



COMPOSITIO MATHEMATICA

Finiteness properties of cubulated groups

G. C. Hruska and Daniel T. Wise

Compositio Math. **150** (2014), 453–506

[doi:10.1112/S0010437X13007112](https://doi.org/10.1112/S0010437X13007112)



FOUNDATION
COMPOSITIO
MATHEMATICA



LONDON
MATHEMATICAL
SOCIETY



Finiteness properties of cubulated groups

G. C. Hruska and Daniel T. Wise

ABSTRACT

We give a generalized and self-contained account of Haglund–Paulin’s wallspaces and Sageev’s construction of the CAT(0) cube complex dual to a wallspace. We examine criteria on a wallspace leading to finiteness properties of its dual cube complex. Our discussion is aimed at readers wishing to apply these methods to produce actions of groups on cube complexes and understand their nature. We develop the wallspace ideas in a level of generality that facilitates their application. Our main result describes the structure of dual cube complexes arising from relatively hyperbolic groups. Let H_1, \dots, H_s be relatively quasiconvex codimension-1 subgroups of a group G that is hyperbolic relative to P_1, \dots, P_r . We prove that G acts relatively cocompactly on the associated dual CAT(0) cube complex C . This generalizes Sageev’s result that C is cocompact when G is hyperbolic. When P_1, \dots, P_r are abelian, we show that the dual CAT(0) cube complex C has a G -cocompact CAT(0) truncation.

Contents

1	Introduction	454
1.1	The goal of this paper	455
1.2	Why cubulate?	457
1.3	Survey of known results towards cubulation of groups	458
2	Wallspaces	459
2.1	Wallspaces	459
2.2	Geometric wallspaces	460
2.3	Tracks and transverse walls	461
2.4	Codimension-1 subgroups	463
2.5	Wallspaces from codimension-1 subgroups	464
2.6	A wallspace structure for a group	465
3	The dual CAT(0) cube complex	466
3.1	CAT(0) cube complexes	466
3.2	The cube complex dual to a wallspace	466
3.3	Cubes	470
3.4	Hemiwallspaces and convex subcomplexes	473
3.5	The dimension of the dual cube complex and uniform local finiteness	475

Received 12 March 2011, accepted 15 October 2012, received in final form 14 October 2013, published online 10 March 2014.
2010 Mathematics Subject Classification 20F65, 20F67 (primary).
Keywords: cube complex, wallspace, relatively hyperbolic group.

The first author’s research supported by NSF grant DMS-0808809. The second author’s research supported by NSERC.

This journal is © Foundation Compositio Mathematica 2014.

4	CAT(0) spaces with convex walls	477
4.1	Main results illustrated in a highly geometric setting	477
4.2	Examples of CAT(0) 2-complexes with convex walls	478
5	Properness and local finiteness	479
5.1	Properness of the action	479
5.2	Local finiteness	482
6	Review of relative hyperbolicity and quasiconvexity	483
7	Relative cocompactness	486
7.1	The hyperbolic case	486
7.2	Various hemiwallsaces induced by a peripheral subspace	486
7.3	Structure of cubulation with peripheral subspaces	488
7.4	Application towards relatively hyperbolic groups	488
7.5	The structure of the dual when G acts properly and cocompactly on X	492
8	Local finiteness and properness in the relatively hyperbolic case	494
8.1	Additional finiteness properties in the presence of a proper action on $C(X)$	494
8.2	Background on saturations and relatively thin triangles	496
8.3	Wall–wall separation and relative hyperbolicity	498
8.4	Structure and local finiteness of $C(X)$ in the relatively hyperbolic case with wall–wall separation	500
9	CAT(0) truncations	502
	Acknowledgements	503
	References	503

1. Introduction

One of the most important themes permeating combinatorial group theory during the past century has been splittings of groups as free products, amalgamated products, and HNN extensions. This topic finally matured when Serre advanced the viewpoint of a group acting on a tree in the 1970s: a group G acts essentially on a tree if and only if G splits.

Gromov introduced CAT(0) cube complexes in his seminal essay on hyperbolic groups in the 1980s. As trees are one-dimensional CAT(0) cube complexes, there was already a plethora of interesting group actions. In the 1990s, Sageev showed how to obtain an action on a CAT(0) cube complex from codimension-1 subgroups in a manner that naturally, but surprisingly, generalized the way Serre’s trees arise as the graphs dual to an embedded collection of surfaces in a 3-manifold.

While only free groups can act freely on a one-dimensional CAT(0) cube complex, the variety of groups acting freely in higher dimensions is staggering. When G acts freely on a CAT(0) cube complex C , the quotient $G \backslash C$ provides a natural $K(G, 1)$ for G , and many properties of G are revealed by examining combinatorial and geometric properties of $G \backslash C$ or the action on C . For instance, applications towards automaticity and the lack of Property (T) are given in [NR97, NR98, NiR98].

While a minimal action of a finitely generated group G on a tree T immediately provides a compact ‘graph of groups’, the situation is substantially more complex in higher dimensions, even when the action is free. For instance, in general the action of G on the cube complex C provided

by Sageev's construction may fail to be cocompact; C may even fail to be finite dimensional. However, Sageev proved that G acts cocompactly on C when the codimension-1 subgroups are quasiconvex and G is hyperbolic, and generalizing this theorem is a primary motivation of this paper. There has been much recent research finding appropriate codimension-1 subgroups of a group G , and then using the dual cube complex to illuminate G . A notable success is the theorem on the 'structure of groups with a quasiconvex hierarchy' ([Wis09] and in a forthcoming paper by Wise). Our main result plays a critical role there in understanding the dual cube complex arising for a cusped hyperbolic 3-manifold.

1.1 The goal of this paper

The main goal of this paper is to examine the finiteness properties of the dual cube complex obtained by applying Sageev's construction to a group together with a collection of codimension-1 subgroups. We have especially concentrated on the case where the group is relatively hyperbolic, as there are powerful results to be gleaned in this case that require some care.

The following is a simplified version of our main result, which is Theorem 7.10. We emphasize that this result plays an important role in the relatively hyperbolic case of the work in a forthcoming paper by Wise.

THEOREM 1.1 (Relative cocompactness). *Let G be hyperbolic relative to P_1, \dots, P_j , and let H_1, \dots, H_k be relatively quasiconvex subgroups. Choose an H_i -wall for each i . Let C be the cube complex dual to the wallspace consisting of these H_i -walls and their G -translates. The induced action of G on C is cocompact relative to P_i -invariant subcomplexes C_i (see Definition 1.2).*

An H_i -wall is an H_i -invariant way to 'cut G in half' along H_i . This is of particular interest when H_i is a codimension-1 subgroup of G (see §2.4). Sometimes there is coarsely a unique 'natural' partition, for instance when G and the H_i are fundamental groups of closed aspherical manifolds [KK05]. Since a subgroup H_i might 'cut G in half' in more than one way, we are unable to associate a specific dual cube complex with a collection of subgroups. Instead we must specify the additional data of the chosen partitions of G , which are the ' H_i -walls' discussed in Definition 2.10.

DEFINITION 1.2 (Relative cocompactness). Let (G, \mathbb{P}) be a group together with a collection of subgroups $\mathbb{P} = \{P_1, \dots, P_r\}$. We say G acts *relatively cocompactly* on the cube complex C if there is a compact subcomplex K and P_i -invariant subcomplexes C_i such that:

- (1) $C = GK \cup \bigcup GC_i$;
- (2) $gC_i \cap g'C_j \subset GK$ unless $gC_i = g'C_j$ and $i = j$.

In many cases, the dual cube complex C of Theorem 1.1 can be 'truncated' by cutting off the infinite 'ends' of the subcomplexes C_i . This is analogous to the classical case, where one truncates the cusps of a finite volume hyperbolic manifold. We illustrate this in the following simplified version of Corollary 9.5, which can be applied to complete the proof that $B(6)$ groups are CAT(0) and similarly to complete an alternate proof of the Alibegović–Bestvina result that limit groups are CAT(0) [AB06].

THEOREM 1.3 (CAT(0) truncation). *Let G be hyperbolic relative to virtually abelian subgroups P_1, \dots, P_j . Let H_1, \dots, H_k be relatively quasiconvex subgroups with associated H_i -walls. Suppose G acts properly on the dual cube complex C . Then C contains a convex cocompact G -invariant subspace. In particular, G acts properly and cocompactly on a CAT(0) space.*

We emphasize that the vast majority of our effort and the exposition is focused around the proof of Theorem 1.1 as well as its relation to the Bounded Packing Property as developed in [HrW09]. However we have also written the paper with the intention of introducing the methods to new researchers as well as providing a self-contained reference for future work in this area. We have included together fundamental results, straightforward results, and some folk results in a unified language.

While the fundamental tool is Sageev's dual cube complex construction, the central unifying notion is a wallspace. These were first introduced by Haglund–Paulin, but we have developed their notion so that it is both more flexible and has greater reach in a manner that maintains its natural geometric motivation.

Wallspaces. Haglund–Paulin's original definition of a wallspace has an underlying set X whose walls are partitions $X = U \sqcup V$. We soften this definition by relaxing their requirement that $U \cap V = \emptyset$. We do not believe there is a theoretical difference in the level of applicability. However, practically this more general definition facilitates applications by avoiding awkward choices. For instance, in a geometric setting when X is not discrete, a wall often naturally arises from a separating subspace W , and it is convenient to define U, V to be the two closed halfspaces intersecting in W . One already saw the arbitrary inclusion of W on one side but not the other in Sageev's thesis, where W arises as a neighborhood of a codimension-1 subgroup.

When the wallspace (X, \mathcal{W}) is independently a metric space (X, d) , we say (X, \mathcal{W}) has the *Linear Separation Property* if there exist constants $\kappa > 0$ and ϵ such that for all points $x, y \in X$ we have $\#(x, y) \geq \kappa d(x, y) - \epsilon$. Here $\#(x, y)$ denotes the number of walls separating x and y . The following theorem, which is restated and proven as Theorem 5.2, has been observed repeatedly in various forms by researchers applying Sageev's construction. It is easy, but very useful.

THEOREM 1.4 (Linear separation \implies proper). *Suppose G acts on a wallspace (X, \mathcal{W}) , and the action on the underlying metric space (X, d) is metrically proper. Then the Linear Separation Property for (X, \mathcal{W}) implies that G acts metrically properly on the dual cube complex $C = C(X, \mathcal{W})$.*

More generally G acts metrically properly on C provided that for some $x \in X$ we have: $\#(x, gx) \rightarrow \infty$ as $g \rightarrow \infty$.

Although Theorem 1.4 holds under the more general hypothesis stated second, we have deliberately chosen to emphasize the linear separation hypothesis, as this stronger property on (X, \mathcal{W}) is often more natural to verify.

In practice it is often more difficult to verify the properness of an action on the dual cube complex than the (relative) cocompactness. This is even true in a δ -hyperbolic setting, where quasiconvexity is a robust phenomenon, but the property of having sufficiently many separating walls is not locally discernible in general.

Metric properness translates directly into a very natural 'filling' condition on the wallspace, but verifying this condition is treacherous because a collection of walls that locally appears to be highly filling can be very ineffective globally. A suggestive (messy) example is given by a highly self-intersecting essential curve σ in a closed surface S so that the complementary components of $S - \sigma$ are all simply connected. When σ is homotopic to a simple curve, then the cube complex dual to the wallspace consisting of translates of $\tilde{\sigma}$ in the universal cover \tilde{S} is quasi-isometric to a tree. In this case, while the action is cocompact it cannot be proper.

We illustrate some of the main results in a simple and suggestive framework in §4. For many of these examples, linear separation is obvious because the simplicity of the walls makes the

necessary computation transparent. In settings that arise ‘in the wild’, these computations can be opaque. For instance, linear separation is proven for Gromov’s random groups at density $<1/6$ by Ollivier–Wise in [OW11], but at density $<1/5$ although a wallspace was constructed and the walls are quasiconvex it remains unknown whether these random groups can act properly, i.e., with finite stabilizers, on a cube complex.

We offer a second closely related criterion for verifying properness that is particularly suited to a hyperbolic setting, where each infinite order element has an axis.

The properness conclusion in the following is a special case of Theorem 5.5. The cocompactness follows from Sageev’s hyperbolic special case of Theorem 1.1.

THEOREM 1.5 (Axis separation). *Let G be a word-hyperbolic group. Let H_1, \dots, H_k be quasiconvex subgroups with associated H_i -walls. Suppose that for each infinite order element $g \in G$ its axis A_g is cut by some translate $g'H_j$. Then G acts properly and cocompactly on the dual cube complex.*

In Bergeron–Wise [BW12], the properness and cocompactness are treated simultaneously in a rather soft fashion that works effectively when there is a rich collection of codimension-1 subgroups to draw from. The idea in [BW12] is to use ∂G to choose finitely many quasiconvex codimension-1 subgroups so that each geodesic (and hence each axis) is separated by a wall.

When G acts properly and cocompactly on the dual cube complex $C = C(X, \mathcal{W})$ then we obtain local finiteness of C by a simple argument. We formulate a characterization of the local finiteness of C that does not require a group action. The following is a simplified statement of Theorem 5.7, which is related to the Parallel Walls Separation Condition of Brink–Howlett as applied by Niblo–Reeves to prove local finiteness of their cubulation of a finitely generated Coxeter group [BH93, NR03].

THEOREM 1.6 (Local finiteness). *Let (X, d) be a metric space with a locally finite collection \mathcal{W} of walls. Then the dual cube complex $C = C(X, \mathcal{W})$ is locally finite if and only if for each compact set $K \subseteq X$ there exists a constant $f(K)$ such that whenever $d(K, W) \geq f(K)$ there is a wall W' separating K from W .*

In § 8 we examine the finiteness properties of $C(X, \mathcal{W})$ in much greater detail in our principal setting where the relatively hyperbolic group G is acting on X .

1.2 Why cubulate?

We describe below a sample of the information one obtains about a group admitting cubulations with various levels of finiteness properties.

The earliest motivation for defining nonpositively curved cube complexes is simply that: a proper/cocompact action on a CAT(0) cube complex yields a proper/cocompact action on a CAT(0) metric space. Interestingly enough, this remains a fascinating albeit simple application, especially in view of the methodology promoted in this paper, which often surprisingly provides a CAT(0) metric from codimension-1 subgroups. As mentioned earlier, the most desirable goal of a cocompact CAT(0) space can sometimes be obtained by truncating.

If G acts essentially on a CAT(0) cube complex then G cannot have Property (T) by Niblo–Roller [NiR98]. Moreover, it is implicit in their argument that if G acts metrically properly then G is a-T-menable (see the discussion in [CCJJV01, § 7.4.1]). If G acts properly and cocompactly on a CAT(0) cube complex then G is biautomatic by Niblo–Reeves [NR98]. If G has bounded torsion and acts properly on a finite dimensional CAT(0) cube complex then G satisfies a strong form of the Tits Alternative by Sageev–Wise [SW05].

Right-angled Artin groups arise as the fundamental groups of well-known nonpositively curved cube complexes. Exploiting their combinatorial geometry, a connection has been developed by Haglund–Wise in [HW08], which has the potential to virtually embed a nicely cubulated group G into a right-angled Artin group, and thus expose many algebraic features of G .

If G acts properly on a CAT(0) cube complex C , then the asymptotic dimension of G is bounded above by the dimension of C by [Wri12]. One obtains a rich family of quasimorphisms of G from irreducible actions on a CAT(0) cube complex by recent work of Caprace–Sageev in [CS11]. When G acts properly on a CAT(0) cube complex, then Hagen has shown that G is weakly hyperbolic relative to the hyperplane stabilizers [Hag13].

1.3 Survey of known results towards cubulation of groups

There has been an increasing amount of work towards obtaining actions of groups on CAT(0) cube complexes.

Although there has certainly been earlier work, the Bass–Serre theory of groups acting on trees is a milestone here [Ser77]. While splittings of groups as amalgamated products or HNN extensions had been studied for a long time previously, great perspective was added to this subject through the theory of groups acting on trees. The edge groups of these splittings had already been identified as codimension-1 subgroups by Scott [Sco77], but these ideas were only put into their natural context with Sageev’s thesis [Sag95]. Sageev gave his dual cube complex construction and showed that a group G acts essentially on a CAT(0) cube complex if and only if G contains a codimension-1 subgroup. The meaning of ‘essential’ was subsequently refined by Gerasimov and Niblo–Roller [Ger97, NiR98] to mean that G acts with no global fixed point. A beautiful exposition of these ideas in a crisp abstract setting was given by Roller in [Ro198].

Two early applications of Sageev’s construction were given by Niblo–Reeves [NR03] towards cubulation of Coxeter groups and by Wise [Wis04] towards cubulation of certain small-cancellation groups. The role played by wallspaces in Sageev’s construction was subsequently made explicit by Nica and Chatterji–Niblo in [CN05, Nic04].

Further applications are towards cubulating one-relator groups with torsion [LW11], cubulating Gromov’s random groups [OW11], cubulating graphs of free groups with cyclic edge groups [HsW10], cubulating arithmetic hyperbolic groups of simple type [BHW11], cubulating rhombus groups [JW13], cubulating certain malnormal amalgams of cubulated groups (in a forthcoming paper by Hsu and Wise), and most recently cubulating the group of Formanek–Procesi [Gau12].

The variety of cubulations that one can obtain by applying Sageev’s construction to a fixed group G is reflected in [Wis12], where it is shown that for each finitely generated infinite index subgroup H of a finitely generated free group G , there is a free G -action on a CAT(0) cube complex with one orbit of hyperplane, such that the stabilizer of each hyperplane equals H .

There has been some other work obtaining a cube complex more directly without an application of Sageev’s construction. In particular, we note the work of Weinbaum using Dehn’s presentation to cubulate prime alternating link complements (see e.g., [Wis06]), the work of Aitchison–Rubinstein [AR90] cubulating certain 3-manifolds, the work of Brady–McCammond [McC09] who recognized how to thicken a complex to obtain a cube complex and applied this to certain 3-manifolds and to $C'(1/4) - T(4)$ complexes, and finally with a very different flavor, the work of Farley [Far03] cubulating the diagram groups of Guba–Sapir [GS06]. At this point, it appears that the strategy of using Sageev’s construction affords the greatest variety of applications.

2. Wallspaces

2.1 Wallspaces

Let X be a nonempty set. A *wall* of X is a pair of subsets $\{U, V\}$ called (closed) *halfspaces* such that $X = U \cup V$. We shall not assume that $U \cap V = \emptyset$ (see Remark 2.1). The *open halfspaces* associated to the wall $\{U, V\}$ are the sets $U - (U \cap V)$ and $V - (U \cap V)$.

Points x, y of X are *separated* by a wall W if they lie in distinct open halfspaces of W . The notation $\#(x, y)$ denotes the number of walls separating x and y . We say that a point x and a wall $\{U, V\}$ *betwixt each other* if $x \in U \cap V$.

A *wallspace* is a set X together with a collection of walls \mathcal{W} on X with the following two finiteness properties:

- (1) $\#(x, y) < \infty$ for all $x, y \in X$; and
- (2) each point x betwixts only finitely many walls.

We do allow our ‘collection’ to contain duplicates (see Remark 2.3) but we do not allow duplicate walls that are genuine partitions (see Remark 2.2).

Distinct walls $W = \{U, V\}$ and $W' = \{U', V'\}$ are *transverse* if all four of the following intersections are nonempty:

$$U \cap U', \quad U \cap V', \quad V \cap U', \quad V \cap V'.$$

The central example of a wallspace arises from a CAT(0) cube complex C by letting $X = C^0$ and letting \mathcal{W} consist of the partitions of C^0 induced by hyperplanes. Note that the walls are genuine partitions here. A similar example is a wallspace on C whose walls are pairs of closed halfspaces associated to hyperplanes of C . We discuss this type of example further in § 2.2.

A group G *acts* on the wallspace (X, \mathcal{W}) if G acts on X and stabilizes \mathcal{W} . Consequently $\#(x, y) = \#(gx, gy)$ for each $g \in G$.

Remark 2.1. We note that in Haglund–Paulin’s original definition [HP98] of ‘*espaces à murs*’ the walls are partitions of X ; in other words $U \cap V = \emptyset$. We have chosen to work in a slightly more flexible setting.

Remark 2.2 (Duplicated walls). In Remark 2.3, we describe a more general definition, essentially using indexed sets, that allows ‘multiplicity’ in our collection of walls, in the sense that it is possible for $W_1 = \{U, V\}$ and $W_2 = \{U, V\}$. However,

- (†) we do not allow duplicate walls that are *genuine partitions*

in the sense that $U \cap V = \emptyset$ and U and V are both nonempty. This will be important in the proof of connectedness of the canonical component of the dual cube complex given in Proposition 3.8. Duplicate *vacuous walls* of the form $\{X, \emptyset\}$ do not create difficulties.

The definition of transversality given below implies that duplicated walls W_1, W_2 are transverse if and only if they are not partitions. It is conceivable that our treatment could be modified in the exceptional partition case so that all duplicated walls are transverse.

One natural way that duplicated walls arise is when (X, \mathcal{W}) is a wallspace and $Y \subset X$ and we give Y the induced *subwallspace* structure whose walls have the form $(Y \cap U, Y \cap V)$, see Definition 2.5. We refer the reader to the right-hand side of Figure 1, where the discrete subset Y has a duplicated wall that is not a partition. In contrast, on the left-hand side of the figure, we see that Y has a duplicated wall that is a genuine partition, which we do not allow. The reader can imagine thick walls arising from two infinite cyclic subgroups of a surface group and Y is a third subgroup.



FIGURE 1. On the left, the discrete subset Y has a duplicated wall that is a genuine partition, which we do not allow. On the right, the discrete subset Y has a duplicated wall that is not a partition.

The following more formal approach correctly permits duplicated walls.

Remark 2.3 (Wallspaces using indexed collections of halfspaces). A *wallspace* is a pair of sets (X, \mathcal{W}) where each element $W \in \mathcal{W}$ is a pair of indexed subsets $\{H_{-W}, H_{+W}\}$ of X called the *closed halfspaces* of W . Note that the subsets are indexed by the elements of $\{-, +\} \times \mathcal{W}$. Moreover we require that the following properties are satisfied:

- (0) $H_{-W} \cup H_{+W} = X$ for all $W \in \mathcal{W}$;
- (1) for all $x, y \in X$, there are finitely many $W \in \mathcal{W}$ such that $x \in (H_{-W}) - (H_{+W})$ and $y \in (H_{+W}) - (H_{-W})$;
- (2) for each $x \in X$, there are finitely many $W \in \mathcal{W}$ such that $x \in H_{-W} \cap H_{+W}$;
- (3) if $\{H_{-W}, H_{+W}\} = \{H_{-W'}, H_{+W'}\}$ and $H_{-W} \cap H_{+W} = \emptyset = H_{-W'} \cap H_{+W'}$ then $W = W'$.

Note that (3) excludes the duplicate walls that are genuine partitions, as in (\dagger) . However, for walls that are not partitions, it is now sensible for W, W' to be distinct walls with the same closed halfspaces. For though $H_{-W}, H_{-W'}$ and $H_{+W}, H_{+W'}$ have the same underlying sets, their indices differ.

2.2 Geometric wallspaces

A *geometric wallspace* is a metric space (X, d) with a collection \mathcal{W} of closed subspaces called *geometric walls* such that each geometric wall W is connected, and $X - W$ has exactly two components.

Letting U and V be the components of $X - W$, one obtains closed halfspaces $W \cup U$ and $W \cup V$, and hence a wall $\{(W \cup U), (W \cup V)\}$ associated to W .

Our earlier definition of separation is motivated by the observation that a geometric wall W separates x and y if they lie in distinct components of $X - W$.

Note that $\#(x, y) < \infty$ provided that x, y are connected by a path intersecting only finitely many walls of \mathcal{W} . Thus when X is path connected, the local finiteness of the collection of geometric walls is the critical ingredient ensuring that $\#(x, y) < \infty$ always holds.

To ensure that geometric wallspaces are associated to underlying (abstract) wallspaces, we will therefore insist that they are path connected and that the collection of geometric walls is locally finite in the sense that compact subsets of X intersect only finitely many geometric walls.

Many natural examples of geometric wallspaces are geodesic metric spaces whose walls are convex subspaces. For instance, this was the case for hyperplanes in a CAT(0) cube complex C .

LEMMA 2.4. *Suppose (X, \mathcal{W}) is a geometric wallspace. Two distinct abstract walls are transverse if and only if their associated geometric walls intersect.*

Proof. If two geometric walls W and W' intersect, then it is clear that their corresponding abstract walls are transverse, since each geometric wall lies in both of its associated halfspaces.

Now suppose W and W' are disjoint geometric walls. Let U and V be the components of $X-W$, and let U' and V' be the components of $X-W'$.

Since W is connected and disjoint from W' , either $W \subset U'$ or $W \subset V'$. Without loss of generality, assume $W \subset V'$. Thus $W \cap U' = \emptyset$. Similarly, we can assume that $W' \subset V$. Thus $W' \cap U = \emptyset$. Since $W \cup U$ is a connected set in the complement of W' that intersects V' , we must have $W \cup U \subseteq V'$. Therefore $(W \cup U) \cap (W' \cup U')$ is empty.

Thus the abstract walls associated to W and W' are not transverse. □

DEFINITION 2.5 (Subwallspace). Let (X, \mathcal{W}_X) be a wallspace and $Y \subset X$. The induced *subwallspace* structure on Y is the wallspace (Y, \mathcal{W}_Y) whose *induced walls* have the form $(Y \cap U, Y \cap V)$, for $(U, V) \in \mathcal{W}_X$. We will ignore the vacuous partitions (Y, \emptyset) in a subwallspace whenever convenient, as they play no role in the dual cube complex.

Note that we allow duplicate walls, but because of the difficulties discussed in Remark 2.2, we will only consider subwallspaces when there do not exist distinct walls (U, V) and (U', V') of X such that $(Y \cap U, Y \cap V)$ and $(Y \cap U', Y \cap V')$ are the same nonvacuous partition.

When X is a geometric wallspace and Y is a connected subspace, there do not exist (nonvacuous) induced walls that are genuine partitions. Indeed, if $Y \cap W = \emptyset$ then the connected set Y lies in one of the open halfspaces of W , so the induced wall is vacuous.

Example 2.6. Let $X = \mathbb{R}^2$ and let \mathcal{W}_X be the x -axis and y -axis. These two walls are transverse. Indeed, they intersect as guaranteed by Lemma 2.4. We now examine the transversality of walls in induced subwallspaces on subsets of X .

Let $Y = S^1 \subset \mathbb{R}^2$, and let (Y, \mathcal{W}_Y) be the induced subwallspace. Note that this is not a geometric wallspace. Observe that the two induced walls of Y are transverse but do not intersect. Indeed the same holds when $Y = \{\pm 1\}^2$.

Now however, let us consider $Z = \{(+1, -1), (+1, +1), (-1, +1)\}$, and let (Z, \mathcal{W}_Z) be the induced subwallspace. In this case, the induced walls of Z are no longer transverse.

2.3 Tracks and transverse walls

Consider a group G that acts on a wallspace (Y, \mathcal{W}) and also acts on a 2-complex X . A G -equivariant map $X^0 \rightarrow Y$ induces a wallspace structure on X^0 . Of course we could extend this to a (noncontinuous) G -equivariant map $X \rightarrow Y$ to make X into a wallspace. But this might not be geometrically satisfying. Instead, the following proposition explains how to enlarge X to X' so that X' and Y have consistent wall structures, and X' is a geometric wallspace whose walls are tracks in the sense of Dunwoody. It is proven by combining variants of Dunwoody resolutions and Stallings binding ties.

PROPOSITION 2.7 (Walls as tracks). *Let G act freely on a simply connected 2-complex X and suppose X^0 is a wallspace with no betwixting and the wall system is locally finite, with finitely many orbits of walls each of which has finitely generated stabilizer. Then there exists a finite G -equivariant expansion of X to a 2-complex X' with the same 0-skeleton and the same set of walls and each wall in X' is represented by a track.*

Proof. Let C be the dual cube complex of X^0 . Consider the G -equivariant map $X^0 \rightarrow C$ that sends a 0-cell to its canonical 0-cube. For each 1-cell of X , choose a combinatorial path in C , and in this way obtain a G -equivariant map $X^1 \rightarrow C$. For each G -representative of 2-cell R of X , choose a minimal area disc diagram $D \rightarrow C$ for the path $\partial R \rightarrow C$. We thus obtain a G -equivariant map $X \rightarrow C$. Since D has minimal area, all tracks within D end on ∂D , and hence all tracks intersect X^1 so no ‘new’ tracks have been added.




FIGURE 2. On the left we illustrate the map $R \rightarrow D$. On the right we illustrate the expanded 2-cell and associated disc diagram.

As disc diagrams might be singular, this map might not be combinatorial. However, D is a square diagram since C is cubical, and the 2-cell R can be subdivided so that the map $R \rightarrow D$ is well-behaved in the sense that the preimage of a ‘dual curve’ (a hyperplane) of D is an embedded arc in R . To see this, for each cut-vertex of D , we choose a tree in R projecting to it. For each cut-edge of D we choose a rectangle in R projecting to it; see the left of Figure 2. Finally squares of D which have a 0-cube on a cut-point can be replaced by pentagons. The result is that the preimage of dual curves of D are tracks in R .

In conclusion, the preimage of each hyperplane of C is a *pattern* in X , meaning that it is a disjoint union of tracks.

We now use that the stabilizers of walls are finitely generated to equivariantly connect each such pattern by equivariantly expanding along 2-cells. Indeed, let p and q be points in X^1 that lie in distinct tracks mapping to the same hyperplane H of C . Let σ be a combinatorial path whose first 1-cell contains p and whose last 1-cell contains q . Replacing σ with a subpath if necessary, we can assume without loss of generality that p, q lie in distinct tracks but that no point of σ between p, q maps to H . We add a new 1-cell e_σ whose endpoints are p, q and add a new 2-cell R_σ between σ and e_σ . We will now choose a disc diagram D_σ that R_σ will map to, similarly to the method used as above.

We now refer the reader to the right of Figure 2. Let L be a geodesic ladder in C joining the 1-cubes containing the images \bar{p} and \bar{q} . Let α and β be the two rails of L . The image $\bar{\sigma}$ of σ is a concatenation $\bar{\sigma}_p \bar{\sigma}' \bar{\sigma}_q$, where the last edge of $\bar{\sigma}_p$ contains \bar{p} and the first edge of $\bar{\sigma}_q$ contains \bar{q} . Let D' be a minimal area diagram between β and $\bar{\sigma}'$. Let D_σ be the disc diagram built from $\bar{\sigma}_p \cup \bar{\sigma}_q \cup L \cup D'$. We map e_σ to the concatenation $\bar{\sigma}_p \alpha \bar{\sigma}_q$. Finally observe that $\partial R_\sigma \rightarrow C$ equals the boundary path of D_σ . This allows us to map R_σ to D_σ .

The ladder L connects the track containing p to the track containing q . No new tracks are introduced. Indeed the minimal area of D' ensures that each dual curve of D_σ intersects the boundary. The geodesicity of L ensures that dual curves crossing L emerge on $\bar{\sigma}'$. And thus since e_σ maps to $\bar{\sigma}_p \alpha \bar{\sigma}_q$, no new tracks emerge on e_σ .

Using that the stabilizers are finitely generated, we must repeat this procedure only finitely many times to connect all of the components. \square

We now give a technical result asserting that transverse walls coarsely intersect. The result is stated using the language of H_i -walls, and we refer the reader to Definition 2.10 for the requisite terminology. Since we will only use this result one time, in the proof of Theorem 7.10, we suggest that this technical point be skipped on first reading. We have chosen to position this lemma together with Lemma 2.4 and Proposition 2.7 in order to emphasize that transverse walls do coarsely intersect in an appropriate sense.

LEMMA 2.8 (Transverse \implies close). *Let G be a group with finite generating set and Cayley graph Γ . Let H_i and H_j be finitely generated subgroups of G . Let $\{U_i, V_i\}$ be an H_i -wall and let $\{U_j, V_j\}$ be an H_j -wall. Then there exists D such that for all $g, g' \in G$, if the walls $g\{U_i, V_i\}$ and $g'\{U_j, V_j\}$ are transverse then $d_\Gamma(gH_i, g'H_j) < D$.*

Technically ‘transverse’ was defined only in the context of a wallspace, but this relation is obviously definable in the same way here. Indeed Γ can be regarded as a wallspace with two walls.

Proof. We now describe a construction for an arbitrary finitely generated subgroup $H \leq G$ and an H -wall $\{U, V\}$ in G . Let U^1 and V^1 be the closed 1-neighborhoods of U and V in Γ . Then H acts cocompactly on $U^1 \cap V^1$. Here we use that U and V have \ddot{H} -finite frontiers and $U \cap V$ is H -finite.

We now use that H is finitely generated. By adding finitely many H -orbits of vertices and edges to U^1 and V^1 we obtain U^2 and V^2 that are connected and where $W^2 = U^2 \cap V^2$ is nonempty. By adding finitely many H -orbits of vertices and edges to W^2 , we obtain an H -cocompact H -invariant connected graph W^3 containing H . Finally let $U^3 = U^2 \cup W^3$ and let $V^3 = V^2 \cup W^3$.

We note that U^3 and V^3 are connected and H -invariant, since each is an intersecting union of connected H -invariant subspaces. Moreover $U^3 \cap V^3$ is connected, since it equals W^3 by construction.

We apply the above construction to $\{U_i, V_i\}$ and $\{U_j, V_j\}$ to obtain geometric walls $\{\bar{U}_i, \bar{V}_i\}$ and $\{\bar{U}_j, \bar{V}_j\}$ in Γ . We emphasize that $\bar{W}_i = \bar{U}_i \cap \bar{V}_i$ is a (connected) H_i -cocompact geometric wall, and likewise for $\bar{W}_j = \bar{U}_j \cap \bar{V}_j$ and H_j .

Observe that transversality persists after thickening. Now apply Lemma 2.4 to the transverse geometric walls $g\{\bar{U}_i, \bar{V}_i\}$ and $g'\{\bar{U}_j, \bar{V}_j\}$. Therefore $g\bar{W}_i \cap g'\bar{W}_j$ is nonempty. The result follows by letting $D = \text{diam}(H_i \setminus \bar{W}_i) + \text{diam}(H_j \setminus \bar{W}_j)$. □

2.4 Codimension-1 subgroups

Let G be a finitely generated group with Cayley graph $\Gamma = \Gamma(G, S)$. A subgroup $H \subset G$ is *codimension-1* if the coset graph $\bar{\Gamma} = H \backslash \Gamma$ is a *multi-ended* graph in the following sense: $\bar{\Gamma} - \Lambda$ has two or more infinite components for some compact subgraph $\Lambda \subset \bar{\Gamma}$. Note that this definition is independent of the finite generating set S .

For example, it is well known that any copy of \mathbb{Z}^n in \mathbb{Z}^{n+1} is codimension-1. A frequently encountered example arises when G is isomorphic to a nontrivial amalgamated product $A *_C B$ or an HNN extension $A *_C$, in which case C is a codimension-1 subgroup by a result of Scott [Sco77]. Another suggestive example is any infinite cyclic subgroup of a closed surface group.

Every finitely generated infinite index subgroup of a free group is codimension-1. However the same cannot be said for a closed surface group. The reader is urged to consider a covering space with a compact core having exactly one boundary circle; the corresponding subgroup is not codimension-1, but it does have a property we now briefly turn to.

Observe that if H is codimension-1, then H is a *divisive* subgroup of G in the sense that for some $d > 0$ the complement $\Gamma - \mathcal{N}_d(H)$ has more than one component K that is *deep*, meaning that K does not lie in $\mathcal{N}_c(H)$ for any $c > 0$.

Returning to the case when G splits, we note that A, B , and C are all divisive subgroups in a nontrivial amalgam $A *_C B$ and likewise A and C are divisive in $A *_C$.

A divisive subgroup is codimension-1 precisely when there exists $d > 0$ such that $\Gamma - \mathcal{N}_d(H)$ has more than one H -orbit of deep components.

The notions of codimension-1 and divisive subgroups are not equivalent but are occasionally confused in the literature. Let n denote the number of deep components in $\Gamma - \mathcal{N}_d(H)$. In the special case when $1 < n < \infty$, the divisive subgroup H has a finite index subgroup H' that is a codimension-1 subgroup of G . Indeed H acts on the collection of deep components and we can let H' be the kernel of the resulting permutation homomorphism.

It is often the case that n increases with an increase in d . As mentioned in the introduction, n is stably equal to 2 for the situation arising from a codimension-1 submanifold. Recently Caprace–Przytycki explored ‘bipolar’ Coxeter groups, which have the property that n is stably equal to 2 for the divisive subgroups stabilizing the walls of the generating reflections [CP11].

There is a more sensitive notion called the *number of ends of the group pair* (G, H) , which takes a value between 0 and ∞ (see [Sco77]). A group is codimension-1 precisely when this end invariant takes the value $e(G, H) \geq 2$.

The property that H is divisive in G is equivalent to saying that the pair (G, H) has more than one ‘filtered end’, indicated in the literature by $\tilde{e}(G, H) \geq 2$ (see [Geo08, KR89]).

2.4.1. *Codimension-1 subgroups on the boundary of divisive subgroups.* We now show that if G has a divisive subgroup H , then G has one or more codimension-1 subgroups, which we refer to as *boundary subgroups* of H . A case of great interest here is when H is divisive but not codimension-1, in which case these boundary subgroups are all conjugate to each other in H .

For example, when $G = A *_C B$ the codimension-1 subgroup C is a boundary subgroup of the divisive subgroup A . When $G = A *_C *_D$ the divisive subgroup A has boundary subgroups C and D . When $G = \pi_1(S)$ for a surface S , and $\hat{S} \rightarrow S$ is an infinite degree cover with a core T having a unique boundary circle ∂T , the divisive subgroup $\pi_1(\hat{S})$ has boundary subgroup $\pi_1(\partial T)$.

THEOREM 2.9. *Let G be a finitely generated group. Every divisive subgroup of G contains a codimension-1 subgroup of G . Consequently, G contains a divisive subgroup if and only if G contains a codimension-1 subgroup.*

The *frontier* ∂S of a subset $S \subseteq G$ is the set of vertices of S that are adjacent to a vertex of $G - S$ in Γ .

Proof. Choose a neighborhood $\mathcal{N}_r(H)$ such that $\Gamma - \mathcal{N}_r(H)$ has at least two deep components. For each deep component Λ , let $K = \text{Stab}_H(\Lambda)$. Let E denote the union of all the edges that join Λ to $\mathcal{N}_r(H)$. Note that E is contained in the 1-neighborhood of the frontier of Λ , and $\Gamma - E$ has at least two K -orbits of deep components: namely Λ and $\Gamma - (\Lambda \cup E)$. The latter is K -deep since it is H -deep and $K \subset H$.

Observe that H acts cocompactly on the frontier $\partial(\Gamma - \mathcal{N}_r(H))$, and H preserves the partition of the frontier into $\sqcup \partial\Lambda_i$ where $\{\Lambda_i\}$ is the collection of all components of $\Gamma - \mathcal{N}_r(H)$. Consequently each $K_i = \text{Stab}_H(\Lambda_i)$ acts cocompactly on $\partial\Lambda_i$. □

2.5 Wallspaces from codimension-1 subgroups

Sageev gave a construction in [Sag95] which produces a G -action on a CAT(0) cube complex from a collection of codimension-1 subgroups H_i of the finitely generated G . His construction has two steps: one first produces a system of walls associated to the subgroups, and one then produces a cube complex dual to the system of walls.

A codimension-1 subgroup H can be used to produce a wall in several ways: for instance, Sageev chose the wall $\{HK, \Gamma - HK\}$ where K is a fixed deep complementary component of $\mathcal{N}_d(H)$. Alternately one could decompose $\Gamma - \mathcal{N}_d(H)$ as a disjoint union $K \sqcup K'$ of disjoint

H -invariant subspaces that are unions of components. We then obtain the wall:

$$\{\mathcal{N}_d(H) \cup K, \mathcal{N}_d(H) \cup K'\}. \quad (\ddagger)$$

While it is of greatest interest to insist that both K and K' are deep, the alternate decomposition in (\ddagger) leads to a functional wall, even without assuming that K or K' are deep (in particular, H need not be codimension-1). We will examine such decompositions in §2.6.

A collection H_1, \dots, H_k of subgroups of G together with choices K_i, K'_i for each H_i leads in this way to a collection of walls in Γ , which we extend to a G -equivariant collection. This transforms Γ , and hence G , into a wallspace.

2.6 A wallspace structure for a group

In this subsection, we provide a framework for endowing a group G with a wallspace structure associated to finitely many ways of cutting G along subgroups H_i . This provides important examples of wallspaces generalizing the construction in §2.5. We will use this material in §7.

DEFINITION 2.10 (H -wall). Let H be a subgroup of G . An H -wall $\{U, V\}$ is a pair of subsets satisfying the following conditions:

- (1) $G = U \cup V$;
- (2) $\{U, V\}$ is H -invariant, in the sense that $\{hU, hV\} = \{U, V\}$ for each $h \in H$; note that an element might interchange the subsets;
- (3) $U \cap V$ is H -finite;
- (4) U and V have \ddot{H} -finite frontiers, where \ddot{H} is the subgroup of H (of index at most two) stabilizing both U and V .

The *stabilizer* of a wall $\{U, V\}$ is the set $\{g \in G \mid \{gU, gV\} = \{U, V\}\}$. Let \overline{H} denote the stabilizer of the H -wall $\{U, V\}$. Note that $H \leq \overline{H}$ is finite index when $\{U, V\}$ is nonvacuous. (The exceptional vacuous H -wall $\{G, \emptyset\}$ has $\overline{H} = G$.) Indeed, \overline{H} stabilizes the union of the frontiers of U and V , and hence stabilizes a collection of finitely many left cosets of \ddot{H} . When $\{U, V\}$ is nonvacuous, this is a *nonempty* collection of left cosets. It thus follows that \overline{H} is commensurable with H .

DEFINITION 2.11 (Wallspace from H -walls). Let H_1, \dots, H_r be subgroups of a finitely generated group G . For each i , choose an H_i -wall $\{U_i, V_i\}$. (We do allow repeated subgroups in the sense that H_i might equal H_j for $i \neq j$. However, we do not allow repeated walls that are nonvacuous partitions. So if $U_i \cap V_i = \emptyset$ and $\{U_i, V_i\} = \{U_j, V_j\}$ then either $i = j$ or $U_i, V_i \neq \emptyset$.)

These choices determine a G -wallspace whose underlying set is G and whose walls are $\{gU_i, gV_i\}_i$. Walls $\{gU_i, gV_i\}_i$ and $\{g'U_j, g'V_j\}_j$ are considered to be equal if their underlying pairs of halfspaces are equal and $i = j$. The group G acts on the wallspace by $g\{U, V\}_i = \{gU, gV\}_i$. Note that the way in which walls are indexed by their associated subgroups remains invariant by the G -action. We shall sometimes suppress these indices.

Let us now verify that this construction produces a wallspace. We first check that each $x \in G$ is betwixed by finitely many walls. If gx and $g'x$ are in the same H_i -orbit, then $gx = hg'x$ for some $h \in H_i$. Thus $H_i g = H_i g'$. Since $U_i \cap V_i$ contains finitely many H_i -orbits, there are finitely many right cosets $H_i g$ such that $gx \in U_i \cap V_i$. If x is betwixed by a wall $\{g^{-1}U_i, g^{-1}V_i\}_i$ then $gx \in U_i \cap V_i$. Therefore the elements g lie in only finitely many right cosets, i.e., H_i -orbits. Since $\{U_i, V_i\}$ is H_i -invariant, for each i there are only finitely many walls betwixting x .

We now verify that $\#(x, y) < \infty$ for all $x, y \in G$. Here we use that G is finitely generated. Let γ be a path in the Cayley graph from x to y . Since U_i has \ddot{H}_i -finite frontier, we see that



FIGURE 3. On the left is a cube complex. On the right the link of a vertex is shown.

finitely many $\ddot{H}_i g$ -translates of γ intersect the frontier of U_i . If x and y are separated by a wall $\{g^{-1}U_i, g^{-1}V_i\}_i$, then $g\gamma$ intersects the frontier of U_i . As above there are finitely many such possibilities for each i . Consequently $\#(x, y) < \infty$ for all x, y .

There are two ways for distinct walls arising in this construction to correspond to an identical pair of halfspaces of G . Suppose $\{gU_i, gV_i\} = \{g'U_j, g'V_j\}$ and $i \neq j$. If $\{U_i, V_i\} \neq \{G, \emptyset\}$ then \overline{H}_i and \overline{H}_j are conjugate, and hence H_i and H_j are commensurable within G . If $\{U_i, V_i\} = \{G, \emptyset\}$ then H_i and H_j can be arbitrary subgroups of G .

Remark 2.12. H -walls are related to ‘ \ddot{H} -almost-invariant subsets’ (see e.g., [NiR98]). A nonvacuous H -wall $\{U, V\}$ is the same as a pair of \ddot{H} -almost-invariant sets U and V that cover G and whose intersection is \ddot{H} -finite. If U is an H -almost-invariant subset of G , then $\{U, G - U\}$ is an H -wall. Conversely, if $\{U, V\}$ is a nonvacuous H -wall then U and V are \ddot{H} -almost-invariant sets whose intersection is \ddot{H} -finite.

3. The dual CAT(0) cube complex

3.1 CAT(0) cube complexes

DEFINITION 3.1. An n -cube is a copy of $[-1, 1]^n$. Restricting i of its coordinates to ± 1 yields a *subcube* which can be identified with an $(n - i)$ -cube in various ways. A *cube complex* is obtained by gluing cubes along subcubes. (The actual gluing maps are modeled on isometries but are determined completely by the combinatorial data.)

A *flag complex* is a simplicial complex with the property that $n + 1$ vertices span an n -simplex if and only if they are pairwise adjacent.

The *link* of a 0-cube $c \in C^0$ is the complex associated to the ‘ ϵ -sphere’ about c in C , as illustrated in Figure 3. More precisely, $\text{link}(c)$ has an n -simplex for each corner of an $(n + 1)$ -cube at c , where such corners are glued together precisely according to the way their associated cubes are glued together.

A cube complex C is *nonpositively curved* if $\text{link}(c)$ is a flag complex for each $c \in C^0$.

Finally, C is CAT(0) if C is simply connected and nonpositively curved. Using a slightly different definition, Gromov introduced nonpositively curved cube complexes as a source of examples of CAT(0) metric spaces. It is a fact that when C is simply connected and satisfies the local nonpositive curvature condition, then C has a CAT(0) metric where each n -cube is isometric to the standard Euclidean n -cube.

This fact was verified by Moussong and Gromov in the finite dimensional case [Gro87, Mou88] and by Leary in the general case [Lea10]. We will not call upon this fact until § 9.

3.2 The cube complex dual to a wallspace

Following Sageev, cube complexes dual to wallspaces were constructed by Niblo–Reeves in [NR03] where the walls are the ‘reflection walls’ naturally considered in the study of Coxeter groups,




FIGURE 4. A wallspace (X, \mathcal{W}) and its dual cube complex $C(X, \mathcal{W})$ are indicated. The set $X = \{a, b, c, d, e, f\}$ and \mathcal{W} consists of the five walls indicated above. Note that a betwixts the second wall, and e betwixts the third and fifth walls. The fourth wall is a genuine partition, while the others are not. The third and fifth walls are duplicates. We have drawn $C(X, \mathcal{W})$ superimposed on (X, \mathcal{W}) so that the reader can see the dual relationship. In fact, X can be regarded as a set of points in $C(X, \mathcal{W})$. Namely a, b are centers of 1-cubes, c, d, f are 0-cubes, and e is the center of a 3-cube. The pairs of closed halfspaces of $C(X, \mathcal{W})$ induce the given wallspace structure on X .

and it was exploited by the second author in [Wis04] where the walls are ‘hypergraphs’, which are tracks cutting across the 2-cells in a 2-complex. However the language of wallspaces had not yet been adopted. Subsequently, the procedure was described and carefully examined by Nica and Chatterji–Niblo in the context of cubulating the spaces with walls of Haglund–Paulin in [CN05, Nic04]. We shall now review the construction in the slightly more general context of wallspaces, as defined in § 2.1.

Construction 3.2. We now define the CAT(0) cube complex C dual to a wallspace (X, \mathcal{W}) . An orientation $\sigma(W)$ of a wall $W = \{U, V\}$ is a choice of one of the two ordered pairs: (U, V) or (V, U) .¹ We use the notation $\sigma(W) = (\overleftarrow{\sigma}(W), \overrightarrow{\sigma}(W))$.

An orientation σ of the wallspace \mathcal{W} is a choice of orientation $\sigma(W)$ for each wall $W \in \mathcal{W}$. We emphasize that duplicated walls need not be oriented in the same way by σ .

A 0-cube c^σ of C is an orientation of the wallspace that satisfies the following conditions:

- (1) $\overleftarrow{c^\sigma}(W) \cap \overleftarrow{c^\sigma}(W') \neq \emptyset$ for all $W, W' \in \mathcal{W}$;
- (2) for each $x \in X$, we have $x \in \overleftarrow{c^\sigma}(W)$ for all but finitely many $W \in \mathcal{W}$.

Two 0-cubes are connected by a 1-cube c^1 if there is a unique wall W to which they assign opposite orientations. In this case, c^1 is dual to the wall W .

For $n \geq 2$, we add an n -cube whenever its $(n - 1)$ -skeleton is present.

We refer the reader to Figure 4 for an example of a wallspace and its dual cube complex.

Remark 3.3. It is sometimes the case that there are no ‘duplicate halfspaces’ in the sense that no halfspace is associated to more than one wall. Indeed, this holds in the setting of Haglund–Paulin, where the walls are distinct partitions of X . In this case, a 0-cube is equivalent to a certain subcollection of the set of all halfspaces, namely a choice of one halfspace from each wall. (This is a separate issue from duplicate walls, which were discussed in Remark 2.2.)

Consideration of adjacent 0-cubes leads to the following definition, which plays a role in a number of fundamental proofs.

¹The reader might wish to review the formalism advanced in Remark 2.3. In that setting, an orientation of W is an ordered pair of indexed halfspaces. This corresponds to a choice of $(-, +)$ or $(+, -)$. Following this formalism, a wall W whose two halfspaces H_{-W} and H_{+W} are both equal to X actually has two orientations.

DEFINITION 3.4. We define a partial order \preceq on $2^X \times 2^X$ as follows: $(U, V) \preceq (U', V')$ if either: $U \subsetneq U'$ or: $U = U'$ and $V \supseteq V'$. As usual the notation \prec indicates \preceq but not equal.

Let c be a 0-cube, and consider the set

$$S = \{c(W) \mid W \in \mathcal{W}\} = \{(\overleftarrow{c}(W), \overrightarrow{c}(W)) \mid W \in \mathcal{W}\}.$$

If the ordered pair $c(W)$ is \preceq -minimal in S , then W is dual to a 1-cube adjacent to the 0-cube c . However the converse does not hold in general, e.g., when $X = \{1, 2, 3\}$ and $\mathcal{W} = \{\{\{1, 2\}, \{1, 2, 3\}\}, \{\{2, 3\}, \{1\}\}\}$. We frequently use the contrapositive of this fact, namely: if a wall W is not dual to an adjacent 1-cube, then $c(W)$ is not \preceq -minimal.

Remark 3.5 (Boundary components). If we were to omit condition (2) from the definition of a 0-cube, Construction 3.2 would produce a disjoint union of CAT(0) cube complexes, of which the dual cube complex C is one connected component. The other components are *boundary components* and can be associated with a boundary at infinity in a manner that has not yet been fully investigated (see for instance Roller [Rol98] and Guralnik [Gur06a, Gur06b]). Some interesting applications towards the Poisson boundary are explored by Nevo–Sageev in [NS11]. We emphasize that Propositions 3.9 and 3.10 apply to each boundary component.

We recommend that the reader consider the example of a wallspace (X, \mathcal{W}) arising from two infinite systems of parallel lines in \mathbb{R}^2 . When we omit condition (2), Construction 3.2 provides a cube complex consisting of nine components, each of which is a CAT(0) cube complex. One of these components is the dual cube complex $C(X, \mathcal{W})$.

DEFINITION 3.6 (Canonical cubes). Let $x_0 \in X$. Let $\{W_j \mid j \in J\}$ be the set of walls of \mathcal{W} that betwixt x_0 , and let $\{W_i \mid i \in I\}$ be the set of walls of \mathcal{W} that do not betwixt x_0 . For each $i \in I$, we orient W_i so that its left halfspace contains x_0 . Observe that J is finite, by the definition of wallspace. The $2^{|J|}$ distinct orientations of $\{W_j \mid j \in J\}$ determine a $|J|$ -cube. This cube is the *canonical cube* corresponding to x_0 .

In the case where all walls are genuine partitions of X , all canonical cubes are 0-cubes, and these have been termed ‘canonical vertices’ in the literature.

THEOREM 3.7. *The cube complex C is a CAT(0) cube complex.*

Proof. We follow Sageev [Sag95] here. We must show that C is connected, that $\pi_1(C) = 1$, and that C is nonpositively curved. These are verified in the following three propositions. \square

PROPOSITION 3.8. *The cube complex C is connected.*

Proposition 3.8 would be false if there were duplicate nonvacuous walls that were genuine partitions $\{U, X - U\}$ (see Remark 2.2). Indeed if $u \in U$ and $v \in X - U$ and there were two walls of the form $\{U, X - U\}$, then there would be no path between the canonical cubes corresponding to u and v . Therefore C would be disconnected. Vacuous walls must always be oriented towards X . Thus their presence has no effect on C .

Proof of Proposition 3.8. It is critical that there are no duplicate walls that are genuine partitions (see Remark 2.2). Choose a basepoint $x_0 \in X$. Let d be an arbitrary 0-cube of C . Let W_1, \dots, W_n be the walls such that $x_0 \notin \overleftarrow{d}(W_i)$.

Consider the ordered pairs $\{d(W_1), \dots, d(W_n)\}$. One of these, say $d(W_1)$, must be \preceq -minimal in $\{d(W_1), \dots, d(W_n)\}$ in the sense of Definition 3.4. As claimed below, the orientation of W_1 can be reversed to obtain another 0-cube adjacent to d . Repeating this procedure n times, we

eventually reach a 0-cube c such that $x_0 \in \overleftarrow{c}(W)$ for all $W \in \mathcal{W}$. Finally we note that c lies on the canonical cube corresponding to x_0 .

Claim. $\overrightarrow{d}(W_1) \cap \overleftarrow{d}(W) \neq \emptyset$ for all $W \in \mathcal{W}$.

We first note that if $W \notin \{W_1, \dots, W_n\}$ then clearly $\overrightarrow{d}(W_1) \cap \overleftarrow{d}(W) \neq \emptyset$ since both halfspaces contain x_0 .

We now show that $\overleftarrow{d}(W_i) \cap \overrightarrow{d}(W_1) \neq \emptyset$ for $i \in \{2, \dots, n\}$. There are two cases.

Case 1: If $\overleftarrow{d}(W_i) \not\subseteq \overleftarrow{d}(W_1)$, then $\overleftarrow{d}(W_i)$ intersects $(X - \overleftarrow{d}(W_1))$ which lies in $\overrightarrow{d}(W_1)$. Hence $\overleftarrow{d}(W_i) \cap \overrightarrow{d}(W_1) \neq \emptyset$.

Case 2: If $\overleftarrow{d}(W_i) \subseteq \overleftarrow{d}(W_1)$, then we must have $\overleftarrow{d}(W_i) = \overleftarrow{d}(W_1)$ since $d(W_i) \neq d(W_1)$. Arguing by contradiction, suppose that $\overrightarrow{d}(W_1) \cap \overleftarrow{d}(W_i) = \emptyset$. Then $\overrightarrow{d}(W_1) \cap \overleftarrow{d}(W_1) = \emptyset$; in other words, W_1 is a genuine partition. Observe that W_1 is not a vacuous wall $\{X, \emptyset\}$ since both of its halfspaces are nonempty. Indeed, $x_0 \in \overrightarrow{d}(W_1)$ by hypothesis, and $\overrightarrow{d}(W_1) \cap \overleftarrow{d}(W_i) \neq \emptyset$ by the definition of 0-cube. It is impossible for W_i to be a duplicate of W_1 , since by hypothesis there are no duplicate genuine partitions (with the possible exception of vacuous walls). Therefore W_i is not (also) a genuine partition; i.e., $\overleftarrow{d}(W_i) \cap \overrightarrow{d}(W_i) \neq \emptyset$. Therefore

$$\overrightarrow{d}(W_1) = X - \overleftarrow{d}(W_1) = X - \overleftarrow{d}(W_i) \subsetneq \overrightarrow{d}(W_i).$$

Consequently $d(W_i) \prec d(W_1)$, which contradicts the minimality of $d(W_1)$. Thus we must have $\overleftarrow{d}(W_i) \cap \overrightarrow{d}(W_1) \neq \emptyset$. □

PROPOSITION 3.9. *The cube complex C is simply connected.*

Proof. Consider a closed edge path $P = e_1 \cdots e_t$. Choose an ‘innermost’ pair of edges e_p, e_q dual to the same wall W_1 . By this we mean that there does not exist a pair $(r, s) \neq (p, q)$ with $p \leq r < s \leq q$ such that e_r and e_s are dual to the same wall W .

If they are consecutive then $e_p = e_q^{-1}$, and we homotope by removing a backtrack.

Otherwise let W_2 be the wall dual to e_{p+1} . We show below that W_1 and W_2 are transverse. Consequently the orientations of W_1, W_2 can be independently reversed at the terminal 0-cube c_{++} of e_p to provide a 4-cycle $e_p e_{p+1} (e'_p)^{-1} (e'_{p+1})^{-1}$ in the 1-skeleton, which bounds a 2-cube by construction. Specifically, its four 0-cubes are $c_{\pm\pm}$, where c_{+-} is the initial 0-cube of e_p and c_{-+} is the terminal 0-cube of e_{p+1} and c_{--} is obtained from c_{++} by reversing the orientations of both W_1 and W_2 .

We are thus able to homotope $e_p e_{p+1}$ to $e'_p e'_{p+1}$. The resulting path

$$e_1 \cdots e_{p-1} e'_p e'_{p+1} e'_{p+2} \cdots e_t$$

contains a shorter innermost pair of edges dual to W_1 . Continuing this process and removing backtracks, we eventually arrive at a trivial path.

To see that W_1 and W_2 are transverse, observe that since the edges e_{p+1}, \dots, e_{q-1} are dual to distinct walls, the edge e_{p+1} is the only one dual to W_2 . Let x and y be the initial and terminal 0-cubes of e_p . Let z and w be the initial and terminal 0-cubes of e_q . Then w, x, y, z assign orientations to the walls W_1 and W_2 in all four possible combinations. By the definition of 0-cube, all four intersections of halfspaces are therefore nonempty. Therefore W_1 and W_2 are transverse. By the definition of the dual cube complex C , the 1-cubes e_p, e_{p+1} dual to W_1, W_2 form the corner of a 2-cube at y . The path $e'_p e'_{p+1}$ is at the opposite corner. □

PROPOSITION 3.10. *The cube complex C is nonpositively curved.*



FIGURE 5. A collection of walls that corresponds to a 3-cube in $C(X, \mathcal{W})$. The figure shows three independent walls and five dependent walls. Note that this 3-cube lies in a 4-cube.

Sketch. We first observe that $\text{link}(c)$ is a *simplicial* complex. The main point is that only a single i -cube is added for each i -fold collection of pairwise crossing walls.

Consider n pairwise adjacent vertices in $\text{link}(c)$. Their associated 1-cubes are dual to n pairwise crossing walls. These n vertices span an $(n - 1)$ -simplex that is a corner of the cube associated to these n walls (and this 0-cube). \square

Remark 3.11 (The induced action on $C(X, \mathcal{W})$). In conclusion, we note that if G acts on a wallspace (X, \mathcal{W}) , there is an induced action on $C(X, \mathcal{W})$. The action on the 0-skeleton is given by $\overleftarrow{g}c(gW) = g\overleftarrow{c}(W)$ or equivalently $\overleftarrow{g}c(W) = g\overleftarrow{c}(g^{-1}W)$. We note that the action preserves the correspondence between walls and hyperplanes.

3.3 Cubes

In this subsection we examine the cubes of the dual cube complex in more detail. It is easiest to describe the maximal cubes of $C(X, \mathcal{W})$, which correspond to finite cardinality maximal collections of pairwise transverse walls. The description of an arbitrary cube is a bit more subtle since an arbitrary cube is not always contained in a maximal cube. Instead we find that an arbitrary cube corresponds to a certain collection of halfspaces. Much of what we discuss holds for the cubes in the boundary components mentioned in Remark 3.5. Indeed one is led to these boundary components by considering infinite cardinality maximal collections of pairwise transverse walls.

Consider a k -dimensional cube c in $C(X, \mathcal{W})$, and let d be a 0-cube of c . The k distinct 1-cubes at the corner of c containing d are dual to pairwise transverse walls $\{W_1, \dots, W_k\}$. Traveling from 0-cube to 0-cube in the 1-skeleton of c corresponds to reversing the orientations of these walls, whilst preserving the orientations of all other walls in \mathcal{W} . We refer to $\{W_1, \dots, W_k\}$ as the *independent* walls of c and the remaining walls as *dependent* walls. See Figure 5 for an example. In § 3.4 we introduce the notion of a hemiwallspace, which generalizes this viewpoint on cubes to interpret arbitrary convex subcomplexes.

We note that a 0-cube has no independent walls; indeed each wall in \mathcal{W} is assigned a fixed orientation. A 1-cube has a single independent wall, namely the wall dual to it.

We emphasize that the data determining the k -cube c requires a fixed choice of orientation for each dependent wall, but allows the orientation to vary for the k independent walls to reach all 2^k of its 0-cubes. The fixed orientations of the dependent walls must be chosen to satisfy conditions (1) and (2) of the definition of 0-cube. There is always at least one way to make these fixed choices (see Remark 3.12). But often there is more than one way; in particular, there is more than one way if and only if $\{W_1, \dots, W_k\}$ is not a maximal collection of pairwise transverse walls.

An orientation d of a wall W is *towards* a point $x \in X$ if $x \in \overleftarrow{d}(W)$. More generally d orients W *towards* a subset $S \subseteq X$ if d orients W towards some point of S . We say that d orients W *towards* a wall W' if d orients W towards each halfspace of W' .

If c is a cube of $C(X, \mathcal{W})$, then each dependent wall W of c has a well-defined orientation by c . This dependent wall is oriented by c towards each independent wall of c and towards the chosen halfspace of each dependent wall. The reader may wish to consider this language in conjunction with Figure 5.

Remark 3.12. Given a finite collection of pairwise transverse walls W_1, \dots, W_k and a point $p \in X$, there is an associated cube c in $C(X, \mathcal{W})$. The independent walls of c are precisely the walls W_1, \dots, W_k together with those walls W that betwixt p and are transverse to every W_i . If W is a dependent wall that is not transverse with some W_i , then c orients W towards W_i . If W is transverse with every W_i and does not betwixt p , then c orients W towards p . We leave it to the reader to verify that the choices above do indeed determine a cube in $C(X, \mathcal{W})$. A similar conclusion is explicitly verified in Lemma 3.17.

Remark 3.12 leads to the following corollary.

COROLLARY 3.13. *The dual cube complex $C(X, \mathcal{W})$ is finite dimensional if and only if there is a finite upper bound on the size of collections of pairwise transverse walls.*

Having interpreted arbitrary cubes in $C(X, \mathcal{W})$ in terms of independent and dependent walls, we now examine the (finite-dimensional) maximal cubes.

PROPOSITION 3.14. *Maximal cubes in the dual cube complex are in one-to-one correspondence with finite maximal collections of pairwise transverse (nonvacuous) walls in a wallspace.*

The proof of Proposition 3.14 is broken into Lemmas 3.16 and 3.17 below.

LEMMA 3.15 (Descending chain condition). *Let c be a 0-cube of the dual cube complex $C(X, \mathcal{W})$. Then the set $\{c(W) \mid W \in \mathcal{W}\}$ does not contain an infinite properly descending chain with respect to \preceq .*

Proof. Suppose there is an infinite properly descending chain:

$$(\overleftarrow{c}(W_1), \overrightarrow{c}(W_1)) \succ (\overleftarrow{c}(W_2), \overrightarrow{c}(W_2)) \succ \dots$$

The first coordinates cannot be an infinite properly descending chain

$$\overleftarrow{c}(W_1) \supsetneq \overleftarrow{c}(W_2) \supsetneq \dots$$

because such a chain would contradict condition (2) from the definition of a 0-cube indicated in Construction 3.2. Indeed if $x \in \overleftarrow{c}(W_1) - \overleftarrow{c}(W_2)$, then $x \notin \overleftarrow{c}(W_i)$ for all $i \geq 2$. This contradicts the fact that c satisfies condition (2).

It follows that if there is an infinite properly descending chain, then there must be one of the form:

$$(H, \overrightarrow{c}(W_1)) \succ (H, \overrightarrow{c}(W_2)) \succ \dots$$

where $H = \overleftarrow{c}(W_i)$ for all i . In particular $\overrightarrow{c}(W_1) \subsetneq \overrightarrow{c}(W_2) \subsetneq \dots$. Letting $x \in H \cap \overrightarrow{c}(W_2)$, we see that x betwixts W_i for all $i \geq 2$. This contradicts the definition of a wallspace. \square

LEMMA 3.16. *For each maximal cube c of the dual cube complex $C(X, \mathcal{W})$, the collection of independent walls is a maximal family of pairwise transverse walls.*




FIGURE 6. A maximal collection of pairwise transverse walls in the plane $X = \mathbb{R}^2$ that is not associated to an infinite increasing sequence of cubes in the canonical component.

Proof. Let c be a maximal cube. Observe that if W is a dependent wall of c , then the orientation $c(W)$ is well defined. Consider the set

$$\mathcal{S} = \{c(W) \mid W \text{ is dependent in } c \text{ and transverse to every independent wall of } c\}.$$

Suppose by way of contradiction that \mathcal{S} is nonempty. Applying Lemma 3.15 to any 0-cube d of c , we see that \mathcal{S} must have a minimal element $c(W_0)$ with respect to \preceq .

Since c is maximal, the orientation of W_0 cannot be reversed, which means there is another wall W_1 so that $\vec{c}(W_0) \cap \overleftarrow{c}(W_1) = \emptyset$. Observe that W_0 is transverse to each independent wall, but is not transverse to W_1 . Therefore W_1 is dependent. Since $X - \overleftarrow{c}(W_0) \subseteq \vec{c}(W_0)$, it follows that $\overleftarrow{c}(W_1) \subseteq \overleftarrow{c}(W_0)$. If $\overleftarrow{c}(W_1) = \overleftarrow{c}(W_0)$, then $W_0 = \{\overleftarrow{c}(W_0), \vec{c}(W_0)\}$ is a partition of X . Now our hypothesis that there are no duplicate genuine partitions implies that $\vec{c}(W_0) \subsetneq \vec{c}(W_1)$. In particular, we have shown that $(\overleftarrow{c}(W_1), \vec{c}(W_1)) \prec (\overleftarrow{c}(W_0), \vec{c}(W_0))$. In other words, $c(W_1) \prec c(W_0)$.

We will now show that W_1 is transverse to every independent wall of c . Indeed let W be independent. Then both halfspaces of W intersect $\overleftarrow{c}(W_1)$. Since $\vec{c}(W_0) \cap \overleftarrow{c}(W_1) = \emptyset$, we have $\vec{c}(W_0) \subseteq X - \overleftarrow{c}(W_1) \subseteq \vec{c}(W_1)$. But W is transverse with W_0 , therefore both halfspaces of W intersect $\vec{c}(W_0)$, and hence $\vec{c}(W_1)$. Therefore W is transverse to W_1 . Since W is arbitrary, it follows that W_1 is transverse to every independent wall of c , contradicting the fact that W_0 was the minimal such wall. \square

LEMMA 3.17. *Each finite maximal collection $\{W_1, \dots, W_k\}$ of pairwise transverse nonvacuous walls is associated to a maximal cube in the dual cube complex $C(X, \mathcal{W})$. Specifically this cube is the unique cube of $C(X, \mathcal{W})$ whose independent walls are $\{W_1, \dots, W_k\}$.*

It is interesting that the associated maximal cube actually arises in $C(X, \mathcal{W})$ and *not* in one of the boundary components. This contrasts with the infinite case: an infinite maximal collection of pairwise transverse walls could be associated with an infinite increasing sequence of cubes in many ways. One of these might lie in $C(X, \mathcal{W})$ but uncountably many lie in distinct boundary components. While $C(X, \mathcal{W})$ may contain some infinite increasing sequences of cubes, there could be maximal collections of pairwise transverse walls that are only associated with boundary components and not associated at all with $C(X, \mathcal{W})$ (see Figure 6).

Proof of Lemma 3.17. Suppose W_1, \dots, W_k is a finite collection of pairwise transverse nondegenerate walls. We consider these to be independent and all other walls to be dependent. Let $x \in X$ be a fixed point. We define an orientation c of \mathcal{W} as follows. Orient the independent walls towards x . (Due to betwixting, there may be more than one way of doing this.) For

each i choose $x_i \in \overrightarrow{c}(W_i)$. (This can be done since all halfspaces are nonempty.) Let W be a dependent wall. Then there exists at least one independent wall W_i that is not transverse to W . In particular, there are orientations $\sigma(W_i)$ and $\sigma(W)$ so that $\overleftarrow{\sigma}(W_i) \cap \overrightarrow{\sigma}(W) = \emptyset$. In particular, we have $\overleftarrow{\sigma}(W_i) \subseteq \overleftarrow{\sigma}(W)$. Furthermore $\sigma(W)$ is the unique orientation of W towards the independent wall W_i . We define $c(W) = \sigma(W)$ for each dependent wall W .

We now verify that c is a 0-cube. Let us first check that each dependent wall is oriented by c towards the chosen halfspace of each other dependent wall. Let W and W' be dependent walls. Then there are independent walls $W_i = \{U_i, V_i\}$ and $W_j = \{U'_j, V'_j\}$ such that $U_i \subseteq \overleftarrow{c}(W)$ and $U'_j \subseteq \overleftarrow{c}(W')$. Since W_i and W_j are transverse, it follows that $U_i \cap U'_j \neq \emptyset$. Therefore $\overleftarrow{c}(W) \cap \overleftarrow{c}(W') \neq \emptyset$.

If $\overleftarrow{c}(W)$ does not contain x , then $\overleftarrow{c}(W_i)$ also does not contain x . Consequently, $\overleftarrow{c}(W_i)$ contains x_i , and therefore so does $\overleftarrow{c}(W)$. Therefore W is one of the finitely many walls separating x from x_i . Since there are only finitely many independent walls, all but finitely many dependent walls are oriented towards x . \square

3.4 Hemiwallspaces and convex subcomplexes

In this section we introduce and study hemiwallspaces, which are associated to convex subcomplexes of the dual cube complex. Hemiwallspaces generalize our association of cubes in $C(X, \mathcal{W})$ with collections of independent and dependent walls in \mathcal{W} .

To accommodate this study, we will shift gears towards a viewpoint of a wallspace as a collection of halfspaces of X . When equipped with an involution ι (generalized complementation), such a collection of halfspaces leads to a very elegant theory, and Roller has adopted this as the main viewpoint [Rol98]. However our own viewpoint will be slightly different and aimed at natural applications in group theory.

When there are no duplicated halfspaces,² a wallspace is equivalent to a collection of subsets \mathcal{V} together with a pairing $\iota: \mathcal{V} \rightarrow \mathcal{V}$ that is a fixed point free involution such that for each $V \in \mathcal{V}$, we have $V \cup \iota(V) = X$, each x lies in finitely many $V \cap \iota(V)$, and for each $x, y \in X$ there are finitely many V such that $x \in V - \iota(V)$ and $y \in \iota(V) - V$. To emphasize that we are thinking of a collection of halfspaces \mathcal{V} rather than a collection of walls \mathcal{W} , we will often use the notation (X, \mathcal{V}) . Technically we have in mind (X, \mathcal{V}, ι) , but we shall suppress mentioning ι , which is sometimes even uniquely determined by \mathcal{V} .

We will use this halfspace terminology throughout this section, and our arguments and definitions will implicitly adhere to this convention. We will later return to this viewpoint in §7. The general case requires few modifications and a slightly different language: ι simply exchanges the two indexed halfspaces of each wall.

Observe that for Haglund–Paulin wallspaces complementation is the only possible way of defining the pairing ι on the collection of halfspaces \mathcal{V} , and so the wallspace is uniquely determined by \mathcal{V} . More generally, we say a collection of halfspaces is *unambiguous* if there is only one possible way of defining ι so that $V \cup \iota(V) = X$ for all $V \in \mathcal{V}$. In practice however, there are two commonly encountered ways that \mathcal{V} can fail to be unambiguous. Firstly, if \mathcal{V} contains distinct halfspaces U_1, U_2, V_1, V_2 such that each $U_i \cup V_j = X$ then ι might not be uniquely determined by \mathcal{V} . Secondly, there are scenarios where it is natural to consider a collection \mathcal{V} where certain halfspaces are repeated (we would then regard \mathcal{V} as an indexed collection of sets). We emphasize that this second type of failure is more severe than the first, as it is not simply a matter of suppressing ι , but rather a matter of the difficulty of defining ι .

² To handle duplicated halfspaces, one must work with indexed collections of halfspaces; cf. Remark 2.3.

DEFINITION 3.18 (Hemiwallspace). For a wallspace (X, \mathcal{V}) , an associated *hemiwallspace* is a subcollection \mathcal{U} of halfspaces of \mathcal{V} such that for each wall $\{V, \iota(V)\}$ at least one of V or $\iota(V)$ is contained in \mathcal{U} .³ We note that there is a partial pairing of \mathcal{U} induced by ι .

Remark 3.19 (Dependent and independent walls). An alternate approach to defining hemiwallspaces involves restricting orientations. We earlier defined an orientation σ on a wallspace as a choice $\sigma(W) \in W = \{U, V\}$ and used this to define a 0-cube as a particular type of orientation whose nature was restricted only by two simple axioms. An associated *hemiwallspace* is equivalent to a specific way of fixing the orientations of certain walls. These are the *dependent* walls, and the other walls are the *independent* walls. More precisely, a hemiwallspace is an orientation of the set of dependent walls.

As we saw in §3.3, a finite set of pairwise transverse walls will provide us with a set of independent walls. Such a set, together with an appropriate choice of orientations for the remaining dependent walls, leads to a cube in the dual cube complex. Similarly a hemiwallspace leads to a certain convex subcomplex of the dual cube complex as we shall see in Lemma 3.24.

Example 3.20. For a subset $P \subset X$ we let

$$\mathcal{U}_P = \{U \in \mathcal{V} \mid P \cap U \neq \emptyset\},$$

and we regard (X, \mathcal{U}_P) as a hemiwallspace. We discuss related but more elaborate hemiwallspaces in Definition 7.4.

Another structure associated to $P \subseteq X$ is the induced subwallspace described in Definition 2.5. We urge the reader to consider the differences between these.

Construction 3.21 (Cubulating a hemiwallspace). Let \mathcal{U} be a hemiwallspace in a wallspace (X, \mathcal{V}) . The cube complex $C(\mathcal{U})$ dual to (X, \mathcal{U}) is defined as follows: 0-cubes are subcollections $c \subset \mathcal{U}$ that have nonempty pairwise intersection such that for each halfspace $V \in \mathcal{V}$ exactly one of V or $\iota(V)$ is in c . Two 0-cubes are connected by a 1-cube if they differ on two complementary halfspaces. For $n \geq 2$ we add an n -cube if its $(n - 1)$ -skeleton appears.

In practice, the hemiwallspaces we consider are induced from subspaces of X , for instance as in Example 3.20. In that case, there is always at least one 0-cube in the dual cube complex $C(\mathcal{U})$. (We shall always assume that this is the case for purposes of the lemmas below, but their statements and proofs hold for boundary components as well.) It is then a consequence of Lemma 3.24 that there is a unique connected component $C(\mathcal{U})$ with the property that for one and hence any $x \in X$ almost all halfspaces of c contain x .

In the more general setting where a hemiwallspace corresponds to restricted orientations, our 0-cubes are simply chosen so that they conform with the restriction determining the hemiwallspace, and higher cubes are added as before.

Remark 3.22. A hemiwallspace (X, \mathcal{U}) naturally yields a wallspace $(X, \ddot{\mathcal{U}})$ by removing all dependent walls; that is, ignoring all unpaired halfspaces. The cube complex $C(\ddot{\mathcal{U}})$ dual to $(X, \ddot{\mathcal{U}})$ is isomorphic to the cube complex $C(\mathcal{U})$ associated above to (X, \mathcal{U}) . In §7, $C(\ddot{\mathcal{U}})$ will be the cubulation of the peripheral subspace and $C(\mathcal{U})$ naturally corresponds to a convex subcomplex of $C(\mathcal{V})$.

LEMMA 3.23. *The cube complex $C(\ddot{\mathcal{U}})$ embeds naturally in $C(\mathcal{V})$.*

³ In Construction 3.21, we shall also adopt the convention that some subcollection of \mathcal{U} is a (canonical) 0-cube of X . Hence hemiwallspaces correspond precisely with convex subcomplexes of $C(X)$.

Proof. By Remark 3.22 there is a natural isomorphism $C(\mathcal{U}) \rightarrow C(\ddot{\mathcal{U}})$ (forgetting the unpaired walls, i.e., the dependent walls). The natural embedding $C(\mathcal{U}) \rightarrow C(\mathcal{V})$ is a feature of Construction 3.21 (unrestricting the orientations, i.e., considering all walls to be independent). □

LEMMA 3.24. *The cube complex $C(\mathcal{U})$ is a convex subcomplex of $C(\mathcal{V})$. Specifically, combinatorial geodesics in $C(\mathcal{V})$ that start and end in $C(\mathcal{U})$ are actually contained in $C(\mathcal{U})$.*

Proof. Consider a geodesic edge path γ in $C(\mathcal{V})$ whose endpoints a and c lie in $C(\mathcal{U})$. Observe that each 0-cube of γ lies in $C(\mathcal{U})$. Indeed for each 0-cube b of γ each halfspace of b is either a halfspace of a or a halfspace of c (or both). □

We note that $C(\mathcal{U})$ is also convex in $C(\mathcal{V})$ with respect to the CAT(0) metric since it is an intersection of ‘shifted halfspaces’. A *shifted halfspace* is the closure of a component of the complement of $\mathcal{N}_\epsilon(H)$ for some hyperplane H and some $\epsilon < 1$. Shifted halfspaces are slightly smaller than ordinary halfspaces and are convex in the CAT(0) metric.

3.5 The dimension of the dual cube complex and uniform local finiteness

DEFINITION 3.25. A wallspace has the *k-Plane Intersection Property* if each collection of $k + 1$ walls contains a nontransverse pair. In other words, every collection of walls that is pairwise transverse has cardinality at most k .

Equivalently, a wallspace has the *k-Plane Intersection Property* if and only if the corresponding CAT(0) cube complex C has dimension at most k .

The *k-Plane Intersection Property* generalizes a property introduced by Scott [Sco83] in the setting of surfaces in 3-manifolds and studied further by Hass–Scott [HS92] and Rubinstein–Sageev [RS99]. Let M be a closed, orientable, irreducible 3-manifold. An *essential surface* $S \rightarrow M$ is a map from a closed, connected, orientable, aspherical surface S to a 3-manifold M such that $\tilde{S} \rightarrow \tilde{M}$ is an embedding. Note that the collection of all translates of \tilde{S} provides a geometric wallspace on \tilde{M} . We thus say that an essential surface $S \rightarrow M$ has the *k-Plane Intersection Property* if this wallspace satisfies the *k-Plane Intersection Property* as defined above.

Embedded essential surfaces in 3-manifolds obviously have the 1-Plane Intersection Property. More generally, if a cover of an essential surface S in a 3-manifold M lifts to an embedding in a finite cover \hat{M} of M , then S has the *k-Plane Intersection Property* where k is the degree of any regular covering $\bar{M} \rightarrow M$ factoring through \hat{M} .

Rubinstein–Wang [RW98] constructed an example of an essential surface in a graph manifold that does not have the *k-Plane Intersection Property* for any k . In particular, the corresponding CAT(0) cube complex is infinite dimensional.

On the other hand, Rubinstein–Sageev [RS99] showed that if M is a finite volume hyperbolic 3-manifold, then every acylindrical essential surface in M has the *k-Plane Intersection Property* for some k . (If M is compact and hyperbolic, then every essential surface in M is acylindrical.) They also proved a similar result for geometrically finite essential surfaces in 3-manifolds M whose JSJ decomposition contains only hyperbolic pieces glued along tori.

DEFINITION 3.26 (Bounded packing). If G is a group with finite generating set \mathcal{A} , a pair of subsets $A, B \subseteq G$ are *D-close* if $d_{\mathcal{A}}(A, B) := \min\{d_{\mathcal{A}}(a, b) \mid a \in A, b \in B\} < D$. A subgroup $H \leq G$ has *bounded packing* in G if for each $D > 0$ there exists k such that every collection of $k + 1$ left cosets gH of H contains a pair that is not *D-close*. In other words, every collection of left cosets of H that is pairwise *D-close* has cardinality at most k . Note that while the specific

value of k depends on the metric, having the bounded packing property is independent of the generating set.

Sageev established bounded packing for quasiconvex subgroups of hyperbolic groups in [Sag97]. We prove the following generalization in [HrW09].

THEOREM 3.27. *Let H be a relatively quasiconvex subgroup of a relatively hyperbolic group G . Suppose $H \cap gPg^{-1}$ has bounded packing in gPg^{-1} for each conjugate of each peripheral subgroup P . Then H has bounded packing in G .*

DEFINITION 3.28. Let (X, \mathcal{W}) be a wallspace. Generalizing the notion of a wall separating two points, we say W separates subsets P, Q of X if P and Q lie in distinct open halfspaces of W . A wall W'' separates W, W' if $W'' \neq W$, $W'' \neq W'$, and there are closed halfspaces U of W and V' of W' such that W'' separates U and V' .

Distinct walls W, W' osculate if they are not transverse and there is no wall separating them.

Remark 3.29. Walls W, W' osculate in X if and only if the corresponding hyperplanes H, H' of $C(X)$ are dual to 1-cubes e, e' that intersect at a 0-cube but do not lie in a 2-cube. Indeed, let $W = \{U, V\}$ and $W' = \{U', V'\}$ be osculating walls with $V \cap V' = \emptyset$. Choose a minimal length path γ in $C(X)$ between $C(\mathcal{U}_V)$ and $C(\mathcal{U}_{V'})$. Reasoning as in the proof of Proposition 3.8 one reaches the conclusion that γ is the concatenation of two edges dual to W and W' respectively. Conversely, walls W, W' corresponding to hyperplanes dual to adjacent 1-cubes e, e' not in a 2-cube cannot be separated by any other wall since it would then be impossible for both W, W' to be flippable.

The following theorem explains how the bounded packing property applies to ensure the k -Plane Intersection Property, and hence the finite dimensionality of the dual cube complex and similarly its uniform local finiteness. We emphasize that the additional necessary hypothesis for finite dimensionality is usually immediate in a geometric setting. However the additional hypothesis for uniform local finiteness is far from transparent because it entails a more global condition related to the Wall–Wall Separation Property (see §§ 8.3 and 8.4).

THEOREM 3.30. *Let G be a finitely generated group that acts on a wallspace (X, \mathcal{W}) with finitely many G -orbits of walls.*

Suppose there are finitely many G -orbits of pairs of transverse walls and that the stabilizer of each wall has bounded packing in G . Then the wallspace has the k -Plane Intersection Property. Consequently the dual cube complex $C(X, \mathcal{W})$ is finite dimensional.

Suppose in addition that there are finitely many G -orbits of pairs of osculating walls. Then $C(X, \mathcal{W})$ is uniformly locally finite, in the sense that there is a uniform upper bound on the number of 1-cubes incident to each 0-cube.

Proof. Choose representatives W_1, \dots, W_r of the finitely many G -orbits of walls, and let H_1, \dots, H_r denote their stabilizers. Observe that each wall $W \in \mathcal{W}$ is a translate gW_i of one of our representatives, and is associated with the coset gH_i .

Choose finitely many pairs of walls representing the finitely many G -orbits of pairs of transverse walls. Let D be an upper bound on the distance between their associated cosets. As there are finitely many distinct G -orbits of walls in X and by hypothesis the bounded packing property holds for their stabilizers, we can choose k to be an upper bound on the size of a collection of pairwise D -close cosets.

By Corollary 3.13, the k -Plane Intersection Property for (X, \mathcal{W}) implies the finite dimensionality of $C(X, \mathcal{W})$.

Similarly, a collection of 1-cubes at a 0-cube corresponds to a collection of walls that pairwise either osculate or are transverse. Hence the result follows from the bounded packing property in the same manner as above. \square

If we assume the subgroups in question are finitely generated, then by Lemma 2.8 there are finitely many orbits of pairs of transverse walls. We thus deduce the following corollary.

COROLLARY 3.31. *Suppose the finitely generated group G has bounded packing with respect to finitely generated subgroups H_1, \dots, H_r . Let $C(X)$ be the dual CAT(0) cube complex with respect to a system of H_i -walls in $X = \Gamma(G)$. Then $C(X)$ is finite dimensional.*

The following converse to the second part of Theorem 3.30 was proven in [HrW09, Theorem 3.2(1)]. (The statement there incorrectly omits the word ‘uniformly’, but uniform local finiteness is explicitly used in the proof.)

THEOREM 3.32. *Suppose a finitely generated group G acts on a CAT(0) cube complex C and $H \leq G$ is the stabilizer of a hyperplane. If C is uniformly locally finite then H has bounded packing in G .*

4. CAT(0) spaces with convex walls

In § 4.1 we describe our main theorems in the setting of a CAT(0) space with convex walls. In § 4.2 we give simple examples of 2-dimensional CAT(0) spaces with convex walls. The discussion of these examples incorporates the subsequent finiteness results on dual cube complexes from §§ 5 and 7 that were stated in the introduction. We hope that the reader will be able to navigate § 4 to have concrete examples in mind and then revisit the section later.

4.1 Main results illustrated in a highly geometric setting

The following is a motivating example of a geometric wallspace similar to the Coxeter and cubical settings.

Example 4.1. Let X be a CAT(0) space. Let \mathcal{W} be a locally finite collection of convex 2-sided subspaces.

We will describe how some such examples arise in certain naturally occurring CAT(0) 2-complexes in § 4.2. Higher dimensional versions of this scenario can be described analogously but appear to be scarcer. In a hyperbolic high dimensional instance, X is \mathbb{H}^n and the convex walls are copies of \mathbb{H}^{n-1} (see for instance Bergeron–Haglund–Wise [BHW11]).

A cocompact action. Assume that a group G acts properly and cocompactly on (X, \mathcal{W}) . Observe that stabilizers of walls are quasi-isometrically embedded in G . According to Theorem 7.1, if G is word-hyperbolic, then G acts cocompactly on the dual cube complex.

A relatively cocompact action. Let us now consider a motivating case of the relatively hyperbolic situation; namely suppose that X is also a CAT(0) space with isolated flats. According to Theorem 7.12, G acts relatively cocompactly on the dual cube complex (see Definition 1.2), and in fact the dual cube complex looks like a quasi-copy of G together with various cubulations of free abelian groups hanging off. The simplest such example arises from a piecewise-Euclidean 2-complex all of whose 2-cells are regular even-sided polygons with at least 6 sides (see § 4.2). Note that the same result holds without assuming that the polygons have an even number of sides [Wis04], but there is no longer an obvious natural choice of isometrically embedded walls.




FIGURE 7. The figure on the left shows a portion of the dual tree in the case when all 2-cells are regular polygons with an even number of sides. The figure in the middle illustrates the case when some 2-cells have an odd number of sides, and the dual tree is forced to pass through 0-cells. On the right is a \ominus -graph and several sets of midpoints that are π -separating, provided that the edges have length $\geq \pi/3$.

A proper action. Assume now that a group G acts properly on (X, \mathcal{W}) . According to Theorem 5.2, G acts metrically properly on the dual cube complex provided that there exists $L > 0$ such that any two points at distance at least L are separated by some wall.

Problem 4.2. Is there an example of a CAT(0) space X with a proper cocompact G -action and a locally finite G -invariant system \mathcal{W} of convex walls such that the dual cube complex $C(X, \mathcal{W})$ is not finite dimensional?

4.2 Examples of CAT(0) 2-complexes with convex walls

In this section we will describe simple conditions on a CAT(0) 2-complex X that endow it with a locally finite collection \mathcal{W} of convex 2-sided trees. Moreover this collection *fills* X in the sense that the Linear Separation condition holds (see Definition 5.1).

We first describe a condition yielding trees that are disjoint from the 0-cells.

Example 4.3. Suppose X is a CAT(0) 2-complex such that each 2-cell is a regular Euclidean polygon with an even number of sides. For each 1-cell e , the tree T dual to e passes through the midpoint m of e as a perpendicular bisector. If e lies on r 2-cells then m is a valence r vertex of T . The tree bisects any 2-cell c that it passes through; $T \cap c$ is an edge of T , which ends at the midpoints of opposite 1-cells of c , as shown on the left in Figure 7. Similar walls have been examined in a slightly more general context in [HP98].

We now consider the more general setting where polygons are not regular and may not have an even number of sides. Conditions guaranteeing the existence of such trees are trickier to state, especially considering that there are such examples where no isometric 2-sided trees exist.

The way in which a tree T locally cuts X in half is modeled by two situations, illustrated in Figure 7:

- (1) perpendicularly cutting through a 1-cell at an internal point; and
- (2) perpendicularly cutting through a 0-cell v .

In the second case, v is actually a vertex of T , and we now describe what we mean by ‘perpendicularly cutting through v ’. There is an embedding $\text{Lk}_T(v) \hookrightarrow \text{Lk}_X(v)$. Note that the edges of $\text{Lk}_X(v)$ are assigned lengths according to the angles of the corresponding corners of 2-cells in X . The *angular metric* on $\text{Lk}_X(v)$ is the induced path metric on $\text{Lk}_X(v)$.

Perpendicularly cutting through v requires that:

- (1) points of $\text{Lk}_T(v)$ are π -separated in $\text{Lk}_X(v)$, in the sense that for each distinct $a, b \in \text{Lk}_T(v)$, the angular distance in $\text{Lk}_X(v)$ between a and b is at least π ; and
- (2) $\text{Lk}_T(v)$ separates $\text{Lk}_X(v)$ into two complementary components.

When $Lk_T(v)$ is a \ominus -graph whose three arcs have length π , then there are two π -separated separating sets. But there are more possibilities when the lengths of the arcs are increased, as shown on the right of Figure 7.

Patching together these local models to yield a locally finite collection of 2-sided trees can be tricky. A simple condition that guarantees that the local models can be patched together requires that:

- (1) all polygons are regular;
- (2) the embedding $Lk_T(v) \subset Lk_X(v)$ lies at the centers of edges; and
- (3) the center of each edge lies in exactly one $Lk_T(v)$ subset.

An example of this is illustrated on the right of Figure 7.

The $Lk_T(v)$ subsets make $Lk_X(v)$ into a wallspace. This type of local wallspace structure is exploited in a forthcoming paper by Wise to create many global wallspaces in a similar fashion.

The following example arose recently in the work of Barré–Pichot [BP10]. Barré–Pichot use these walls to see that $\pi_1 X$ is a-T-menable, by applying a criterion of Haglund–Paulin [HP98].

Example 4.4. Let X be the standard 2-complex of the following presentation, and let x be the 0-cell of X :

$$\langle a, b, c \mid abc, acb, ad^{-1}bd^{-1}cd^{-1} \rangle.$$

The three 2-cells and $\text{link}(x)$ are illustrated in Figure 8. Observe that X has a nonpositively curved piecewise-Euclidean 2-complex whose 2-cells are equilateral triangles and a regular hexagon. The walls in \tilde{X} are ‘generated’ by geodesic segments that cut through the 2-cells as illustrated in Figure 8. The three such segments in each triangle are the perpendicular bisectors of its sides. There are nine such segments in the hexagon: three of these are the perpendicular bisectors of the hexagon, and the other six cut across the hexagon to form a triangle with consecutive sides of the hexagon.

The graph $\text{link}(x)$ is illustrated on the right of Figure 8. There are three walls that perpendicularly cut through $\text{link}(x)$. We have illustrated one of these walls, namely the one corresponding to the four bold vertices of the wall segments on the left.

The linear separation property for the geometric wallspace on \tilde{X} generated by these walls is immediate. Indeed, let L denote the union of all the walls, and consider $\tilde{X} - L$. Each component of $\tilde{X} - L$ is a subspace of the union of cells containing a single 0-cell, and hence has uniformly bounded diameter.

5. Properness and local finiteness

5.1 Properness of the action

DEFINITION 5.1 (Linear separation). Let G act on a geometric wallspace (X, \mathcal{W}) . Recall that the underlying space (X, d) is a metric space. The geometric wallspace has the *Linear Separation Property* if there exist constants $\kappa, \epsilon > 0$ such that for all points $x, y \in X$ we have

$$\#(x, y) \geq \kappa d(x, y) - \epsilon.$$

The reader may wish to verify that Linear Separation holds for the examples in §4.2.

An inequality of the form $\#(x, y) \leq \kappa d(x, y) + \epsilon$ is of less interest here. For instance this inequality holds whenever the following ubiquitous condition is satisfied: X is a geodesic metric space, G acts on (X, \mathcal{W}) , there exists a bounded set B such that $X = GB$, and there are finitely many distinct walls that separate points of B .




FIGURE 8. The walls in \tilde{X} are generated by segments cutting through the 2-cells; some of these are indicated on the left. Each wall perpendicularly cuts $\text{link}(x)$. One of the three types of such cuts is illustrated within $\text{link}(x)$ on the right.

An action of a group G on a topological space X is *proper* if $\{g \in G \mid gK \cap K \neq \emptyset\}$ is finite for each compact subspace $K \leq X$. We note that an action G on a cube complex C is proper precisely when all 0-cubes of C have finite stabilizers.

An action of a group G on a metric space X is *metrically proper* if for each $r > 0$ and each $x \in X$ the set $\{g \mid gB_r(x) \cap B_r(x) \neq \emptyset\}$ is finite. Note that if G acts metrically properly then G acts properly. The two notions coincide when X is *proper*, which means that every closed ball in X is compact.

In particular, observe that a proper action on a locally finite cube complex is always metrically proper.

THEOREM 5.2. *Suppose G acts on a geometric wallspace (X, \mathcal{W}) , and the action on the underlying metric space (X, d) is metrically proper. Then the Linear Separation Property for (X, \mathcal{W}) implies that G acts metrically properly on $C(X, \mathcal{W})$.*

More generally G acts metrically properly on $C(X, \mathcal{W})$ provided that for some $x \in X$ we have: $\#(x, gx) \rightarrow \infty$ as $g \rightarrow \infty$.

We have in mind the case when G has a proper metric (e.g., a word metric), in which case a sequence of distinct elements in G always tends to infinity.

Although Theorem 5.2 holds under the more general hypothesis stated second, we have deliberately chosen to emphasize the Linear Separation hypothesis as this stronger property on (X, \mathcal{W}) is often the natural one to verify.

Proof. Suppose (g_i) is a sequence of distinct elements in G , and let $x \in X$ be any basepoint. Since G acts metrically properly on X , we have $d(x, g_i x) \rightarrow \infty$ as $i \rightarrow \infty$. Thus $\#(x, g_i x) \rightarrow \infty$. We shall now verify that $\#(x, y) = d_C(c_x, c_y)$, where c_x denotes the canonical cube containing x . Consequently $d_C(c_x, g_i c_x) = d_C(c_x, c_{g_i x}) \rightarrow \infty$.

Recall that the independent walls of c_x are precisely those that betwixt x . Observe that if W separates x from y then it is betwixted by neither of these points. Hence W is a dependent wall oriented towards x in the cube c_x and towards y in c_y . Consequently each such separating W contributes a value of 1 to $d_C(c_x, c_y)$. Moreover each hyperplane separating c_x from c_y corresponds to such a separating wall. □

A wall W *crosses* a set S if S intersects each closed halfspace of W .

LEMMA 5.3 (Ball–Ball separation). *Let X be a metric space and also a wallspace such that each finite radius ball in X is crossed by finitely many walls. Suppose G acts cocompactly on*

(X, \mathcal{W}) and acts properly on the dual cube complex $C(X)$. For each r there exists m such that if $d_X(x_1, x_2) > m$ then there is a wall separating $\mathcal{N}_r(x_1)$ from $\mathcal{N}_r(x_2)$.

‘Ball–Ball Separation’ is considered among other separation properties in Section 8.3.

Proof. By increasing r , we can assume without loss of generality that each $x_i = g_i x$ for some $g_i \in G$ and a chosen point x . The hemiwallspace associated with $\mathcal{N}_r(x)$ corresponds to a compact convex subcomplex $C_r(\{x\})$ of $C(X)$. By properness of the G action on $C(X)$, the set of elements $S = \{g \mid gC_r(\{x\}) \cap C_r(\{x\}) \neq \emptyset\}$ is finite. Disjoint convex subcomplexes of $C(X)$ are separated by a hyperplane (see for instance [Rol98, Theorem 2.7]). Consequently $\mathcal{N}_r(gx)$ can be separated from $\mathcal{N}_r(x)$ provided $g \notin S$. Let $m = \max_{g \in S} d_X(x, gx)$. Then $\mathcal{N}_r(g_1x)$ is separated from $\mathcal{N}_r(g_2x)$ provided that $d(g_1x, g_2x) > m$. □

Remark 5.4. A wallspace is Hausdorff if $\#(x, y) \neq 0$ whenever $x \neq y$. Note that $(X, \#)$ is a metric space when (X, \mathcal{W}) is a Hausdorff wallspace. The condition that $\#(x, gx) \rightarrow \infty$ as $g \rightarrow \infty$ can then be restated as: G acts metrically properly on $(X, \#)$. Theorem 5.2 is well known in this setting.

The following non-Hausdorff example shows that $\#$ can fail to satisfy the triangle inequality. Let $X = \{x, y, z\}$. Let \mathcal{W} consist of a single wall given by $(\{x, z\}, \{y, z\})$. Then $\#(x, y) = 1$ but $\#(x, z) = \#(y, z) = 0$. Nevertheless $\#$ is related to a distance function since, as explained in the proof above, $\#(x, y) = d_C(c_x, c_y)$.

We motivate Theorem 5.5 by the following scenario: let g act on X by an isometry preserving an axis A isometric to the real line. Let W be a wall in X that separates the ends $\pm\infty$ of A in the sense that $(-\infty, -n)$ and $(n, +\infty)$ lie in distinct open halfspaces of W . We think of W as ‘cutting’ g .

A more extreme scenario is where $W = \{U, V\}$ and we have $g(U) \subsetneq U$ and $V \subsetneq g(V)$.

THEOREM 5.5 (Axis separation). *Let $g \in G$ be an infinite order element. Suppose an infinite order element $g \in G$ is cut by a wall $W = \{U, V\}$ in the following sense:*

- (1) $W \neq g^n W$ when $n \neq 0$;
- (2) $U \cap g^n U \neq \emptyset$ for all n ;
- (3) $V \cap g^n V \neq \emptyset$ for all n .

If g fixes a 0-cube c of $C(X, \mathcal{W})$, then (X, \mathcal{W}) has an infinite collection of pairwise transverse walls, and hence $C(X, \mathcal{W})$ is infinite dimensional. In particular, if $C(X, \mathcal{W})$ is finite dimensional, then G acts with torsion stabilizers.

Proof. Suppose $gc = c$ for some $c \in C^0$, and hence $g^{-m}c \stackrel{\leftarrow}{\leftarrow} c$ for each m . This means that $g^{-m}c$ and c orient each wall W in the same way, so we have $g^{-m}c(W) = \overleftarrow{c}(W)$. As mentioned in Remark 3.11, $\overleftarrow{g^{-m}c}(W) = g^{-m} \overleftarrow{c}(g^m W)$. Thus $\overleftarrow{c}(g^m W) = g^m \overleftarrow{c}(W)$ for all m .

We will see that the walls $\{g^m W \mid m \in \mathbb{Z}\}$ are pairwise transverse. Suppose by way of contradiction that W and $g^n W$ are not transverse for some n . By hypothesis the quarterspaces $\overleftarrow{c}(W) \cap g^n \overleftarrow{c}(W)$ and $\overrightarrow{c}(W) \cap g^n \overrightarrow{c}(W)$ are nonempty. It follows that at least one of the quarterspaces $\overleftarrow{c}(W) \cap \overrightarrow{c}(g^n W)$ and $\overrightarrow{c}(W) \cap \overleftarrow{c}(g^n W)$ is empty. Therefore one of the halfspaces $\overleftarrow{c}(W)$ and $\overleftarrow{c}(g^n W)$ is contained in the other. Replacing n with $-n$ if necessary, we can arrange that $c(g^n W) \prec c(W)$. This yields an infinite properly descending sequence

$$c(W) \succ c(g^n W) \succ c(g^{2n} W) \succ \dots$$

which is impossible by Lemma 3.15.

We thus conclude that W and g^mW are transverse for each m , and thus g^kW and g^mW are transverse for each $k \neq m$. By Corollary 3.13, the dual cube complex is infinite dimensional. \square

5.2 Local finiteness

The investigation of local finiteness of C is often simplified in the presence of a group action on (X, \mathcal{W}) because if G acts properly and cocompactly on C then C is locally finite.

More generally suppose G acts properly on C . Assume there is a collection of subcomplexes $C_i \subset C$ and a compact subcomplex $K \subset C$ such that $C = GK \cup \bigcup_i GC_i$. Suppose each 0-cube has a neighborhood that either lies in some gC_i or lies in GK . Then C is locally finite provided that each C_i is locally finite. We emphasize that the subcomplexes gC_i might intersect, and this intersection could be unrelated to GK .

We will show in §7.4 that the above decomposition of C arises naturally when cubulating certain relatively hyperbolic groups. In that setting we obtain the stronger conclusion that points outside of GK lie in a *unique* gC_i . In other words, the subcomplexes gC_i can intersect only inside GK . We then examine the local finiteness of C in that setting in §8.4.

In this subsection we describe a simple criterion on a wallspace to recognize local finiteness of its cubulation. We give conditions similar to the Parallel Wall Separation Condition of Brink–Howlett that was used by Niblo–Reeves to prove local finiteness of their cubulations of Coxeter groups [BH93, NR03]. Guralnik has undertaken a study of local finiteness of the cube complex dual to a system of convex walls in a CAT(0) space obtaining a result similar to one implication of Theorem 5.7 [Gur06b].

DEFINITION 5.6 (Compact-wall separation). Let X be a wallspace that is also a metric space. Then X has *compact-wall separation* if for each compact set $K \subseteq X$ there exists a constant $f(K)$ such that whenever $d(K, W) \geq f(K)$ there is a wall W' separating K from W . By this we mean that K lies in one open halfspace of W' and a closed halfspace of W lies in the other.

In particular, this condition holds when X has *ball-wall separation at $x \in X$* , in the sense that for each $r \geq 0$ there exists $f(r, x)$ such that if $d(x, W) \geq f$ then there is a wall W' separating $\mathcal{N}_r(x)$ from W . Note that ball-wall separation with respect to a particular $x \in X$ implies ball-wall separation with respect to every $y \in X$.

THEOREM 5.7. *Suppose X is a metric space and also a wallspace (with no vacuous walls), and that each finite radius ball is crossed by finitely many walls. Then (X, \mathcal{W}) has compact-wall separation if and only if $C(X, \mathcal{W})$ is locally finite.*

Proof. First assume that (X, \mathcal{W}) has compact-wall separation. Let c be a 0-cube of $C(X, \mathcal{W})$. In order to bound the valence of c , we need to bound the number of oriented walls $c(W)$ that are *reversible*, in the sense that reversing only the orientation of W yields a new 0-cube.

Let $x \in X$ be a basepoint. Consider the finitely many walls W_1, \dots, W_n that are *oriented away* from x by c , in the sense that $x \notin \overleftarrow{c}(W_j)$ for all j . For each j choose $x_j \in \overleftarrow{c}(W_j)$. (Here we use that there are no vacuous walls.) Let $K = \{x, x_1, \dots, x_n\}$.

Note that each W_j crosses K since $x \in \overrightarrow{c}(W_j)$ but $x_j \in \overleftarrow{c}(W_j)$. Note that by hypothesis there are finitely many walls crossing $\mathcal{N}_f(K)$, where $f = f(K)$ is also given by hypothesis.

We claim that each reversible oriented wall of c is one of these finitely many walls crossing $\mathcal{N}_f(K)$. Indeed, if W is reversible and doesn't cross $\mathcal{N}_f(K)$ then by compact-wall separation, there exists W' separating K and W . Observe that all walls not crossing K are oriented towards x . Since W' does not cross K , it follows that $\overleftarrow{c}(W')$ contains K . For the same reason $\overleftarrow{c}(W)$ contains K . As W' separates W from K , it follows that $\overleftarrow{c}(W')$ is disjoint from $\overrightarrow{c}(W)$, which contradicts the reversibility of W .

We now show that if $C(X, \mathcal{W})$ is locally finite, then (X, \mathcal{W}) has compact-wall separation. Let K be a compact set in X .

The local finiteness of walls in X implies that the hemiwallspace induced by K has finitely many independent walls. This ensures the compactness of the dual $C(K)$, which is a convex subcomplex of $C(X, \mathcal{W})$. The local finiteness of $C(X, \mathcal{W})$ implies that there are finitely many hyperplanes dual to 1-cubes with a 0-cube on $C(K)$.

Suppose there is a sequence of walls W_i in \mathcal{W} such that $d(K, W_i) \geq i$, but each W_i is not separated from K . We claim that each W_i corresponds to a hyperplane H_i in $C(X, \mathcal{W})$ dual to a 1-cube with a 0-cube on $C(K)$, which is impossible. To verify the claim, consider a minimal (combinatorial) geodesic γ between $C(K)$ and $C(W_i)$, which is the cubical neighborhood of H_i . One can verify that each edge of γ corresponds to a hyperplane separating $C(K)$ and $C(W_i)$. This is associated to a wall separating K and W_i . \square

We close this section with the following example of a wallspace (X, \mathcal{W}) , which has compact-wall separation even though its dual cube complex $C(X, \mathcal{W})$ is not locally finite.

Example 5.8. Let $X = \mathbb{R}$, and let the walls in \mathcal{W} be of the following two types:

- (1) $\{(-\infty, n], [n, \infty)\}$ for each $n \in \mathbb{Z}$; and
- (2) $\{\{r\}, \mathbb{R} - \{r\}\}$ for each $r \in \mathbb{R}$.

The dual cube complex $C(X, \mathcal{W})$ is the union of a horizontal line L consisting of infinitely many horizontal edges with an uncountable set of vertical edges attached at each vertex of L and a square freely attached at each horizontal edge of L .

The n th square arises since $\{(-\infty, n], [n, \infty)\}$ is transverse to $\{\{n\}, \mathbb{R} - \{n\}\}$. The vertical edges between the n th and $(n + 1)$ th square correspond to $\{\{r\}, \mathbb{R} - \{r\}\}$ with $n \leq r \leq n + 1$.

6. Review of relative hyperbolicity and quasiconvexity

In this section we review the definition of a relatively hyperbolic group and a relatively quasiconvex subgroup of such a group.

A geodesic metric space (\widehat{X}, d) is δ -hyperbolic if every geodesic triangle with vertices in \widehat{X} is δ -thin in the sense that each side lies in the δ -neighborhood of the union of the other two sides. If \widehat{X} is a δ -hyperbolic space, the *boundary* (or *boundary at infinity*) of \widehat{X} , denoted $\partial\widehat{X}$, is the set of all equivalence classes of geodesic rays in \widehat{X} , where two rays γ, γ' are equivalent if the distance $d(\gamma(t), \gamma'(t))$ remains bounded as $t \rightarrow \infty$. The set $\partial\widehat{X}$ has a natural topology, which is compact and metrizable (see, for instance, Bridson–Haefliger [BH99] for details).

Increasing the constant δ if necessary, we will also assume that every geodesic triangle with vertices in $\widehat{X} \cup \partial\widehat{X}$ is δ -thin.

Let $\Delta = \Delta(x, y, z)$ be a triangle with vertices $x, y, z \in \widehat{X} \cup \partial\widehat{X}$. A *center* of Δ is a point $w \in \widehat{X}$ such that the ball $B(w, \delta)$ intersects all three sides of the triangle. It is clear that each side of Δ contains a center of Δ .

If $\xi \in \partial\widehat{X}$, a function $h: \widehat{X} \rightarrow \mathbb{R}$ is a *horofunction centered at ξ* if there is a constant D_0 such that the following holds. Let $\Delta(x, y, \xi)$ be a geodesic triangle, and let $w \in \widehat{X}$ be a center of the triangle. Then

$$|(h(x) + d(x, w)) - (h(y) + d(y, w))| < D_0.$$

A subset $B \subseteq \widehat{X}$ is a *horoball centered at ξ* if there is a horofunction h centered at ξ and a constant D_1 such that $h(x) \geq -D_1$ for all $x \in B$ and $h(x) \leq D_1$ for all $x \in \widehat{X} - B$.

We remark that every horoball or horofunction has a unique center, and every $\xi \in \partial\widehat{X}$ is the center of a horofunction (see Gromov [Gro87, § 7.5] for details).

The following is a well-known result about the interaction between geodesics and horoballs. We leave the proof as an exercise for the reader.

LEMMA 6.1. *Let B and B' be two horoballs of \widehat{X} centered at the same point ξ such that $B \subset B'$. There is a constant $M = M(B, B')$ such that the following holds. Let γ be a geodesic of \widehat{X} with endpoints $x, y \in B' - B$. Then*

$$\gamma \cap (\widehat{X} - B) \subseteq \mathcal{N}_M(\{x, y\}).$$

In particular, if $d(x, y) > 2M$, then γ must intersect B .

Suppose G has a proper, isometric action on a proper δ -hyperbolic space \widehat{X} . An element $g \in G$ is *loxodromic* if it has infinite order and fixes exactly two points of $\partial\widehat{X}$. A subgroup $P \leq G$ is *parabolic* if it is infinite and contains no loxodromic element. A parabolic subgroup P has a unique fixed point in $\partial\widehat{X}$, called a *parabolic point*. The stabilizer of a parabolic point is always a maximal parabolic group. Thus there is a natural one-to-one correspondence between parabolic points and maximal parabolic subgroups.

The following is Gromov’s definition of relative hyperbolicity from [Gro87], as elaborated by Bowditch in [Bow12].

DEFINITION 6.2 (Relatively hyperbolic). Let G be a countable group with a finite collection \mathbb{P} of infinite subgroups. Suppose G acts properly on a proper δ -hyperbolic space \widehat{X} , and \mathbb{P} is a set of representatives of the conjugacy classes of maximal parabolic subgroups. Suppose also that there is a G -invariant collection of disjoint open horoballs centered at the parabolic points of G , with union \mathcal{B} , such that the quotient of $\widehat{X} - \mathcal{B}$ by the action of G is compact. Then (G, \mathbb{P}) is *relatively hyperbolic*.

In this case, we say that the action of (G, \mathbb{P}) on \widehat{X} is *cuspid uniform*, and the space $X = \widehat{X} - \mathcal{B}$ is a *truncated space* for the action. We refer to the horoball components B of \mathcal{B} as *horoballs of X* . The subgroups $P \in \mathbb{P}$ are called *peripheral subgroups*. When $\mathbb{P} = \emptyset$ the truncated space X is equal to \widehat{X} , and the group G is word-hyperbolic.

The hypothesis that all parabolic subgroups are infinite is a technical luxury and not at all necessary for the theory.

In [Dah03], Dahmani defined a subgroup $H \leq G$ to be quasiconvex relative to the peripheral subgroups if the action of G as a geometrically finite convergence group restricts to a geometrically finite convergence group action of H on its limit set. Subsequently Osin [Osi06] gave the following alternate formulation of relative quasiconvexity. In [Hru10], Hruska showed that these definitions are equivalent.

DEFINITION 6.3 (Relatively quasiconvex (A)). Let (G, \mathbb{P}) be relatively hyperbolic with a finite generating set \mathcal{S} . Let $\mathcal{P} = \bigcup_{P \in \mathbb{P}} P - \{e\}$. A subgroup $H \leq G$ is *relatively quasiconvex* if there is a constant κ such that the following condition holds: for each geodesic γ in $\text{Cayley}(G, \mathcal{S} \cup \mathcal{P})$ connecting two points of H , every vertex of γ lies within a distance κ of H in $\text{Cayley}(G, \mathcal{S})$.

We occasionally find the following formulation of relative quasiconvexity useful. This is also equivalent to the above by [Hru10].

DEFINITION 6.4 (Relatively quasiconvex (B)). Let (G, \mathbb{P}) be relatively hyperbolic. A subset $H \leq G$ is κ -*relatively quasiconvex* if the following holds. Let (\widehat{X}, d) be some proper δ -hyperbolic space on which (G, \mathbb{P}) has a cuspid uniform action. Let $\widehat{X} - \mathcal{B}$ be some truncated space for G acting on \widehat{X} . For some basepoint $x \in \widehat{X} - \mathcal{B}$ there is a constant κ such that whenever γ is a geodesic in \widehat{X} with endpoints in the orbit Hx , we have

$$\gamma \subseteq \mathcal{N}_\kappa(Hx) \cup \mathcal{B}.$$

We say that H is *relatively quasiconvex* if it is κ -relatively quasiconvex for some κ .

Relative quasiconvexity of H is independent of the choice of \widehat{X} , \mathcal{B} , and x . However the value of κ depends on these choices. See [Hru10] for details.

DEFINITION 6.5 (Quasiconvex). Let Z be a geodesic metric space and $\kappa > 0$. A subset $A \subseteq Z$ is κ -quasiconvex if $\gamma \subseteq \mathcal{N}_\kappa(A)$ whenever γ is a geodesic in Z whose endpoints lie on A . We say A is quasiconvex if it is κ -quasiconvex for some $\kappa > 0$.

We often refer to quasiconvex subsets of a word-hyperbolic group G , i.e., those that are quasiconvex as subsets of $\text{Cayley}(G, S)$ for a finite generating set S . Note that quasiconvexity is independent of the particular choice of S when G is word-hyperbolic, but this can fail for general G .

When G is relatively hyperbolic, quasiconvexity is a stronger property than relative quasiconvexity.

We refer the reader to §8.2 for the notion of saturation, denoted Sat.

PROPOSITION 6.6 [HrW09, Proposition 6.5]. Let (G, \mathbb{P}) be relatively hyperbolic. For each $\tau > 0$ there is a constant $\lambda = \lambda(\tau) > 0$ such that the following holds. Let γ and γ' be geodesics in $\text{Cayley}(G, \mathcal{S} \cup \mathcal{P})$, such that the endpoints u_0 and u_1 of γ lie within an \mathcal{S} -distance τ of $\text{Sat}_\tau(\text{Vert}(\gamma'))$. Then each vertex v of γ lies within an \mathcal{S} -distance λ of either u_0, u_1 , or a vertex of γ' .

LEMMA 6.7. Let H be a κ -relatively quasiconvex subgroup of (G, \mathbb{P}) and let $P \in \mathbb{P}$ be a peripheral subgroup. Assume G is finitely generated, and fix a word metric on G . Suppose cosets gP and aH have infinite coarse intersection, meaning that $\mathcal{N}_r(gP) \cap aH$ has infinite diameter for some $r < \infty$. Then $\mathcal{N}_s(gP) \cap aH$ has infinite diameter, where s is a constant depending on H and P but not depending on r, g , and a .

Proof. Choose a cusp uniform action of (G, \mathbb{P}) on a δ -hyperbolic space (\widehat{X}, d) , with truncated space $\widehat{X} - \mathcal{B}$, and basepoint $x \in \widehat{X} - \mathcal{B}$. Although we claim that k is independent using the word metric, we shall actually work with a (pseudo)metric coming from \widehat{X} . Note that the result holds for any proper, left-invariant metric on G if and only if it holds for a particular such metric (although the exact value of k will typically change). For convenience, we will identify G with the orbit Gx and use the induced metric from \widehat{X} .

The maximal parabolic group gPg^{-1} stabilizes $gP(x)$ and has a fixed point $\xi \in \partial\widehat{X}$. Let B be the horoball component of \mathcal{B} centered at ξ , and let $h: \widehat{X} \rightarrow \mathbb{R}$ be an associated horofunction. Let D_0 and D_1 be the constants given by the definition of horofunction and horoball. Let S denote the ‘horosphere’ obtained as the topological boundary of B . Without loss of generality, we can assume that B and h are invariant under the action of gPg^{-1} . In other words, the horofunction h is constant on the orbit $gP(x)$, and this orbit lies in some horoball B' centered at ξ and containing B .

Choose a sequence of distinct points $ah_i(x)$ in $\mathcal{N}_r(gP(x)) \cap aH(x)$. Relative quasiconvexity of H , implies that a geodesic between any two of these points lies in $\mathcal{N}_\kappa(aH(x)) \cup \mathcal{B}$. On the other hand, if $d(ah_i(x), ah_j(x))$ is sufficiently large, then a geodesic γ joining them is forced to penetrate the horoball B by Lemma 6.1. This yields a point of the horosphere S that lies within a distance κ of $aH(x)$. Since gPg^{-1} acts cocompactly on both $gP(x)$ and S , the Hausdorff distance between $gP(x)$ and S is finite, say η . Thus $\mathcal{N}_\eta(gP) \cap \mathcal{N}_\kappa(aH(x))$ is nonempty, since it has a nontrivial intersection with S . Observe that η is a uniform constant which we could have chosen at the beginning relating the geometry of horoballs to their corresponding peripheral cosets; in particular η is independent of our choice of coset gP and even independent of H .

The first and last points y and z of $\gamma \cap \overline{B}$ lie within a distance M of $ah_i(x)$ and $ah_j(x)$ respectively, where M is the constant given by Lemma 6.1. We thus find that the diameter of $\mathcal{N}_\eta(gP) \cap \mathcal{N}_\kappa(aH)$ is infinite, as it exceeds $d(ah_i(x), ah_j(x)) - 2M$. The result therefore follows with $s = \eta + \kappa$. □

7. Relative cocompactness

7.1 The hyperbolic case

In [Sag97], Sageev proved the following result. We note that his actual statement focuses on the case where each H_i is a codimension-1 subgroup and he produced his walls as described in §2.5.

THEOREM 7.1. *Let H_1, \dots, H_k be quasiconvex subgroups of a word-hyperbolic group G . Choose an H_i -wall for each i . Then G acts cocompactly on the associated dual cube complex.*

As a prelude to the work in this section, let us examine Sageev’s theorem in more detail. The following lemma is an easy consequence of Corollary 3.13.

LEMMA 7.2. *Let G act on a wallspace (X, \mathcal{W}) . Suppose there are finitely many orbits of collections of pairwise transverse walls in X . Then G acts cocompactly on $C(X, \mathcal{W})$.*

Note that the above condition implies in particular that there are only finitely many orbits of walls. Lemma 7.2 applies naturally when X is δ -hyperbolic and walls are quasiconvex. Note that pairwise transverse walls correspond to D -close cosets. The hypothesis of Lemma 7.2 follows from the following lemma, which is implicit in Sageev [Sag97] where it is deduced from results of Gitik–Mitra–Rips–Sageev [GMRS98].

LEMMA 7.3. *Let \mathcal{H} be a finite collection of κ -quasiconvex subgroups of the δ -hyperbolic group G with Cayley graph $\Gamma = \Gamma(G, S)$. For each $D \geq 0$, there exists $M = M(D, \delta, \kappa)$ such that the following holds. Let \mathcal{A} be any collection of left cosets aH with $a \in G$ and $H \in \mathcal{H}$ such that all pairs of cosets $aH, a'H' \in \mathcal{A}$ are D -close. Then the following intersection is nonempty:*

$$\bigcap_{aH \in \mathcal{A}} \mathcal{N}_M(aH) \neq \emptyset.$$

As clarified by Niblo–Reeves [NR03], a proof of Lemma 7.3 involves two steps. Firstly, for any collection of n distinct D -close κ -quasiconvex subspaces of a δ -hyperbolic space there is a point g that lies within a distance $M(n)$ of each of these spaces. Secondly, there is an upper bound N on the maximal number of D -close left cosets $g_j H_i$.

It was this second step that motivated us to introduce and explore the notion of bounded packing described in Definition 3.26.

7.2 Various hemiwalls induced by a peripheral subspace

DEFINITION 7.4 (Hemiwalls induced by a subspace). Let (X, \mathcal{V}) be a wallspace, where \mathcal{V} denotes the collection of halfspaces with pairing ι arising from a collection of walls \mathcal{W} . For a nonempty subset $P \subset X$, we list below several natural ways of inducing a hemiwallspace, the first of which was already seen in Example 3.20. The other ways presume that X is a metric space, and P is an infinite diameter subspace. Indeed, if $\text{diam}(P) < \infty$ then items (3) to (5) do not lead to a hemiwallspace, as there is a violation of the requirement that at least one halfspace of each wall be included:

- (1) $\mathcal{U}_0 = \{U \in \mathcal{V} \mid U \cap P \neq \emptyset\}$;
- (2) $\mathcal{U}_r = \{U \in \mathcal{V} \mid U \cap \mathcal{N}_r(P) \neq \emptyset\}$;
- (3) $\mathcal{U}_\infty = \{U \in \mathcal{V} \mid \text{diam}(U \cap P) = \infty\}$;




FIGURE 9. A wallspace (X, \mathcal{V}) and a subset Y . The left-hand side shows the halfspaces of \mathcal{U}_0 . In the center, we see the halfspaces of \mathcal{U}_∞ . The right-hand side illustrates \mathcal{U}_* .

- (4) $\mathcal{U}_* = \{U \in \mathcal{V} \mid \text{diam}(U \cap \mathcal{N}_r(P)) = \infty \text{ for some } r > 0\}$;
 (5) $\mathcal{U}_{r*} = \{U \in \mathcal{V} \mid \text{diam}(U \cap \mathcal{N}_r(P)) = \infty\}$.

Here and below we shall use the symbol \mathcal{U}_\heartsuit for a chosen hemiwallspace from the above list. Several of these hemiwallspaces are illustrated in Figure 9. In each case it is easily verified that for each halfspace in \mathcal{V} , either it, its complement, or both lie in \mathcal{U}_\heartsuit . Note that if a group J acts on (X, \mathcal{V}) and stabilizes P , then J acts on \mathcal{U}_\heartsuit as well.

Although \mathcal{U}_r is similar to \mathcal{U}_0 , it will be appropriate for applications, since sometimes P needs to be ‘thickened up’ to capture the properties of its interaction with walls. In addition, \mathcal{U}_* is a compromise that is appropriate for an application towards relatively hyperbolic groups we shall give later.

When $\mathcal{N}_r(P)$ is a finite set, the axioms imply that each \mathcal{U}_r has finitely many independent walls, and hence $C(\mathcal{U}_r)$ is finite. As mentioned above, we do not consider \mathcal{U}_∞ in this case, as it contains *no* halfspaces and is not a hemiwallspace. However, when $\text{diam}(P) = \infty$ each of \mathcal{U}_∞ and \mathcal{U}_* and \mathcal{U}_{r*} is a hemiwallspace. Indeed, just as either $U \cap P$ or $(X - U) \cap P$ is nonempty, when $\text{diam}(P) = \infty$ either $U \cap P$ or $(X - U) \cap P$ has infinite diameter.

Remark 7.5. Another natural way to associate a wallspace to a subspace $P \subset X$, where X is itself a wallspace, is to consider the collection of all partitions of P that arise from partitions of X . Namely, for each wall $(U, X - U)$ we obtain a wall $(U \cap P, P - U)$. This is the induced *subwallspace* of Definition 2.5. Subwallspaces and hemiwallspaces play different roles. We will return to subwallspaces in §8.

There are several subtleties. A minor issue is that several different partitions of X might induce the same partition of P . In particular, many partitions of X could induce the vacuous partition $\{P, \emptyset\}$ of P . There are ways to disregard this nuisance. A more serious issue is that we would like to have a natural way of embedding the cube complex dual to the P wallspace in the cube complex dual to the X wallspace. This approach would require some additional work to suggest such an embedding. The most important pitfall to this approach is that walls that are transverse in X might induce walls that are not transverse in P . If one were to use the subwallspace approach and aim for an embedding of the dual cube complexes, then one would have to declare walls of P to cross by decree. It is for this reason that we have adopted the approach of taking a hemiwallspace, consisting of a subcollection of halfspaces as it appears to naturally avoid all of these difficulties.

7.3 Structure of cubulation with peripheral subspaces

We will now examine the structure of a cube complex dual to a wallspace X that has a system \mathcal{P} of peripheral subspaces. The strategy is to understand the structure of the cube complex $C(X)$ dual to (X, \mathcal{V}) by recognizing that important parts of it are convex subcomplexes $C(P)$ where $P \in \mathcal{P}$. We call elements of \mathcal{P} *peripheral subspaces* or *peripheries* because of our later relatively hyperbolic applications, but the initial results do not employ relative hyperbolicity.

Technically speaking we should use the notation $C(X, \mathcal{V})$ for the cube complex dual to (X, \mathcal{V}) and $C(X, \mathcal{U}_\heartsuit(P))$ for the cube complex dual to the hemiwallspace induced by some $P \in \mathcal{P}$. We have elected to adopt the notation $C(X)$ for the former and $C_\heartsuit(P)$ for the latter, as it highlights the method of inducing hemiwallspace.

Recall that each n -cube in $C(X)$ corresponds to a collection of unpaired halfspaces, together with n walls, which are paired halfspaces, such that any orientation of these walls yields a 0-cube, as explained in § 3.3.

Having chosen a way of inducing a hemiwallspace \mathcal{U}_\heartsuit from a subspace P , we say a cube c is *represented* in a peripheral subspace P if all these halfspaces lie in \mathcal{U}_\heartsuit . Technically c is a ‘subhemiwallspace’; yikes!

The *depth* of a cube $c \in C(X)$ is its distance to a nearest canonical cube.

We say the peripheries in \mathcal{P} are *isolated* with respect to \mathcal{V} if there exists m such that each cube of depth $\geq m$ is represented in a unique periphery.

It is an interesting problem to find geometric hypotheses on \mathcal{P} that imply isolation. Our main results in the relatively hyperbolic setting depend upon the fundamental point that the coarse geometric isolation of the peripheral subspaces \mathcal{P} implies their isolation with respect to \mathcal{V} .

The following tautology sets a framework for our main results. In its strongest form, it concludes that G acts relatively cocompactly on $C(X)$.

PROPOSITION 7.6. *Let (X, \mathcal{V}) be a wallspace with a collection \mathcal{P} of peripheral subspaces. Suppose G acts on X stabilizing \mathcal{V} and \mathcal{P} . Let \mathbb{P} be a collection of representatives of the G -orbits of peripheries in \mathcal{P} . Let \mathcal{U}_\heartsuit be one of the five ways of inducing hemiwallspace described in Definition 7.4.*

Suppose there exists m such that each (maximal) cube of depth $\geq m$ is represented in some periphery. Then $C(X) = GK \cup \bigcup_{P \in \mathbb{P}} GC_\heartsuit(P)$ where K is a collection of depth $\leq m$ cubes.

Suppose the family of all peripheries gP is isolated with respect to \mathcal{V} . Then $gC_\heartsuit(P) \cap g'C_\heartsuit(P') \subset GK$ unless $gP = g'P'$.

Proof. Let K denote a collection of cubes representing all the G -orbits of cubes of depth $\leq m$. By hypothesis, each cube c either has depth $\leq m$ so lies in gK for some $g \in G$, or c has depth $> m$ and so lies in $gC_\heartsuit(P)$ for some $g \in G$ and $P \in \mathbb{P}$.

If depth $> m$ cubes are represented in unique peripheries, then each intersection $gC_\heartsuit(P) \cap g'C_\heartsuit(P')$ consists of cubes of depth $\leq m$ and hence $gC_\heartsuit(P) \cap g'C_\heartsuit(P') \subset GK$, unless $gP = g'P'$. □

7.4 Application towards relatively hyperbolic groups

The following result is obtained in [HrW09].

THEOREM 7.7 [HrW09, Theorem 8.9]. *Let (G, \mathbb{P}) be relatively hyperbolic, choose a finite generating set \mathcal{S} for G , and let \mathcal{H} be a finite collection of κ -relatively quasiconvex subgroups of G . For each constant D , there are constants R and M such that the following holds. Let \mathcal{A} be any set of left cosets aH with $a \in G$ and $H \in \mathcal{H}$ such that for all $aH, a'H' \in \mathcal{A}$ we have*

$d_{\mathcal{S}}(aH, a'H') < D$. Suppose the following intersection is empty:

$$\bigcap_{aH \in \mathcal{A}} \mathcal{N}_M(aH) = \emptyset.$$

Then we have the following:

- (1) there is a unique peripheral coset bP such that for all $aH \in \mathcal{A}$ the intersection $\mathcal{N}_R(bP) \cap \mathcal{N}_R(aH)$ is nonempty;
- (2) $\mathcal{N}_R(bP) \cap \mathcal{N}_R(aH) \cap \mathcal{N}_R(a'H')$ is nonempty for all $aH, a'H' \in \mathcal{A}$;
- (3) $\mathcal{N}_R(bP) \cap \mathcal{N}_R(aH)$ is unbounded for all $aH \in \mathcal{A}$;
- (4) $bPb^{-1} \cap aHa^{-1}$ is infinite for all $aH \in \mathcal{A}$.

The next result follows from the proof of [HrW09, Proposition 7.8]. Specifically, for the transition point z that was proven to lie in $\bigcap_{A \in \mathcal{A}} \mathcal{N}_{M_1}(A)$, it was also shown in the proof that $z \in \mathcal{N}_{M_1}(g_0P_0)$.

PROPOSITION 7.8 [HrW09, Proposition 7.8]. *Let (G, \mathbb{P}) be relatively hyperbolic with a finite generating set \mathcal{S} . For each $M > 0$ and $\kappa > 0$ there exists $M_1 = M_1(M, \kappa)$ such that the following holds. Suppose \mathcal{A} is a collection of κ -relatively quasiconvex subsets and $g_0P_0 \neq g_1P_1$ are distinct peripheral cosets. If $d_{\mathcal{S}}(g_iP_i, A) < M$ for all $i \in \{0, 1\}$ and for all $A \in \mathcal{A}$, then the following intersection is nonempty:*

$$\mathcal{N}_{M_1}(g_0P_0) \cap \bigcap_{A \in \mathcal{A}} \mathcal{N}_{M_1}(A).$$

Let G be a finitely generated group with Cayley graph Γ . The *frontier* of a subset $S \subseteq G$ is the set of vertices of S that are adjacent to a vertex of $G - S$ in Γ . When S is an H -invariant subset of a subgroup H , we say the frontier of S is H -finite if it consists of finitely many H -orbits. While the frontier of S obviously depends on Γ , its H -finiteness does not. Indeed let Γ, Γ' be two Cayley graphs of G with respect to finite generating sets A, A' . An edge in Γ between S and $G - S$ yields a finite path in Γ' whose length is bounded above by a constant D . Some vertex on that path must be in the frontier of S with respect to Γ' . Therefore every vertex v in the frontier with respect to Γ lies within a distance D in Γ' of a vertex v' in the frontier with respect to Γ' .

LEMMA 7.9. *Let (G, \mathbb{P}) be relatively hyperbolic with a finite generating set \mathcal{S} , and let H be a subgroup that is relatively quasiconvex and finitely generated. Let U be an H -invariant subset whose frontier F is H -finite. Then U is relatively quasiconvex.*

Proof. Finite neighborhoods in G of relatively quasiconvex subspaces are relatively quasiconvex. Therefore the frontier F of U is relatively quasiconvex. But any set with a frontier that is relatively quasiconvex is itself relatively quasiconvex. Indeed, it suffices to consider the case when \mathcal{S} is a finite generating set for G that contains generating sets for each $P \in \mathbb{P}$. Let $\mathcal{P} = \bigcup_{P \in \mathbb{P}} P - \{e\}$. Let γ be a geodesic in $\text{Cayley}(G, \mathcal{S} \cup \mathcal{P})$ whose endpoints u_0, u_1 lie in U but whose internal vertices are not in U . If the first edge of γ is labeled by an element of \mathcal{S} then $u_0 = f_0$ for some $f_0 \in F$. If the first edge e of γ is labeled by an element $p_0 \in P_0$ for some $P_0 \in \mathbb{P}$ then e connects u_0 with an element $u_0p_0 \in G - U$. In this case, the coset u_0P_0 intersects F since it intersects both U and $G - U$; i.e., $u_0 \in f_0P_0$ for some $f_0 \in F$. Similarly either $u_1 = f_1$ or $u_1 \in f_1P_1$ for some $f_1 \in F$ and $P_1 \in \mathcal{P}$.

Let γ' be the geodesic in $\text{Cayley}(G, \mathcal{S} \cup \mathcal{P})$ from f_0 to f_1 . By Proposition 6.6 each vertex of γ is within an \mathcal{S} -distance λ of $\{u_0, u_1\} \cup \text{Vert}(\gamma')$. Since F is relatively quasiconvex, each vertex of γ' is within an \mathcal{S} -distance κ of F , which is a subset of U . Thus each vertex of γ is within an \mathcal{S} -distance $\lambda + \kappa$ of U . □

We now state and prove the main result of this paper.

THEOREM 7.10 (Relative cocompactness for relatively hyperbolic groups). *Let (G, \mathbb{P}) be relatively hyperbolic. Let H_1, \dots, H_r be a collection of relatively quasiconvex finitely generated subgroups. Suppose an H_i -wall $\{U_i, V_i\}$ is chosen for each i . Let (G, \mathcal{V}) be the associated wallspace.*

Let $\mathcal{P} := \{gP \mid g \in G, P \in \mathbb{P}\}$. Let $C = C(G)$ be the cube complex dual to (G, \mathcal{V}) . Then we have the following conclusions.

(1) *There exists a compact K_* such that $C = GK_* \cup \bigcup_P GC_*(P)$ and $gC_*(P) \cap g'C_*(P')$ is contained in GK_* unless $gP = g'P'$.*

(2) *There exists r_0 such that for any $r \geq r_0$ there exists a compact K_r such that $C = GK_r \cup \bigcup_P GC_r(P)$ and $gC_r(P) \cap g'C_r(P')$ is contained in GK_r unless $gP = g'P'$.*

(3) *There exists r_0 such that for any $r \geq r_0$ there exists a compact K_{r*} such that $C = GK_{r*} \cup \bigcup_P GC_{r*}(P)$ and $gC_{r*}(P) \cap g'C_{r*}(P')$ is contained in GK_{r*} unless $gP = g'P'$.*

Note that the three conclusions of the theorem are nearly identical except for the method of inducing hemiwalls from peripheral subgroups. The key point enabling these conclusions is that the resulting peripheral cubulations should be large enough to nearly cover all of C . Shifting the method of inducing hemiwalls only really affects the size of the compact subcomplex K that is necessary for isolation and covering the rest of C . Recall our convention that \mathcal{P} is a collection of *infinite* subgroups.

Remark 7.11. A special case of relative cocompactness is that $C(G)$ is ‘relatively finite dimensional’. In other words, if each $C_{\heartsuit}(P)$ is finite dimensional then so is $C(G)$.

Proof of Theorem 7.10. Let $\mathcal{H} = (H_1, \dots, H_r)$, let $\overline{\mathcal{H}} = (\overline{H}_1, \dots, \overline{H}_r)$, and note that we are viewing this as an ordered list, which could have duplicates.

Consider the map from walls to cosets defined by $\{gU_i, gV_i\}_i \mapsto g\overline{H}_i$. This is well defined since $\{g\overline{h}_iU_i, g\overline{h}_iV_i\}_i$ maps to $g\overline{h}_i\overline{H}_i = g\overline{H}_i$. If we view these as ‘indexed cosets’ then this map is actually a bijection. We shall now forget the indexing, and view two indexed cosets as equal if they are identical subsets of G .

The map from walls gW_i to associated cosets $g\overline{H}_i$ is finite-to-one. Indeed if $g\overline{H}_i = g'\overline{H}_j$, then $\overline{H}_i = \overline{H}_j$. The finiteness of $\overline{\mathcal{H}}$ gives an upper bound r on the number of walls mapping to a given coset. We emphasize that commensurable subgroups H_i and H_j could lead to walls with the same underlying halfspace pairs, depending upon our choices. This is the source of the noninjectivity of the map from walls to cosets.

By Lemma 2.8 there is a constant D such that the following holds: if $W_i = \{U_i, V_i\}$ and $W_j = \{U_j, V_j\}$ and the walls gW_i and $g'W_j$ are transverse, then $d(g\overline{H}_i, g'\overline{H}_j) < D$. Henceforth we use the notation $g\overline{H}$ to indicate a coset, and drop the subscripts which no longer play an indexing role.

A subtle aspect of the argument is that C could be infinite dimensional. Let c be a cube of C and let d be a 0-cube of c . By Zorn’s lemma the collection of hyperplanes dual to 1-cubes of c extends to a maximal collection of pairwise crossing hyperplanes dual to 1-cubes at d . This maximal collection is finite when c lies in a maximal cube of C .

In the first part of the argument we examine the maximal cubes of C . Some of the maximal cubes will be contained in GK ; others will be contained in some $gC_{\heartsuit}(P)$. Each cube that is not contained in a maximal cube will be shown to lie in some $gC_{\heartsuit}(P)$ as well.

By Proposition 3.14, maximal cubes c in C are in one-to-one correspondence with finite cardinality maximal collections of pairwise transverse walls. We have already indicated a finite-to-one map from collections of walls to collections of cosets $g\overline{H}$ with $g \in G$ and $\overline{H} \in \overline{\mathcal{H}}$. Moreover

a collection of pairwise transverse walls leads to a collection \mathcal{A} of pairwise D -close cosets. The dimension of c is equal to the cardinality of \mathcal{A} .

Each \overline{H} is relatively quasiconvex since its finite index subgroup H is relatively quasiconvex. Let R and M be the constants of Theorem 7.7. Suppose there exists

$$x \in \bigcap_{a\overline{H} \in \mathcal{A}} \mathcal{N}_M(a\overline{H}).$$

Translating c by the element x^{-1} gives a cube $x^{-1}c$ whose corresponding cosets $x^{-1}a\overline{H}$ all intersect the ball of radius M centered at the identity. There are only finitely many such cosets counted with multiplicities, since each coset occurs with multiplicity at most r and the ball is a finite set.

On the other hand, if $\bigcap_{a\overline{H} \in \mathcal{A}} \mathcal{N}_M(a\overline{H})$ is empty, then Theorem 7.7 gives a (unique) peripheral coset bP such that for all $a\overline{H} \in \mathcal{A}$ the intersection $\mathcal{N}_R(bP) \cap \mathcal{N}_R(a\overline{H})$ is nonempty. Furthermore, $\mathcal{N}_R(bP) \cap \mathcal{N}_R(a\overline{H})$ is unbounded for all $a\overline{H} \in \mathcal{A}$.

Therefore, if aW is a wall associated with $a\overline{H}$, then both halfspaces of aW have infinite intersection with the S -neighborhood of bP , for some $S = S(R, \overline{H}, W)$. Indeed $\mathcal{N}_R(a\overline{H})$ is contained in an $(S - R)$ -neighborhood of each halfspace of aW for some S . By maximality of the collection of pairwise transverse walls, each dependent wall W' of c is not transverse with some independent wall aW of c . The halfspace corresponding to the c -orientation of W' contains one of the halfspaces of aW , which has infinite intersection with $\mathcal{N}_S(bP)$. In particular, c is in $C_{r*}(bP)$ for each $r \geq S$. Let $r_0 = \max\{S(R, \overline{H}, W)\}$ where the maximum is taken over the finitely many choices for \overline{H} and W .

Suppose now that c does not lie in a maximal cube of C . As mentioned above, the independent walls of c extend to an infinite cardinality maximal collection of pairwise transverse walls dual to 1-cubes at a 0-cube d of c . As above, this family of walls leads to a collection \mathcal{A} of pairwise D -close cosets. Any ball of radius M can intersect only finitely many such cosets. Therefore Theorem 7.7 gives a (unique) peripheral coset bP such that for all $a\overline{H} \in \mathcal{A}$ the intersection $\mathcal{N}_R(bP) \cap \mathcal{N}_R(a\overline{H})$ is nonempty. Again we conclude that $\mathcal{N}_R(bP) \cap \mathcal{N}_R(a\overline{H})$ is unbounded for all $a\overline{H} \in \mathcal{A}$. We conclude that c is in $C_{r*}(bP)$ for each $r \geq S$ as before.

Let us first verify the almost uniqueness of peripheral representations of cubes in the C_r case (as well as the C_{r*} case). This almost uniqueness holds for any $r > 0$. Indeed suppose c is a cube (not necessarily maximal) represented in two peripheries. Thinking of c as a collection of halfspaces, we see that each of its halfspaces intersects the r -neighborhoods of these two peripheries. Observe that \ddot{H} is finitely generated and relatively quasiconvex since H is. The halfspaces of c are κ -relatively quasiconvex for some κ by Lemma 7.9 applied to (U, \ddot{H}) and (V, \ddot{H}) . Thus by Proposition 7.8 there exists $g \in G$ such that each halfspace comes within a distance $M_1 = M_1(r, \kappa)$ of g . As before, after translating c by g^{-1} we can assume that every halfspace in c intersects the ball of radius M_1 centered at the identity. There are only finitely many possibilities for such a cube c , since the finitely many walls intersecting this ball can be oriented in only finitely many ways, and the remaining walls must all be oriented towards the identity. Thus the original cube c lies in one of finitely many orbits.

To verify almost uniqueness in the C_* case we observe that by Lemma 6.7 there exists s such that whenever $g\overline{H}, g'P$ have infinite coarse intersection then the distance between $g\overline{H}$ and $g'P$ is at most s . It follows that each $C_*(gP)$ is contained in $C_s(gP)$ and therefore almost uniqueness follows from the above argument. \square

7.5 The structure of the dual when G acts properly and cocompactly on X

The previous subsection focused on the ‘pure case’ where G is itself a wallspace. The main result of this subsection is Theorem 7.12, which applies Theorem 7.10 to the more general setting of G acting on a general wallspace (X, \mathcal{W}) . Theorem 7.12 provides additional information about the structure of $C(X)$ by explaining that each $C_\heartsuit(Y)$ is attached to GK along PK .

The stabilizer H of a wall $W = \{U, V\}$ acts cocompactly on W if H acts cocompactly on the intersection of closed neighborhoods $\overline{\mathcal{N}}_r(U) \cap \overline{\mathcal{N}}_r(V)$ for all r .

THEOREM 7.12. *Let (X, \mathcal{W}) be a wallspace such that X is also a length metric space. Suppose G acts properly and cocompactly on X preserving both its metric and wallspace structures, and the action on \mathcal{W} has only finitely many G -orbits of walls. Suppose G is hyperbolic relative to \mathbb{P} , and each $H = \text{Stab}(W)$ acts cocompactly on W and H is relatively quasiconvex. Choose $\heartsuit \in \{r, *, r*\}$. For each $P \in \mathbb{P}$ let $Y = Y(P) \subset X$ be a nonempty P -invariant P -cocompact subspace.*

Then there exists a compact K with GK connected such that $C(X)$ has the following structure:

- (1) $C(X) = GK \cup \bigcup_P GC_\heartsuit(Y)$;
- (2) $gC_\heartsuit(Y) \cap g'C_\heartsuit(Y')$ is contained in GK unless $gP = g'P'$;
- (3) $C_r(Y) \cap GK = PK'$ for some compact K' . In particular, $C_r(Y) \cap GK$ lies in a finite neighborhood of PK within GK with respect to the path metric on the 1-skeleton GK^1 . The same conclusion holds for $C_*(Y) \cap GK$.

Proof. Our strategy will be to use (X, \mathcal{W}) to produce a wallspace $(G, \overline{\mathcal{W}})$ such that there is a one-to-one correspondence between \mathcal{W} and $\overline{\mathcal{W}}$, and transverse walls of \mathcal{W} map to transverse walls of $\overline{\mathcal{W}}$.

Let $\{\{U_i, V_i\}\}$ be a set of representatives of the finitely many G -orbits of walls in \mathcal{W} . Let H_i be the stabilizer of $\{U_i, V_i\}$.

Let $x \in X$, and choose D such that $X = \mathcal{N}_D(Gx)$. Let $\phi: G \rightarrow X$ be the map $g \mapsto gx$. For each wall $W_i = \{U_i, V_i\}$ in X we obtain an H_i -wall $\overline{W}_i = \{\overline{U}_i, \overline{V}_i\} = \{\phi^{-1}(\mathcal{N}_D(U_i)), \phi^{-1}(\mathcal{N}_D(V_i))\}$ in G . As usual this extends equivariantly to a collection $\overline{\mathcal{W}}$ of walls in G .

Note that $\overline{\mathcal{W}}$ has no walls that are genuine partitions, with the possible exception of vacuous walls $\{\emptyset, G\}$ arising from $\{\emptyset, X\}$. A reader uncomfortable with duplicate walls can avoid them by varying the ‘thickening’ parameter D to D_i dependent on H_i , so that distinct H_i -walls will be distinguished by varying amounts of overlaps (i.e., betwixtness).

Transverse walls in \mathcal{W}_X map to transverse walls in \mathcal{W}_G . More specifically, any pair of intersecting halfspaces of X maps to a pair of intersecting halfspaces of G . Indeed, if $p \in U \cap U'$ in X , then there exists g such that $d(gp, p) < D$, so $gp \in \mathcal{N}_D(U) \cap \mathcal{N}_D(U')$. Thus $g \in \overline{U} \cap \overline{U}'$.

Observe that there is a G -equivariant cubical map from $C(X)$ to $C(G)$. Indeed since intersecting halfspaces map to intersecting halfspaces, a 0-cube in $C(X)$ naturally maps to a 0-cube in $C(G)$. Since transverse walls map to transverse walls, this embedding extends naturally to the higher dimensional cubes.

The embedding $C(X) \rightarrow C(G)$ is an isometric embedding. Indeed, the distance between two 0-cubes of $C(X)$ equals the number of walls that they orient differently. These map to 0-cubes of $C(G)$ that orient differently the corresponding walls (and no others).

Since G acts properly, cocompactly on X , it is generated by a finite set A . We let Γ denote the Cayley graph of G with respect to A and use the word metric on G . Furthermore, there are constants α, β so that for all $g, g' \in G$:

$$\frac{1}{\alpha} d_G(g, g') - \beta \leq d_X(gx, g'x) \leq \alpha d_G(g, g') + \beta.$$

For each $P \in \mathbb{P}$ let $Y = Y(P) \subset X$ be a periphery stabilized by P . For each choice of $\heartsuit \in \{r, *, r*\}$ we will make a corresponding choice $\overline{\heartsuit} \in \{\bar{r}, *, \bar{r}*\}$. Each hemiwallspace \mathcal{U}_{\heartsuit} maps to a hemiwallspace $\mathcal{U}_{\overline{\heartsuit}}$. There are consequently embeddings $C_{\heartsuit}(Y) \subseteq C_{\overline{\heartsuit}}(P)$.

As G acts cocompactly on $C(G)$ relative to the $C_{\overline{\heartsuit}}(P)$ by Theorem 7.10, G also acts cocompactly on the subspace $C(X)$ relative to the $C_{\overline{\heartsuit}}(Y)$. Indeed, specifically if $C(G) = G\bar{K} \cup \bigcup_P GC_{\overline{\heartsuit}}(P)$ with distinct $gC_{\overline{\heartsuit}}(P)$ intersecting in $G\bar{K}$, then $C(X) = GK \cup \bigcup_Y GC_{\heartsuit}(Y)$, where $K = C(X) \cap \bar{K}$. Notice that the intersection of distinct $gC_{\heartsuit}(Y)$ lies in the intersection of the corresponding distinct $gC_{\overline{\heartsuit}}(P)$, which therefore lies in $C(X) \cap G\bar{K} = GK$.

To ensure that the peripheral hemiwallspaces of X map to peripheral hemiwallspaces of G , we choose D sufficiently large that $Y = Y(P) \subset \mathcal{N}_D(Px)$ for each $P \in \mathbb{P}$. For each $r < \infty$ there exists $\bar{r} = \alpha r + \beta$ so that $\mathcal{U}_r(Y)$ maps to $\mathcal{U}_{\bar{r}}(P)$, since if a halfspace U intersects $\mathcal{N}_r(Y)$ then $\mathcal{N}_D(U)$ necessarily intersects $\mathcal{N}_r(Px)$. Hence $\phi^{-1}(\mathcal{N}_D(U))$ intersects $\mathcal{N}_{\alpha r + \beta}(P)$, so $\bar{U} \cap P \neq \emptyset$. The cases $\heartsuit = *$ and $\heartsuit = r*$ are similar. We leave the details for the reader.

The third claim is proved separately in Lemma 7.13. □

The following two lemmas are within the framework and hypotheses of Theorem 7.12.

LEMMA 7.13 (Footprint of $C(Y)$ is PK). *Suppose G acts cocompactly on the wallspace X . Then $C_r(Y) \cap GK = PK'$ for some compact K' . In particular, if GK is connected, then $C_r(Y) \cap GK$ lies in a finite neighborhood of PK within GK with respect to the path metric on the 1-skeleton GK^1 . The same conclusion holds for $C_*(Y) \cap GK$.*

Proof. Recall that a d -cube u corresponds to a hemiwallspace with d independent walls (that are pairwise transverse). Since K contains finitely many cubes, there exists $B > 0$ such that for each cube u in GK there is a point $z \in X$ such that the B -neighborhood of each halfspace of u contains z . If u also lies in $C_r(Y)$ then all halfspaces of u intersect $\mathcal{N}_r(Y)$.

By Lemma 7.14 there exists a uniform A with the following property: for any collection of halfspaces that intersect $\mathcal{N}_r(Y)$ and also intersect $\mathcal{N}_B(z)$, there exists a point $y \in Y$ such that they all intersect $\mathcal{N}_A(y)$. Consequently, each cube u in $C_r(Y) \cap GK$ lies in some cube complex $C_A(y) = C(\mathcal{U}_0(\mathcal{N}_A(y)))$.

Choose L such that for some $y_0 \in Y$ the P -translates of the ball $\mathcal{N}_L(y_0)$ cover Y . Thus for each $y \in Y$ there exists $p \in P$ so that $\mathcal{N}_A(y) \subset p\mathcal{N}_{L+A}(y_0)$. Moreover, we can assume that L is chosen large enough so that $\mathcal{N}_{L+A}(y_0)$ intersects every halfspace of each cube of K . Observe that $K \subseteq K' = C_{L+A}(y_0)$.

In summary, we have shown that u lies in the subcomplex $C_A(y)$, which is contained in a P -translate of $K' = C_{L+A}(y_0)$. Consequently $C_r(Y) \cap GK$ lies in PK' , which contains only finitely many P -orbits of cubes since K' is finite.

The result also holds for $C_*(Y)$ since $C_*(Y) \subseteq C_r(Y)$ when r equals the constant s of Lemma 6.7. □

LEMMA 7.14 (Nearby point). *Let $\{V_i\}$ denote a collection of halfspaces of X whose stabilizers are uniformly relatively quasiconvex, and let Y denote a peripheral subspace. There exists $A = A(\kappa, \lambda, \mu, X)$ such that: if each $V_i \cap \mathcal{N}_\mu(Y) \neq \emptyset$ and each $V_i \cap \mathcal{N}_\lambda(x) \neq \emptyset$ for some $x \in X$, then there exists $y \in Y$ such that each $V_i \cap \mathcal{N}_A(y) \neq \emptyset$.*

Proof. Consider the quasi-isometry $\phi: G \rightarrow X$ induced by $g \mapsto gx_0$ for some $x_0 \in X$. Observe that Y is coarsely equal to $\phi(P_0)$ for some peripheral subgroup P_0 . Let $U_i = \phi^{-1}(V_i)$. By Lemma 7.9 each U_i is relatively quasiconvex in G . Choose g_1 coarsely equal to $\phi^{-1}(x)$. For consistency with the statement of Proposition 7.8, we add the trivial subgroup, denoted by P_1 , to the peripheral structure of G . Note that $\{g_1\} = g_1P_1$. Since nearness in X provides nearness of preimages in

G , Proposition 7.8 implies that $\{U_i\}$ are uniformly near a point in P_0 . Applying ϕ , we see that $\{V_i\}$ are uniformly near a point $y \in Y$. □

The following corollary to Theorem 7.12 is aimed at situations like finite volume hyperbolic lattices with sufficiently many regular hyperplanes.

COROLLARY 7.15. *Let (G, \mathbb{P}) be relatively hyperbolic. Suppose G acts on a proper δ -hyperbolic space \widehat{X} , each $P \in \mathbb{P}$ stabilizes a quasiconvex subspace $Y = Y(P)$, and $\widehat{X} = GK \cup \bigcup GY$ where K is compact and assume that $gY \cap g'Y'$ is contained in GK unless $gP = g'P'$.*

Suppose $(\widehat{X}, \mathcal{W})$ is a wallspace with a G -action, the halfspaces are quasiconvex, and each wall separates GK . Let Z be a G -invariant G -finite subset of X . Assume that distinct halfspaces of X have distinct intersection with Z . Assume that $Z \cap U_i \cap V_j \neq \emptyset$ whenever two walls $\{U_1, V_1\}$ and $\{U_2, V_2\}$ of X are transverse then all four ‘quarterspaces’ intersect Z ; e.g., $U_1 \cap V_2 \cap Z \neq \emptyset$.

Then all but finitely many orbits of cubes are represented in a unique periphery (that is, $C_\heartsuit(gY)$). Consequently, G acts relatively cocompactly on the dual cube complex $C(\widehat{X})$.

Proof. Remove the cusps of \widehat{X} and apply Theorem 7.12. □

8. Local finiteness and properness in the relatively hyperbolic case

The goal of this section is to show that, within the framework of Theorem 7.12, additional assumptions lead to further finiteness properties of the action of G on $C(X)$. The main result of this section is Theorem 8.13, which can play an important role in supporting the proof of virtual specialness of certain relatively hyperbolic groups. In §8.1 we collect those properties that are consequences of the additional assumption that G acts properly on $C(X)$. These are listed in Proposition 8.1. Section 8.2 provides some background on the geometry of relatively hyperbolic groups. Section 8.3 explains that the Wall–Wall Separation property holds for X provided that various analogous properties hold for each periphery Y . Section 8.4 focuses on finiteness properties of $C(X)$ that are consequences of hypotheses on the nature of the walls in each periphery Y of X . These consequences are stated in Theorem 8.13. The results of §§8.2 and 8.3 are used to support the proof of Theorem 8.13.

8.1 Additional finiteness properties in the presence of a proper action on $C(X)$

We again emphasize that we are working under the hypothesis of Theorem 7.12. In this case when we also know that G acts properly on $C(X)$, i.e., with finite stabilizers, then there are numerous additional finiteness properties, which we collect together in the following statement. These statements are proven separately in this subsection.

PROPOSITION 8.1. *Let $G, X, C(X), P, Y = Y(P), K, \heartsuit, \dots$ be as in Theorem 7.12. Suppose G acts properly on $C(X)$.*

- (1) *For each $P \in \mathbb{P}$ and corresponding Y , the subcomplex $C_*(Y)$ is superconvex in $C(X)$.*
- (2) *Each 0-cube of GK lies in uniformly finitely many $gC(Y_i)$.*
- (3) *If each P acts metrically properly on $C(Y)$ then G acts metrically properly on $C(X)$.*
- (4) *If each subwallspace Y has linear separation (with respect to the subspace metric on Y), then X has linear separation.*
- (5) *If each $C_\heartsuit(Y)$ is locally finite then $C(X)$ is locally finite, and G acts metrically properly on $C(X)$.*

The convex subcomplex $A \subset B$ is *superconvex* if there is a uniform upper bound on the diameter d of any *hanging flap* $[0, 1] \times [0, d] \subset B$ that is a subcomplex not contained in A but with $\{0\} \times [0, d] \subset A$ isometrically embedded.

LEMMA 8.2. *For each $P \in \mathbb{P}$ with periphery Y the subcomplex $C_*(Y)$ is superconvex in $C(X)$.*

Proof. Let K be a compact subcomplex of $C(X)$ satisfying the conclusion of Theorem 7.12. By Lemma 7.13, there exists a compact subcomplex K' such that for each of the finitely many $P \in \mathbb{P}$, we have $C_*(Y) \cap GK \subset PK'$. By hypothesis G acts properly on $C(X)$, so P also acts properly on $C_*(Y) \cap GK$.

By almost malnormality of the peripheral subgroups, there exists an upper bound M on the diameter $P_i K' \cap g P_j K'$ for all $P_i, P_j \in \mathcal{P}$ with $P_i \neq g P_j$. Consequently for a hanging flap F whose base lies on $C_*(Y) \cap GK$, the intersection between F and any $g C_*(Y')$ has diameter at most M . Since $g_i C_*(Y_i) \cap g_j C_*(Y_j) \subset GK$, we deduce that each sufficiently large diameter flap $F = [0, 1] \times [0, d]$ has some internal 1-cube $[0, 1] \times \{e\}$ in GK .

Suppose there were arbitrarily large diameter hanging flaps. Since P acts properly and cocompactly on $C_*(Y) \cap GK$, there is a sequence of hanging flaps whose base lines $\{0\} \times [-d, d]$ in $C_*(Y) \cap GK$ are an increasing sequence with union a line $\{0\} \times \mathbb{R}$. For each flap F_ℓ there exists $0 < e_\ell < M$ such that the internal 1-cube $[0, 1] \times \{e_\ell\}$ of F_ℓ lies in GK . Passing to a subsequence, we can assume $e_\ell = e$ is independent of ℓ . As GK is locally finite, we can pass to a further subsequence so that the 1-cubes $[0, 1] \times \{e\}$ of each F_ℓ coincide. This 1-cube is dual to a hyperplane V of $C(X)$.

We now show $\text{Stab}(V) \cap P$ is infinite. Our infinite sequence of hanging flaps shows that GK contains infinitely many 1-cubes dual to V . Each of these lies in the 1-neighborhood of PK . Since GK is locally finite, the 1-neighborhood of PK in GK is P -cocompact. Consequently infinitely many elements of P map dual 1-cubes of V to dual 1-cubes of V , and thus stabilize V . We conclude that $\text{Stab}(V) \cap P$ is infinite.

Let W denote the wall of X corresponding to V . Since $\text{Stab}(V) = \text{Stab}(W)$, we see that $\text{Stab}(W) \cap P$ is infinite. It follows that both halfspaces of W have unbounded intersection with some $\mathcal{N}_r(Px)$. But then $W \in \mathcal{U}_*$ and hence V intersects C_*P which is a contradiction. \square

COROLLARY 8.3 (Local finiteness of $\{C(Y_i)\}$). *Each 0-cube of GK lies in uniformly finitely many $gC(Y_i)$.*

Proof. Since G acts properly on $C(X)$ it acts properly on GK . There are finitely many peripheral subgroups $P \in \mathbb{P}$, so the collection of subcomplexes of the form gPK' arising in Lemma 7.13 is a locally finite collection. \square

COROLLARY 8.4 (Inheriting metrical properness). *If each P_i acts metrically properly on $C(Y_i)$ then G acts metrically properly on $C(X)$. If each subwallspace Y_i has linear separation (with respect to the subspace metric on Y_i), then X has linear separation.*

Proof. We prove the latter statement. First observe that if x, x' are points of X with corresponding canonical cubes d, d' , then $d_{C(X)}(d, d')$ equals the number of walls of X separating x, x' . Enlarge K slightly to ensure that GK is connected and contains all canonical cubes. Since G acts properly, cocompactly on both GK and on X , we see that X is quasi-isometric to GK with its induced path metric. In order to verify linear separation, it suffices to show that $d_{GK}(d, d') \leq \kappa d_{C(X)}(d, d')$ for some constant κ .

By hypothesis, each subwallspace Y_i has linear separation. As above if y, y' are points of Y_i with corresponding canonical cubes c, c' , then $d_{C(X)}(c, c')$ equals the number of walls of X

separating y, y' , which also equals the number of walls of the subwallspace Y_i separating y, y' . By Lemma 7.13 the footprint $GK \cap C_r(Y_i)$ consists of finitely many P_i -orbits of cubes. Each peripheral subgroup P_i is finitely generated since G is finitely generated [Osi06]. Enlarge K further to ensure that each $GK \cap C_r(Y_i)$ is connected. Since P_i acts properly, cocompactly both on $GK \cap C_r(Y_i)$ and on Y_i , we see that Y_i is quasi-isometric to $GK \cap C_r(Y_i)$ with its induced path metric. By linear separation, there is a constant κ_i such that the distance in the footprint $GK \cap C_r(Y_i)$ is bounded above by κ_i times the distance between c and c' in $C_r(Y_i)$ and hence in $C(X)$.

To complete the proof, consider a geodesic γ in $C(X)$ from d to d' . For each subpath σ of γ in some $C_r(Y_i) - GK$ with endpoints in GK , we replace it by a path σ' in $GK \cap C_r(Y_i)$ with $|\sigma'| \leq \kappa |\sigma|$ for some uniform κ .

To prove the metric properness, we would use functions $f_i: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ instead of the linear factors κ_i . □

COROLLARY 8.5. *If each $C_\heartsuit(Y)$ is locally finite then $C(X)$ is locally finite, and G acts metrically properly on $C(X)$.*

Proof. Since G acts properly on $C(X)$, it also acts properly on GK . Therefore GK is locally finite, as it admits a proper, cocompact group action. If we endow GK with its induced path metric, then Lemma 7.13 implies that each finite ball in GK intersects only finitely many of the subcomplexes $gC_\heartsuit(Y)$. Since each $gC_\heartsuit(Y)$ is itself locally finite, by hypothesis, it follows immediately that $C(X)$ is locally finite.

Finally, a proper action on a locally finite complex is always metrically proper. □

8.2 Background on saturations and relatively thin triangles

In this subsection we consider a relatively hyperbolic group (G, \mathbb{P}) acting properly, cocompactly on a space X . For each $P_i \in \mathbb{P}$ there is an associated P_i -cocompact nonempty subspace Y_i . The peripheries of X are the G -translates of the various Y_i .

The J -saturation of a subspace $A \subseteq X$, denoted $\text{Sat}_J(A)$, is the union of $\mathcal{N}_J(A)$ together with all peripheries gY_i intersecting $\mathcal{N}_J(A)$.

LEMMA 8.6. *Let $H = \text{Stab}(A)$ act cocompactly on A , and suppose H is relatively quasiconvex in G . For all sufficiently large J , the saturation $\text{Sat}_J(A)$ has the following properties:*

- (1) $\text{Sat}_J(A)$ is κ -quasi-isometrically embedded in X ;
- (2) for any δ there exists $M = M(\delta)$ such that $\text{Sat}_J(A)$ is ‘isolated’ in the sense that for any periphery Y , either $\mathcal{N}_\delta(Y) \cap \text{Sat}_J(A)$ has diameter $\leq M$, or $Y \subset \text{Sat}_J(A)$.

Proof. The first claim follows from the quasiconvexity of saturations due to Druţu–Sapir [DS05]. The key property of $\text{Sat}_J(A)$ is that there exists J' such that the intersection of distinct peripheries lies within $\mathcal{N}_{J'}(A)$.

The second claim holds when J exceeds the constant s provided by Lemma 6.7 (after applying an appropriate quasi-isometry). Indeed if $\mathcal{N}_\delta(Y) \cap \text{Sat}_J(A)$ has large diameter, then since there is a uniform bound on the diameters of coarse intersection between distinct peripheries, we see that either Y equals one of the gY' added to form the saturation, or using the key property, $\mathcal{N}_\delta(Y) \cap \mathcal{N}_{J'}(A)$ must have large diameter. But a sufficiently large diameter implies an infinite diameter, which implies that Y comes s -close to A , and hence $Y \subset \text{Sat}_J(A)$. □

The following property of quasigeodesic triangles is a simplification of a result from [DS05, §8.1.3].




FIGURE 10. A quasigeodesic triangle in a relatively hyperbolic group with a peripheral coset at the center.

THEOREM 8.7. *Let (G, \mathbb{P}) be relatively hyperbolic. For each ϵ there is a constant δ such that if Δ is an ϵ -quasigeodesic triangle with sides c_0, c_1 and c_2 , then there is either:*

- (1) a point p such that $\mathcal{N}_{\delta/2}(p)$ intersects each side of Δ ; or
- (2) a peripheral coset gP such that $\mathcal{N}_{\delta}(gP)$ intersects each side of Δ .

In the second case, illustrated in Figure 10, each side c_i of Δ has a subpath c'_i that lies in $\mathcal{N}_{\delta}(gP)$ such that the terminal endpoint of c'_i and the initial endpoint of c'_{i+1} are mutually within a distance δ (indices modulo 3).

LEMMA 8.8. *Let (X, \mathcal{W}) be a wallspace where X is a geodesic metric space. Let (G, \mathbb{P}) be relatively hyperbolic. Suppose G acts properly and cocompactly on X , and each $P \in \mathbb{P}$ acts cocompactly on a chosen connected subspace $Y = Y(P)$. Finally suppose each wall of \mathcal{W} has a relatively quasiconvex stabilizer.*

Let W' be a connected subspace with relatively quasiconvex cocompact stabilizer. Choose J sufficiently large as in Lemma 8.6. Let γ be a geodesic from a point p to $\text{Sat}_J(W')$. Let q' be the first point on γ that lies in $\text{Sat}_J(W')$. There exists $U > 0$ such that each wall V that crosses both pq' and W' must also cross $\mathcal{N}_U(q')$.

We will explain the proof in the special case when X is a geometric wallspace. The general case is almost identical. One applies the same line of reasoning to each closed halfspace of V ; these are also relatively quasiconvex in an appropriate sense generalizing Definition 6.4.

We will also explain the proof in the case where each wall V of X is quasi-isometrically embedded. This happens precisely when the parabolic subgroups of $\text{Stab}(V)$ are quasi-isometrically embedded in the peripheral subgroups of G by [Hru10, Theorem 10.5].

To obtain the proof in general, one would have to consider a quasigeodesic in the saturation of V and observe that the points a and c in Figure 11 are within a bounded distance of V .

Proof of Lemma 8.8. Let s be a point in $V \cap W'$. Let t be a point in $V \cap \gamma$. Consider a quasigeodesic triangle $\Delta(stq')$ whose three sides are in $V, \text{Sat}_J(W')$, and pq' .

There exists M such that $\text{Sat}_J(W')$ is ‘isolated’ in the sense that for any periphery Y , either $\mathcal{N}_{\delta}(Y) \cap \text{Sat}_J(W')$ has diameter $\leq M$, or Y lies entirely within $\text{Sat}_J(W')$.

We will prove the statement for $U = 3\delta + M$. The proof breaks into two cases according to whether the central periphery Y within $\Delta(stq')$ has large or small overlap with $\text{Sat}_J(W')$ in the sense that $d(b, \bar{c}) \leq M$ or $d(b, \bar{c}) \geq M$. (The small overlap case includes the situation where the central periphery consists of a singleton.) We refer the reader to Figure 11.

If $d(b, \bar{c}) \leq M$ then

$$d(V, q') \leq d(c, q') \leq d(c, \bar{c}) + d(\bar{c}, b) + d(b, \bar{b}) + d(\bar{b}, q') \leq \delta + M + \delta + \delta \leq 3\delta + M.$$




FIGURE 11. The triangle $\Delta(stq')$ contains a central periphery Y , whose overlap with $\text{Sat}_J(W')$ is approximated by $d(b, \bar{c})$.

If $d(b, \bar{c}) \geq M$ then $Y \subset \text{Sat}_J(W')$. Thus $d(\bar{a}, q') \leq d(\bar{a}, Y) \leq \delta$ since q' is a closest point to $\text{Sat}_J(W')$. We thus have:

$$d(V, q') \leq d(a, q') \leq d(a, \bar{a}) + d(\bar{a}, q') \leq \delta + \delta. \quad \square$$

8.3 Wall–wall separation and relative hyperbolicity

Let (X, \mathcal{W}) be a wallspace where (X, d) is a metric space. We say (X, \mathcal{W}) has the *Wall–Wall Separation Property* if there exists a constant $D > 0$ such that any two walls W, W' in X with $d(W, W') > D$ are separated by another wall W'' in the sense that they lie in distinct open halfspaces of W'' . The Wall–Wall Separation Property was termed ‘Strong Parallel Walls’ (see [NR03]). We say (X, \mathcal{W}) has the *Ball–Ball Separation Property* if for each r there exists m such that if $d(x_1, x_2) > m$ then $\mathcal{N}_r(x_1)$ and $\mathcal{N}_r(x_2)$ are separated by a wall.

Let Y be a connected subspace of the wallspace X . We say Y has *Ball–WallNbd Separation* if for each $r > 0$ there exists s such that if $q \in Y$ and W is a wall of X and $d(\mathcal{N}_r(q) \cap Y, \mathcal{N}_r(W) \cap Y) > s$, then there exists a wall $W'' \cap Y$ of Y separating the sets $\mathcal{N}_r(q) \cap Y$ and $\mathcal{N}_r(W) \cap Y$.


We say Y has *WallNbd–WallNbd Separation (for osculating walls)* if for each $r > 0$ there exists s such that if W, W' are (osculating) walls of X and $d(\mathcal{N}_r(W) \cap Y, \mathcal{N}_r(W') \cap Y) > s$, then there exists a wall $W'' \cap Y$ of Y separating the sets $\mathcal{N}_r(W) \cap Y$ and $\mathcal{N}_r(W') \cap Y$. Note that when one or both of these sets are empty, then any wall separates them. Note also that all infinite intersections arise for some sufficiently large critical value of r ; specifically, there exists r_0 such that if $\mathcal{N}_r(W) \cap Y$ is infinite for some r then $\mathcal{N}_{r_0}(W) \cap Y$ is infinite.

THEOREM 8.9. *Let G act properly and cocompactly on a wallspace (X, \mathcal{W}) . Suppose (G, \mathbb{P}) is relatively hyperbolic. Suppose for each wall W , the stabilizer H is relatively quasiconvex. For each $P \in \mathbb{P}$ let Y be a P -cocompact connected subspace. Suppose X has Ball–Ball Separation and each Y has WallNbd–WallNbd Separation for osculating walls and Ball–WallNbd Separation.*

Then (X, \mathcal{W}) has Wall–Wall Separation.

COROLLARY 8.10. *Under the hypotheses of Theorem 8.9, if Y has WallNbd–WallNbd Separation for arbitrary walls, it follows that X has WallNbd–WallNbd Separation for arbitrary walls.*

Proof. Given a wallspace (X, \mathcal{W}) , for each $t \geq 0$ there is a new wallspace (X, \mathcal{W}_t) whose walls are $\{\mathcal{N}_t(U), \mathcal{N}_t(V)\}$ for each wall $\{U, V\}$ of \mathcal{W} . In particular, for geometric wallspaces the walls of \mathcal{W}_t have the form $W_t = \mathcal{N}_t(W)$ for each geometric wall $W \in \mathcal{W}$. Observe that X has Ball–Ball Separation and Y has WallNbd–WallNbd Separation and Ball–WallNbd Separation in (X, \mathcal{W}_t) if


 FIGURE 12. In each case, there exists W'' separating W, W' .

the corresponding properties hold in (X, W) . For instance, if $d(\mathcal{N}_{r+2t}(W) \cap Y, \mathcal{N}_{r+2t}(W') \cap Y) > s = s(r + 2t)$, then there exists a wall $W'' \cap Y$ of Y separating the sets $\mathcal{N}_{r+2t}(W_t) \cap Y$ and $\mathcal{N}_{r+2t}(W'_t) \cap Y$, and hence W'' separates $\mathcal{N}_r(W_t) \cap Y$ and $\mathcal{N}_r(W'_t) \cap Y$. We apply Theorem 8.9 to conclude that (X, W_t) has Wall–Wall Separation for every $t \geq 0$. We infer that (X, W) has WallNbd–WallNbd Separation. \square

Proof of Theorem 8.9. Choose J sufficiently large that Lemma 8.6 holds when A is any of the finitely many types of walls of X . We will show that if W, W' osculate then

$$d(W, W') \leq \max\{s + 2r_1, t + 2(u + 2r)\}$$

where the constants are chosen in the cases below.

We first consider two osculating walls W, W' of X with the property that $\text{Sat}_J(W)$ and $\text{Sat}_J(W')$ contain a common periphery Y . By Lemma 8.11 we can choose r_1 such that if some wall W'' intersects Y , and W'' crosses a wall W of X , and W crosses the J -neighborhood of Y , then W'' crosses $\mathcal{N}_{r_1}(W) \cap Y$. Moreover we shall assume that $r_1 > J$ to ensure that $\mathcal{N}_{r_1}(W) \cap Y$ and $\mathcal{N}_{r_1}(W') \cap Y$ are nonempty. Let $s = s(r_1)$ be the WallNbd–WallNbd constant. If $d(W, W') > s + 2r_1$, then $d(\mathcal{N}_{r_1}(W) \cap Y, \mathcal{N}_{r_1}(W') \cap Y) > s$. By hypothesis there exists a wall W'' separating $\mathcal{N}_{r_1}(W) \cap Y$ and $\mathcal{N}_{r_1}(W') \cap Y$, as shown on the left of Figure 12. Since W, W' osculate, W'' must cross either W or W' , say W . But then our choice of r_1 guarantees that W'' crosses $\mathcal{N}_{r_1}(W) \cap Y$, which is impossible.

We now consider the case where $\text{Sat}_J(W)$ and $\text{Sat}_J(W')$ do not contain a common periphery. We may also assume that $\mathcal{N}_J(W) \cap \mathcal{N}_J(W') = \emptyset$, since otherwise $d(W, W') < 2J$ and so the conclusion holds. Consider a shortest geodesic qq' from $\text{Sat}_J(W)$ to $\text{Sat}_J(W')$. Choose $r > \max\{J, r_1, r_2\}$, where r_1, r_2 are the constants from Lemmas 8.11 and 8.12. Let $u = u(r)$ be the Ball–WallNbd constant.

Suppose $d(q, W) > u + 2r$. Since $d(q, W) > J$, we have $q \in Y$ for some periphery $Y \subset \text{Sat}_J(W)$. It follows that $d(\mathcal{N}_r(q) \cap Y, \mathcal{N}_r(W) \cap Y) > u$ and so Ball–WallNbd Separation in Y yields a wall W'' that separates $\mathcal{N}_r(q) \cap Y$ from $\mathcal{N}_r(W) \cap Y$. Since W, W' osculate, W'' must cross either W or W' . But if W'' crosses W then Lemma 8.11 guarantees that W'' crosses $\mathcal{N}_r(W) \cap Y$, which is impossible. Similarly if W'' crosses W' (as shown on the right side of Figure 12) then Lemma 8.12 guarantees that W'' crosses $\mathcal{N}_r(q) \cap Y$, which is impossible. An analogous argument works if $d(q', W') > u + 2r$.

Now suppose $d(q, W)$ and $d(q', W')$ are both $\leq u + 2r$. By Ball–Ball Separation, there exists $t = t(r)$ such that whenever $d(q, q') > t$, there exists a wall W'' that separates $\mathcal{N}_r(q)$ and $\mathcal{N}_r(q')$. If $d(W, W') > t + 2(u + 2r)$, then $d(q, q') > t$. Consequently there exists a wall W'' that separates $\mathcal{N}_r(q)$ and $\mathcal{N}_r(q')$. Hence W'' separates W, W' by Lemma 8.12. This contradicts that W, W' osculate. \square

LEMMA 8.11. *Let Y be a periphery and W, W'' be walls in X . Suppose W'' crosses W , and W'' crosses Y , and $d(W, Y) < J$. Then $\mathcal{N}_{r_1}(W) \cap W'' \cap Y$ is nonempty for $r_1 = r_1(J)$.*

Proof. For simplicity we assume that X is a geometric wallspace. By [HrW09, Proposition 7.5(2)] there exists $A = A(J)$ such that the intersection $\mathcal{N}_A(W) \cap \mathcal{N}_A(W'') \cap Y$ contains a point p . Let p'' denote a point in $W'' \cap \mathcal{N}_A(p)$. For each $q \in W'' \cap Y$ there exists a path $p''q$ from p'' to q in $W'' \cap \mathcal{N}_B(Y)$ where B is a uniform constant. Indeed observe that $\text{Stab}(W'')$ is a finitely generated relatively hyperbolic group and $\text{Stab}(W'') \cap \text{Stab}(Y)$ is one of its peripheral subgroups. There are only finitely many conjugacy classes of such subgroups and each is finitely generated by [Osi06]. Let q denote a point of $W'' \cap Y$ where the path $p''q$ is within 1 of the infimal length. We claim that the length of $p''q$ is uniformly bounded. Indeed there exists $s = s(A)$ so that if $p''q$ is longer than s , a standard argument using the pigeonhole principle and cocompactness allows us to chop out a segment and effectively translate q to a point $q'' \in W'' \cap Y$ connected to p'' by a shorter path $p''q''$. The result follows with $r_1 = s + 2A$. \square

LEMMA 8.12. *Assume $\mathcal{N}_J(W) \cap \mathcal{N}_J(W') = \emptyset$ and $\text{Sat}_J(W) \cap \text{Sat}_J(W')$ does not contain a periphery. There exists $r_2 = r_2(J)$ such that the following holds: let qq' be a shortest geodesic between $\text{Sat}_J(W)$ and $\text{Sat}_J(W')$. Any wall W'' crossing both saturations must also intersect both $\mathcal{N}_{r_2}(q)$ and $\mathcal{N}_{r_2}(q')$. Similarly, any wall W'' crossing both $\text{Sat}_J(W)$ and qq' must intersect $\mathcal{N}_{r_2}(q)$.*

Sketch. We describe the proof in the case when walls are μ -quasi-isometrically embedded. The union $U = \text{Sat}_J(W) \cup qq' \cup \text{Sat}_J(W')$ is uniformly quasiconvex by the methods of [DS05]. Moreover for each $s > 0$ there exists r_2 such that W, W' lie in distinct components of $\mathcal{N}_s(U) - \mathcal{N}_{r_2}(q)$ and also of $\mathcal{N}_s(U) - \mathcal{N}_{r_2}(q')$. If there exists a quasigeodesic γ in W'' from $p \in \text{Sat}_J(W)$ to $p' \in \text{Sat}_J(W')$, then $\gamma \subset \mathcal{N}_s(U)$ so γ must come r_2 -close to both q and q' . Similarly if there exists a quasigeodesic γ' in W'' from $p \in \text{Sat}_J(W)$ to $p'' \in qq'$, then $\gamma' \subset \mathcal{N}_s(U)$ so γ' must come r_2 -close to q (see the middle of Figure 12).

In general, when the walls are μ -relatively quasiconvex, one uses an argument similar to the proof of [Hru10, Theorem 10.5]. \square

8.4 Structure and local finiteness of $C(X)$ in the relatively hyperbolic case with wall-wall separation

The purpose of Theorem 8.13 is to assert that $C(X)$ has very strong hyperplane finiteness properties under the additional assumption of appropriate peripheral finiteness properties. It is most desirable that $C(X)$ be cocompact, but this often fails to hold outside of a hyperbolic setting. In many cases having finitely many orbits of hyperplanes is a suitable generalization of cocompactness, and this is an immediate consequence of having finitely many orbits of walls in X . The conclusion of Theorem 8.13 ensures that there are finitely many orbits of ‘ways’ in which hyperplanes of $C(X)$ can interact, i.e., osculate or cross. This finiteness property enables one to prove virtual specialness in the presence of corresponding double hyperplane coset separability [HW10].

THEOREM 8.13. *Let G act properly, cocompactly on a geometric wallspace (X, \mathcal{W}) . Suppose (G, \mathbb{P}) is relatively hyperbolic. Let W_1, \dots, W_ℓ be representatives of the finitely many G -orbits of walls in X . Let $H_j = \text{Stab}_G(W_j)$, and assume each H_j is finitely generated and relatively quasiconvex. For each $P \in \mathbb{P}$ let $Y = Y(P)$ be a P -cocompact subspace of X . Suppose:*

- (1) *each Y has WallNbd-WallNbd and Compact-WallNbd Separation;*
- (2) *G acts properly on $C(X)$;*
- (3) *$H_j \cap gPg^{-1}$ has bounded packing in gPg^{-1} for each $g \in G, P \in \mathbb{P}$, and each j .*

Then:

- (1) there are finitely many G -orbits of pairs of osculating hyperplanes and also of pairs of transverse hyperplanes in $C(X)$;
- (2) $C(X)$ is uniformly locally finite;
- (3) G acts metrically properly on $C(X)$;
- (4) for each hyperplane L in $C(X)$ there are finitely many $\text{Stab}_G(L)$ -orbits of hyperplanes L' in $C(X)$ that are transverse with or osculate with L ;
- (5) for each hyperplane M in $C_\heartsuit(Y)$ there are finitely many $\text{Stab}_P(M)$ -orbits of hyperplanes M' in $C_\heartsuit(Y)$ that are transverse with or osculate with M . Here $\heartsuit \in \{r, *, r*\}$.

There are two natural settings in which Theorem 8.13 applies: the first is where X is a truncated finite volume hyperbolic 3-manifold; the second is where X is a ‘mixed’ 3-manifold [PW12].

In applications, verifying the hypotheses of Theorem 8.13 requires some care, but is primarily confined to features of the peripheral subgroups. For instance hypothesis (3) holds in the finite volume hyperbolic case since all subgroups of $\mathbb{Z} \times \mathbb{Z}$ have bounded packing. In the mixed case, the peripheral subgroups are fundamental groups of graph manifolds, so the walls must be chosen carefully to ensure hypothesis (3).

We have discussed conditions that imply hypothesis (2) in §5.

The consequences of hypothesis (1) are the focus of §8.3. As a hemiwallspace on Y is somewhat extrinsic and depends on the features of the embedding $Y \rightarrow X$, we decided to couch things in terms of WallNbd–WallNbd and Compact–WallNbd Separation, which depend only on interactions between various subsets of Y .

Proof of Theorem 8.13. We first prove conclusion (1). Since G acts properly on $C(X)$, Lemma 5.3 implies that X has Ball–Ball Separation. By Theorem 8.9, X has Wall–Wall Separation, which gives a uniform upper bound A on the distance in X between any two osculating walls. By Lemma 2.4 any two transverse walls are uniformly close (they intersect). Therefore there are only finitely many G -orbits of osculating and/or transverse pairs of walls.

We now prove conclusion (2). Our hypothesis that each H_j is relatively quasiconvex, together with hypothesis (3), allows us to conclude that each H_j has bounded packing in G by Theorem 3.27. Combining this with conclusion (1), we can apply Theorem 3.30 to conclude that $C(X)$ is uniformly locally finite.

Conclusion (3) follows from conclusion (2) together with our hypothesis that G acts properly on $C(X)$.

To see conclusion (4), as in the proof of conclusion (1) there is a uniform bound A on the distance between pairs of osculating and/or transverse walls of X . Since G acts properly and cocompactly on X , there are only finitely many H_j -orbits of walls intersecting $\mathcal{N}_A(W_j)$.

Let us now verify conclusion (5). Let W denote the wall of X corresponding to M . Any hyperplane M' of $C_\heartsuit(Y)$ that is transverse with or osculates with M corresponds to a wall W' of X that is transverse with or osculates with W . By the definition of $C_\heartsuit(Y)$ there exists a constant $t = t(\heartsuit)$ such that W and W' intersect $\mathcal{N}_t(Y)$. By [HrW09, Proposition 7.5(2)] there exists $s = s(t)$ such that $(\mathcal{N}_s(W) \cap \mathcal{N}_t(Y)) \cap \mathcal{N}_s(W')$ is nonempty. The result follows since $\text{Stab}_P(M)$ acts cocompactly on $\mathcal{N}_s(W) \cap \mathcal{N}_t(Y)$. □

9. CAT(0) truncations

The goal of this section is to explain that the cube complex obtained within the relatively hyperbolic situation studied in §7 can often be ‘truncated’ to obtain a cocompact CAT(0) space.

If G acts on a space B , a subset $A \subset B$ is G -cocompact provided that A contains a compact set D such that $A = GD$. Note that outside of a hyperbolic situation it is not always possible to obtain a G -cocompact convex subcomplex. For instance, let G be any diagonal action of an infinite cyclic group on the square tiling of \mathbb{E}^2 . Then there is no G -cocompact convex subcomplex but there is a G -cocompact subspace that is convex with respect to the CAT(0) metric. The possibility of obtaining G -cocompact subspaces is poorly understood.

DEFINITION 9.1. The G action on X has the *convex core property* if each compact subset $D \subset X$ lies in a closed, convex, G -cocompact subset $E \subset X$.

Example 9.2. If G is virtually \mathbb{Z}^n and X is a proper CAT(0) space with a semi-simple G -action then (G, X) has the convex core property.

Indeed, let $F \subset X$ be a flat such that $GF = F$ and F is G -cocompact. Then any closed neighborhood $\mathcal{N}_r(F)$ has the properties that:

- (1) $\mathcal{N}_r(F)$ is a convex subspace of X ;
- (2) $\mathcal{N}_r(F)$ is G -stable and G -cocompact (since X is proper and $\mathcal{N}_r(F)$ is cobounded).

There are many examples of groups G that act properly and cocompactly on a CAT(0) cube complex X , such that G contains a subgroup H that does not act properly and cocompactly on a CAT(0) space. For instance, H might not have a quadratic isoperimetric function, or H might not have finite homotopy type. In this case the H -action on X cannot have the convex core property. Typical free subgroups of $F_2 \times \mathbb{Z}$ do not have the convex core property for the action on the Cartesian product of a tree and a line.

THEOREM 9.3. *Let G be a finitely generated group acting properly on the complete CAT(0) space C . Let $\{C_i \mid i \in I\}$ be a collection of closed convex subspaces of C , and let $G_i = \text{Stab}(C_i)$ for each i . Assume that the family of translates $\{gC_i\}$ is locally finite, so any finite ball intersects only finitely many of the translates.*

Suppose that $C = GK \cup \bigcup_i GC_i$ for some compact $K \subset C$. And suppose that $g_i C_i \cap g_j C_j \subset GK$ unless $C_i = C_j$ and $g_j^{-1}g_i \in G_i$.

Then (G, C) has the convex core property provided that (G_i, C_i) has the convex core property for each i .

Proof. Since G is finitely generated, we can assume that the stated compact set K is large enough that GK is also connected.

Let D be a compact subspace of C . For each i , let $D_i = C_i \cap (GK \cup GD)$, and note that D_i is G_i -cocompact.

Indeed if $A \subset C$ is compact then $C_i \cap GA$ is clearly G -cocompact. We will show that it is also G_i -cocompact. Let $\{g_j A \mid j \in J\}$ be a minimal family of G -translates of A that covers $C_i \cap GA$. Consider the family of subspaces $\{g_j^{-1}C_i \mid j \in J\}$. This family is finite, since each of its elements is a subspace intersecting the bounded set A . If $g_j^{-1}C_i = g_k^{-1}C_i$ then $g_j g_k^{-1} \in G_i$. In other words, g_j and g_k lie in the same right coset $G_i g_j = G_i g_k$. So the family $\{g_j A \mid j \in J\}$ lies in finitely many G_i -orbits. Let $A' = A_1 \cup \dots \cup A_\ell$ be a union of representatives of these orbits. Then $G_i A'$ contains $C_i \cap GA$.

By hypothesis, we can choose $E_i \subset C_i$ to be a G_i -invariant, G_i -cocompact convex subset containing $\mathcal{N}_1(D_i)$. The union

$$E = GK \cup \bigcup_{i \in I, g \in G} gE_i$$

is clearly G -cocompact. Observe that E is connected because GK is connected, each gE_i is connected, and each gE_i intersects GK . Furthermore E is closed and locally convex because GK is in the interior of E and each point of $E - GK$ lies in a unique gC_i and hence has a convex open neighborhood in $gC_i - GK$ not meeting any other $g'C_j$.

The convexity holds because it is a closed, connected, locally convex subspace of a complete CAT(0) space. \square

The following is obtained by combining Proposition 7.6 with Theorem 9.3.

COROLLARY 9.4. *Suppose the finitely generated group G acts on a wallspace X with an isolated collection \mathcal{P} of peripheries. Suppose the $\text{Stab}(P)$ action on $C(P)$ has the convex core property for each periphery P . Then the G action on $C(X)$ has the convex core property.*

In particular as a consequence of Theorem 7.10 we have the following corollary.

COROLLARY 9.5. *Suppose G is hyperbolic relative to virtually abelian subgroups. Let H_1, \dots, H_r be a collection of quasi-isometrically embedded subgroups with chosen H_i -walls so that G acts properly on the associated wallspace. Then G acts properly and cocompactly on a CAT(0) space.*

We also obtain the following result, which was claimed without proof in [Wis04].

COROLLARY 9.6. *Every finitely presented $B(6)$ group acts properly and cocompactly on a CAT(0) space.*

Proof. It can be proven directly that a $B(6)$ group G is hyperbolic relative to virtually abelian subgroups. Thus the corollary follows by combining Corollary 9.5 with the main result of [Wis04].

Alternatively, the main result of [Wis04] shows that G acts properly on a CAT(0) cube complex $C = GK \cup \bigcup_i C(P_i)$ satisfying the hypothesis of Theorem 9.3. Moreover each $G_i = \text{Stab}(P_i)$ is virtually abelian. Consequently C has a G -cocompact convex core \bar{C} . \square

ACKNOWLEDGEMENTS

We are grateful to Mark Hagen, Frédéric Haglund, Piotr Przytycki, Hung Tran, and Wenyuan Yang for corrections and especially grateful to the referees for many helpful corrections that improved this paper.

REFERENCES

- AR90 I. R. Aitchison and J. H. Rubinstein, *An introduction to polyhedral metrics of nonpositive curvature on 3-manifolds*, in *Geometry of low-dimensional manifolds: 2. Symplectic manifolds and Jones–Witten theory Durham, 1989*, London Mathematical Society Lecture Note Series, vol. 151, eds S. K. Donaldson and C. B. Thomas, (Cambridge University Press, Cambridge, 1990), 127–161.
- AB06 E. Alibegović and M. Bestvina, *Limit groups are CAT(0)*, J. Lond. Math. Soc. (2) **74** (2006), 259–272.
- BP10 S. Barré and M. Pichot, *Removing chambers in Bruhat–Tits buildings*, Preprint (2010), [arXiv:1003.4614](https://arxiv.org/abs/1003.4614) [math.MG].
- BHW11 N. Bergeron, F. Haglund and D. T. Wise, *Hyperplane sections in arithmetic hyperbolic manifolds*, J. Lond. Math. Soc. (2) **83** (2011), 431–448.

- BW12 N. Bergeron and D. T. Wise, *A boundary criterion for cubulation*, Amer. J. Math. **134** (2012), 843–859.
- Bow12 B. H. Bowditch, *Relatively hyperbolic groups*, Internat. J. Algebra Comput. **22** (2012), doi:[10.1142/S0218196712500166](https://doi.org/10.1142/S0218196712500166).
- BH99 M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature* (Springer, Berlin, 1999).
- BH93 B. Brink and R. B. Howlett, *A finiteness property and an automatic structure for Coxeter groups*, Math. Ann. **296** (1993), 179–190.
- CP11 P.-E. Caprace and P. Przytycki, *Bipolar Coxeter groups*, J. Algebra **338** (2011), 35–55.
- CS11 P.-E. Caprace and M. Sageev, *Rank rigidity for $CAT(0)$ cube complexes*, Geom. Funct. Anal. **21** (2011), 851–891.
- CN05 I. Chatterji and G. Niblo, *From wall spaces to $CAT(0)$ cube complexes*, Internat. J. Algebra Comput. **15** (2005), 875–885.
- CCJJV01 P.-A. Cherix, M. Cowling, P. Jolissaint, P. Julg and A. Valette, *Groups with the Haagerup property: Gromov’s a - T -menability*, Progress in Mathematics, vol. 197 (Birkhäuser, Basel, 2001).
- Dah03 F. Dahmani, *Combination of convergence groups*, Geom. Topol. **7** (2003), 933–963.
- DS05 C. Druţu and M. Sapir, *Tree-graded spaces and asymptotic cones of groups*, Topology **44** (2005), 959–1058; With an appendix by Denis Osin and Mark Sapir.
- Far03 D. S. Farley, *Finiteness and $CAT(0)$ properties of diagram groups*, Topology **42** (2003), 1065–1082.
- Gau12 F. Gautero, *A non-trivial example of a free-by-free group with the Haagerup property*, Groups Geom. Dyn. **6** (2012), 677–699.
- Geo08 R. Geoghegan, *Topological methods in group theory*, Graduate Texts in Mathematics, vol. 243 (Springer, New York, 2008).
- Ger97 V. N. Gerasimov, *Semi-splittings of groups and actions on cubings*, in *Algebra, geometry, analysis and mathematical physics (Russian) Novosibirsk, 1996*, Izdat. Ross. Akad. Nauk Sib. Otd. Inst. Mat., eds Yu. G. Reshetnyak, L. A. Bokut’, S. K. Vodop’yanov and I. A. Taïmanov (1997), 91–109; Engl. transl. *Fixed-point-free actions on cubings*, Siberian Adv. Math. **8** (1998), 36–58.
- GMRS98 R. Gitik, M. Mitra, E. Rips and M. Sageev, *Widths of subgroups*, Trans. Amer. Math. Soc. **350** (1998), 321–329.
- Gro87 M. Gromov, *Hyperbolic groups*, in *Essays in group theory*, Mathematical Sciences Research Institute Publications, vol. 8, ed S. M. Gersten, (Springer, New York, 1987), 75–263.
- GS06 V. S. Guba and M. V. Sapir, *Diagram groups and directed 2-complexes: homotopy and homology*, J. Pure Appl. Algebra **205** (2006), 1–47.
- Gur06a D. Guralnik, *Coarse decompositions of boundaries for $CAT(0)$ groups*, Preprint (2006), [arXiv:math/0611006](https://arxiv.org/abs/math/0611006) [math.GR].
- Gur06b D. Guralnik, *Local finiteness for cubulations of $CAT(0)$ groups*, Preprint (2006), [arXiv:math/0610950](https://arxiv.org/abs/math/0610950) [math.GR].
- Hag13 M. F. Hagen, *Weak hyperbolicity of cube complexes and quasi-arboreal groups*, J. Topology, published online 23 August 2013, doi:[10.1112/jtopol/jtt027](https://doi.org/10.1112/jtopol/jtt027), 34 pages.
- HP98 F. Haglund and F. Paulin, *Simplicité de groupes d’automorphismes d’espaces à courbure négative*, in *The Epstein birthday schrift*, Geometry & Topology Monographs, vol. 1, eds I. Rivin, C. Rourke and C. Series (Geometry & Topology Publications, Coventry, 1998), 181–248.
- HW08 F. Haglund and D. T. Wise, *Special cube complexes*, Geom. Funct. Anal. **17** (2008), 1551–1620.

- HW10 F. Haglund and D. T. Wise, *Coxeter groups are virtually special*, Adv. Math. **224** (2010), 1890–1903.
- HS92 J. Hass and P. Scott, *Homotopy equivalence and homeomorphism of 3-manifolds*, Topology **31** (1992), 493–517.
- Hru10 G. C. Hruska, *Relative hyperbolicity and relative quasiconvexity for countable groups*, Algebr. Geom. Topol. **10** (2010), 1807–1856.
- HrW09 G. C. Hruska and D. T. Wise, *Packing subgroups in relatively hyperbolic groups*, Geom. Topol. **13** (2009), 1945–1988.
- HsW10 T. Hsu and D. T. Wise, *Cubulating graphs of free groups with cyclic edge groups*, Amer. J. Math. **132** (2010), 1153–1188.
- JW13 D. Janzen and D. T. Wise, *Cubulating rhombus groups*, Groups Geom. Dyn. **7** (2013), 419–442.
- KK05 M. Kapovich and B. Kleiner, *Coarse Alexander duality and duality groups*, J. Differential Geom. **69** (2005), 279–352.
- KR89 P. H. Kropholler and M. A. Roller, *Relative ends and duality groups*, J. Pure Appl. Algebra **61** (1989), 197–210.
- LW11 J. Lauer and D. T. Wise, *Cubulating one-relator groups with torsion*, submitted, 2011.
- Lea10 I. J. Leary, *A metric Kan–Thurston theorem*, Preprint (2010), [arXiv:1009.1540](https://arxiv.org/abs/1009.1540) [math.GR].
- McC09 J. P. McCammond, *Constructing non-positively curved spaces and groups*, in *Geometric and cohomological methods in group theory Durham, 2003*, London Mathematical Society Lecture Note Series, vol. 358, eds M. R. Bridson, P. H. Kropholler and I. J. Leary, (Cambridge University Press, Cambridge, 2009), 162–224.
- Mou88 G. Moussong, *Hyperbolic coxeter groups*. PhD thesis, Ohio State University, (1988).
- NS11 A. Nevo and M. Sageev, *The Poisson boundary of $CAT(0)$ cube complex groups*, Preprint (2011), [arXiv:1105.1675](https://arxiv.org/abs/1105.1675) [math.GT].
- Nic04 B. Nica, *Cubulating spaces with walls*, Algebr. Geom. Topol. **4** (2004), 297–309.
- NR97 G. Niblo and L. Reeves, *Groups acting on $CAT(0)$ cube complexes*, Geom. Topol. **1** (1997), 1–7.
- NR98 G. A. Niblo and L. D. Reeves, *The geometry of cube complexes and the complexity of their fundamental groups*, Topology **37** (1998), 621–633.
- NR03 G. A. Niblo and L. D. Reeves, *Coxeter groups act on $CAT(0)$ cube complexes*, J. Group Theory **6** (2003), 399–413.
- NiR98 G. A. Niblo and M. A. Roller, *Groups acting on cubes and Kazhdan’s property (T)*, Proc. Amer. Math. Soc. **126** (1998), 693–699.
- OW11 Y. Ollivier and D. T. Wise, *Cubulating random groups at density less than 1/6*, Trans. Amer. Math. Soc. **363** (2011), 4701–4733.
- Osi06 D. V. Osin, *Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems*, Mem. Amer. Math. Soc. **179** (2006).
- PW12 P. Przytycki and D. T. Wise, *Mixed 3-manifolds are virtually special*, Preprint (2012), [arXiv:1205.6742](https://arxiv.org/abs/1205.6742) [math.GR].
- Rol98 M. A. Roller, *Poc sets, median algebras and group actions: an extended study of Dunwoody’s construction and Sageev’s theorem*, Preprint (1998), University of Southampton.
- RS99 H. Rubinstein and M. Sageev, *Intersection patterns of essential surfaces in 3-manifolds*, Topology **38** (1999), 1281–1291.
- RW98 J. H. Rubinstein and S. Wang, π_1 -*injective surfaces in graph manifolds*, Comment. Math. Helv. **73** (1998), 499–515.
- Sag95 M. Sageev, *Ends of group pairs and non-positively curved cube complexes*, Proc. Lond. Math. Soc. (3) **71** (1995), 585–617.

- Sag97 M. Sageev, *Codimension-1 subgroups and splittings of groups*, J. Algebra **189** (1997), 377–389.
- SW05 M. Sageev and D. T. Wise, *The Tits alternative for $CAT(0)$ cubical complexes*, Bull. Lond. Math. Soc. **37** (2005), 706–710.
- Sco77 P. Scott, *Ends of pairs of groups*, J. Pure Appl. Algebra **11** (1977), 179–198.
- Sco83 P. Scott, *There are no fake Seifert fibre spaces with infinite π_1* , Ann. of Math. (2) **117** (1983), 35–70.
- Ser77 J.-P. Serre, *Arbres, amalgames, SL_2* , Astérisque, vol. 46 (Société Mathématique de France, Paris, 1977); written in collaboration with H. Bass.
- Wis04 D. T. Wise, *Cubulating small cancellation groups*, Geom. Funct. Anal **14** (2004), 150–214.
- Wis06 D. T. Wise, *Subgroup separability of the figure 8 knot group*, Topology **45** (2006), 421–463.
- Wis09 D. T. Wise, *Research announcement: the structure of groups with a quasiconvex hierarchy*, Electron. Res. Announc. Amer. Math. Soc. **16** (2009), 44–55.
- Wis12 D. T. Wise, *Recubulating free groups*, Israel J. Math. **191** (2012), 337–345.
- Wri12 N. Wright, *Finite asymptotic dimension for $CAT(0)$ cube complexes*, Geom. Topol. **16** (2012), 527–554.

G. C. Hruska chruska@uwm.edu

Department of Mathematical Sciences, University of Wisconsin–Milwaukee, PO Box 413,
Milwaukee, WI 53201, USA

Daniel T. Wise wise@math.mcgill.ca

Department of Mathematics & Statistics, McGill University, Montreal, QC, Canada H3A 0B9