta$, i.e. $d_Y(0, b) > t_0 - \delta$. Hence $d_Y(b, 0) \geq t_0$, as required. At the same time, $\alpha = \alpha_c$ is a parametrized geodesic ray and $\alpha(t_0)$ is a point of $K(v)$ at distance exactly t_0 from 0, hence $d(0, K(v)) = t_0$.

The projection ϖ and the open set $U'(v)$: Let $\varpi : Y \setminus \{0\} \rightarrow \partial_T X$ be the radial projection. Recall that $v_\rho(0)$ is the open ρ -neighborhood of 0. We claim that for every $0 < \rho \leq 1$, the restriction of ϖ to $Y \setminus v_\rho(0)$ is $1/\rho$ -Lipschitz. Indeed, by [7, Proposition I.5.10(1)], it suffices to check this claim on a sector of \mathbb{R}^2 . In polar coordinates on \mathbb{R}^2 , the distance element ds satisfies $ds^2 = dr^2 + r^2 d\theta^2$, hence the projection ϖ rescales ds by a factor of at most r^{-1} , hence ϖ is $1/\rho$ -Lipschitz when $r \geq \rho$.

Recall the constant $\epsilon(v)$ defined in Definition 9.13, and define a positive constant

$$\epsilon' = \epsilon'(v) = \frac{\epsilon(v)}{1 + \epsilon(v)} t_0 \in (0, t_0), \quad \text{which satisfies} \quad \frac{\epsilon'}{t_0 - \epsilon'} = \epsilon(v). \quad (9.15)$$




Fig. 7 Construction of $U' = U'(v)$ as a small-radius neighborhood of $K = B \cap Z$, where $Z \subset Y$ is the cone on $Q \subset \partial_T X$ and B is an appropriately chosen horoball determined by α . The radius is chosen so that U' avoids the $(t_0 - \epsilon')$ -ball around 0

Let $U'(v) = v_{\epsilon'}(K(v))$ be the open ϵ' -neighborhood of $K(v)$ in Y . Since $K(v) = B \cap Z$ is convex and Y is CAT(0), it follows that $U'(v)$ is convex. Furthermore, $U'(v)$ is open and disjoint from the ball of radius $t_0 - \epsilon'$ about 0. See Fig. 7.

Definition and properties of $U(v)$: Let $U(v) = \varpi(U'(v)) \subset \partial_T X$. We claim that this has all of the properties claimed in the lemma statement. First, since $Q(v) \subset K(v)$ by construction, and ϖ is the identity map on $\partial_T X \subset Y$, it follows that $Q(v) = \varpi(K(v)) \subset U(v)$. On the other hand, since $d(0, U'(v)) \geq t_0 - \epsilon'$, the projection ϖ is $1/(t_0 - \epsilon')$ -Lipschitz on $U'(v)$. Therefore, (9.15) implies

$$U(v) = \varpi(v_{\epsilon'}(K(v))) \subset v_{\epsilon(v)}(\varpi(K(v))) = v_{\epsilon(v)}(Q(v)),$$

proving the first bullet of the lemma.

Second, recall from Definition 9.13 that $\epsilon(v) \leq r_v/4$. Thus Proposition 8.12.(4) implies

$$\text{diam}(U(v)) \leq \text{diam}(Q(v)) + 2\epsilon(v) \leq \frac{\pi}{2} + \frac{r_v}{2} < \frac{\pi}{2} + \frac{\pi}{4} = \frac{3\pi}{4}.$$

Third, observe that ϖ is an open map. Indeed, $Y \setminus \{0\}$ is homeomorphic to $\partial_T X \times (0, \infty)$, and ϖ is projection to the first factor, which is an open map. Thus $U(v) = \varpi(U'(v))$ is open.

Fourth, it follows from [7, Proposition I.5.10.(1)] and the convexity of $U'(v)$ that $U(v)$ is convex in $\partial_T X$. Since $U(v)$ has diameter less than π , it is uniquely geodesic, hence contractible.

Fifth, observe that all of the ingredients in the definition of $U(v)$ are $\text{Aut}(X)$ -equivariant. Indeed, the intrinsic radius r_v is invariant under $\text{Aut}(X)$, hence the number $t_0 = \cos(r_v)$ is as well. Thus the definitions of the sub-cone $Z(v)$, the horoball $B(v)$, and the set $K(v) = Z(v) \cap B(v)$ are all $\text{Aut}(X)$ -equivariant. The constant $\epsilon'(v)$, defined in (9.15) is invariant under $\text{Aut}(X)$, because both t_0 and $\epsilon(v)$ are. It follows that $U(v) = \varpi(v_{\epsilon'}(K(v)))$ is also $\text{Aut}(X)$ -equivariant. \square

Recall that in Sect. 9.1, we have defined a simplicial complex \mathcal{N}_T . By Corollary 9.5, \mathcal{N}_T coincides with the nerve of the cover of $\partial_T X$ by the closed sets $Q(v)$ for $v \in \text{MaxVis}(X)$. By Lemmas 9.3 and 9.14, the open sets $U(v)$ for $v \in \text{MaxVis}(X)$ also cover $\partial_T X$.

By analogy with Definition 9.2, we define a simplicial complex \mathcal{C}_T with vertex set $\text{MaxVis}X$, where vertices $v_0, \dots, v_n \in \text{MaxVis}X$ span an n -simplex if and only if $\bigcap_{i=0}^n U(v_i) \neq \emptyset$. The content of the following lemma is that \mathcal{C}_T is isomorphic to \mathcal{N}_T .

Lemma 9.16 *For Roller classes $v_0, \dots, v_n \in \text{MaxVis}(X)$, we have $\bigcap_{i=0}^n U(v_i) \neq \emptyset$ if and only if $\bigcap_{i=0}^n Q(v_i) \neq \emptyset$. Consequently, the identity map $\text{MaxVis}(X) \rightarrow \text{MaxVis}(X)$ induces an $\text{Aut}(X)$ -equivariant simplicial isomorphism $\mathcal{C}_T \rightarrow \mathcal{N}_T$.*

Proof Observe that $\bigcap_{i=0}^n Q(v_i) \neq \emptyset$ trivially implies $\bigcap_{i=0}^n U(v_i) \neq \emptyset$. For the other direction, we assume $\bigcap_{i=0}^n U(v_i) \neq \emptyset$ and aim to show $\bigcap_{i=0}^n Q(v_i) \neq \emptyset$.

We will argue by induction on n . In the base case, $n = 1$, suppose that $U(v_0) \cap U(v_1) \neq \emptyset$. Since $U(v_i) \subset \mathcal{N}_{\epsilon(v_i)}(Q(v_i))$, and $\epsilon(v_i) \leq \epsilon_0/2$ by Definition 9.13, it follows that there are points $a_i \in Q(v_i)$ such that $d_T(a_0, a_1) < \epsilon_0$. But then Lemma 9.6 implies $Q(v_0) \cap Q(v_1) \neq \emptyset$, proving the base case.

For the inductive step, assume that $\bigcap_{i=0}^n U(v_i) \neq \emptyset$ and $\bigcap_{i=0}^{n-1} Q(v_i) \neq \emptyset$. We claim that there is a Roller class $w \in \text{Vis}(X)$ such that $\bigcap_{i=0}^{n-1} Q(v_i) = Q(w)$.

By Lemma 9.7, we have convex subcomplexes $\mathcal{J}_0, \dots, \mathcal{J}_{n-1} \subset X$, such that $Q(v_i) = \partial_T \mathcal{J}_i$ and $\bigcap_{i=0}^{n-1} Q(v_i) = \partial_T \bigcap_{i=0}^{n-1} \mathcal{J}_i$. By Lemma 6.20, the hyperplane collection $\mathcal{W}(\mathcal{J}_i)$ is a UBS. Furthermore, since each \mathcal{J}_i is convex, we have $\mathcal{W}(\bigcap \mathcal{J}_i) = \bigcap \mathcal{W}(\mathcal{J}_i)$, which is a UBS because $\bigcap \mathcal{J}_i$ is unbounded. By Lemma 6.18, the UBS $\mathcal{W}(\bigcap \mathcal{J}_i)$ represents a Roller class w' , which means that $Q(w') = \partial_T \bigcap_{i=0}^{n-1} \mathcal{J}_i = \bigcap_{i=0}^{n-1} Q(v_i)$. Now, we let $w = \psi(\varphi(w'))$. Then Definition 8.19 says that $w \in \text{Vis}(X)$, and Proposition 8.18.(3) says that $Q(w) = Q(w') = \bigcap_{i=0}^{n-1} Q(v_i)$, proving the Claim.

Next, we claim that $Q(v_n) \cap Q(w) \neq \emptyset$. Since $Q(w) = \bigcap_{i=0}^{n-1} Q(v_i)$, this will complete the inductive step and prove the Lemma.

Suppose for a contradiction that $Q(v_n) \cap Q(w) = \emptyset$. Then, by Lemma 9.6, we have $d_T(Q(v_n), Q(w)) > \epsilon_0$. Let ϵ_2 be as in Definition 9.13. Fix $a \in \bigcap_{i=0}^n U(v_i) \subset \bigcap_{i=0}^n \mathcal{N}_{\epsilon(v_i)}(Q(v_i))$, which ensures that $d_T(a, Q(v_n)) < \epsilon_2$. Since $\epsilon(v_i) \leq \epsilon_2$ for each v_i , Lemma 9.9 implies $d_T(a, Q(w)) < f(\epsilon_2)$. Thus $d_T(Q(v_n), Q(w)) < \epsilon_2 + f(\epsilon_2) \leq \epsilon_0$, a contradiction. This proves the Claim and the equivalence $\bigcap_{i=0}^n U(v_i) \neq \emptyset \Leftrightarrow \bigcap_{i=0}^n Q(v_i) \neq \emptyset$.

This gives the simplicial isomorphism $\mathcal{C}_T \rightarrow \mathcal{N}_T$. This isomorphism is $\text{Aut}(X)$ -equivariant, by Corollary 8.8 and Proposition 9.14. □

9.3 Homotopy equivalence of $\partial_T X$ and \mathcal{N}_T

Theorem 9.17 *There is an $\text{Aut}(X) \sim$ homotopy equivalence from the simplicial complex \mathcal{N}_T to $\partial_T X$.*

Proof. By Corollary 9.5, \mathcal{N}_T is the nerve of the covering of $\partial_T X$ by the closed sets $Q(v)$, where v varies over $\text{MaxVis}(X)$.

By Proposition 9.14, $\partial_T X$ admits an open covering $\{U(v) : v \in \text{MaxVis}(X)\}$, with each $U(v)$ a convex, uniquely geodesic subspace of $\partial_T X$. Hence, if $v_0, \dots, v_n \in \text{MaxVis}(X)$, then either $\bigcap_{i=0}^n U(v_i) = \emptyset$, or $\bigcap_{i=0}^n U(v_i)$ is contractible. Let \mathcal{L}_T be the nerve of this open covering. By the Equivariant Open Nerve Theorem 2.7, there is an $\text{Aut}(X) \sim$ homotopy equivalence $\mathcal{L}_T \rightarrow \partial_T X$.

As above, let \mathcal{C}_T be the simplicial complex with one vertex for each $v \in \text{MaxVis}(X)$, with v_0, \dots, v_n spanning an n -simplex whenever $\bigcap_{i=0}^n U(v_i) \neq \emptyset$. Note that \mathcal{C}_T is not necessarily isomorphic to \mathcal{L}_T , since we may have $U(v_i) = U(v_j)$ for some pair v_i, v_j with $i \neq j$. But \mathcal{C}_T is homotopy-equivalent to \mathcal{L}_T , which can be seen as follows. The $\text{Aut}(X)$ -equivariant assignment $v \mapsto U(v)$ (i.e. the surjection $\mathcal{C}_T^{(0)} \rightarrow \mathcal{L}_T^{(0)}$) determines an equivariant simplicial map $\mathcal{C}_T \rightarrow \mathcal{L}_T$ such that the preimage of each 0-simplex $U(v)$ is the full subcomplex of \mathcal{C}_T spanned by the set of w with $U(v) = U(w)$. Since any finite set of such w span a simplex of \mathcal{C}_T , this preimage is contractible: it is homeomorphic to the cone on the link of any of its vertices. Now, let σ be a simplex of \mathcal{C}_T . Then by the definition of \mathcal{L}_T , the preimage of σ is the join of the preimages of the vertices of σ , each of which we have just shown to be contractible, so the preimage of σ is contractible. Hence the map $\mathcal{C}_T \rightarrow \mathcal{L}_T$ is a homotopy equivalence, by Quillen’s fiber theorem (see e.g. [11, Theorem 10.5]).

To conclude, recall that by Lemma 9.16, we have an equivariant simplicial isomorphism $\mathcal{C}_T \rightarrow \mathcal{N}_T$. Hence we have a chain of $\text{Aut}(X)$ -homotopy equivalences

$$\mathcal{N}_T \xrightarrow{\cong} \mathcal{C}_T \xrightarrow{\sim} \mathcal{L}_T \xrightarrow{\sim} \partial_T X. \quad \square$$

9.4 Cuboid generalization

All of the results of this section have a straightforward generalization to the cuboid setting. In the following description, we refer to the notation introduced in Sect. 8.4.

Definition 9.1 generalizes immediately to give a set of maximal ℓ^2 visible Roller classes, namely $\text{MaxVis}^\rho(X) \subset \text{Vis}^\rho(X)$. Following Definition 9.2, we get a simplicial complex \mathcal{N}_T^ρ with vertex set $\text{MaxVis}^\rho(X)$, where vertices v_0, \dots, v_n span an n -simplex if and only if $\bigcap_{i=1}^n Q^\rho(v_i) \neq \emptyset$. Then, Lemmas 9.3 and 9.4 generalize immediately because their proofs are assembled from the results of Sect. 8, and we have already checked that those results generalize to cuboids. Thus Corollary 9.5 generalizes as well, and we learn that \mathcal{N}_T^ρ is the covering of $\partial_T^\rho X$ by the closed sets $\{Q^\rho(v) : v \in \text{MaxVis}^\rho(X)\}$.

(Recall that Remark 8.22 outlined an argument that $\text{Vis}(X) = \text{Vis}^\rho(X)$ for every admissible rescaling ρ . It follows that $\text{MaxVis}(X) = \text{MaxVis}^\rho(X)$ as well. The nerves \mathcal{N}_T and \mathcal{N}_T^ρ can then be shown to coincide, as mentioned in Remark 8.22. However, we do not use this.)

Moving ahead to Sect. 9.2, the results generalize as follows. The existence statement of Lemma 9.6 extends, with the modification that the constant ϵ_0 depends on $\dim(X)$ and ρ . This is because the proof of Lemma 9.6 uses the constants λ_0, λ_1 of Lemma 3.1, which needs to be replaced by the constants $\lambda_0^\rho, \lambda_1^\rho$ of Lemma 4.8. Lemma 9.7 extends to cuboids as well, because the proof of that lemma uses combinatorial arguments in combination with the d_X^ρ -convexity of cubical convex hulls. Similarly, Lemma 9.9 extends: its proof works perfectly well in the cuboid metric d_X^ρ , once we substitute Lemma 4.8 for Lemma 3.1. As above, the outputs of that lemma depend on both $\dim(X)$ and ρ . Combining these ingredients, we can follow Definition 9.13 to define a thickening constant $\epsilon^\rho(w) > 0$ for every Roller class w . This constant depends only ρ and the $\text{Aut}^\rho(X)$ -orbit of w .

Proposition 9.14 extends verbatim to cuboids. Indeed, its proof is a CAT(0) argument that takes place in the Euclidean cone on $\partial_T X$, which works equally well in the Euclidean cone on $\partial_T^\rho X$. (The proof also uses Proposition 8.12, which works for cuboids.) At the end of the lemma, we get a collection of open sets $U^\rho(v)$ for $v \in \text{Vis}^\rho(X)$, which is invariant under $\text{Aut}^\rho(X)$.


By analogy with Definition 9.2, we define a simplicial complex \mathcal{C}_T^ρ , where the simplices are defined by intersection patterns of the open sets $U(v_i)$ for $v_i \in \text{MaxVis}^\rho(X)$. Now,

Lemma 9.16 extends verbatim to give an $\text{Aut}^\rho(X)$ -equivariant simplicial isomorphism $C_T^\rho \rightarrow \mathcal{N}_T^\rho$. The proof of Theorem 9.17 also extends verbatim, because all of its ingredients extend. Thus we obtain the following cuboid extension of Theorem 9.17:

Theorem 9.18 *Let ρ be a G -admissible hyperplane rescaling of X . Then there is a G -homotopy homotopy equivalence $\mathcal{N}_T^\rho \rightarrow \partial_T^\rho X$.*

10 Proof of the Main Theorems

In this section, we prove Theorem A, which was stated in the Introduction. That is, we construct the following commutative diagram of $\text{Aut}(X)$ -homotopy equivalences:



The homotopy equivalence $R^S: \partial_\Delta X \rightarrow \mathfrak{R}_\Delta X$ will be proved in Proposition 10.12, while the homotopy equivalence $T^R: \mathfrak{R}_\Delta X \rightarrow \partial_T X$ will be proved in Corollary 10.11. At the end of the section, we check that all of the arguments extend to cuboids, proving Theorem B.

To prove Corollary 10.11, we first study a subcomplex of $\mathfrak{R}_\Delta X$.

Definition 10.1 (*Visible subcomplexes of $\mathfrak{R}_\Delta X$*) Let $\mathfrak{R}_\Delta^\triangleleft X$ be the subcomplex of $\mathfrak{R}_\Delta X$ consisting of all simplices corresponding to chains $v_0 \leq \dots \leq v_n$ such that $v_i \in \text{Vis}(X)$ for $0 \leq i \leq n$.

For each $v \in \text{Vis}(X)$, let Σ_v be the connected subcomplex of $\mathfrak{R}_\Delta^\triangleleft X$ consisting of all simplices corresponding to chains $v_0 \leq \dots \leq v_n$ for which:

- each $v_i \in \text{Vis}(X)$;
- $v_n = v$.

10.1 Closed cover of $\mathfrak{R}_\Delta^\triangleleft X$

Next, we show that the complexes Σ_v yield a closed covering of $\mathfrak{R}_\Delta^\triangleleft X$, and use this covering to build a nerve.

Lemma 10.2 *The set $\{\Sigma_v : v \in \text{MaxVis}(X)\}$ is a closed covering of $\mathfrak{R}_\Delta^\triangleleft X$. Moreover, for all $v, v_0, \dots, v_n \in \text{Vis}(X)$, we have the following:*

- (1) Σ_v is finite.
- (2) $\Sigma_{v_i} = \Sigma_{v_j}$ only if $v_i = v_j$.
- (3) $\bigcap_{i=0}^n \Sigma_{v_i}$ is either empty or contractible.
- (4) For every $g \in \text{Aut}(X)$, we have $g\Sigma_v = \Sigma_{gv}$.

Proof Let $w \in \text{Vis}(X)$. Then there exists $v \in \text{MaxVis}(X)$ with $w \leq v$, and hence $w \in \Sigma_v^{(0)}$. Thus $\{\Sigma_v : v \in \text{MaxVis}(X)\}$ is a closed covering of $\mathfrak{R}_\Delta^\triangleleft X$.

Let v be a (visible) Roller class. By Corollary 6.30, there are finitely many other classes w such that $w \leq v$. This proves assertion (1).

Now suppose that v_i, v_j are (visible) Roller classes satisfying $\Sigma_{v_i} = \Sigma_{v_j}$. Since $v_j \in \Sigma_{v_i}$, we have $v_j \leq v_i$. Likewise, $v_i \leq v_j$. Hence $v_i = v_j$, proving assertion (2).

Next, let $v_0, \dots, v_n \in \text{Vis}(X)$ and suppose that $\bigcap_{i=0}^n \Sigma_{v_i} \neq \emptyset$. First, note that $\Sigma_{v_i} = v_i \star L_{v_i}$, where L_{v_i} (the link of v_i in Σ_{v_i}) is a subcomplex of $\mathfrak{R}_\Delta^{\leq} X$ and \star denotes the simplicial join operation. Hence Σ_{v_i} is topologically a cone, and is thus contractible.

For each i , let \mathcal{V}_i be a UBS representing v_i (recall Definition 6.21). Then $\bigcap_{i=0}^n \mathcal{V}_i$ is unidirectional and contains no facing triple, since those properties hold for each \mathcal{V}_i and are inherited by subsets. Second, $\bigcap_{i=0}^n \mathcal{V}_i$ is inseparable, because each \mathcal{V}_i is.

Now, if $u \in \bigcap_{i=0}^n \Sigma_{v_i}$ is a 0-simplex, then $u \leq v_i$ for all i , so u is represented by a UBS of the form $\mathcal{U} \leq \bigcap_{i=0}^n \mathcal{V}_i$. Hence $\bigcap_{i=0}^n \mathcal{V}_i$ is infinite. Combined with the above discussion, this shows that $\bigcap_{i=0}^n \mathcal{V}_i$ is a UBS. Let v be the corresponding Roller class. Since $\mathcal{U} \leq \bigcap_{i=0}^n \mathcal{V}_i$, we have $u \leq v$.

Since each such u is ℓ_2 -visible, the set of visible Roller classes w with $w \leq v$ is nonempty. Hence there is a unique Roller class m such that $m \leq v$, and m is ℓ_2 -visible, and m is \leq -maximal with those properties, by Lemma 8.14 and Definition 8.19. Since each $u \in \bigcap_{i=0}^n \Sigma_{v_i}$ satisfies $u \leq m \in \mathfrak{R}_\Delta^{\leq} X$, there is a subcomplex $L \subset \mathfrak{R}_\Delta^{\leq} X$ such that $\bigcap_{i=0}^n \Sigma_{v_i} = m \star L$, so $\bigcap_{i=0}^n \Sigma_{v_i}$ is contractible in $\mathfrak{R}_\Delta^{\leq} X$.

Finally, statement (4) follows from the equivariance of $\text{Vis}(X)$ and the partial order \leq . □

Definition 10.3 (Simplicial complex \mathcal{N}_Δ) Let \mathcal{N}_Δ be the nerve of the closed covering $\{\Sigma_v : v \in \text{MaxVis}(X)\}$ of $\mathfrak{R}_\Delta^{\leq} X$. By Lemma 10.2.(2), \mathcal{N}_Δ has vertex-set $\text{MaxVis}(X)$, and v_0, \dots, v_n span an n -simplex if and only if $\bigcap_{i=0}^n \Sigma_{v_i} \neq \emptyset$. By Lemma 10.2.(4), $\text{Aut}(X)$ acts by simplicial automorphisms on \mathcal{N}_Δ .

Lemma 10.4 Let $v_0, \dots, v_n \in \text{Vis}(X)$. Then

$$\bigcap_{i=0}^n \Sigma_{v_i} \neq \emptyset \Leftrightarrow \bigcap_{i=0}^n Q(v_i) \neq \emptyset.$$

Consequently, there is an $\text{Aut}(X)$ -equivariant simplicial isomorphism $\mathcal{N}_\Delta \rightarrow \mathcal{N}_T$.

Proof Suppose that there exists $a \in \bigcap_{i=0}^n Q(v_i)$. Then Lemma 8.9 says $\psi(a) \leq v_i$ for all i . Hence, by Definition 10.1, we have $\psi(a) \in \Sigma_{v_i}$ for all i .

Conversely, suppose that there exists a vertex (i.e. a Roller class) $w \in \bigcap_{i=0}^n \Sigma_{v_i}$. Then $w \leq v_i$ for all i , by definition. Hence, by Lemma 8.11, $Q(w) \subset Q(v_i)$ for all i , from which it follows that $\varphi(w) \in \bigcap_{i=0}^n Q(v_i)$.

Finally, the simplicial isomorphism $\mathcal{N}_\Delta \rightarrow \mathcal{N}_T$ comes from identifying the 0-skeleta of \mathcal{N}_Δ and \mathcal{N}_T with $\text{MaxVis}(X)$. □

10.2 Homotopy equivalence between $\mathfrak{R}_\Delta X$ and $\partial_T X$

We can now assemble the proof that $\mathfrak{R}_\Delta X$ is homotopy equivalent to $\partial_T X$. We do this in two propositions:

Proposition 10.5 There is an $\text{Aut}(X)$ -homotopy equivalence $\mathfrak{R}_\Delta^{\leq} X \rightarrow \partial_T X$.

Proof. Consider the covering $\{\Sigma_v : v \in \text{MaxVis}(X)\}$ of $\mathfrak{R}_\Delta^{\leq} X$ by the subcomplexes Σ_v , coming from Lemma 10.2. By Definition 10.3, the nerve of this cover is \mathcal{N}_Δ . By Lemma 10.2.(3), the intersection of any finite collection of the Σ_v is either empty or contractible. Hence, by

the Equivariant Simplicial Nerve Theorem 2.8, there is an $\text{Aut}(X)\sim$ homotopy equivalence $\mathfrak{R}_\Delta^{\triangleleft} X \rightarrow \mathcal{N}_\Delta$.

By Lemma 10.4, there is a simplicial isomorphism $\mathcal{N}_\Delta \rightarrow \mathcal{N}_T$. Finally, by Theorem 9.17, there is an $\text{Aut}(X)\sim$ homotopy equivalence $\mathcal{N}_T \rightarrow \partial_T X$. Putting it all together, we obtain a chain of $\text{Aut}(X)\sim$ homotopy equivalences

$$\mathfrak{R}_\Delta^{\triangleleft} X \xrightarrow{\sim} \mathcal{N}_\Delta \xrightarrow{\cong} \mathcal{N}_T \xrightarrow{\sim} \partial_T X. \quad \square$$

Proposition 10.6 *The inclusion $\mathfrak{R}_\Delta^{\triangleleft} X \hookrightarrow \mathfrak{R}_\Delta X$ is an $\text{Aut}(X)$ -equivariant homotopy equivalence. Its homotopy inverse is an $\text{Aut}(X)\sim$ deformation retraction $\mathfrak{R}_\Delta X \rightarrow \mathfrak{R}_\Delta^{\triangleleft} X$.*

Proof First, observe that since $\mathfrak{R}_\Delta X$ is $\text{Aut}(X)$ -invariant by Definition 5.7, and the set $\text{Vis}(X)$ is $\text{Aut}(X)$ -invariant by Definition 8.19, the inclusion $\mathfrak{R}_\Delta^{\triangleleft} X \hookrightarrow \mathfrak{R}_\Delta X$ is $\text{Aut}(X)$ -equivariant. In the remainder of the proof, we will construct a deformation retraction $\mathfrak{R}_\Delta X \rightarrow \mathfrak{R}_\Delta^{\triangleleft} X$ that serves as a homotopy inverse to the inclusion $\mathfrak{R}_\Delta^{\triangleleft} X \hookrightarrow \mathfrak{R}_\Delta X$.

Let w be an invisible Roller class (viewed as a 0-simplex of $\mathfrak{R}_\Delta X$). By Lemma 8.21, w is not minimal, so there exists a minimal Roller class m with $m < w$. Let $n(w) \geq 1$ be the maximum length of a chain of the form $m < \dots < w$ with m a minimal Roller class.

Let D be the maximal number such that there exists an invisible Roller class w with $n(w) = D$. By Remark 5.8, every simplex in $\mathfrak{R}_\Delta X$ has at most $\dim X$ vertices, hence $D \leq \dim X$.

Let I_n be the set of invisible Roller classes v with $n(v) = n$. Then

$$(\mathfrak{R}_\Delta X)^{(0)} = (\mathfrak{R}_\Delta^{\triangleleft} X)^{(0)} \sqcup \bigsqcup_{n=1}^D I_n.$$

Let $(\mathfrak{S}X)_0 = \mathfrak{R}_\Delta^{\triangleleft} X$. For $1 \leq N \leq D$, let $(\mathfrak{S}X)_N$ be the subcomplex spanned by $(\mathfrak{R}_\Delta^{\triangleleft} X)^{(0)} \sqcup \bigsqcup_{n=1}^N I_n$, so

$$\mathfrak{R}_\Delta^{\triangleleft} X = (\mathfrak{S}X)_0 \subset (\mathfrak{S}X)_1 \subset \dots \subset (\mathfrak{S}X)_D = \mathfrak{R}_\Delta X.$$

Given a Roller class w and a positive integer $n \leq D$, define the *open star* $st_n(w)$ to be the union of $\{w\}$ and all of the open simplices of $(\mathfrak{S}X)_n$ whose closures contain w . Let $St_n(w)$ be the union of all (closed) simplices of $(\mathfrak{S}X)_n$ containing w . Define the (*downward*) *link* $L_n(w) = St_n(w) \setminus st_n(w)$.

Claim 10.7 *For $w \in I_n$, we have $L_n(w) \subset (\mathfrak{S}X)_{n-1}$.*

Let $w \in I_n \subset (\mathfrak{S}X)_n$. Let σ be a maximal simplex of $(\mathfrak{S}X)_n$ containing w and lying in $(\mathfrak{S}X)_n$. Let $v_0 < \dots < v_m$ be the corresponding chain, with $w = v_i$ for some i . Since $n(w) = n$, the chain $v_0 < \dots < v_i$ has length at most n .

First, consider v_j for $j < i$. If v_j is visible, then $v_j \in (\mathfrak{S}X)_0 \subset (\mathfrak{S}X)_{n-1}$. If v_j is invisible, then since $v_0 < \dots < v_j$ is a chain of length less than n , and σ is maximal, $n(v_j) < n$, hence $v_j \in (\mathfrak{S}X)_{n-1}$.

Next, consider v_j for $j > i$. If v_j is visible, then $v_j \in (\mathfrak{S}X)_0 \subset (\mathfrak{S}X)_{n-1}$. If v_j is invisible, then $v_j \notin I_n$, since there is a chain $v_0 < \dots < v_j$ of length more than $n = n(v_i)$, which would contradict the containment $\sigma \subset (\mathfrak{S}X)_n$.

Hence the simplex σ' corresponding to the chain $v_0 < \dots < v_{i-1} < v_{i+1} < \dots < v_m$ lies in $(\mathfrak{S}X)_{n-1}$. This proves the claim.

Contractibility of links: Now, fix $w \in I_n$. We will show that $L_n(w)$ is contractible. Establishing this involves several claims, leading up to Claim 10.10.

Claim 10.8 For a minimal Roller class with $m \in L_n(w)$, let $St(m) = \bigcup_n St_n(m)$ denote the star of m in $\mathfrak{R}_\Delta X$. Then:

- (1) $L_n(w) \cap St(m)$ is topologically a cone with cone-point m .
- (2) $L_n(w)$ is the union of the subcomplexes $L_n(w) \cap St(m)$, as m varies over the finitely many minimal Roller classes with $m \leq w$.

Let $L(m)$ denote the link of m in $\mathfrak{R}_\Delta X$. Then, by definition, $St(m)$ decomposes as a join: $St(m) = m \star L(m)$. Moreover, $m \in L_n(w)$, so

$$L_n(w) \cap St(m) = L_n(w) \cap (m \star L(m)) = m \star (L_n(w) \cap L(m)),$$

which proves the first assertion.

By Corollary 6.30, there are finitely many minimal Roller classes $m \leq w$. For each such m , we clearly have $L_n(w) \cap St(m) \subset L_n(w)$. Conversely, if $v \in L_n(w)$, then there exists a minimal class m such that $m \leq v$ and $m \leq w$, hence $v \in L_n(w) \cap St(m)$. This proves the claim.

Claim 10.9 Let m_0, \dots, m_k be distinct minimal Roller classes with $m_i < w$ for all i . Then $\bigcap_{i=0}^k St(m_i) \cap L_n(w)$ is topologically a cone, and is in particular nonempty.

By Lemma 8.21, each m_i is ℓ^2 -visible, so we can choose CAT(0) geodesic rays $\xi_i : [0, \infty) \rightarrow X$ such that $\xi_i(0) = \xi_j(0)$ for all i , and such that $\xi_i(\infty) = \varphi(m_i)$ for each i . For each i , define $\mathcal{M}_i = \mathcal{W}(\xi_i)$, a UBS representing m_i . Since m_i is minimal for all i , Lemma 6.8 implies the intersection $\mathcal{M}_i \cap \mathcal{M}_j$ is finite whenever $i \neq j$. Moving the common basepoint $\xi_1(0) = \dots = \xi_k(0)$ ensures that $\mathcal{M}_i \cap \mathcal{M}_j = \emptyset$ for all $i \neq j$, and that each \mathcal{M}_i is the inseparable closure of an infinite descending chain. Indeed, by moving the basepoint across each of the finitely many hyperplanes appearing in some $\mathcal{M}_i \cap \mathcal{M}_j$, we arrange the first property, and the inseparable closure property holds by Lemma 7.6, Lemma 7.7, and minimality of the m_i .

Let \mathcal{A} be a UBS representing w . Since $m_i \leq w$ for all i , we have $\mathcal{M}_i \leq \mathcal{A}$ for each i . Thus it is readily checked that $\bigcup_{i=0}^k \mathcal{M}_i$ is a UBS.

Consider a pair of indices $i \neq j$. Then since $\bigcup_{i=0}^k \mathcal{M}_i$ is unidirectional and $\mathcal{M}_i \cap \mathcal{M}_j = \emptyset$, any pair of hyperplanes $\hat{h}_i \in \mathcal{M}_i$ and $\hat{h}_j \in \mathcal{M}_j$ must cross. Indeed, $\hat{h}_i \neq \hat{h}_j$, and \hat{h}_i is the base of a chain in \mathcal{M}_i , and \hat{h}_j is the base of a chain in \mathcal{M}_j . If \hat{h}_i, \hat{h}_j did not cross, then the union of this pair of chains would violate unidirectionality of $\mathcal{M}_i \cap \mathcal{M}_j$.

Thus, by Lemma 4.6, the convex hull of $\bigcup_{i=0}^k \xi_i$ is isometric to $\prod_{i=0}^k Y_i$, where Y_i is the cubical convex hull of ξ_i . For each $t > 0$, let x_t be the image of $(\xi_0(t), \dots, \xi_k(t))$ under the isometric embedding $\prod_{i=0}^k Y_i \rightarrow X$. Then segments joining the basepoint to the points x_t converge uniformly on compact sets to a geodesic ray representing a point $b \in \partial_T X$ with the property that $\psi(b) = u$ is represented by $\mathcal{U} = \bigsqcup_{i=0}^k \mathcal{M}_i$. Hence u is visible by Definition 8.19. Moreover, $m_i \leq u \leq w$ for all i . To complete the proof of the claim, we will show that u is the promised cone point of $\bigcap_{i=0}^k St(m_i) \cap L_n(w)$.

For each i , the fact that $m_i \leq u$ implies $u \in St(m_i)$. On the other hand, $u \neq w$ since u is visible but w is not. Thus $u < w$, so $u \in L_n(w)$. Thus far, we have shown that $\bigcap_{i=0}^k St(m_i) \cap L_n(w) \neq \emptyset$. Now, suppose $v \in L_n(w) \cap St(m_i)$ for all i . To complete the claim, we must show that u and v are \leq -comparable. By definition, $m_i \leq v$ for all i , hence $\mathcal{M}_i \leq \mathcal{V}$ for all i , where \mathcal{V} is a UBS representing v . Therefore, $\mathcal{U} = \bigsqcup_{i=0}^k \mathcal{M}_i \leq \mathcal{V}$, which implies $u \leq v$. Thus $\bigcap_{i=0}^k St(m_i) \cap L_n(w)$ is a cone with cone point u , completing Claim 10.9.

Claim 10.10 $L_n(w)$ is contractible.

By Claim 10.8, $L_n(w)$ is a finite union of subcomplexes $L_n(w) \cap St(m)$, where m varies over the minimal Roller classes satisfying $m \leq w$. By Claim 10.9, any collection of the $L_n(w) \cap St(m)$ intersect in a contractible subcomplex. Hence, by the Simplicial Nerve Theorem 2.8, $L_n(w)$ is homotopy equivalent to the nerve of this covering. (We do not need to check equivariance at this step.) But, since any collection of the $L_n(w) \cap St(m)$ have nonempty intersection, this nerve is a finite simplex and thus contractible. Hence $L_n(w)$ is contractible.

Conclusion: We have shown that the following hold for $1 \leq n \leq D$:

- Claim 10.7 says that for each vertex $v \in (\mathfrak{S}X)_n \setminus (\mathfrak{S}X)_{n-1}$, the link in $(\mathfrak{S}X)_n$ of v is contained in $(\mathfrak{S}X)_{n-1}$;
- Claim 10.10 says that for each vertex $v \in (\mathfrak{S}X)_n \setminus (\mathfrak{S}X)_{n-1}$, the link in $(\mathfrak{S}X)_n$ of v is contractible.

The first fact says that the open stars $st_n(v)$ of the vertices $v \in (\mathfrak{S}X)_n \setminus (\mathfrak{S}X)_{n-1}$ are pairwise disjoint. Together with the second fact, this implies that $(\mathfrak{S}X)_n$ is homotopy equivalent to

$$(\mathfrak{S}X)_{n-1} = (\mathfrak{S}X)_n \setminus \left(\bigcup_{v \in (\mathfrak{S}X)_n^{(0)} \setminus (\mathfrak{S}X)_{n-1}^{(0)}} st_n(v) \right).$$

It follows that we can independently deformation retract the various open stars $st_n(v)$ to the corresponding links $L_n(v)$ to get a deformation retraction $(\mathfrak{S}X)_n \rightarrow (\mathfrak{S}X)_{n-1}$ that is a homotopy inverse for the inclusion. Composing these retractions (for $1 \leq n \leq D$) gives the desired deformation retraction $\mathfrak{R}_\Delta X \rightarrow \mathfrak{R}_\Delta^< X$.

Finally, recall from the discussion after Definition 1.1 that every homotopy inverse of a $\text{Aut}(X)$ -equivariant inclusion $\mathfrak{R}_\Delta^< X \rightarrow \mathfrak{R}_\Delta X$ is itself an $\text{Aut}(X)$ -homotopy equivalence. Thus we have an $\text{Aut}(X)$ -deformation retraction $\mathfrak{R}_\Delta X \rightarrow \mathfrak{R}_\Delta^< X$. □

Combining Propositions 10.5 and 10.6 gives a proof of half of Theorem A.

Corollary 10.11 *There is an $\text{Aut}(X)$ -homotopy equivalence $\text{T}^{\mathbb{R}}: \mathfrak{R}_\Delta X \rightarrow \partial_T X$.*

10.3 Homotopy equivalence between $\partial_\Delta X$ and $\mathfrak{R}_\Delta X$

The following proposition proves the homotopy equivalence of the simplicial boundary and the simplicial Roller boundary, completing the proof of Theorem A.

Proposition 10.12 *There is an $\text{Aut}(X)$ -homotopy equivalence $\text{R}^{\mathbb{S}}: \partial_\Delta X \rightarrow \mathfrak{R}_\Delta X$.*

Proof Let σ be a maximal simplex of $\partial_\Delta X$, corresponding to a class $[\mathcal{V}_\sigma] \in \text{UBS}(X)$. Then, by Theorem 6.27, $v_\sigma = \text{R}^{\text{U}}([\mathcal{V}_\sigma])$ is a maximal Roller class.

Let \mathcal{A} be the nerve of the covering of $\partial_\Delta X$ by maximal simplices. Now, since any collection of maximal simplices intersect in \emptyset or a simplex, and simplices are contractible, the Equivariant Simplicial Nerve Theorem 2.8 provides an $\text{Aut}(X)$ -homotopy equivalence $\partial_\Delta X \rightarrow \mathcal{A}$.

Now, for each maximal Roller class v , consider the subcomplex Υ_v of $\mathfrak{R}_\Delta X$ consisting of all simplices corresponding to chains in which v is the maximal element. Note that the set of such Υ_v is a cover of $\mathfrak{R}_\Delta X$.

A simpler version of the proof of Lemma 10.2 implies that for all finite collections $\{v_1, \dots, v_n\}$ of maximal Roller classes, $\bigcap_{i=1}^n \Upsilon_{v_i}$ is empty or contractible. Indeed, suppose

this intersection is nonempty. Exactly as in the proof of Lemma 10.2, each v_i corresponds to a commensurability class $[\mathcal{V}_i]$ of UBSes whose intersection is a UBS \mathcal{V} , which in turn determines a \leq -maximal Roller class v such that $v \leq v_i$ for all i . As in Lemma 10.2, this implies that $\bigcap_{i=1}^n \Upsilon_{v_i}$ is a cone with cone-point v . (In the context of Lemma 10.2, there is an additional step to check that v is visible, but that is unnecessary here since we are working in all of $\mathfrak{R}_\Delta X$ rather than in $\mathfrak{R}_\Delta^< X$.)

Let \mathcal{B} be the nerve of the covering of $\mathfrak{R}_\Delta X$ by the subcomplexes Υ_v , as v varies over the maximal Roller classes. Observe that the assignment $v \mapsto \Upsilon_v$ is a bijection from the set of maximal Roller classes to $\mathcal{B}^{(0)}$. By Theorem 2.8, as above, there is an $\text{Aut}(X) \sim$ homotopy equivalence $\mathfrak{R}_\Delta X \rightarrow \mathcal{B}$.

To conclude, we will show that \mathcal{A} and \mathcal{B} are equivariantly isomorphic. The vertex set of \mathcal{A} is the set of maximal simplices σ of $\partial_\Delta X$, so we may define a function $f: \mathcal{A}^{(0)} \rightarrow \mathcal{B}^{(0)}$ via the composition

$$\sigma \mapsto [\mathcal{V}_\sigma] \mapsto v_\sigma \mapsto \Upsilon_{v_\sigma},$$

where the middle arrow is the equivariant map R^U of Theorem 6.27. Since both σ and v_σ are maximal by definition, the map f is a bijection by Theorem 6.27.(3). Every arrow is $\text{Aut}(X)$ -equivariant by construction.

To extend f to a simplicial isomorphism $f: \mathcal{A} \rightarrow \mathcal{B}$, it suffices to check that for all maximal simplices $\sigma_0, \dots, \sigma_n$ of $\partial_\Delta X$, we have $\bigcap_{i=0}^n \sigma_i \neq \emptyset$ if and only if $\bigcap_{i=0}^n \Upsilon_{v_{\sigma_i}} \neq \emptyset$; that this is sufficient follows since \mathcal{A} is the nerve of the covering of $\partial_\Delta X$ by maximal simplices and \mathcal{B} is the nerve of the covering of $\mathfrak{R}_\Delta X$ by the subcomplexes Υ_v .

Suppose that $\bigcap_{i=0}^n \sigma_i \neq \emptyset$. Let \mathcal{V}_i be a UBS representing σ_i . So, \mathcal{V}_i also represents the corresponding Roller class $f(\sigma_i)$. Let $\mathcal{V} = \bigcap_{i=0}^n \mathcal{V}_i$, and note that \mathcal{V} is infinite. Then by Remark 9.8, \mathcal{V} is an ℓ^1 -visible UBS. Hence \mathcal{V} represents a Roller class w such that $w \leq v_{\sigma_i}$ for all i . Thus $\bigcap_{i=0}^n \Upsilon_{v_{\sigma_i}}$ contains w . The converse is similar. Hence f extends to an isomorphism, and it follows that

$$R^S: \partial_\Delta X \xrightarrow{\sim} \mathcal{A} \xrightarrow{f} \mathcal{B} \xrightarrow{\sim} \mathfrak{R}_\Delta X$$

is an $\text{Aut}(X) \sim$ homotopy equivalence. □

10.4 Cuboid generalization

We can now conclude the proof of Theorem B.

Proof of Theorem B Let G be a group acting on X by cubical automorphisms, and let ρ be a G -admissible rescaling of X , with rescaled metric d_X^ρ . Recall from Definition 4.4 that the action of G factors through $\text{Aut}^\rho(X)$. In Sect. 8.4, we have checked that all of the constructions and results about the Tits boundary $\partial_T X$ also apply to $\partial_T^\rho X$, the Tits boundary of the rescaled metric (X, d_X^ρ) , in a G -equivariant way.

In Sect. 9.4, we have checked that the constructions of open and closed nerves for $\partial_T X$ also work for $\partial_T^\rho X$. In particular, by Theorem 9.18, there is a $G \sim$ homotopy equivalence between $\partial_T^\rho X$ and the nerve \mathcal{N}_T^ρ of the closed cover $\{Q^\rho(v) : v \in \text{MaxVis}^\rho(X)\}$.

Now, we inspect the results of this section. In Lemma 10.2, one simply needs to replace $\text{MaxVis}(X)$ by $\text{MaxVis}^\rho(X)$ and $\text{Aut}(X)$ by $\text{Aut}^\rho(X)$; the same exact proof then applies. In Definition 10.3, replacing $\text{MaxVis}(X)$ by $\text{MaxVis}^\rho(X)$ yields a nerve \mathcal{N}_Δ^ρ . The proof of Lemma 10.4 extends verbatim to give an $\text{Aut}^\rho(X)$ -equivariant simplicial isomorphism $\mathcal{N}_\Delta^\rho \rightarrow \mathcal{N}_T^\rho$. Proposition 10.5 extends verbatim, because its proof is a top-level assembly of previous

results. Similarly, Proposition 10.6 extends immediately to cuboids, because its proof is a topological argument about simplicial complexes. (The proof of Proposition 10.6 does use several lemmas from Sects. 7 and 8, particularly in Claim 10.9, but all of those lemmas have been extended to cuboids. Compare Remark 7.8.) Combining the cuboid versions of Propositions 10.5 and 10.6 gives a G -homotopy equivalence $\mathfrak{R}_\Delta X \rightarrow \partial_T^\rho X$.

Finally, note that the $\text{Aut}(X)$ -homotopy equivalence $\partial_\Delta X \rightarrow \mathfrak{R}_\Delta X$ established in Proposition 10.12 is also a G -homotopy equivalence, because G acts by cubical automorphisms. Thus we have both of the G -homotopy equivalences claimed in the theorem.

Acknowledgements We are grateful to Craig Guilbault, Jingyin Huang, Dan Ramras, and Kim Ruane for some helpful discussions. In particular, we thank Ramras for a correction, for pointing out the reference [Bjo81], and for the observation about Borel G -homotopy equivalence (Remark 1.2). We thank the organizers of the conference “Nonpositively curved groups on the Mediterranean” in May of 2018, where the three of us began collaborating as a unit. Fernós was partially supported by NSF grant DMS-005640. Futer was partially supported by NSF grant DMS-907708. Hagen was partially supported by EPSRC New Investigator Award EP/R042187/1.

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. Agol, I.: The virtual Haken conjecture. *Doc. Math.* **18**, 1045–1087 (2013). (With an appendix by Agol, Daniel Groves, and Jason Manning)
2. Bandelt, H.-J., Chepoi, V.: Surveys on Discrete and Computational Geometry. In: Goodman, J.E., Pach, J., Pollack, R. (eds.) *Metric Graph Theory and Geometry: A Survey*. Contemporary Mathematics, vol. 453, pp. 49–86. American Mathematical Society, Providence, RI (2008)
3. Brodzki, J., Campbell, S.J., Guentner, E., Niblo, G.A., Wright, N.J.: Property A and CAT(0) cube complexes. *J. Funct. Anal.* **256**(5), 1408–1431 (2009)
4. Bestvina, M.: Local homology properties of boundaries of groups. *Mich. Math. J.* **43**(1), 123–139 (1996)
5. Beyrer, J., Fioravanti, E.: Cross ratios and cubulations of hyperbolic groups. *Math. Ann.* **384**(3–4), 1547–1592 (2022)
6. 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)
7. 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)
8. Behrstock, J., Hagen, M.F.: Cubulated groups: thickness, relative hyperbolicity, and simplicial boundaries. *Groups Geom. Dyn.* **10**(2), 649–707 (2016)
9. Behrstock, J., Hagen, M.F., Sisto, A.: Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups. *Geom. Topol.* **21**(3), 1731–1804 (2017)
10. Björner, A.: Homotopy type of posets and lattice complementation. *J. Combin. Theory Ser. A* **30**(1), 90–100 (1981)
11. Björner, A.: *Topological Methods*. Handbook of Combinatorics, vol. 1, 2, pp. 1819–1872. Elsevier, Amsterdam (1995)
12. Balsler, A., Lytchak, A.: Centers of convex subsets of buildings. *Ann. Global Anal. Geom.* **28**(2), 201–209 (2005)
13. Borsuk, K.: On the imbedding of systems of compacta in simplicial complexes. *Fund. Math.* **35**, 217–234 (1948)

14. Bourgin, D.G.: The paracompactness of the weak simplicial complex. *Proc. Nat. Acad. Sci. USA* **38**, 305–313 (1952)
15. Bowditch, B.H.: Some properties of median metric spaces. *Groups Geom. Dyn.* **10**(1), 279–317 (2016)
16. Bowditch, B.H.: Quasiflats in coarse median spaces. <https://homepages.warwick.ac.uk/~masgak/papers/quasiflats.pdf> (2019)
17. Bowers, P.L., Ruane, K.: Boundaries of nonpositively curved groups of the form $G \times Z^n$. *Glasgow Math. J.* **38**(2), 177–189 (1996)
18. Buyalo, S.V.: Geodesics in Hadamard spaces. *Algebra i Analiz* **10**(2), 93–123 (1998)
19. 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)
20. Chepoi, V.: Graphs of some CAT(0) complexes. *Adv. Appl. Math.* **24**(2), 125–179 (2000)
21. Croke, C.B., Kleiner, B.: Spaces with nonpositive curvature and their ideal boundaries. *Topology* **39**(3), 549–556 (2000)
22. Croke, C.B., Kleiner, B.: The geodesic flow of a nonpositively curved graph manifold. *Geom. Funct. Anal.* **12**(3), 479–545 (2002)
23. Caprace, P.-E., Lytchak, A.: At infinity of finite-dimensional CAT(0) spaces. *Math. Ann.* **346**(1), 1–21 (2010)
24. Chatterji, I., Niblo, G.: From wall spaces to CAT(0) cube complexes. *Int. J. Algebra Comput.* **15**(5–6), 875–885 (2005)
25. Caprace, P.-E., Sageev, M.: Rank rigidity for CAT(0) cube complexes. *Geom. Funct. Anal.* **21**(4), 851–891 (2011)
26. Druţu, C., Kapovich, M.: *Geometric Group Theory*, vol. 63. American Mathematical Society, Providence (2018). (With an appendix by Bogdan Nica)
27. Dowker, C.H.: Homology groups of relations. *Ann. Math.* **2**(56), 84–95 (1952)
28. Fernós, T.: The Furstenberg-Poisson boundary and CAT(0) cube complexes. *Ergod. Theory Dyn. Syst.* **38**(6), 2180–2223 (2018)
29. Fioravanti, E.: The Tits alternative for finite rank median spaces. *Enseign. Math.* **64**(1–2), 89–126 (2018)
30. Fioravanti, E.: Roller boundaries for median spaces and algebras. *Algebr. Geom. Topol.* **20**(3), 1325–1370 (2020)
31. Fioravanti, E., Hagen, M.: Deforming cubulations of hyperbolic groups. *J. Topol.* **14**(3), 877–912 (2021)
32. Fernós, T., Lécureux, J., Mathéus, F.: Random walks and boundaries of CAT(0) cubical complexes. *Comment. Math. Helv.* **93**(2), 291–333 (2018)
33. Fernós, T., Lécureux, J., Mathéus, F.: Contact graphs, boundaries, and a central limit theorem for CAT(0) cubical complexes. [arXiv:2112.10141](https://arxiv.org/abs/2112.10141) (2021)
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. Guillbault, C., Mooney, C.: Cell-like equivalences and boundaries of CAT(0) groups. *Geom. Dedicata* **160**, 119–145 (2012)
36. Gromov, M.: *Hyperbolic Groups*. In: Gersten, S.M. (ed.) *Essays in Group Theory*. Mathematical Sciences Research Institute Publications, vol. 8, pp. 75–263. Springer, New York (1987)
37. Guralnik, D.: Coarse decompositions of boundaries for CAT(0) groups. [arXiv:math/0611006](https://arxiv.org/abs/math/0611006) (2006)
38. Haglund, F.: Isometries of CAT(0) cube complexes are semi-simple. *Ann. Math. Qué.* **47**(2), 249–261 (2023)
39. Hagen, M.F.: The simplicial boundary of a CAT(0) cube complex. *Algebr. Geom. Topol.* **13**(3), 1299–1367 (2013)
40. Hagen, M.F.: Corrigendum to the article “The simplicial boundary of a CAT(0) cube complex”. *Algebr. Geom. Topol.* **18**(2), 1251–1256 (2018)
41. Hagen, M.F.: Large facing tuples and a strengthened sector lemma. *Tunis. J. Math.* **4**(1), 55–86 (2022)
42. Hatcher, A.: *Algebraic Topology*. Cambridge University Press, Cambridge (2002)
43. Huang, J., Kleiner, B., Stadler, S.: Morse quasiflats II. *Adv. Math.* **425**(41), Paper No. 109075 (2023)
44. Hosaka, T.: On equivariant homeomorphisms of boundaries of CAT(0) groups and Coxeter groups. *Differ. Geom. Appl.* **43**, 68–94 (2015)
45. Hagen, M.F., Susse, T.: On hierarchical hyperbolicity of cubical groups. *Israel J. Math.* **236**(1), 45–89 (2020)
46. Huang, J.: Top-dimensional quasiflats in CAT(0) cube complexes. *Geom. Topol.* **21**(4), 2281–2352 (2017)
47. Kleiner, B., Lang, U.: Higher rank hyperbolicity. *Invent. Math.* **221**(2), 597–664 (2020)
48. Kar, A., Sageev, M.: Ping pong on CAT(0) cube complexes. *Comment. Math. Helv.* **91**(3), 543–561 (2016)
49. Miesch, B.: Injective metrics on cube complexes. [arXiv:1411.7234](https://arxiv.org/abs/1411.7234) (2014)

50. Mooney, C.P.: Examples of non-rigid CAT(0) groups from the category of knot groups. *Algebr. Geom. Topol.* **8**(3), 1666–1689 (2008)
51. Munkres, J.R.: *Elements of Algebraic Topology*. Addison-Wesley Publishing Company, Menlo Park (1984)
52. Nica, B.: Cubulating spaces with walls. *Algebr. Geom. Topol.* **4**, 297–309 (2004)
53. Niblo, G.A., Reeves, L.D.: The geometry of cube complexes and the complexity of their fundamental groups. *Topology* **37**(3), 621–633 (1998)
54. Nevo, A., Sageev, M.: The Poisson boundary of CAT(0) cube complex groups. *Groups Geom. Dyn.* **7**(3), 653–695 (2013)
55. Ontaneda, P.: Cocompact CAT(0) spaces are almost geodesically complete. *Topology* **44**(1), 47–62 (2005)
56. Qing, Y.: Boundary of CAT(0) groups with right angles. Thesis (Ph.D.)—Tufts University, ProQuest LLC, Ann Arbor, MI (2013)
57. Ramras, D.A.: Variations on the nerve theorem. [arxiv:2305.04794](https://arxiv.org/abs/2305.04794) (2023)
58. Roller, M.: Poc sets, median algebras and group actions. [arXiv:1607.07747](https://arxiv.org/abs/1607.07747) (2016)
59. Ruane, K.E.: Boundaries of CAT(0) groups of the form $\Gamma = G \times H$. *Topol. Appl.* **92**(2), 131–151 (1999)
60. Sageev, M.: Ends of group pairs and non-positively curved cube complexes. *Proc. Lond. Math. Soc.* (3) **71**(3), 585–617 (1995)
61. Wilson, J.M.: A CAT(0) group with uncountably many distinct boundaries. *J. Group Theory* **8**(2), 229–238 (2005)
62. Wise, D.T.: *From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry*. CBMS Regional Conference Series in Mathematics, vol. 117, Published for the Conference Board of the Mathematical Sciences, Washington, DC (2012)
63. Wise, D.T.: *The Structure of Groups with a Quasiconvex Hierarchy*. Annals of Mathematics Studies, vol. 209. Princeton University Press, Princeton (2021)
64. Xie, X.: The Tits boundary of a CAT(0) 2-complex. *Trans. Am. Math. Soc.* **357**(4), 1627–1661 (2005)

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