Annals of Mathematics **176** (2012), 1427–1482
http://dx.doi.org/10.4007/annals.2012.176.3.2

A combination theorem for special cube complexes

By Frédéric Haglund and Daniel T. Wise

Abstract

We prove that certain compact cube complexes have special finite covers. This means they have finite covers whose fundamental groups are quasiconvex subgroups of right-angled Artin groups. As a result we obtain linearity and the separability of quasiconvex subgroups for the groups we consider. Our result applies, in particular, to a compact negatively curved cube complex whose hyperplanes do not self-intersect. For a cube complex with word-hyperbolic fundamental group, we show that it is virtually special if and only if its hyperplane stabilizers are separable. In a final application, we show that the fundamental groups of every simple type uniform arithmetic hyperbolic manifolds are cubical and virtually special.

Contents

Research supported by NSERC.

1. **Introduction**

In this paper we give a high-dimensional generalization of the 1-dimensional work in [Wis02]. The main result there can be loosely reformulated as follows: the 2-complex built by amalgamating two graphs along a malnormal immersed graph is "virtually special." Recall that a subgroup $H \subset G$ is *malnormal* if $gHg^{-1} \cap H = \{1\}$ for each $g \notin H$, and an immersed subgraph is malnormal if its π_1 maps to a malnormal subgroup. More precisely, we consider two combinatorial graph immersions $A \leftarrow M$ and $M \to B$ and form a nonpositively curved square complex $X = A \cup_M B$ by attaching a copy of $M \times [-1, 1]$ to A and B using the maps $A \leftarrow M \times \{-1\}$ and $M \times \{1\} \to B$. The main result of [Wis02] can be reformulated as

Proposition 1.1. *If $\pi_1 M$ is malnormal in $\pi_1 A$ and $\pi_1 B$, then $X = A \cup_M B$ is virtually special.*

From a group theoretical viewpoint, two particularly salient features of a graph Γ are that Γ retracts onto its connected subgraphs and that any finite immersed subgraph embeds in a finite cover of Γ. These features lead to notions of canonical completion and retraction for graphs that were studied in [Wis02]. The *special cube complexes* introduced in [HaW08] are higher dimension spaces that also admit canonical completion and retraction. Simple aspects of these notions and their peculiar properties were verified and extended to higher dimensions in [HaW08], and more difficult such aspects are treated in this paper. Using these, we prove the following statement, which is reformulated and proven in Theorem 8.2.

Theorem 1.2. *Let A and B be compact virtually special cube complexes with word-hyperbolic π_1, and let $A \leftarrow M$ and $M \to B$ be local isometries of cube complexes such that $\pi_1 M$ is quasiconvex and malnormal in $\pi_1 A$ and $\pi_1 B$. Let $X = A \cup_M B$ be the cube complex obtained by gluing A and B together with $M \times [-1, 1]$. Then X is virtually special.*

The specificness of Proposition 1.1 is misleading; in fact, it was surprisingly widely applicable to many 2-dimensional groups appearing in combinatorial group theory. We expect Theorem 1.2 to be even more powerful and dynamic, both because it can reach higher-dimensional groups and also because sometimes 2-dimensional groups are not fundamental groups of 2-dimensional special cube complexes, but require the greater flexibility of higher-dimensions.

A $\mathcal{VH}$*-complex* X is (up to double cover) a square 2-complex with the property that the link of each vertex is bipartite. Furthermore, X is $n$egatively curved if each link has girth ≥ 5. An immediate application of Proposition 1.1 is that negatively curved $\mathcal{VH}$-complexes are virtually special. In higher dimensions, this naturally extends to negatively curved "foldable complexes" as defined in [BŚ99]. The cube complex X is *foldable* if there is a combinatorial map $X \to Q$ onto a cube Q. X is *negatively curved* if the link of each vertex is a flag complex, and moreover, any 4-cycle in the link bounds the union of two 2-simplices.

THEOREM 1.3. *Let C be a compact negatively curved foldable cube complex. Then C has a finite cover $\widehat{C}$ such that $\widehat{C}$ is a special cube complex.*

The work in [Wis02] led to connections between negative curvature and residual finiteness. Our paper extends this connection considerably, since it is now understood that surprisingly many of the groups studied in combinatorial group theory actually act properly and cocompactly on CAT(0) cube complexes and are thus approachable through these results.

A subgroup H is *separable* in a group G if H is the intersection of finite index subgroups of G. In particular, G is *residually finite* if $\{1_G\}$ is separable.

THEOREM 1.4. *Let C be a compact nonpositively curved cube complex, and suppose that $\pi_1 C$ is word-hyperbolic. Then C is virtually special if and only if $\pi_1 D$ is separable in $\pi_1 C$ for each immersed hyperplane $D \to C$.*

Using the already known properties of virtually special cube complexes we get the following corollary.

THEOREM 1.5. *Let C be a compact nonpositively curved cube complex. Suppose that $\pi_1 C$ is word-hyperbolic and that $\pi_1 D$ is separable in $\pi_1 C$ for each immersed hyperplane $D \to C$. Then $\pi_1 C \subset GL(n, \mathbb{Z})$ for some large n and every quasiconvex subgroup of $\pi_1 C$ is separable.*

In a final application we apply Theorem 1.4 to obtain an interesting structural result about certain arithmetic hyperbolic lattices.

THEOREM 1.6. *Let G be a uniform arithmetic hyperbolic lattice of "simple type." Then G has a finite index subgroup F that is the fundamental group of a compact special cube complex.*

Hidden inside Theorem 1.6 is the claim that every such lattice acts properly and cocompactly on a CAT(0) cube complex, which was not known previously. But combining with results from [HaW08], we obtain the following subgroup separability consequence.

Corollary 1.7. *Let G be a uniform arithmetic hyperbolic lattice of "simple type." Then each quasiconvex subgroup of G is a virtual retract and is thus closed in the profinite topology.*

Partial results were made towards separability of arithmetic hyperbolic lattices in low dimensions in [ALR01]. Their method is similar to ours in that they virtually embed such groups into right-angled hyperbolic Coxeter groups. The results have been substantially extended by Ian Agol to deal with various lattices in up to eleven dimensions that satisfy an orthogonality condition on certain of their hyperplanes [Ago06]. Perhaps the paucity of hyperbolic reflection groups has limited the scope of Scott's method.

The application towards uniform arithmetic lattices was not the original intention of this research, but it emerged as a consequence of our combination theorem. This application does not require the full strength of our main theorem, and in a future paper, using a method more specific to the situation, we will give an account of the virtual specialness of nonuniform simple arithmetic hyperbolic lattices.

As an application of the cubulation of uniform lattices of the real hyperbolic space we get the following result.

Theorem 1.8. *Every word-hyperbolic group is quasi-isometric to a uniformly locally finite* CAT(0) *cube complex.*

Proof. Let Γ be a word-hyperbolic group. Then by the work of Bonk and Schramm (see [BS00]) there exists a quasi-isometric embedding $\Gamma \to \mathbb{H}^n$ for n large enough. Consider any standard arithmetic uniform lattice G of H^n. We let G act freely cocompactly on a locally finite CAT(0) cube complex X. Then $\mathbb{H}^n$ is quasi-isometric to G, which is quasi-isometric to X.

We thus get a quasi-isometric embedding of Γ into X. The image of Γ in X is a quasiconvex subset $Y \subset X$. Now X is a uniformly locally finite, Gromov-hyperbolic CAT(0) cube complex and $Y \subset X$ is quasiconvex. Thus by Theorem 4.2 the combinatorial convex hull Z of Y inside X stays at finite hausdorff distance of Y. $\square$

It is not very difficult to prove that a Gromov-hyperbolic uniformly locally finite CAT(0) cube complex embeds in a product of finitely many trees (see, for example, [Hag]), and the embedding is isometric on the 1-skeleton equipped with the combinatorial distance. Thus, as a corollary of Theorem 1.8, we see that every word-hyperbolic group has a quasi-isometric embedding in a product of finitely many trees, a result which was first proved by Buyalo and Schroeder (see [BS05]).

2. Special cube complexes

In this section we review definitions related to special and nonpositively curved cube complexes. See [Sag95], [HaW08].

2.A. *Nonpositively curved cube complex.*

Definition 2.1. A 0-cube is a single point. A 1-*cube* is an isometric copy of $[-1, 1]$ and has a cell structure consisting of 0-cells $\{\pm 1\}$ and a single 1-cell. An n-*cube* is an isometric copy of $[-1, 1]^n$, and has the product cell structure, so each closed cell of $[-1, 1]^n$ is obtained by restricting some of the coordinates to $+1$ and some to -1.

A *cube complex* is obtained from a collection of cubes of various dimensions by isometrically identifying certain subcubes. We shall often call 0-cubes *vertices* and 1-cubes *edges*. A map between cube complexes is *combinatorial* if it isometrically maps open cubes to open cubes. Note that we are using the Euclidean metric on open cubes only to rigidly specify our maps between cubes.

A *flag complex* is a simplicial complex with the property that every finite set of pairwise adjacent vertices spans a simplex.

Let X be a cube complex. The *link* of a vertex v in X is a complex built from simplices corresponding to the corners of cubes adjacent to v. One can think of $\mathrm{link}(v)$ as being the "ε-sphere" about v in X.

The cube complex X is *nonpositively curved* if $\mathrm{link}(v)$ is a flag complex for each $v \in X^0$. A (finite dimensional) simply-connected nonpositively curved cube complex has a CAT(0) metric in which each n-cube is isometric to the subspace $[-1, 1]^n \subset \mathbb{E}^n$ (see [Gro87]). We thus refer to simply-connected nonpositively curved cube complexes as CAT(0) *cube complexes*. Similarly, a (finite dimensional) nonpositively curved cube complex admits a locally CAT(0) metric, and hence the choice of terminology for the combinatorial flag complex condition.

A combinatorial map $f : X \to Y$ of nonpositively curved cube complexes is a *local isometry* if $\mathrm{link}(v, X)$ maps injectively to a full subcomplex of $\mathrm{link}(f(v), Y)$ for each $v \in X^0$. Recall that a subcomplex of a simplicial complex is *full* if it is spanned by its set of vertices.

2.B. *Hyperplanes.*

Definition 2.2. A *midcube* D in a cube C is the subspace obtained by restricting exactly one coordinate to 0. For instance, $[-1, 1] \times [-1, 1] \times \{0\} \times [-1, 1]$ is a midcube in $[-1, 1]^4$.

Given a cube complex X, consider the disjoint union of all midcubes of X, and let Y be the cube complex obtained by identifying lower-dimensional midcubes with their images as subcubes of higher-dimensional midcubes. The connected components of Y are the *hyperplanes* of X.

It is not difficult to check that when X is nonpositively curved, then each hyperplane of X is nonpositively curved. Moreover, using the metrics of nonpositive curvature, the natural map $H \to X$ is a local isometry. When this local isometry is injective we say that H *embeds*.

When X is CAT(0), each hyperplane of X is itself CAT(0) and embeds.

2.C. *Right-angled Artin groups.*

Definition 2.3. The *right-angled Artin group* presentation associated to a simplicial graph Γ is defined to be:

$$\Big\langle a : a \in \mathrm{V}(\Gamma) \mid [a,b] : \{a,b\} \in \mathrm{E}(\Gamma)\Big\rangle.$$

The associated right-angled Artin group will be denoted by $A = A(\Gamma)$.

For each clique of n pairwise adjacent vertices in Γ, we add an n-cube to the standard 2-complex of the presentation above to obtain a nonpositively curved cube complex, which we shall denote by $R = R(\Gamma)$. We call any such complex R an *Artin cube complex.*

Specifically, for each vertex a of $\mathrm{V}(\Gamma)$, let S_a denote a graph with a single vertex and a single edge. Then R is a subcomplex of the combinatorial torus $\prod_{a\in \mathrm{V}(\Gamma)} S_a$. Namely, R equals the union of all subtori corresponding to complete subgraphs of Γ. To obtain an isomorphism $A \to \pi_1 R$ we choose an orientation for the edges of the circles S_a.

2.D. *Special cube complex.*

Definition 2.4. A *special cube complex* is a cube complex C such that there exists a combinatorial local isometry $C \to R$ where $R = R(\Gamma)$ for some simplicial graph Γ.

Observe that if $C \to D$ is a local isometry of nonpositively curved cube complexes and D is special, then so is C (by composing local isometries).

We give below an intrinsic combinatorial characterization of the special property. We first introduce the adapted notations and definitions.

We denote with an arrow the oriented edges of a cube complex. The (unoriented) edge associated with an oriented edge $\overrightarrow{a}$ will always be denoted by a. We denote by $\iota(\overrightarrow{a})$ and $\tau(\overrightarrow{a})$ the initial vertex and terminal vertex of the oriented edge $\overrightarrow{a}$.

Definition 2.5 (parallelism). Two oriented edges of a Euclidean unit square are *parallel* if the associated unit vectors are the same. Then the parallelism of oriented edges in a cube complex C is the equivalence relation generated by parallelism inside a square. By forgetting orientation we also get a parallelism relation on (unoriented) edges of C. Note that two edges of C are parallel if and only if their midpoints belong to the same hyperplane.

An embedded hyperplane Y in X is 2-*sided* if its open cubical neighborhood is isomorphic to the product $Y \times (-1,1)$, where we identify Y with $Y \times \{0\}$. For each $y \in Y^0$, the 1-cell in X whose open 1-cell corresponds to $y \times (-1,1)$ is *dual* to Y. An oriented edge $\overrightarrow{a}$ is *dual* to Y if a is dual to Y. When Y is 2-sided, the projection map $Y \times (-1,1) \to (-1,1)$ allows us to choose an orientation on each 1-cell dual to Y so that all corresponding oriented edges are parallel. In other words, no oriented edge dual to Y is parallel to its opposite edge.

The 2-sided hyperplane Y *directly self-osculates* if there are distinct oriented dual edges $\overrightarrow{a}$ and $\overrightarrow{b}$ such that $\iota(\overrightarrow{a}) = \iota(\overrightarrow{b})$. The hyperplane Y *self-osculates* if there are distinct dual edges a and b that share a vertex. Thus self-osculation of 2-sided hyperplanes consists of direct self-osculation, and also indirect self-osculation, where there are distinct oriented dual edges $\overrightarrow{a}$ and $\overrightarrow{b}$ such that $\iota(\overrightarrow{a}) = \tau(\overrightarrow{b})$.

Consider two distinct oriented edges $\overrightarrow{a}, \overrightarrow{b}$ with origin a given vertex v. Identify $\overrightarrow{a}, \overrightarrow{b}$ with vertices of $\mathrm{link}(v)$. When $\overrightarrow{a}, \overrightarrow{b}$ span an edge of $\mathrm{link}(v)$ we say that $\overrightarrow{a}, \overrightarrow{b}$ are *perpendicular* at v. When $\overrightarrow{a}, \overrightarrow{b}$ are not joined in $\mathrm{link}(v)$ we say that $\overrightarrow{a}, \overrightarrow{b}$ *osculate* at v. We say that two edges a, b are perpendicular [osculate] when there are orientations $\overrightarrow{a}, \overrightarrow{b}$ such that $\overrightarrow{a}, \overrightarrow{b}$ are perpendicular [osculate] at some vertex.

Two hyperplanes A, B *intersect* (i.e., *cross*) if they have perpendicular dual edges a, b. The *graph of hyperplanes* of the cube complex X is the graph Γ_X whose vertices are the hyperplanes of X, and where an edge joins A, B precisely when A, B intersect.

Two hyperplanes A, B *osculate* if they have osculating dual edges a, b. Two hyperplanes A and B *inter-osculate* if they both intersect and osculate. See Figure 1.

LEMMA 2.6. *A nonpositively curved cube complex is special if and only if*

(S_1) *Each hyperplane embeds.*
(S_2) *Each hyperplane is two-sided.*
(S_3) *No hyperplane directly self-osculates.*
(S_4) *No two hyperplanes interosculate.*

For a nonpositively curved cube complex C satisfying $(S_1), (S_2), (S_3), (S_4)$, consider the graph of hyperplanes Γ_C. We let $R(C)$ denote the Artin cube complex associated with Γ_C.

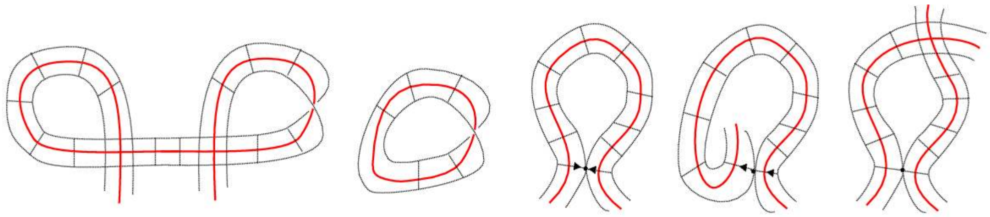

Figure 1. From left to right, the diagrams above correspond to the excluded pathologies enumerated in Lemma 2.6. The third and fourth diagrams illustrate direct self-osculation and indirect self-osculation.

Sketch. Let $R = R(C)$. Choose arbitrary orientations for the edges of R. Then there is a map $C \to R$ induced by sending an oriented edge $\overrightarrow{a}$ of C to the oriented edge of R corresponding to the hyperplane that $\overrightarrow{a}$ is dual to. It is not difficult to verify that conditions $(S_1), (S_2), (S_3), (S_4)$ imply that this map is a local isometry.

Conversely, if $C \to D$ is a local isometry of nonpositively curved cube complexes, then it is easy to verify that if D satisfies conditions $(S_1), (S_2), (S_3), (S_4)$, then C does. The lemma follows since it is easy to check that an Artin cube complex $R(C)$ always satisfies conditions $(S_1), (S_2), (S_3), (S_4)$. We refer to [HaW08] for details. $\square$

If X is a special cube complex, then the cube complex obtained by subdividing X along a hyperplane is still special. This is an easy consequence of Lemma 2.6.

3. Canonical completion and retraction and wall projections

3.A. *Canonical completion and retraction.* In this section we recall how to factorize some combinatorial immersions $X \to Y$ as the composition of an inclusion and a covering map. The key point here is to give a canonical construction that can be used as an elementary machine in more elaborate constructions. Due to its naturality, the construction will enjoy many nice formal properties.

The possibility of lifting an immersion $X \to Y$ to an inclusion $X \to Y'$ in a finite cover $Y' \to Y$ is related to the separability of $\pi_1 X < \pi_1 Y$. In a residually finite group, any virtual retract is separable (indeed retracts of hausdorff topological spaces are closed). We show below that when Y is a special cube complex it is possible to lift a local isometry $X \to Y$ to an inclusion $X \to \mathsf{C}(X, Y)$, where $\mathsf{C}(X, Y) \to Y$ is a "canonically" defined covering *and, furthermore,* $\mathsf{C}(X, Y)$ "*canonically*" *retracts to* X. This construction has been made for graph immersions in [Wis02] and generalized to arbitrary local isometries of special cube complexes in [HaW08]. We first recall the case of graphs.

Definition 3.1 (canonical completion and retraction for immersion of graph to bouquet). Let $A_b \to b$ be an immersion of graphs where b consists of a single loop and A_b is finite. Each component of A_b is either already a cover of b or can be completed to a cover of b by the addition of a single edge. We define $\mathsf{C}(A_b, b) \to b$ to be the resulting covering space. Note that $A_b \subset \mathsf{C}(A_b, b)$ and that there is a retraction map $\mathsf{C}(A_b, b) \to A_b$ defined by sending each new open edge e to the component $I(e)$ of A_b that it was attached along. We choose the map $e \to I(e)$ to have constant speed.

Let $A \to B$ be an immersion of graphs where B is a bouquet of circles and A is finite. For each loop b in B, let A_b denote the preimage of b in A. Observe that $A^0 = A_b^0 = (\mathsf{C}(A_b, b))^0$, and consider the natural embedding of A^0 into the graph $\mathsf{C}(A_b, b)$ defined above. We define $\mathsf{C}(A, B)$ to be the quotient of $\sqcup_{b\in \mathrm{edges}(B)}\mathsf{C}(A_b, b)$ obtained by identifying the embedded copies of A^0. The induced map $\mathsf{C}(A, B) \to B$ is a covering space. Note that $A \subset \mathsf{C}(A, B)$ and that the retraction maps $\mathsf{C}(A_b, b) \to A_b$ induce a retraction map $\mathsf{C}(A, B) \to A$. In conclusion, by adding edges in a "canonical" manner, we have completed A to a finite degree covering of B that furthermore retracts onto A.

Definition 3.2 (canonical completion and retraction for local isometries of cube complexes in an Artin complex). Let $X \to R$ be a local isometry of cube complexes where R is an Artin cube complex and X is finite. We have already defined canonical completions and retractions of graphs to obtain $X^1 \leftarrow \mathsf{C}(X^1, R^1) \to R^1$. Using the local isometry assumption, a case-by-case inspection shows that the boundary of a square in R always lifts to a closed curve in $\mathsf{C}(X^1, R^1)$. More specifically, there is a local isometry $[0, n] \times [0, m] \to X$ such that the images of the four sides $[0, n] \times \{0\}$, $\{n\} \times [0, m]$, $[0, n] \times \{m\}$, $\{0\} \times [0, m]$ coincide with the images of the four sides of the lift under the retraction map. See [HaW08] for a written argument and Figure 2, where we depict possible scenarios. Thus we can extend the previous covering $\mathsf{C}(X^1, R^1) \to R^1$ to a covering map of square complexes $\mathsf{C}(X, R)^2 \to R^2$. The 2-skeleton of each higher-dimensional cube of R lifts to $\mathsf{C}(X, R)^2$, and we attach all such corresponding cubes. The resulting space $\mathsf{C}(X, R)$ covers R and contains X.

The graph retraction $X^1 \leftarrow \mathsf{C}(X^1, R^1)$ extends naturally to a retraction $X^2 \leftarrow \mathsf{C}(X, R)^2$ but, in general, this retraction cannot be made cellular; see Figure 2 for a square that maps onto a rectangle. We remedy this later in Definition 3.5 by subdividing $\mathsf{C}(X, R)^2$. Note that when X and R are compact then so is $\mathsf{C}(X, R)$.

Definition 3.3 (fiber-product). Given a pair of combinatorial maps $X \to W$ and $Y \to W$ (between cube complexes), we define their *fiber-product* $X \otimes_W Y$

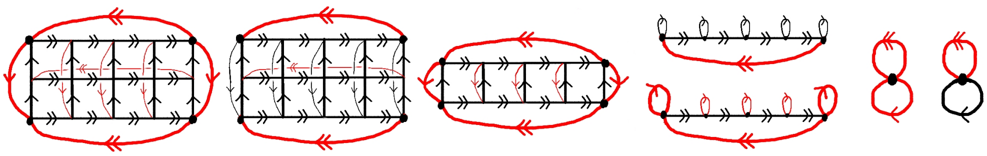

Figure 2. Each new square in $\mathsf{C}(X^1, R^1)$ is associated with a local isometry $Y \to T$ where T is a torus of R. Typically, Y is an $m \times n$-grid like the 4×2-grid on the left, but in degenerate cases, Y may be a cylinder, or m, n may be zero. Canonical retraction sends the new square to Y.

to be a cube complex, whose i-cubes are pairs of i-cubes in X, Y that map to the same i-cube in W. There is a commutative diagram

$$\begin{array}{ccc} X \otimes_W Y & \to & Y \\ \downarrow & & \downarrow \\ X & \to & W. \end{array}$$

Note that $X \otimes_W Y$ is the complex in $X \times Y$ that is the preimage of the diagonal $D \subset W \times W$ under the map $X \times Y \to W \times W$. Note that D is naturally a cube complex since for any cube Q, the diagonal of Q^2 is isomorphic to Q by either of the projections.

We remark that when $X \to W$ and $Y \to W$ are covering maps, then so is $X \otimes_W Y \to W$. Moreover, in this case $\pi_1\big(X \otimes_W Y, (x, y)\big) = \pi_1(X, x) \cap \pi_1(Y, y)$. We emphasize that $X \otimes_W Y$ may not be connected even when X and Y are connected.

We will use the universal property of $X \otimes_W Y$, which is that any commutative diagram as displayed below is the pull-back under some combinatorial map $C \to X \otimes_W Y$ of the diagram with $X \otimes_W Y$ (which is thus minimal):

$$\begin{array}{ccc} C & \to & Y \\ \downarrow & & \downarrow \\ X & \to & W. \end{array}$$

Definition 3.4 (canonical completion and retraction of local isometries). Let $A \to B$ be a local isometry of cube complexes, where A is finite, B is special, and let $R = R(B)$. By composition we get a local isometry $A \to R$, and we already know how to complete A to a cover $\mathsf{C}(A, R) \to R$.

Our plan is to pull back $\mathsf{C}(A, R) \to R$ through $B \to R$ to obtain $\mathsf{C}(A, B) \to B$, and we note that the map $A \to B$ lifts to an inclusion in $\mathsf{C}(A, B)$. In order to investigate the formal properties of this construction, we will use the fiber-product language.

We thus set $\mathsf{C}(A, B) := B \otimes_R \mathsf{C}(A, R)$. The local isometry $A \to B$ and the inclusion map coincide on R and thus define an inclusion $A \to \mathsf{C}(A, B)$.

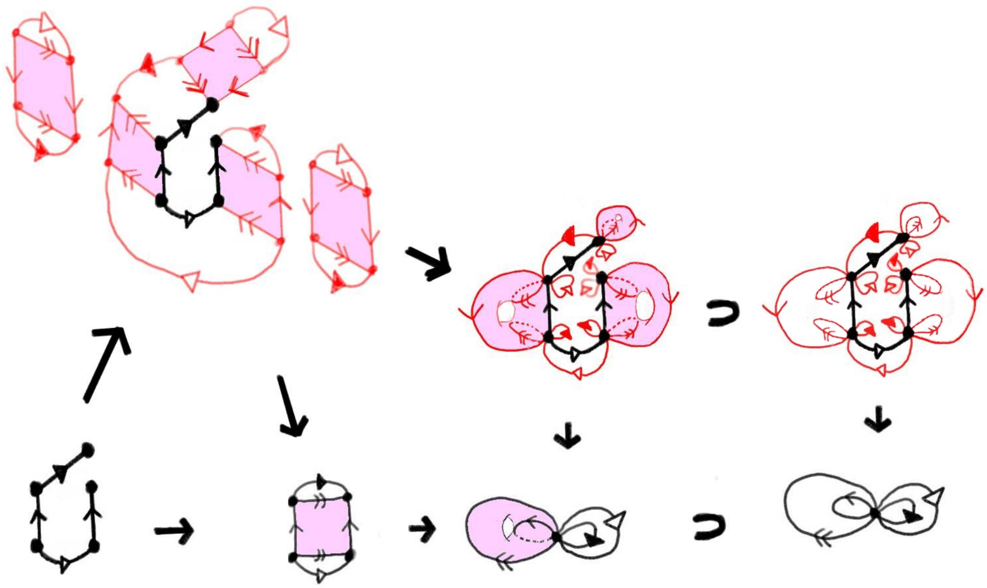

Figure 3. The figure corresponds to the following commutative diagram:

$$\begin{array}{ccccccc} & & \mathsf{C}(A,B) & \to & \mathsf{C}(A,R) & \supset & \mathsf{C}(A^1,R^1) \\ & \nearrow & \downarrow & & \downarrow & & \downarrow \\ A & \to & B & \to & R & \supset & R^1. \end{array}$$

The map $\mathsf{C}(A,B) \to B$ is a covering because $\mathsf{C}(A,R) \to R$ is a covering. We define the canonical retraction map $\mathsf{C}(A,B) \to A$ to be the composition $\mathsf{C}(A,B) \to \mathsf{C}(A,R) \to A$.

We will now describe a cubical subdivision $\mathsf{C}_{\boxminus}(A,B)$ of $\mathsf{C}(A,B)$. It is the minimal subdivision so that $\mathsf{C}_{\boxminus}(A,B) \to B$ is cellular, mapping cubes to cubes by possibly collapsing some dimensions. This subdivision is used here only to support the proof of Corollary 3.11. We have $\mathsf{C}(A,B) = \mathsf{C}_{\boxminus}(A,B)$ precisely when there is no concatenable pair of oriented edges $\overrightarrow{a}_1, \overrightarrow{a}_2$ in A whose images $\overrightarrow{b}_1, \overrightarrow{b}_2$ are parallel oriented edges of B. And when $\mathsf{C}(A,B) \neq \mathsf{C}_{\boxminus}(A,B)$, the covering map $\mathsf{C}_{\boxminus}(A,B) \to B$ is no longer combinatorial.

Definition 3.5 (The subdivision $\mathsf{C}_{\boxminus}(A,B)$). Our choice of subdivision is determined by the 1-skeleton, so let us first revisit the details of the construction given in Definition 3.1: When the interval $I(e)$ consists of $m \geq 1$ edges we subdivide e into m edges and then r induces a combinatorial isomorphism of the subdivided edge e onto $I(e)$. We will denote by $\mathsf{C}_{\boxminus}(A_b,b)$ the corresponding subdivision of $\mathsf{C}(A_b,b)$. The retraction map $\mathsf{C}_{\boxminus}(A_b,b) \to A_b$ is now cellular, and the inclusion map $A_b \to \mathsf{C}_{\boxminus}(A_b,b)$ is an embedding of subgraphs. We then define $\mathsf{C}_{\boxminus}(A,B)$ to be the quotient of $\sqcup_{b\in \text{edges}(B)}\mathsf{C}_{\boxminus}(A_b,b))$ obtained by identifying the copies of A^0. Note that the graph $\mathsf{C}_{\boxminus}(A,B)$ can also be obtained by subdividing certain edges of $\mathsf{C}(A,B)$. The retraction map $\mathsf{C}_{\boxminus}(A,B) \to A$ is cellular, and the inclusion map $A \to \mathsf{C}_{\boxminus}(A_b, b)$ is combinatorial.

We now revisit the details in the construction of Definition 3.2. The attaching map of each square in $\mathsf{C}(X,R)^2$ is an immersion of a 4-cycle in $\mathsf{C}(X^1,R^1)$. We subdivide the square to obtain a rectangle of the form $[0,p] \times [0,q]$ so that the attaching map becomes a combinatorial immersion into $\mathsf{C}_{\boxminus}(X^1,R^1)$. Attaching these rectangles to the graph $\mathsf{C}_{\boxminus}(X^1,R^1)$ yields a square complex $\mathsf{C}_{\boxminus}(X,R)^2$ that is a square subdivision of $\mathsf{C}(X,R)^2$. Since no square of X was subdivided, the inclusion map $X^2 \to \mathsf{C}_{\boxminus}(X,R)^2$ is combinatorial. We will denote by $\mathsf{C}_{\boxminus}(X,R)$ the unique cubical subdivision of $\mathsf{C}(X,R)$ whose 2-skeleton is $\mathsf{C}_{\boxminus}(X,R)^2$.

The same case-by-case inspection as in Definition 3.2 shows that the cellular retraction map $X^1 \leftarrow \mathsf{C}_{\boxminus}(X^1,R^1)$ extends to a cellular retraction $X^2 \leftarrow \mathsf{C}_{\boxminus}(X,R)^2$. More precisely, the image of a square of X^2 is itself, and the image of a new square of $\mathsf{C}(X,R)^2$ corresponding to a rectangle $[0,p]\times[0,q]$ in $\mathsf{C}_{\boxminus}(X,R)^2$ is a rectangle $[0,n]\times[0,m]$ with $n=p$ or $n=0$, and $m=q$ or $m=0$. The cellular retraction $X^2 \leftarrow \mathsf{C}_{\boxminus}(X,R)^2$ extends naturally to $X \leftarrow \mathsf{C}_{\boxminus}(X,R)$. Forgetting the subdivision, we thus get the (original) topological retraction $X \leftarrow \mathsf{C}(X,R)$.

Finally, we revisit the details in the construction of Definition 3.4. We let $\mathsf{C}_{\boxminus}(A,B)$ denote the unique cubical subdivision of $\mathsf{C}(A,B)$ such that the induced map $\mathsf{C}_{\boxminus}(A,B) \to \mathsf{C}_{\boxminus}(A,R)$ is a combinatorial local isometry. The induced retraction map $\mathsf{C}_{\boxminus}(A,B) \to A$ is cellular.

Definition 3.6 (projective). A map $f : X \to Y$ between cube complexes is *projective* if for each cube Q of X, there is a face $Q' < Q$ such that f is combinatorial on Q', and for any $x \in Q$, we have $f(x) = f(x')$, where x' is the orthogonal projection of x onto Q'.

Remark 3.7. Being projective is essentially determined by the 2-skeleton. Indeed if $f : X \to Y$ is projective, then so is $f : X^2 \to Y$. Conversely, if Y is nonpositively curved, then any projective map $f : X^2 \to Y$ extends to a unique projective map $f : X \to Y$.

Note also that projective maps preserve parallelism. More precisely, let $f : X \to Y$ be projective and let a, b denote parallel edges of X. If $f(a)$ is an edge, then so is $f(b)$ and, furthermore, $f(a), f(b)$ are parallel edges of Y.

LEMMA 3.8 (projective retraction). *Let $A \to B$ be a local isometry of cube complexes with A finite and B special. Then the retraction map $A \leftarrow \mathsf{C}_{\boxminus}(A,B)$ is projective.*

Proof. Let $R = R(B)$. We claim that $\mathsf{C}_{\boxminus}(A,R) \to A$ is projective. By Remark 3.7 it suffices to examine 2-skeleta. The squares of $\mathsf{C}(A,R)$ that are already contained in A are not subdivided in $\mathsf{C}_{\boxminus}(A,R)$, and the retraction map is the identity on them. The new squares of $\mathsf{C}_{\boxminus}(A,R)$ are subdivided to Euclidean rectangles $[0,p]\times[0,q]$, on which the retraction map consists of orthogonal projection to one of $\{0\}\times\{0\}$, $[0,p]\times\{0\}$, $\{0\}\times[0,q]$, or $[0,p]\times[0,q]$. Thus $\mathsf{C}_{\boxminus}(A,R)^2\to A^2$ is projective. The retraction $\mathsf{C}_{\boxminus}(A,B)\to B$ is then projective since it is the composition of the local isometry $\mathsf{C}_{\boxminus}(A,B)\to\mathsf{C}_{\boxminus}(A,R)$ with the projective map $\mathsf{C}_{\boxminus}(A,R)\to A$. □

LEMMA 3.9 (retraction of walls). *Let B be a special cube complex, and let $A\to B$ be a local isometry. Then any edge e' of $\mathsf{C}_{\boxminus}(A,B)$ parallel with an edge e of A is retracted onto an edge e'' of A that is parallel to e within A.*

Proof. By Remark 3.7, projective maps preserve parallelism. □

Definition 3.10 (wall-injective). A combinatorial map $D\to C$ of (special) cube complexes induces a map $V_D\to V_C$ between the sets of hyperplanes. We say that the map $D\to C$ is *wall-injective* if the map $V_D\to V_C$ is injective.

COROLLARY 3.11 (wall-injective in completion). *Let C be a special cube complex, and let $D\to C$ be a local isometry. Then D is wall-injective in both $\mathsf{C}_{\boxminus}(D,C)$ and $\mathsf{C}(D,C)$.*

Proof. The wall-injectivity of D in $\mathsf{C}_{\boxminus}(D,C)$ is an immediate consequence of Lemma 3.9. Consider a sequence of edges $e_0,e_1,\ldots,e_n$ with e_0,e_n edges of D and e_i,e_{i+1} opposite in some square of $\mathsf{C}(D,C)$. Since e_0 is an edge of D, it follows that no e_i edge is subdivided in $\mathsf{C}_{\boxminus}(D,C)$. Thus e_0 and e_n are parallel in $\mathsf{C}_{\boxminus}(D,C)$. □

In the sequel we will often consider $\mathsf{C}(D,D)$, which is the canonical completion of the identity map $D\to D$. Note $\mathsf{C}(D,D)$ is a complicated object: it contains a copy of D but also other components, which can be nontrivial covers of D (see Figure 4). This complexity is (part of) the price to pay for the useful functorial properties of the canonical completion.

LEMMA 3.12. *Suppose $D\to C$ is a wall-injective local-isometric embedding of cube complexes with D finite and C special. Then there is a natural embedding of $\mathsf{C}(D,D)$ in $\mathsf{C}(D,C)$ that is consistent with the inclusion, retraction, and covering maps so that we have the following two diagrams:*

$$(\maltese)\qquad \begin{array}{ccc} D & = & D \\ \cap & & \cap \\ \mathsf{C}(D,D) & \subset & \mathsf{C}(D,C) \\ \downarrow & & \downarrow \\ D & \to & C, \end{array} \qquad\qquad \begin{array}{ccc} D & = & D \\ \uparrow & & \uparrow \\ \mathsf{C}(D,D) & \subset & \mathsf{C}(D,C). \end{array}$$

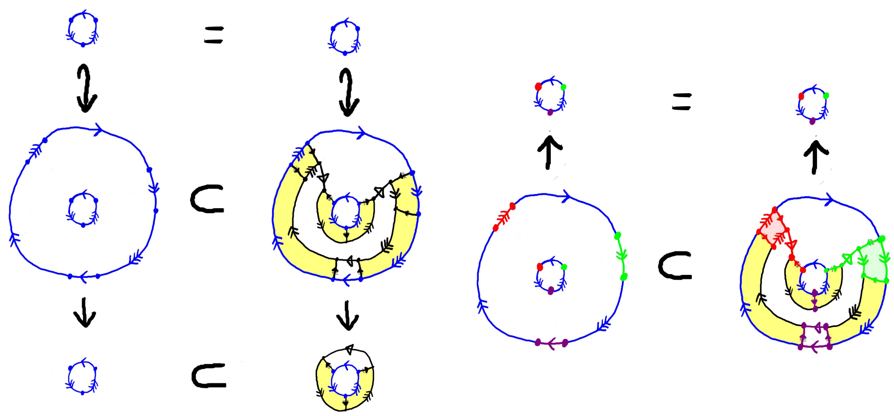

Figure 4. The above two figures illustrate the commutative diagrams in Equation (✠). The complex D is a circle, and the complex C is obtained from $D \times I$ by removing a single square. The fibers of vertices under the retraction map are indicated in different colors on the right.

We refer to Figure 4, which illustrates these commutative diagrams in a simple case.

Proof. The wall-injective local isometry $D \to C$ induces an embedding $V_D \to V_C$ and a combinatorial embedding $R(D) \subset R(C)$ (which is not necessarily a local isometry).

We first check that there is a well-defined map $\mathsf{C}(D, R(D)) \to \mathsf{C}(D, R(C))$ and begin with the 1-skeleta. Note that D^1 is a subgraph of both $\mathsf{C}(D^1, R(D)^1)$ and $\mathsf{C}(D^1, R(C)^1)$. Let a be an edge of $\mathsf{C}(D^1, R(D)^1)$ not contained in D^1, and let b denote the image of a inside $R(D)^1$. We let D_b denote the subgraph of D^1 such that $D_b \cup a$ is a circle. Then b is also an edge of $R(C)$, and D_b is also a connected component of the preimage of b under $D \to R_C$. Thus by construction there is a unique edge a' in $\mathsf{C}(D^1, R(C)^1)$ such that $D_b \cup a'$ is a circle. We then map a to a' by the unique homeomorphism compatible with $a \to b$ and $a' \to b$.

We have now extended the embedding $D^1 \subset \mathsf{C}(D^1, R(C)^1)$ to a combinatorial map $\mathsf{C}(D^1, R(D)^1) \to \mathsf{C}(D^1, R(C)^1)$, which by construction is compatible with retractions onto $D^1$and such that the following diagram commutes:

$$\begin{array}{ccc} \mathsf{C}(D^1, R(D)^1) & \to & \mathsf{C}(D^1, R(C)^1) \\ \downarrow & & \downarrow \\ R(D)^1 & \subset & R(C)^1. \end{array}$$

The map $\mathsf{C}(D^1, R(D)^1) \to \mathsf{C}(D^1, R(C)^1)$ is injective on the 0-skeleton and it is locally injective, thus it is injective. It sends the boundary of a square of $\mathsf{C}(D, R(D))$ onto the boundary of a square of $\mathsf{C}(D, R(C))$, by the very definition of the squares in canonical completion. The same happens for higher-dimensional cubes. Thus we get an injective combinatorial map $\mathsf{C}(D, R(D)) \to \mathsf{C}(D, R(C))$, compatible with retractions onto $D$and such that the following diagram commutes:

$$\begin{array}{ccc} \mathsf{C}(D, R(D)) & \subset & \mathsf{C}(D, R(C)) \\ \downarrow & & \downarrow \\ R(D) & \subset & R(C). \end{array}$$

Now the composition maps

$$\mathsf{C}(D, D) \to D \to C, \quad \mathsf{C}(D, D) \to \mathsf{C}(D, R(D)) \to \mathsf{C}(D, R(C))$$

are compatible with the projections onto $R(C)$. Thus we get a map $\mathsf{C}(D, D) \to \mathsf{C}(D, C)$. Injectivity follows from the injectivity of the maps $\mathsf{C}(D, D) \to D \times \mathsf{C}(D, R(D))$, $D \to C$, and $\mathsf{C}(D, R(D)) \to \mathsf{C}(D, R(C))$. Chasing diagrams, one verifies that this map has the other desired properties. $\square$

LEMMA 3.13. *Let C be a special cube complex, and let $D \subset C$ be a wall-injective locally convex subcomplex. Then the preimage of D in $\mathsf{C}(D, C)$ is* (*isomorphic to*) $\mathsf{C}(D, D)$.

Proof. By definition, the preimage of D in $\mathsf{C}(D, C)$ is $D \otimes_{R(C)} \mathsf{C}(D, R(C))$. The image of a cube $Q \subset D$ in $R(C)$ is in fact a cube of the subcomplex $R(D)$. It follows that $D \otimes_{R(C)} \mathsf{C}(D, R(C)) = D \otimes_{R(D)} \mathsf{C}(D, R(D)) = \mathsf{C}(D, D)$. $\square$

3.B. *Wall projections.* We now study the notion of a wall projection of one subcomplex onto another. This will play the role of the intersection between subgraphs of a graph.

Definition 3.14 (parallel cubes and wall-projection). Let X denote a cube complex. Recall that 1-cubes a, b are *parallel* in X provided they are dual to the same immersed hyperplane.

Let A and B be subcomplexes of X. We define $\mathsf{WProj}_X(A \to B)$, the *wall projection of A onto B in X*, to equal the union of B^0 together with all cubes of B whose 1-cubes are all parallel to 1-cubes of A.

We say the wall projection $\mathsf{WProj}_X(A \to B)$ is *trivial* when any closed loop of $\mathsf{WProj}_X(A \to B)$ is homotopically trivial inside X.

Remark 3.15 (locally convex wall projection). Assume that B is a locally convex subcomplex of a nonpositively curved cube complex X. Let A be any subcomplex. Then $\mathsf{WProj}_X(A \to B)$ is locally convex. Indeed let Q denote a cube of B, and let v be a vertex of Q. Then by definition, $Q \subset \mathsf{WProj}_X(A \to B)$ if and only if each edge of Q at v belongs to $\mathsf{WProj}_X(A \to B)$.

Figure 5. The above examples illustrate a length 3 interval A whose wall projection onto B is all of B. On the left B is circle and on the right it is a 3-cube. The reader can show that for any B, there is a cube complex X containing B and a subcomplex $A \cong I$ such that $\mathsf{WProj}_X(A \to B) = B$.

Lemma 3.16 (wall-projection controls retraction). *Let A and D be subcomplexes of a special cube complex B. Assume A is locally convex. Let $\widehat{D}$ denote the preimage of D in $\mathsf{C}(A, B)$, and let $r : \mathsf{C}(A, B) \to A$ be the canonical retraction map. Then $r(\widehat{D}) \subset \mathsf{WProj}_B(D \to A)$.*

Proof. Consider an edge $\hat{b}$ of $\widehat{D}$. By definition, $\hat{b}$ consists of a pair (b, b') where b is an edge of $D \subset B$ and b' is an edge of $\mathsf{C}(A, R(B))$, and b, b' map to the same edge e in $R(B)$. Let A_e be the linear subgraph of $\mathsf{C}(A, R(B))$ such that $A_e \cup b'$ is a circle and $r(\hat{b}) = A_e$. Assume A_e is not a vertex. By assumption, $A_e \cup b'$ is a connected component of the preimage of the loop e; thus each edge of A_e is parallel to b in B and $r(\hat{b}) \subset \mathsf{WProj}_B(D \to A)$. Let Q be any cube of $\mathsf{C}_\boxminus(A, B)$ contained inside $\widehat{D}$. Since the retraction map $r : \mathsf{C}_\boxminus(A, B) \to A$ is projective, there is a face $Q' < Q$ such that $r(Q) = r(Q')$ is a cube of A isomorphic to Q'. We already know the edges of $r(Q')$ belong to $\mathsf{WProj}_B(D \to A)$. Since $\mathsf{WProj}_B(D \to A)$ contains A^0 and is locally convex, we deduce that $\mathsf{WProj}_B(D \to A)$ contains $r(Q') = r(Q)$, and the lemma follows. □

3.C. *Elevations.* We now indicate some terminology related to covering spaces.

Definition 3.17 (elevations). Let $\bar{X} \to X$ denote a covering map. Let $A \subset X$ denote a connected subspace. An *elevation of A to $\bar{X}$* is a connected component of the preimage of A under $\bar{X} \to X$.

Let $A \to X$ denote a map with A connected. Consider two commutative diagrams:

$$(\mathcal{D}_1): \begin{array}{ccc} \bar{A}_1 & \to & \bar{X} \\ \downarrow & & \downarrow \\ A & \to & X, \end{array} \qquad (\mathcal{D}_2): \begin{array}{ccc} \bar{A}_2 & \to & \bar{X} \\ \downarrow & & \downarrow \\ A & \to & X, \end{array}$$

where $\bar{A}_i \to A$ are connected covers, and write $(\mathcal{D}_1) \leq (\mathcal{D}_2)$ if there is a map $\bar{A}_2 \to \bar{A}_1$ whose composition with $\bar{A}_1 \to \bar{X}, \bar{A}_1 \to A$ gives $\bar{A}_2 \to \bar{X}, \bar{A}_2 \to A$.

Say two diagrams $(\mathcal{D}_1), (\mathcal{D}_2)$ are *equivalent* if $(\mathcal{D}_1) \leq (\mathcal{D}_2)$ and $(\mathcal{D}_2) \leq (\mathcal{D}_1)$. Then, up to equivalence, $\leq$ is a partial order. An *elevation* of $A \to X$ to $\bar{X}$ is a minimal diagram.

Here is a more concrete description. Let $\bar{A}$ be a connected component of the fiber-product $A \otimes_X \bar{X}$. The projections induce two maps, $\bar{A} \to \bar{X}$ and $\bar{A} \to A$. The latter is a covering map, so we get a diagram as above:

$$(\mathcal{E}) : \begin{array}{ccc} \bar{A} & \to & \bar{X} \\ \downarrow & & \downarrow \\ A & \to & X. \end{array}$$

Then any elevation is equivalent to such a diagram $(\mathcal{E})$, and any such $(\mathcal{E})$ is an elevation.

When the map $A \to X$ is injective, the associated map $A \otimes_X \bar{X} \to \bar{X}$ is also injective. More generally, we say that a map $A \to X$ has *embedded elevations with respect to a cover* $\bar{X} \to X$ if the associated map $A \otimes_X \bar{X} \to \bar{X}$ is injective. In this case each elevation is injective, and two distinct components of $A \otimes_X \bar{X}$ have disjoint images so that connected components of $A \otimes_X \bar{X}$ are in 1-to-1 correspondence with equivalence classes of elevations. We may thus identify the elevations with their images inside $\bar{X}$, and we recover the case of connected subspaces $A \subset X$.

For example, let $A \subset X$ denote a locally convex subcomplex, with A compact and X special. Then the subcomplex $A \subset \mathsf{C}(A, X)$ is an elevation of A to $\mathsf{C}(A, X)$. By Lemma 3.13, the other elevations of A are the remaining components of $\mathsf{C}(A, A) \subset \mathsf{C}(A, X)$. We regard $A \subset \mathsf{C}(A, X)$ as the *base elevation* of A. In general, when $A, X, \bar{X}$ have basepoints and $A \to X$, $A \to \bar{X}$ are basepoint preserving maps, the *base elevation* refers to the based space $\bar{A}$ and basepoint preserving map $\bar{A} \to \bar{X}$ of the diagram $\mathcal{E}$.

4. **Connected intersection theorem**

The goal of this section is to prove Theorem 4.25 and its corollary, which will play an important role in Section 5. Given based local isometries $B_j \to X$ where $1 \leq j \leq n$, Theorem 4.25 explains how to choose a based finite cover $\widehat{X}$ such that their injective based elevations have connected intersection in $\widehat{X}$.

4.A. *Some geometric lemmas on cubical complexes.* The *distance* $d(u, v)$ between two vertices in a connected cube complex X is the length of the shortest combinatorial path joining them. A *geodesic between u and v* is a combinatorial path whose length is $d(u, v)$. For subcomplexes U, V, we let $d(U, V)$ denote the length of the shortest geodesic connecting points $u \in U, v \in V$. For a subset $S \subset X$, its cubical neighborhood $N(S)$ is defined to be the union of closed cubes intersecting S.

We recall the following concerning the combinatorial geometry of a CAT(0) cube complex $\widetilde{X}$. (For details see, for example, [Hag08].)

A combinatorial path in $\widetilde{X}$ is a geodesic if and only if the sequence of hyperplanes it crosses has no repetition. A path σ in a nonpositively curved cube complex X whose lift $\widetilde{\sigma}$ is a combinatorial geodesic of the universal cover $\widetilde{X}$ will be called a *local geodesic*.

- — A full subcomplex $Y \subset \widetilde{X}$ is *convex* if and only if it is combinatorially convex, in the sense that any combinatorial geodesic of $\widetilde{X}$ with endpoints in Y has all of its vertices inside Y. A subcomplex A in a cubical complex Z is *full* provided that a cube of Z belongs to A if and only if its vertices do.
- — For any convex subcomplex $Y \subset \widetilde{X}$, its cubical neighborhood $N(Y)$ is again a convex subcomplex.

Let S be a subset of $\widetilde{X}$. We define S^{+0} to be the smallest subcomplex of $\widetilde{X}$ containing S. For $R \geq 1$, we define $S^{+R} = N(S^{+(R-1)})$ to be the *cubical R-thickening of S in $\widetilde{X}$*. When Y is a convex subcomplex, we have $Y^{+0} = Y$, and consequently Y^{+R} is convex for each R. For a 0-cell v, the subcomplex v^{+R} is called the *cubical ball with center v and radius R*. We note that each cubical ball about v is convex whereas combinatorial metric balls are often not convex. Note that $H^{+0} = N(H)$ when H is a hyperplane and, in this case, $N(H)$ is actually a convex subcomplex, and consequently H^{+R} is convex for each R.

- — Every combinatorial path with initial point v and length $\leq R$ is contained in v^{+R}.
- — The *convex hull* of a subcomplex $Y \subset \widetilde{X}$ is the intersection of all convex subcomplexes containing Y.

Remark 4.1. If X has dimension $\leq D$, then any vertex x in v^{+R} is joined to v by an edge-path of length $\leq DR$. In particular, v^{+R} has diameter $\leq 2DR$.

A fundamental result is then

Theorem 4.2 (convex hull). *Let X denote a* CAT(0) *cube complex. Assume that X is uniformly locally finite, in the sense that there is a uniform bound on the number of edges containing a given vertex. Then there exists a function $L \mapsto R(L)$ with the following property.*

For any subcomplex $Y \subset X$, if Y is L-quasiconvex, then the convex hull of Y is contained inside $Y^{+R(L)}$.

Here we say that Y is *L-quasiconvex* if any vertex of a combinatorial geodesic with endpoints inside Y is at combinatorial distance $\leq L$ of some vertex in Y. For a proof of Theorem 4.2 see, for instance, [Hag08].

We will also make use of the following fundamental fact, which the reader can find in [Ger97] and [Rol].

THEOREM 4.3. *Let* $Y_1, \dots, Y_r$ *be convex subcomplexes of the* CAT(0) *cube complex* X. *Suppose that* $Y_i \cap Y_j$ *is nonempty for each* i, j. *Then* $\cap_1^r Y_i$ *is nonempty.*

Definition 4.4 (R-embeddings). Let X be a nonpositively curved, connected cube complex. Let $Y \to X$ denote any local isometry with Y connected. There is an induced equivariant isometric embedding $\widetilde{Y} \to \widetilde{X}$. We have already defined the cubical R-thickening $\widetilde{Y}^{+R}$ of $\widetilde{Y}$ in the simply-connected $\widetilde{X}$. For any integer $R \geq 0$, the *cubical* R*-thickening* of Y is $Y^{+R} := \pi_1(Y)\backslash\widetilde{Y}^{+R}$.

The map $Y \to X$ induces a local isometry $Y^{+R} \to X$, which will play an important role below. When $Y^{+R} \to X$ is an embedding, we say that $Y \to X$ is an R*-embedding*, In particular, a locally convex subcomplex $Y \subset X$ is R*-embedded* when the inclusion $Y \hookrightarrow X$ is R-embedding. In this case, we identify Y^{+R} with its image in X and refer to this locally convex subcomplex as the R*-thickening* of Y in X. Lifting the situation to the universal cover, we see that a locally convex subcomplex $Y \subset X$ fails to be R-embedded precisely when there is a point $\tilde{p} \in \widetilde{Y}^{+R}$ and an element $g \in \pi_1 X - \text{Stabilizer}(\widetilde{Y})$ such that $g\tilde{p} \in \widetilde{Y}^{+R}$.

The *embedding radius of* $Y \subset X$ is the supremum of the integers $R \geq 0$ for which $Y \subset X$ is R*-embedded.* The *injectivity radius of* X is the minimum of the embedding radii of vertices, denoted by $\mathsf{InjRad}(X)$. In other words, $\mathsf{InjRad}(X) \geq R$ if and only if $v^{+R} \to X$ is injective for any vertex v.

When $H \to X$ is a hyperplane, we likewise define the R*-thickening* of $H \to X$ to be $H^{+R} := \pi_1(H)\backslash\widetilde{H}^{+R}$. There is an induced local isometry $H^{+R} \to X$, and when it embeds, we identify H^{+R} with its image and call it the R*-thickening of* H *in* X.

The *hyperplane embedding radius of* X is the minimum of the embedding radii of hyperplane neighborhoods, denoted by $\mathsf{HEmbRad}(X)$. In other words, $\mathsf{HEmbRad}(X) \geq R$ if and only if for any hyperplane H of X with cubical neighborhood $N(H)$, we have that H^{+R} embeds in X.

Note that Y^{+R} is compact provided Y is compact and X is locally finite.

LEMMA 4.5. *Let* A, B *be locally convex connected subcomplexes of the connected nonpositively curved cube complex* X. *Suppose that* A *and* B *are* R*-embedded and that* $A \cap B$ *is connected. Then*

(1) $A \cap B$ *is* R*-embedded.*
(2) $(A \cap B)^{+R}$ *equals the component of* $A^{+R} \cap B^{+R}$ *containing* $A \cap B$.

Proof. We choose a basepoint x in $C = A \cap B$ and let $\widetilde{A}, \widetilde{B}, \widetilde{C}$ be the based elevations at $\tilde{x} \in \widetilde{X}$. It follows from convexity that $\widetilde{C} = \widetilde{A} \cap \widetilde{B}$.

We prove the first statement. Let $\tilde{p} \in \widetilde{C}^{+R}$ and $g \in \pi_1 X$ such that $g\tilde{p} \in \widetilde{C}^{+R}$. We show that $g \in \pi_1 C = \text{Stabilizer}(\widetilde{C})$. Observe that $g\tilde{p} \in \widetilde{A}^{+R}$, and so since A is R-embedded, we see that $g \in \text{Stabilizer}(\widetilde{A})$. Likewise, $g \in \text{Stabilizer}(\widetilde{B})$. Thus $g \in \text{Stabilizer}(\widetilde{C})$ as claimed.

We now prove the second statement. Let σ be a path in $A^{+R} \cap B^{+R}$ starting at $x \in A \cap B$ and ending at some point $p \in A^{+R} \cap B^{+R}$. Let $\tilde{\sigma}$ be the elevation of σ at $\tilde{x}$, and let $\tilde{p}$ be the endpoint of $\tilde{\sigma}$. Observe that $\tilde{\sigma}$ lies entirely in $\widetilde{A}^{+R}$ and, consequently, $\tilde{p} \in \widetilde{A}^{+R}$. Thus $\tilde{p}^{+R} \cap \widetilde{A}$ is nonempty. Similarly, $\tilde{p}^{+R} \cap \widetilde{B}$ is nonempty. Applying Theorem 4.3 we see that $\tilde{p}^{+R} \cap \widetilde{C} = \tilde{p}^{+R} \cap \widetilde{A} \cap \widetilde{B}$ is nonempty, and so taking the images in X, we see that $p \in (A \cap B)^{+R}$. $\square$

Lemma 4.6. *Let X be a nonpositively curved cube complex with $\mathsf{InjRad}(X) \geq R$. Let $Y \subset X$ be any connected subcomplex where any vertex is joined by a path with $\leq R$ edges to some fixed vertex y. Then Y is null-homotopic in X.*

Proof. Indeed the CAT(0) complex y^{+R} embeds in X and contains Y. $\square$

Lemma 4.7 (short self-connections $\Rightarrow$ small embedding radius). *Let $R > 0$ be an integer. Let H denote a hyperplane of a nonpositively curved cube complex X. Let β be a local geodesic of length $k \leq 2R$ between two vertices of $N(H)$. If β is not contained in $N(H)$, then the embedding radius of H is $< R$.*

Proof. Any lift of β to the universal cover $\widetilde{X}$ is a combinatorial geodesic $\widetilde{\beta}$ connecting the cubical neighborhoods of hyperplanes $\widetilde{H}, \widetilde{H'}$ projecting onto H in X. Note that $\widetilde{H} \neq \widetilde{H'}$, otherwise by convexity of cubical neighborhoods of hyperplanes, the path β would stay inside $N(H)$.

By assumption there is a vertex x (on $\widetilde{\beta}$) at distance $\leq R$ of both $N(\widetilde{H})$ and $N(\widetilde{H'})$. Choose $g \in \pi_1 X$ such that $g\widetilde{H'} = \widetilde{H}$, and note that $g \notin \pi_1 H$. Then the vertex $x' = gx$ is not identified with x in $H^{+R} = \pi_1 H \backslash \widetilde{H}^{+R}$, but x, x' are identified in X. $\square$

Using the fundamental group interpretation for elevations and the fact that the natural map $Y \to Y^{+R}$ is a π_1-isomorphism, we get

Lemma 4.8 (elevations of neighborhoods). *Let X denote a connected, nonpositively curved cube complex, and let $Y \to X$ denote a local isometry with Y connected. Assume $X' \to X$ is a cover, and let $R \geq 0$ denote some integer. For any elevation $Y' \to X'$ of $Y \to X$, the map $Y'^{+R} \to X'$ is an elevation of $Y^{+R} \to X$.*

In particular, any elevation of an R-embedding is an R-embedding. For any cover $X' \to X$, we always have $\mathsf{InjRad}(X') \geq \mathsf{InjRad}(X)$ and $\mathsf{HEmbRad}(X') \geq \mathsf{HEmbRad}(X)$.

LEMMA 4.9 (virtually high embedding radius). *Let X denote a compact, connected, nonpositively curved special cube complex. Let $Y \to X$ denote a local isometry of compact cube complexes. Then for each integer $R \geq 0$, there is a finite cover $X_{Y,R} \to X$ such that for any further cover $X' \to X_{Y,R}$, any elevation $Y' \to X'$ of $Y \to X$ is an R-embedding.*

Note that for $R = 0$, Lemma 4.9 provides a finite cover where each elevation is injective.

Proof. We canonically complete the local isometry $Y^{+R} \to X$. We thus get a finite cover $\mathsf{C}(Y^{+R}, X) \to X$ such that $Y \to \mathsf{C}(Y^{+R}, X)$ is an R-embedding. We now let $X_{Y,R} \to X$ denote any regular finite cover factoring through $\mathsf{C}(Y^{+R}, X) \to X$. By Lemma 4.8 some elevation of $Y \subset X$ to $X_{Y,R}$ is an R-embedding. But by regularity all elevations are. Thus the lemma holds for $X' = X_{Y,R}$. Applying again Lemma 4.8, we deduce that the lemma holds for arbitrary $X' \to X_{Y,R}$. □

COROLLARY 4.10 (virtually high (hyperplane) injectivity radius). *Let X denote a compact, connected, nonpositively curved special cube complex. Then for any integer $R \geq 0$, there is a finite cover $X' \to X$ such that for any further cover $X'' \to X'$, we have $\mathsf{InjRad}(X'') \geq R$ and $\mathsf{HEmbRad}(X'') \geq R$.*

Proof. For Y any singleton $\{v\}$ or any hyperplane neighborhood, we consider a finite cover $X_{Y,R} \to X$ as in Lemma 4.9. We then let $X' \to X$ denote any finite cover factoring through the finitely many finite covers $X_{Y,R} \to X$. We conclude as above with Lemma 4.8. □

LEMMA 4.11 (embedding radius of convex amalgams). *Let A, B be connected, locally convex subcomplexes of a nonpositively curved cube complex X such that $A \cup B$ is connected and locally convex. Assume the embedding radii of A, B are $\geq R$ and each component of $A^{+R} \cap B^{+R}$ intersects a component of $A \cap B$. Then the embedding radius of $A \cup B$ is $\geq R$.*

Proof. Set $Y = A \cup B$. Let $\widetilde{Y}$ be some lift of Y to the universal cover $\widetilde{X}$. Let $\gamma \in \pi_1 X$ map $\tilde{p} \in \widetilde{Y}^{+R}$ to $\tilde{p}' \in \widetilde{Y}^{+R}$, and let us show that γ stabilizes $\widetilde{Y}$.

We fix a base vertex $\tilde{x}$ in $\widetilde{Y}$ mapping inside $A \cap B$ and denote by $\widetilde{A}, \widetilde{B}$ the lifts of A, B at $\tilde{x}$. The subspace $\widetilde{Y}$ is covered by the translates $g\widetilde{A}, h\widetilde{B}$ for $g, h \in \pi_1 Y$; thus $\widetilde{Y}^{+R}$ is covered by the translates $g\widetilde{A}^{+R}, h\widetilde{B}^{+R}$. The stabilizer of $\widetilde{Y}$ is identified with $\pi_1 Y$.

We may assume that $\tilde{p} \in (g\widetilde{A})^{+R}$. There are two possibilities for $\tilde{p}'$. The first possibility is that there exists $h \in \pi_1 Y$ such that $\tilde{p}' \in (h\widetilde{A})^{+R}$. In this case we note that $h^{-1}\gamma g\widetilde{A}^{+R}$ intersects $\widetilde{A}^{+R}$. Since the embedding radius of A is $\geq R$ it follows that $h^{-1}\gamma g \in \pi_1 A \subset \pi_1 Y$, and thus $\gamma \in \pi_1 Y$.

The second possibility is that there exists $h \in \pi_1 Y$ such that $\tilde{p}' \in (h\widetilde{B})^{+R}$. Since both $g, h \in \pi_1 Y$, we have $\gamma \in \pi_1 Y \iff h^{-1}\gamma g \in \pi_1 Y$. In other words, we may assume $g = h = 1$. Let p denote the common image in X of $\tilde{p}$ and $\tilde{p}'$, and note that $p \in A^{+R} \cap B^{+R}$. By the relative connectedness assumption, there is a path σ inside $A^{+R} \cap B^{+R}$ connecting p to $q \in A \cap B$. Choose a path $\tilde{\alpha} \subset \widetilde{A}^{+R}$ connecting $\tilde{x}$ to $\tilde{p}$, then compose its image α inside X with σ. Since $A \subset A^{+R}$ is a π_1-isomorphism, the product $\alpha\sigma$ is homotopic with fixed endpoints inside A^{+R} to a path a contained in A. Thus up to translating by an element of Stabilizer$(\widetilde{A})$, the point $\tilde{p}$ is the endpoint of the lift at $\tilde{x}$ of the path $a\sigma^{-1}$. Similarly, up to translating by an element of Stabilizer$(\widetilde{B})$, there is a path $b \subset B$ connecting x to the endpoint of σ such that $\tilde{p}'$ is the endpoint of the lift at $\tilde{x}$ of the path $b\sigma^{-1}$. It follows that $\gamma \in \pi_1 A a^{-1} b \pi_1 B$, and so $\gamma \in \pi_1 Y$. $\square$

4.B. *Quasiconvex amalgams.* In a δ-hyperbolic space X, any local quasigeodesic segment is uniformly near to a geodesic with the same endpoints. This implies that a concatenation of geodesics whose overlaps are sufficiently small yields a quasigeodesic of X. This well-known fact can be used to show that a subspace obtained by combining various quasiconvex subspaces of X is still quasiconvex. Below, we give two precise statements of this type where X is a CAT(0) cube complex.

Lemma 4.12 (quasiisometric line of spaces). *Let X be a* CAT(0) *cube complex that is δ-hyperbolic. There exist $R_0 \geq 0$ and $L_0 \geq 1$ with the following property.*

Let $Y_0, \ldots, Y_m$ be a sequence of convex subcomplexes of X such that $Z_{i+1} := Y_i \cap Y_{i+1}$ is nonempty for each $0 \leq i < m$. Let $Y = Y_0 \sqcup_{Z_1} Y_1 \sqcup_{Z_2} \cdots \sqcup_{Z_m} Y_m$, and note that there is an induced map $\phi : Y \to X$.

If $d(Z_i, Z_{i+1}) > R_0$ for each $0 \leq i < m$, then for $y_0 \in Y_0$, $y_m \in Y_m$, we have $d_X(\phi(y_0), \phi(y_m)) \leq d_Y(y_0, y_m) \leq L_0 d_X(\phi(y_0), \phi(y_m))$.

Proof. This is implicit in the proof of the quasiconvex amalgam assertion proven in [HaW08, Lemma 8.11]. See also [Git99]. $\square$

Lemma 4.13 (quasiisometric tree of spaces). *Let X be a δ-hyperbolic* CAT(0) *cube complex. There exist $R_0 \geq 0$ and $L_0 \geq 1$ with the following property.*

Let Γ be a graph with universal cover $\widetilde{\Gamma}$. For each $v \in \Gamma^0$, let Y_v denote a convex subcomplex of X. For each edge $\{u, v\}$ of Γ, we let $Y_{\{u,v\}} = Y_u \cap Y_v$. Suppose that each such $Y_{\{u,v\}}$ is nonempty.

Let Y_Γ denote the abstract union of the Y_v along their pairwise intersections, and let $Y_{\widetilde{\Gamma}}$ be the corresponding tree of spaces. Consider the map $Y_{\widetilde{\Gamma}} \to X$ obtained by composing the universal cover $Y_{\widetilde{\Gamma}} \to Y_\Gamma$ with the natural map $Y_\Gamma \to X$.

If $d(Y_{\{u,v\}}, Y_{\{v,w\}}) > R_0$ for each pair of distinct adjacent edges $\{u, v\}$, $\{v, w\}$, then $Y_{\widetilde{\Gamma}} \to X$ is an $(L_0, 0)$-quasiisometric embedding (that is, an L_0-bilipschitz embedding). It follows, in particular, that Γ is a tree.

Proof. This follows from Lemma 4.12. □

LEMMA 4.14 (immersed quasiconvex amalgam). *Let X be a compact nonpositively curved cube complex. Assume the universal cover $\widetilde{X}$ is δ-hyperbolic. Then there are constants R_0, K_0 depending only on $\widetilde{X}$ such that the following holds.*

Let A, B, C be (nonempty) connected locally convex subcomplexes of X such that C is a connected component of $A \cap B$. Consider the space $S = A \cup_C B$ and the natural map $f : S \to X$. If both A and B are R_0-embedded, then $S \to X$ factors through an embedding $S \to T$ and a local isometry $T \to X$ such that T is compact, $S \to T$ is a π_1-isomorphism, and every point $p \in T$ is at distance $\leq K_0$ of a point of S.

In particular, $S \to X$ is π_1-injective.

Proof. The universal cover $\widetilde{S}$ consists of a collection of copies of universal covers of $\widetilde{A}$ and $\widetilde{B}$. The nerve of this covering of $\widetilde{S}$ by this collection of subspaces is a tree Γ. The map $S \to X$ induces a map $\widetilde{S} \to \widetilde{X}$, putting us in the framework of Lemma 4.13. We can thus conclude that $\widetilde{S} \to \widetilde{X}$ is an $(L_0, 0)$-quasiisometric embedding.

We now regard $\widetilde{S}$ as a subcomplex of $\widetilde{X}$. Since $(L_0, 0)$-quasigeodesics κ-fellow travel- geodesics for some $K_0 = K_0(L_0, \delta)$, we see that $\widetilde{S}$ is actually K_0-quasiconvex.

Apply Theorem 4.2 to $\widetilde{S} \subset \widetilde{X}$ to obtain a convex $\pi_1 S$-invariant subcomplex $\widetilde{T}$ that is contained in the combinatorial $R(L_0)$-neighborhood of $\widetilde{S}$ and is thus $\pi_1 S$-cocompact. We let $T = \pi_1 S \backslash \widetilde{T}$ and note that $S \to X$ factors as $S \to T \to X$ to satisfy our claim with $K_0 = R(L_0) \times \sqrt{\dim(X)}$, which accounts for the difference between the combinatorial neighborhoods considered here and the CAT(0) metric balls. □

COROLLARY 4.15 (embedded quasiconvex amalgam). *Let X be a compact nonpositively curved special cube complex. Assume the universal cover $\widetilde{X}$ is Gromov-hyperbolic. Then there are constants $R \geq K$ such that the following holds.*

Let A, B, C be (nonempty) connected locally convex subcomplexes of X such that C is a connected component of $A \cap B$. Consider the space $S = A \cup_C B$ and the natural map $f : S \to X$. If the embedding radii of A and B are $\geq R$, then there exists a finite cover $X' \to X$ such that $S \to X$ factors as $S \subset X' \to X$, and there is a connected wall-injective locally convex subcomplex T of X' that contains S, with every point $p \in T$ at distance $\leq K$ of S, and $S \subset T$ is a π_1-isomorphism. Furthermore, T is the union of two locally convex subcomplexes $\mathcal{A}, \mathcal{B}$ such that

(1) $A \subset \mathcal{A} \subset A^{+K}$ *and* $B \subset \mathcal{B} \subset B^{+K}$.
(2) $\mathcal{A} \cap \mathcal{B}$ *is connected and* $\mathcal{A} \cap \mathcal{B} \subset (A \cap B)^{+K} = C^{+K} \subset X'$.

We say that $T \subset X'$ *is an embedded locally convex thickening of the amalgam* $S \to X$.

Remark 4.16. Note that $C \to X'$ is an R-embedding by Lemma 4.5.1 so, in particular, $C^{+K} \subset X'$ is an embedding. Note also that $C \to \mathcal{A} \cap \mathcal{B}$ is a π_1-isomorphism. Indeed composing $C \to \mathcal{A} \cap \mathcal{B}$ with $\mathcal{A} \cap \mathcal{B} \to C^{+K}$ yields a π_1-isomorphism.

Proof of Corollary 4.15. Let R_0, K_0 be the constants of Lemma 4.14. Since X is finite dimensional, let K be an integer such that for any convex subcomplex $\widetilde{Y} \subset \widetilde{X}$, the cubical neighborhood $(\widetilde{Y})^{+K}$ contains each point whose CAT(0) distance to $\widetilde{Y}$ is $\leq K_0$. Let $R = \max(R_0, K)$, and let A, B be subcomplexes of X with embedding radius $\geq R$. We apply Lemma 4.14 to X, A, B, C to extend the amalgam $S \to X$ to a local isometry that we denote by $T \to X$.

For any choice of a lift $\widetilde{A}$ to $\widetilde{T} \subset \widetilde{X}$, the intersection with $\widetilde{T}$ of the cubical K-thickening of $\widetilde{A}$ inside $\widetilde{X}$ is $\pi_1 A$-invariant, and we let $\mathcal{A} = \pi_1 A \backslash \big((\widetilde{A})^{+K} \cap \widetilde{T}\big)$. Since $R \geq K$, the natural local isometry $\mathcal{A} \to T$ is injective. We define $\mathcal{B} \subset T$ similarly. We then have $T = \mathcal{A} \cup \mathcal{B}$ since each point in T is at distance $\leq K_0$ of either A or B.

We now show that $\mathcal{A} \cap \mathcal{B}$ is connected. Let $\mathcal{C}$ denote the connected component of C inside $\mathcal{A} \cap \mathcal{B}$. Form the space $\mathcal{S} := \mathcal{A} \cup_{\mathcal{C}} \mathcal{B}$. Since $S \to T$ is π_1-surjective, the composition $S \to \mathcal{S} \to T$ shows that $\mathcal{S} \to T$ is π_1-surjective. If $\mathcal{A} \cap \mathcal{B}$ is not connected, then there are nontrivial connected covers of T to which $\mathcal{S}$ lifts isomorphically. But this implies that the image of $\pi_1 \mathcal{S}$ inside $\pi_1 T$ is a proper subgroup, which is impossible.

Since we have shown that $\mathcal{A} \cap \mathcal{B}$ is connected, it follows from Lemma 4.5.2 that $(\mathcal{A} \cap \mathcal{B}) \subset (A \cap B)^{+K}$. Finally, $T \subset X'$ is wall-injective by Corollary 3.11. □

4.C. *Virtually connected intersection.* Corollary 4.15 above is a precise statement of the following idea: given two locally convex subcomplexes A, B of a compact special cube complex X and a connected component C of $A \cap B$, we are able to embed A, B in a finite cover $X' \to X$ so that their intersection in X' is just C. Our goal in this section is to adapt this construction to an arbitrary finite collection $B_1, \ldots, B_n$ of locally convex subcomplexes of X that all contain a given vertex. We will do this where X is a compact special cube complex with word-hyperbolic fundamental group, but in the simple case where X is a graph, our statement reads

Proposition 4.17 (connected intersection property for subgraphs). *Let X be a finite graph with base vertex v. Let $B_1, \ldots, B_n$ be connected subgraphs containing v. Then there is a finite cover $X' \to X$ based at v' such that any two of the elevations $B_1', \ldots, B_n'$ of $B_1, \ldots, B_n$ at v' have a connected intersection.*

Moreover, the elevation C' at v' of the connected component C of v inside $B_1 \cap \cdots \cap B_n$ maps isomorphically onto C.

When $n \geq 3$, our elevations of the B_i's to the finite cover $X' \to X$ cannot always remain isomorphic to B_i. (Think of $B_3 = X$ and $B_1 \cap B_2$ not connected.) The proof in the general case may be thought of as a thickening of the graph case.

Lemma 4.18 (à la Helly). *Suppose the nonpositively curved cube complex $\mathcal{C}$ deformation retracts to the locally convex connected subcomplex C. Let $B_1, \ldots, B_n$ be locally convex connected subcomplexes of $\mathcal{C}$ such that $B_i \cap C$ is nonempty for each i. Then*

(1) *$B_i \cap C$ is connected.*
(2) *Each component of $B_1 \cap B_2$ contains a point of C.*
(3) *If $\cap_i(B_i \cap C)$ is connected, then $\cap_i B_i$ is connected.*

Proof. We first prove the first claim. Let γ be a local geodesic in B_i that starts and ends on $B_i \cap C$. Let $\widetilde{\gamma}$ be a lift of γ. Let $\widetilde{B}_i$ be the lift of B_i that contains $\widetilde{\gamma}$. Note that since $\mathcal{C}$ deformation retracts onto C, there is only one component $\widetilde{C}$ in the preimage of C. It follows that the geodesic path $\widetilde{\gamma}$ is contained in the convex subcomplex $\widetilde{C}$. Thus $\widetilde{\gamma} \subset \widetilde{B}_i \cap \widetilde{C}$, and so $\gamma \subset B_i \cap C$.

We now prove the second claim. For $i \in \{1, 2\}$, let σ_i be a path in B_i from $B_i \cap C$ to $p \in B_1 \cap B_2$. Since $\mathcal{C}$ deformation retracts to C, we can let σ be a path in C that is path homotopic to $\sigma_1\sigma_2^{-1}$, and so $\sigma\sigma_2\sigma_1^{-1}$ lifts to a closed path $\widetilde{\sigma}\widetilde{\sigma}_2\widetilde{\sigma}_1^{-1}$ in $\widetilde{\mathcal{C}}$. Let $\widetilde{\Sigma}_i$ be the smallest convex subcomplex of $\mathcal{C}$ containing $\widetilde{\sigma}_i$, and note that $\widetilde{\Sigma}_i \subset \widetilde{B}_i$. Similarly, let $\widetilde{\Sigma}$ be the smallest convex subcomplex containing $\widetilde{\sigma}$ and note that $\widetilde{\Sigma} \subset \widetilde{C}$. The closed path $\widetilde{\sigma}\widetilde{\sigma}_2\widetilde{\sigma}_1^{-1}$ shows that the convex subcomplexes $\widetilde{\Sigma}_1, \widetilde{\Sigma}_2, \widetilde{\Sigma}$ have nonempty pairwise intersection, and so Helly's Theorem 4.3 shows that $\widetilde{\Sigma}_1 \cap \widetilde{\Sigma}_2 \cap \widetilde{\Sigma}$ is nonempty. The image of $\widetilde{\Sigma}_1 \cap \widetilde{\Sigma}_2$ is thus a connected subspace of $B_1 \cap B_2$; moreover, it passes through p and some point of C.

The third claim follows from the following statement:

(3′) Each component of $\cap_i B_i$ contains a point of C.

This follows from (2) by induction on n. Indeed given any family of subcomplexes $B_1, \ldots, B_n, B_{n+1}$ satisfying the assumption of the lemma, let B_n' be any component of $B_n \cap B_{n+1}$. By (2) the inductive hypothesis applies to $B_1, \ldots, B_{n-1}, B_n'$. □

The *profinite topology* on a group G is generated by the basis of finite index cosets. The profinite topology is hausdorff exactly when G is residually finite, which holds exactly when $\{1_G\}$ is separable. A subset $S \subset G$ is *separable* if S is closed in the group G relative to the profinite topology. Equivalently,

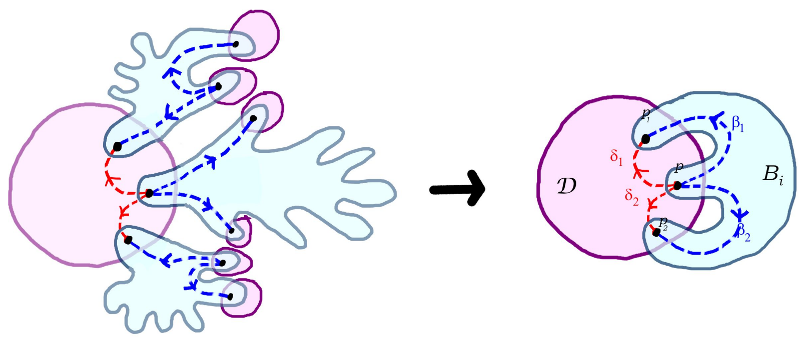

Figure 6. The cover $\widehat{X}_i \to X$ is partially illustrated above.

S is separable provided that for each $g \notin S$, there is a finite quotient $G \to \bar{G}$ such that $\bar{g} \notin \bar{S}$. Typically, S is a subgroup of G or a double coset $AB \subset G$ where A, B are subgroups.

Remark 4.19. When X is compact virtually special and $\pi_1 X$ is word-hyperbolic, the quasiconvex subgroups of $\pi_1 X$ are separable [HaW08], and consequently the double quasiconvex cosets are separable as proven by Minasyan [Min06].

Separability of a subgroup can often be applied to reduce self-intersections of a lifted subspace in a finite cover. Suppose Y is compact and $Y \to X$ lifts to an embedding $Y \hookrightarrow \widehat{X}$ in a covering space $\widehat{X} \to X$ with $\pi_1 \widehat{X}$ separable in $\pi_1 X$. As explained by Scott [Sco78], there then exists a finite intermediate cover $\widehat{X} \to X$ such that Y lifts to an embedding in $\widehat{X}$. Along this vein, the following lemma uses double coset separability to reduce intersections in a cover.

Lemma 4.20 (separating $\mathcal{D}$ from stuff). *Let X be a compact connected nonpositively curved cube complex with basepoint p. Suppose $B_1, \dots, B_n, \mathcal{D}$ are locally convex, connected, subcomplexes of X containing p. Suppose $\pi_1 \mathcal{D} \pi_1 B_i$ is separable in $\pi_1 X$ for each i.*

Then there is a based finite cover $\widehat{X}$ such that, letting $\widehat{\mathcal{D}}, \widehat{B}_1, \dots, \widehat{B}_n$ denote the based elevations, we have $\widehat{\mathcal{D}} \cong D$ and, moreover, each intersection $\widehat{\mathcal{D}} \cap \widehat{B}_i$ is connected.

Proof. It suffices to prove the result for $n = 1$. For then, treating each pair $\mathcal{D}, B_i$ independently, for each i we have a finite based cover $\widehat{X}_i \to X$ in which $\widehat{\mathcal{D}}_i \cong D$ and the based elevation of B_i has connected intersection with $\widehat{\mathcal{D}}_i$. The fiber-product of the covers $\widehat{X}_i \to X$ has the desired property. We write B for B_1 below.

The base component $B \cap \mathcal{D}$ contains the basepoint p. Let p_k denote a point in each other component C_k of $B \cap \mathcal{D}$. For each k, let β_k denote a local geodesic in B from p to p_k, and let δ_k denote a local geodesic in $\mathcal{D}$ from p to p_k.

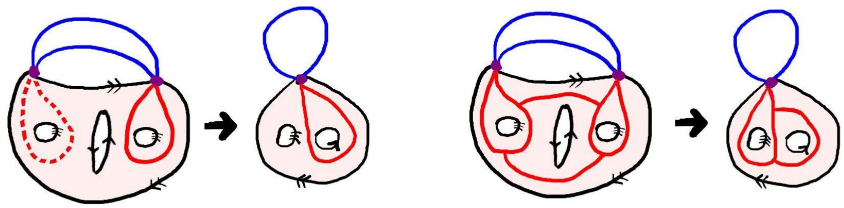

Figure 7. Synchronism fails on the left but succeeds on the right.

Observe that $\delta_k\beta_k^{-1} \notin \pi_1\mathcal{D}\pi_1 B$. Indeed suppose $b\beta_k$ is homotopic to $d\delta_k$ for some $b \in \pi_1 B$ and $d \in \pi_1\mathcal{D}$. Their lifts to $\widetilde{X}$ are homotopic to geodesic paths δ, β. But then $\delta = \beta$ by uniqueness of geodesics, so p_k and p must lie in the same component of $B \cap \mathcal{D}$.

By separability, let N be a finite index normal subgroup of $\pi_1 X$ such that for each k, we have $\delta_k\beta_k^{-1} \notin N\pi_1\mathcal{D}\pi_1 B_i$. Let $(\widehat{X}, \widehat{p}) \to (X, p)$ be the finite based cover corresponding to $N\pi_1\mathcal{D}$, and let $\widehat{\mathcal{D}}, \widehat{B}$ denote the based elevations of $\mathcal{D}, B$. Note that $\widehat{\mathcal{D}} \cong D$.

Suppose $\widehat{\mathcal{D}} \cap \widehat{B}$ has a component not containing $\widehat{p}$. Then there is a path $\hat{\beta}$ inside $\widehat{B}$ from $\widehat{p}$ to some preimage of p_k. The path $\hat{g} = \delta_k\hat{\beta}^{-1}$ is closed in $\widehat{X}$, thus $\hat{g} \in N\pi_1\mathcal{D}$. We also have $\hat{g} = \delta_k\beta_k^{-1}b$ for some $b \in \pi_1 B$, which is a contradiction. $\square$

In our applications of Lemma 4.20 the cube complex X is special and $\pi_1 X$ is word-hyperbolic, and we may thus use Corollary 4.15 to get an alternative argument.

Let X have a basepoint x. We say $Y \subset X$ is a *based subspace* to mean that $x \in Y$ and that x is the basepoint of Y.

Definition 4.21. Let $V, U, \mathcal{U}$ be connected based subspaces of X. We say $V, \mathcal{U}$ is *synchronized with* V, U, if for each based cover $\widehat{X}$ and based elevations $\widehat{V}, \widehat{U}, \widehat{\mathcal{U}}$, the subspace $\widehat{V} \cap \widehat{\mathcal{U}}$ is *connected relative to* $\widehat{V} \cap \widehat{U}$ in the sense that each component of $\widehat{V} \cap \widehat{\mathcal{U}}$ contains a component of $\widehat{V} \cap \widehat{U}$.

By definition, the property of being synchronized is stable under further cover.

Example 4.22. Let $\mathcal{U}$ denote a pair of pants, let V denote a circle, and let $X = \mathcal{U} \vee V$ denote their wedge along a boundary point of $\mathcal{U}$. As illustrated on the left of Figure 7, synchronism can fail if we let U denote a circle, but as illustrated on the right, synchronism succeeds if we let U denote a π_1-surjective graph in $\mathcal{U}$.

LEMMA 4.23 (relatively connected intersection + π_1-surjective $\Rightarrow$ synchronized). *Let $V, U, \mathcal{U}$ be connected based subspaces of X. Suppose $U \subset \mathcal{U}$. Suppose $\pi_1 U \to \pi_1\mathcal{U}$ is surjective and $V \cap \mathcal{U}$ is connected relative to $V \cap U$. Then $V, \mathcal{U}$ is synchronized with V, U.*

Proof. Let $f : \widehat{X} \to X$ be a cover whose basepoint $\widehat{p}$ maps to the basepoint p of X and $\widehat{V}, \widehat{U}, \widehat{\mathcal{U}}$ be the based elevations. Let p_i be a point in each component of $U \cap V$. Note that $\widehat{\mathcal{U}} \cap \widehat{V}$ is the disjoint union of components covering components of $\mathcal{U} \cap V$. So by the relative connectivity of $\mathcal{U} \cap V$, each component of $\widehat{\mathcal{U}} \cap \widehat{V}$ contains a point of $f^{-1}(p_i) \cap \widehat{\mathcal{U}}$. But since $\pi_1 U \to \pi_1 \mathcal{U}$ is surjective, $f^{-1}(p_i) \cap \widehat{\mathcal{U}} = f^{-1}(p_i) \cap \widehat{U}$. Consequently each component of $\widehat{\mathcal{U}} \cap \widehat{V}$ contains a component of $\widehat{U} \cap \widehat{V}$. $\square$

LEMMA 4.24 (intermediate synchronism). *Let V and $U \subset \mathcal{U} \subset U^+$ be connected locally convex based subcomplexes of X, and assume that $\mathcal{U} \to U^+$ is π_1-surjective. If V, U^+ is synchronized with V, U, then $V, \mathcal{U}$ is synchronized with V, U.*

Proof. Let $\widehat{X}$ be a connected cover, and note that $\widehat{V}$ and $\widehat{U} \subset \widehat{\mathcal{U}} \subset \widehat{U}^+$ are connected and locally convex. Let $q \in \widehat{V} \cap \widehat{\mathcal{U}}$. Since V, U^+ is synchronized with V, U, we see that there is a path σ in $\widehat{U}^+ \cap \widehat{V}$ from q to $p \in \widehat{U}$. By π_1-surjectivity σ can be homotoped to a local geodesic γ in $\widehat{\mathcal{U}}$. By local convexity γ lies in $\widehat{V}$. Thus the component of $\widehat{V} \cap \widehat{\mathcal{U}}$ containing q also contains $p \in \widehat{V} \cap \widehat{U}$. $\square$

THEOREM 4.25 (virtually connected intersection). *Let X be a compact special cube complex with $\widetilde{X}$ δ-hyperbolic. Let $(B_0, \ldots, B_n, A)$ be connected locally convex subcomplexes containing the basepoint of X. Suppose that $A \subset \bigcap_{j \in \{0,\ldots,n\}} B_j$. Then there are a based finite cover $\bar{X}$ and based elevations $\bar{B}_0, \ldots, \bar{B}_n$ with $\bar{A} \cong A$ such that $\cap_{j \in J} \bar{B}_j$ is connected for each $J \subset \{0, \ldots, n\}$.*

The *p-component of S* denoted by $[S]_p$, is the component of S containing the point p, and we use the notation $[S] = [S]_b$, where b is the basepoint.

When indices are clear from the context, we will use the notation $B_J = \cap_{j \in J} B_j$ and $\ddot{B}_J = \cap_{j \in J} \ddot{B}_j$, etc. Given a subspace $A \subset X$ containing the basepoint, and a based cover $\widehat{X}$, we will employ the notation $\widehat{A}$ to denote the based elevation of A. Likewise, $\ddot{X}, \ddot{A}$ and $\bar{X}, \bar{A}$, etc. Using these notations we have $[\widehat{B}_J] = \widehat{[B_J]}$.

Proof. The reader may wish to first understand the proof under the simplifying assumption that X is a graph. In that case the proof is easier since no thickening is necessary to ensure local convexity: we take $R = 0$ in the argument below. We thus avoid Steps 0 and 1, and the proof consists of Steps 2–6, with $\mathcal{C} = C, \mathcal{D} = D, \mathcal{B}_0 = \hat{B}_0$.

Step 0: *Ensuring large embedding radii and synchronism.* Let $R \geq K$ be the constants in Corollary 4.15. For each $I \subset \{0, \ldots, n\}$, using the separability of $\pi_1[B_I]$, the local isometry $[B_I]^{+R} \to X$ can be extended to a finite cover $\ddot{X}_I \to X$. The smallest cover $\ddot{X} \to X$ factoring through each $\ddot{X}_I \to X$ has finite degree. Moreover, $\ddot{A} \cong A$ and by Lemma 4.8, each $[\ddot{B}_I]^{+R}$ embeds.

Let $I, J \subset \{0, \dots, n\}$, and apply Lemma 4.20 to $[\ddot{B}_I], [\ddot{B}_J]^{+R}$ to obtain a finite cover $\overset{\bowtie}{X}$ such that $[\overset{\bowtie}{B}_J] \cong [\ddot{B}_J]$ and $[\overset{\bowtie}{B}_I]^{+R}$ have connected intersection. In particular, $\overset{\bowtie}{A} \cong A$. By Lemma 4.23, in any further cover $\dot{X} \to \overset{\bowtie}{X}$, the pair of subspaces $[\dot{B}_I]^{+R}, [\dot{B}_J]$ is synchronized with $[\dot{B}_I], [\dot{B}_J]$, and by Lemma 4.24, the same holds if we replace $[\dot{B}_I]^{+R}$ by a deformation retract containing $[\dot{B}_I]$.

For each $I, J \subset \{0, \dots, n\}$, we apply the above construction to obtain a finite cover $\overset{\bowtie}{X}_{IJ} \to X$. Let $\dot{X}$ be a finite cover that factors through each $\overset{\bowtie}{X}_{IJ}$ with $\dot{A} \cong A$; for instance, the based fiber-product of the $\overset{\bowtie}{X}_{IJ} \to X$. We complete the proof by applying the following claim to $\dot{X}$ with $(\dot{B}_0, \dots, \dot{B}_n, \dot{A})$. The remainder of the proof will focus on verifying this claim.

Claim. Let X be a compact special cube complex with $\widetilde{X}$ δ-hyperbolic. Let $(B_0, \dots, B_n, C)$ be locally convex connected subcomplexes containing the basepoint of X. Suppose that $C \cap B_J$ is connected for each $J \subset \{0, \dots, n\}$. Suppose that each $[B_I]^{+R}$ embeds in X, and suppose that $[B_I], [B_J]^{+R}$ is synchronized with $[B_I], [B_J]$ for each $I, J \subset \{0, \dots, n\}$. Then there is a based cover $\bar{X}$ with $\bar{C} \cong C$ such that $\bar{B}_J = \cap_{j \in J} \bar{B}_j$ is connected for each $J \subset \{0, \dots, n\}$.

The claim holds for $n = 0$ with $\bar{X} = X$, and it will be proven by induction on n.

Step 1: *Embedding neighborhoods of C and ensuring connected intersections with neighborhoods of each $\vec{B}_j$.* We first pass to a finite cover of X such that C^{+R} embeds. We then apply Lemma 4.20 several times to obtain a finite cover $\vec{X}$ such that

(1) $C^{+R} \cong \vec{C}^{+R}$ embeds in $\vec{X}$.
(2) $\vec{C}^{+R} \cap \vec{B}_j$ is connected for each j.
(3) $\vec{C}^{+K} \cap \vec{B}_j^{+K}$ is connected for each j.

As a consequence, the following holds by Lemma 4.5 as we have just enforced that C is R-embedded, and $\vec{C} \cap \vec{B}_0$ is connected and B_0 is R-embedded by hypothesis of the claim that

(4) $(\vec{C} \cap \vec{B}_0)^{+R}$ embeds.

Step 2: *Making intersections connected in B_0.* Consider the subspaces $[\vec{B}_0 \cap \vec{B}_i]$ for $i \in \{1, \dots, n\}$. We claim that $\big(\cap_{j \in J} [\vec{B}_0 \cap \vec{B}_j]\big) \cap \vec{C}$ is connected for any $J \subset \{1, \dots, n\}$. Indeed using the connectivity hypothesis, we have the following inclusions, which are thus equalities:

$$\begin{aligned}\vec{B}_{J \cup \{0\}} \cap \vec{C} \subseteq \big(\cap_{j \in J} [\vec{B}_0 \cap \vec{B}_j]\big) \cap \vec{C} &= \cap_{j \in J} [\vec{B}_0 \cap \vec{B}_j] \cap \vec{C} \\ &\subseteq \cap_{j \in J} \vec{B}_0 \cap \vec{B}_j \cap \vec{C} = \vec{B}_{J \cup \{0\}} \cap \vec{C}.\end{aligned}$$

Observe that $\left[\cap_{j\in J}[\vec{B}_0\cap\vec{B}_j]\right]=[\cap_{j\in J\cup\{0\}}\vec{B}_j]$. Indeed, $\left[\cap_{j\in J}[\vec{B}_0\cap\vec{B}_j]\right]\subset[\cap_{j\in J\cup\{0\}}\vec{B}_j]$ because $\cap_{j\in J}[\vec{B}_0\cap\vec{B}_j]\ \subset\ \cap_{j\in J}(\vec{B}_0\cap\vec{B}_j)\ =\ \cap_{j\in J\cup\{0\}}\vec{B}_j$. The reverse inclusion $[\cap_{j\in J\cup\{0\}}\vec{B}_j]\ \subset\ \left[\cap_{j\in J}[\vec{B}_0\cap\vec{B}_j]\right]$ holds because $\cap_{j\in J\cup\{0\}}\vec{B}_j\ \subset\ \vec{B}_0\cap\vec{B}_j$ for each $j\in J$, and so $[\cap_{j\in J\cup\{0\}}\vec{B}_j]\ \subset\ [\vec{B}_0\cap\vec{B}_j]$, and so $[\cap_{j\in J\cup\{0\}}\vec{B}_j]\ \subset\ \cap_{j\in J}[\vec{B}_0\cap\vec{B}_j]$.

The above two observations show that the hypotheses of our claim hold for the family $\left([\vec{B}_0\cap\vec{B}_1],\ldots,[\vec{B}_0\cap\vec{B}_n],\vec{C}\right)$. Thus, by induction there is a finite covering space $\widehat{X}_o\to\vec{X}$ and an isomorphic based elevation $\widehat{C}\to\vec{C}$ such that the collection of subspaces $\widehat{[\vec{B}_0\cap\vec{B}_j]}=[\widehat{B}_0\cap\widehat{B}_j]$ has the connected multiple intersection property: $\cap_{j\in J}[\widehat{B}_0\cap\widehat{B}_j]$ is connected for each $J\subset\{1,\ldots,n\}$. We will ensure that in each further cover the based elevation of B_0 remains isomorphic with $\widehat{B_0}$.

Step 3: *Formation of* $\mathcal{D}$. Let $D=\widehat{B}_0\cup\widehat{C}$, and since $\widehat{B}_0$ and $\widehat{C}$ are R-embedded, let $\mathcal{D}$ be the locally convex thickening of D provided by Corollary 4.15 (where $D,\mathcal{D}$, $\widehat{X}_o$, $\widehat{X}$ correspond to S, T, X, X', respectively). Furthermore, $\mathcal{D}$ decomposes as $\mathcal{D}=\mathcal{B}_0\cup\mathcal{C}$, where $\widehat{B}_0\subset\mathcal{B}_0\subset\widehat{B}_0^{+K}$ and $\widehat{C}\subset\mathcal{C}\subset\widehat{C}^{+K}$ (Corollary 4.15.1). The inclusions $\widehat{B}_0\subset\mathcal{B}_0\subset\widehat{B}_0^{+K}$ and $\widehat{C}\subset\mathcal{C}\subset\widehat{C}^{+K}$ are homotopy equivalence by local convexity, and we shall soon use this. Moreover, $(\widehat{B}_0\cap\widehat{C})\subset\mathcal{B}_0\cap\mathcal{C}\subset(\widehat{B}_0\cap\widehat{C})^{+K}$ where the locally convex intersection $\mathcal{B}_0\cap\mathcal{C}$ is connected. Since the intermediate subcomplex $\mathcal{B}_0\cap\mathcal{C}$ is locally convex and connected, it deformation retracts to $\widehat{B}_0\cap\widehat{C}$; indeed the inclusion is a π_1-isomorphism by Remark 4.16.

Step 4: *Making* $\mathcal{D}$ *have connected intersection with each of* $\widehat{B}_1,\ldots,\widehat{B}_n$. Apply Lemma 4.20 to pass to a finite cover $\check{X}$ such that $\check{\mathcal{D}}\cong\mathcal{D}$ and such that each $\check{B}_j\cap\check{\mathcal{D}}$ is connected.

Step 5: *Multiple intersections are connected in* $\check{\mathcal{D}}$. Recall that $\check{B}_J$ denotes $\cap_{j\in J}\check{B}_j$. Our goal now is to show that $\check{\mathcal{D}}\cap\check{B}_J$ is connected for each $J\subset\{1,\ldots,n\}$. This intersection can be expressed as the union of two sets containing the basepoint p:

$$\check{\mathcal{D}}\cap\check{B}_J=(\check{\mathcal{C}}\cup\check{\mathcal{B}}_0)\cap\check{B}_J=(\check{\mathcal{C}}\cap\check{B}_J)\cup(\check{\mathcal{B}}_0\cap\check{B}_J),$$

and it therefore suffices to verify the connectivity of both $(\check{\mathcal{C}}\cap\check{B}_J)=(\check{\mathcal{C}}\cap(\cap_{j\in J}\check{B}_j))$ and $(\check{\mathcal{B}}_0\cap\check{B}_J)=(\check{\mathcal{B}}_0\cap(\cap_{j\in J}\check{B}_j))$.

To reach this goal we aim to apply Lemma 4.18.3 to the families $(\check{\mathcal{C}}\cap\check{B}_j)_{j\in J}$ and $(\check{\mathcal{B}}_0\cap\check{B}_j)_{j\in J}$, respectively contained in $\check{\mathcal{C}},\check{\mathcal{B}}_0$, where $\check{\mathcal{C}}$ deformation retracts to $\check{C}$ and $\check{\mathcal{B}}_0$ deformation retracts to $\check{B}_0$ (as we have seen in Step 3).

We will be done after verifying the connectedness of each of the following:

(a) $\check{\mathcal{C}}\cap\check{B}_j$,
(b) $\check{C}\cap\check{B}_J$,

(c) $\check{\mathcal{B}}_0 \cap \check{B}_j$,
(d) $\check{B}_0 \cap \check{B}_J$.

(a) We now show that $\check{\mathcal{C}} \cap \check{B}_j$ is connected. By choice of $\vec{X}$, we know that $\vec{C}^{+K} \cap \vec{B}_j$ is connected. Since $\vec{\mathcal{C}} \subset \vec{C}^{+K}$ is a homotopy equivalence, it follows by Lemma 4.18.1 that $\vec{\mathcal{C}} \cap \vec{B}_j$ is connected, and we are done.

(b) By assumption, $\check{C} \cap \check{B}_J$ is connected. At this point we already have that each $\check{\mathcal{C}} \cap \check{B}_J$ is connected.

(c) We now show that $\check{\mathcal{B}}_0 \cap \check{B}_j$ is connected for each j. Let $q \in \check{\mathcal{B}}_0 \cap \check{B}_j$. By assumption, $\check{B}_j, \check{B}_0^{+R}$ is synchronized with $\check{B}_j, \check{B}_0$. Thus by Lemma 4.24, $[\check{\mathcal{B}}_0 \cap \check{B}_j]_q$ intersects $\check{B}_0$. Since $\check{B}_j \cap \check{\mathcal{D}}$ is connected, we see that $[\check{\mathcal{B}}_0 \cap \check{B}_j]_q$ intersects $\check{\mathcal{B}}_0 \cap \check{\mathcal{C}}$ at some point r. Recall that

$$(\check{\mathcal{B}}_0 \cap \check{\mathcal{C}}) \cong (\mathcal{B}_0 \cap \mathcal{C}) \subset (B_0^{+K} \cap C^{+K})$$

is connected. Apply Lemma 4.18.2 to see that the r-component of $(\check{\mathcal{B}}_0 \cap \check{\mathcal{C}}) \cap [\check{\mathcal{B}}_0 \cap \check{B}_j]_q$ contains a point s of $\check{B}_0 \cap (\check{\mathcal{B}}_0 \cap \check{\mathcal{C}}) \cap [\check{\mathcal{B}}_0 \cap \check{B}_j]_q$. But $s \in \check{B}_0 \cap \check{B}_j \cap \check{\mathcal{C}}$, which we proved is connected and which contains p.

(d) We now show that $\cap_{j\in J}(\check{B}_0 \cap \check{B}_j)$ is connected. Since $\check{\mathcal{B}}_0 \cap \check{B}_j$ is connected, it follows from Lemma 4.18.1 that $\check{B}_0 \cap \check{B}_j$ is connected, and so $[\check{B}_0 \cap \check{B}_j] = \check{B}_0 \cap \check{B}_j$. We thus see that $\cap_{j\in J}(\check{B}_0 \cap \check{B}_j) = \cap_{j\in J}[\check{B}_0 \cap \check{B}_j]$, but the latter is connected since $\cap_{j\in J}[\widehat{B}_0 \cap \widehat{B}_j]$ is connected by choice of $\widehat{X}$.

Step 6: *Applying induction again.* We now apply the inductive assumption to the family $(\check{B}_1, \dots, \check{B}_n, \check{\mathcal{D}})$. There is a finite covering space $\bar{X} \to X$ and an isomorphic based elevation $\bar{\mathcal{D}} \cong \check{\mathcal{D}}$ such that letting $\bar{B}_j$ denote the based elevation of $\check{B}_j$, then for any nonempty subset $J \subset \{1, \dots, n\}$ the intersection $\bar{B}_J$ is connected.

For any subset $J = \{0\} \cup I$ with $I \subset \{1, \dots, n\}$, we have $\bar{B}_J = \cap_{i\in I}(\bar{B}_0 \cap \bar{B}_i)$. Since $\bar{B}_0 \subset \bar{\mathcal{D}}$, the map $\bar{B}_0 \to \check{B}_0$ is an isomorphism, and we have already shown in Step (5d) that $\cap_{i\in I}(\check{B}_0 \cap \check{B}_i)$ is connected. □

Corollary 4.26 (virtually connected intersection for local isometries). *Let X be a compact special cube complex with $\widetilde{X}$ δ-hyperbolic. Let $(B_0 \to X, \dots, B_n \to X, A \to X)$ be local isometries of connected complexes. Suppose that $A \to X$ is injective and factors as $A \to B_j \to X$ for each j. Then there is a finite cover $\bar{X}$ with an elevation $\bar{A} \cong A$ such that the elevations $\bar{B}_0 \to \bar{X}, \dots, \bar{B}_n \to \bar{X}$ at $\bar{A} \to \bar{X}$ are injective, and the intersection $\cap_{j\in J}\bar{B}_j$ of each subcollection of their images is connected.*

Proof. For each i, let $X_i = \mathsf{C}(B_i, X)$, and let $\widehat{X}$ denote the component of the fiber-product of $X_0 \to X, \dots, X_n \to X$ that contains $\widehat{A} \cong A$. Each elevation $\widehat{B}_i$ embeds in $\widehat{X}$ and contains A. We may thus apply Theorem 4.25 to $(\widehat{B}_0, \dots, \widehat{B}_n, \widehat{A})$. □

5. **Trivial wall projections**

5.A. *Introduction.* This section is devoted to proving Corollary 5.8, which is a higher-dimensional generalization of the following

PROPOSITION 5.1. *Let $A \to X, B \to X$ be immersions of finite connected graphs to the connected graph X. Suppose that all conjugates of $\pi_1 A$ and $\pi_1 B$ have trivial intersection in $\pi_1 X$. Then there is a finite cover $\widehat{X} \to X$ such that each pair of distinct elevations $\widehat{A}, \widehat{B}$ embed and intersect in a forest.*

In Corollary 5.8 we have a similar statement for local isometries of complexes of (virtually) special cube complexes (with word hyperbolic fundamental groups) by substituting "trivial wall projection" for "forest intersection." In the case of graphs the wall projection of A onto B equals $B^0 \cup (A \cap B)$ so that the statement in Proposition 5.1 is really about trivial wall projections. While the proof of Proposition 5.1 is rather simple (see, for example, [Wis02]), we found its generalization to special cube complexes to be a very challenging part of this paper. The reader might choose to skip this lengthy section at first reading, after becoming familiar with Corollary 5.8, which of course uses the language of elevations from Definition 3.17.

The conclusion of Proposition 5.1 still holds in any further cover $X' \to \ddot{X}$. Trivial wall projection is also preserved under covering.

LEMMA 5.2. *Suppose $X' \to \dot{X}$ is a covering map of connected cube complexes. Let $\dot{A}, \dot{B} \subset \dot{X}$ denote connected subcomplexes, and let $A', B' \subset X'$ denote elevations of those. If $\mathsf{Proj}_{\dot{X}}(\dot{B} \to \dot{A})$ is trivial, then $\mathsf{Proj}_{X'}(\dot{B}' \to \dot{A}')$ is trivial.*

Proof. This holds as $\mathsf{WProj}_{X'}(B' \to A')$ maps to $\mathsf{WProj}_{\dot{X}}(\dot{B} \to \dot{A})$ under $X' \to \dot{X}$. $\square$

The main work in proving Corollary 5.8 will be to prove the following theorem, which focuses on a single local isometry. In fact this is a special case of Corollary 5.8 when B consist of a single 0-cell. Observe that the following result is immediate when $A \to X$ is an injection of graph. The difficulty in higher dimensions comes from the reach of wall-projections, even when $A \to X$ is injective.

THEOREM 5.3 (trivial wall projections). *Let X be a compact virtually special cube complex, and let $A \to X$ be a local isometry with A compact and $\pi_1 A \subset \pi_1 X$ malnormal. Assume $\pi_1 X$ is word hyperbolic. Then there exists a finite cover $A_0 \to A$ such that any further finite cover $\bar{A} \to A_0$ can be completed to a finite special cover $\bar{X} \to X$ with the following properties*:

(1) *all elevations of $A \to X$ to $\bar{X}$ are injective,*
(2) *$\bar{A}$ is wall-injective in $\bar{X}$,*
(3) *every elevation of A distinct from $\bar{A}$ has trivial wall-projection onto $\bar{A}$.*

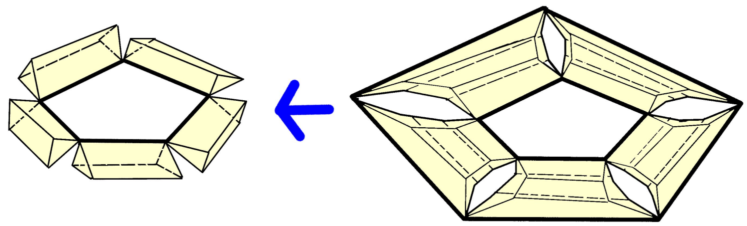

Figure 8. $\widehat{A}_1$ and $\widehat{A}_2$ are the bold pentagons.

One of the difficulties in proving Theorem 5.3 is that the triviality of $\mathsf{WProj}_{\widehat{X}}(B \to \widehat{A})$ for elevations $B, \widehat{A}$ of A is not stable under taking further covers. Lemma 5.2 gives no control in the case that $\dot{A} = \dot{B}$. For then $\mathsf{WProj}_{\dot{X}}(\dot{B} \to \dot{A}) = \dot{A}$ can be nontrivial, and indeed, there can be elevations $B' \neq A'$ with nontrivial wall projection onto A'. This behavior is exhibited in the following example, which illustrates the delicacy of Theorem 5.3. This example shows that, in general, there does not exist a finite cover $\widehat{X} \to X$ with the property that we have trivial wall projections in any further cover.

Example 5.4. Let X denote the standard 2-complex of $\langle a, b, c \mid a^{-1}b^{-1}ac \rangle$ so X is obtained from a cylinder by identifying two points on distinct bounding circles. Let A denote the subcomplex consisting of the 0-cell and the 1-cell labeled by a, and note that $\pi_1 A$ is malnormal in the free group $\pi_1 X$.

Consider any based finite cover $\widehat{X} \to X$. Let $\widehat{A}$ denote the based elevation of A. Observe that for some n, there is an immersion $D \to X$ where D is formed from $\widehat{A}$ by attaching a distinct strip $I_n \times I$ to each 1-cell of $\widehat{A}$ along both $\{0\} \times I$ and $\{n\} \times I$. For instance, we could let n denote the order of the image of $\pi_1 X$ in the left coset representation on $\pi_1 \widehat{X}$.

Let $\widehat{D}$ denote the double cover of D corresponding to the morphism $\pi_1 D \to \mathbb{Z}/2\mathbb{Z}$ that maps the loop corresponding to $\widehat{A}$ to 0 and that is nontrivial on each simple loop contained in one of the annuli attached to $\widehat{A}$. The preimage of $\widehat{A}$ in $\widehat{D}$ consists of two isomorphic components $\widehat{A}_1, \widehat{A}_2$. It is easy to see that $\mathsf{WProj}_{\widehat{D}}(\widehat{A}_1 \to \widehat{A}_2)$ consists of all of $\widehat{A}_2$. Now let $\bar{X}$ denote a cover of $\widehat{X}$ that contains $\widehat{D}$, and we see that $\widehat{A}_2 = \mathsf{WProj}_{\widehat{D}}(\widehat{A}_1 \to \widehat{A}_2) \subset \mathsf{WProj}_{\bar{X}}(\widehat{A}_1 \to \widehat{A}_2) \subset \widehat{A}_2$. This construction is partially illustrated in Figure 8.

5.B. *Narrow wall-projection.* We say hyperplanes H, K of X are *M-close* if $d(N(H), N(K)) \leq M$, where we recall that $d(U, V)$ is the length of the shortest combinatorial path with endpoints on U, V.

Lemma 5.5 (narrow implies trivial). *Suppose X is virtually special and compact. Let $A \to X$ be a local isometry with A compact, and let $M > 0$ be a positive number. There exists a finite cover $X_0 \to X$ with a based elevation A_0*

such that all elevations of A to X_0 are injective and such that for any further cover $(\bar{X}, \bar{A}) \to (X_0, A_0)$, the following holds.

Let $\bar{A}' \neq \bar{A}$ be another elevation of A to $\bar{X}$. If all pairs of hyperplanes intersecting both $\bar{A}$ and $\bar{A}'$ are M-close, then $\bar{A}'$ has trivial wall-projection onto $\bar{A}$.

Remark 5.6. The hypothesis that X is compact can be relaxed to a hypothesis of finitely many hyperplanes. Furthermore, the hypothesis that A is compact can be relaxed to the following hypothesis: $A \to X$ factors through a local isometry $\bar{A} \to \bar{X}$ such that the preimage of $\pi_1\bar{A}$ equals $\pi_1 A$, and $\bar{A}$ is compact, and $\bar{X}$ is virtually special.

Proof of Lemma 5.5. By Lemma 4.9 and Corollary 4.10, there is a finite special cover $X' \to X$ such that all elevations of A to X' are injective and all hyperplanes of X' have embedding radius $> M$. Let L denote the number of hyperplanes in X'. Let $A_0 \to A$ be a finite cover factoring through X' such that the finite CAT(0) complex $v^{+(L+1)} \to A_0$ embeds as a subcomplex for each vertex $v \in A_0$ and such that $v^{+(L+1)}$ is wall-injective in A_0. Let $X_0 = \mathsf{C}(A_0, X')$. Note that each $v^{+(L+1)}$ is then wall-injective in X_0.

Let $(\bar{X}, \bar{A}) \to (X_0, A_0)$ be any further cover. For each $\bar{v}$ in $\bar{A}$, the CAT(0) cubical ball $\bar{v}^{+(L+1)}$ is still a CAT(0) and wall-injective subcomplex of $\bar{X}$.

Consider an edge-path $\bar{\sigma}$ in $\bar{A}$ of length $L+1$, and assume that any two hyperplanes dual to edges of $\bar{\sigma}$ are M-close. Then $\bar{\sigma}$ has two distinct edges dual to hyperplanes $\bar{H}_1, \bar{H}_2$ that project to the same hyperplane H' of X'. Let $\bar{\tau}$ be a path of length $\leq M$ between $N(\bar{H}_1)$ and $N(\bar{H}_2)$. Since $\bar{\tau}$ projects to a short path τ' from $N(H')$ to itself, we see by Lemma 4.7 that τ' is homotopic into $N(H')$ and thus $\bar{H}_1 = \bar{H}_2$. The cubical $(L+1)$-ball $\bar{B}$ of $\bar{A}$ centered at the origin of $\bar{\sigma}$ is wall-injective in $\bar{X}$, and thus there is a $\bar{B}$-hyperplane dual to two distinct edges of $\bar{\sigma}$. Since $\bar{B}$ is CAT(0), it follows that $\bar{\sigma}$ is not a local geodesic.

By the local convexity of wall projections (Remark 3.15), the above argument shows that for any distinct elevation $\bar{A}'$ whose common hyperplanes with $\bar{A}$ are M-close, each component of $\mathsf{WProj}_{\bar{X}}(\bar{A}' \to \bar{A})$ is contained in a cubical L-ball; thus it is trivial. $\square$

While Example 5.4 shows that the property of having trivial wall projections for *arbitrary* elevations is not stable under further covers, the following shows how to obtain this stability for all *nearby* elevations.

LEMMA 5.7 (nearby elevations). *Let $A \to X$ be a local isometry of special cube complexes with A compact. Assume $\pi_1 A \to \pi_1 X$ is malnormal and the universal cover $\widetilde{X}$ is Gromov-hyperbolic. Then for any positive number $D > 0$, there is a finite cover $X_1 \to X$ with a based elevation A_1 with the following properties.*

All elevations of $A \to X$ to X_1 are injective, and for any further cover $(\widehat{X}, \widehat{A}) \to (X_1, A_1)$ and any elevation $\widehat{A}' \neq \widehat{A}$, if $d(\widehat{A}', \widehat{A}) \leq D$, then $\widehat{A}'$ has trivial wall-projection onto $\widehat{A}$.

In the argument below malnormality is employed to ensure that the hyperplanes between A_1 and A_1' are close enough. It then suffices to apply Lemma 5.5.

Proof. We will use the constants R, K of Corollary 4.15. By Lemma 4.9, we may pass to an initial finite cover $\dot{X} \to X$ so that all elevations of $A \to X$ to $\dot{X}$ are embedded and have embedding radius $> D + R$. Let $\dot{A}$ be a fixed elevation of A to $\dot{X}$. Since the embedding radius of $\dot{A}$ is $> D + R$, we see that $\dot{A}^{+D}$ embeds in $\dot{X}$. Let $\dot{A}_i$ be an elevation intersecting $\dot{A}^{+D}$, and let C_{ij} be the various components of $\dot{A}^{+D} \cap \dot{A}_i$. Let $M = \max_{ij}\big(\text{diameter}(C_{ij}^{+K})\big)$, which is obviously finite.

We apply Lemma 5.5 to $\dot{A} \subset \dot{X}$ with the constant M in order to obtain the finite cover $\bar{X}$ and elevation $\bar{A}$ such that if $\bar{A}' \neq \bar{A}$ is an elevation whose common hyperplanes with $\bar{A}$ are M-close, then $\bar{A}'$ has trivial wall projection onto $\bar{A}$ in $\bar{X}$, and this persists in further covers.

Since the embedding radius of $\bar{A}$ is $> D+R$, the space $(\bar{A})^{+D}$ embeds in $\bar{X}$. For any distinct elevation $\bar{A}_k$ within distance D of $\bar{A}$, we choose a connected component $\bar{C}_{k\ell}$ of the nonempty intersection $\bar{A}^{+D} \cap \bar{A}_k$, and we form a space $S_{k\ell}$ by attaching $\bar{A}^{+D}$ with $\bar{A}_k$ along $\bar{C}_{k\ell}$. By malnormality of $\pi_1 A$, each C_{ij} is simply-connected, and so each $\bar{C}_{k\ell}$ factors isomorphically through some C_{ij}. Consequently diameter$(\bar{C}_{k\ell}^{+K}) \leq M$ always holds.

The embedding radius of $\bar{A}_k$ inside $\bar{X}$ is $> D + R \geq R$. The embedding radius of $\bar{A}^{+D}$ inside $\bar{X}$ is $> (D + R) - D = R$. Thus by Corollary 4.15, the natural map $S_{k\ell} \to \bar{X}$ factors as $S_{k\ell} \to T_{k\ell} \hookrightarrow X_{k\ell} \to \bar{X}$, where $X_{k\ell} \to \bar{X}$ is a finite cover, $S_{k\ell} \to T_{k\ell}$ is an injective π_1-isomorphism, $T_{k\ell} \subset X_{k\ell}$ is a connected, wall-injective locally convex subcomplex, and any path in $T_{k\ell}$ connecting $\bar{A}_k$ to $\bar{A}^{+K}$ enters $\bar{C}_{k\ell}^{+K}$.

Since $T_{k\ell} \subset X_{k\ell}$ is wall-injective, it follows that any hyperplane of $X_{k\ell}$ from $\bar{A}_k$ to $\bar{A}$ enters $\bar{C}_{k\ell}^{+K}$, and thus any two such hyperplanes are M-close. Consequently the wall projection from $\bar{A}_k$ to $\bar{A}$ is trivial.

We consider the various covers $X_{k\ell} \to \bar{X}$ associated to the finitely many choices of $C_{k\ell}$. Each finite cover $X_{k\ell}$ contains an isomorphic elevation of $\bar{A}$, and we denote by $(X_1, A_1) \to (\bar{X}, \bar{A})$ the $\bar{A}$-component of the fiber-product of the various covers $\{X_{k\ell} \to \bar{X}\}$. Note that $A_1 \cong \bar{A}$.

The cover $X_1 \to X$ has the desired properties. Indeed consider any further cover $(\widehat{X}, \widehat{A}) \to (X_1, A_1)$ and any elevation $\widehat{A}' \neq \widehat{A}$ with $d(\widehat{A}', \widehat{A}) \leq D$.

In $\bar{X}$ the embedding radius of $\bar{A}$ is $> D$. Since $\widehat{A}' \neq \widehat{A}$ and $d(\widehat{A}', \widehat{A}) \leq D$, it follows that the image of $\widehat{A}'$ inside $\bar{X}$ is distinct from $\bar{A}$ and thus equals one of the $\bar{A}_k$ discussed above. Let A'_1 be the image of $\widehat{A}'$ in X_1, and again note that $A'_1 \neq A'$ — indeed it maps to $\bar{A}_k \neq \bar{A}$ in $\bar{X}$.

Observe that $X_1 \to \bar{X}$ factors through $X_{k\ell}$ and A'_1 maps onto $\bar{A}_k \subset T_{k\ell} \subset X_{k\ell}$. Since the wall-projection of $\bar{A}_k$ onto $\bar{A}$ inside $X_{k\ell}$ is trivial, it follows by Lemma 5.2 that $\mathsf{WProj}_{X_1}(A'_1 \to A_1)$ is trivial, and hence so is $\mathsf{WProj}_{\widehat{X}}(\widehat{A}' \to \widehat{A})$. □

5.C. *Proof of Theorem* 5.3.

Step 1: *Preparation.* Let $R \geq K$ be the constants of Corollary 4.15. Choose constants $D > (3K+1)\dim(X)$, $M \geq 4K\dim(X)$, and $R_1 \geq \max(K+R, 4K+2)$.

We first pass to a finite cover $(X_0, A_0) \to (X, A)$ such that

(1) X_0 is special.

(2) Each elevation of A to X_0 is injective.

(3) Each elevation of A to X_0 has embedding radius $> R_1$. Each hyperplane of X_0 has embedding radius $> R_1$.

(4) Let $(\bar{X}, \bar{A}) \to (X_0, A_0)$ be any further cover. Then for each elevation $\bar{A}' \neq \bar{A}$ with $d(\bar{A}', \bar{A}) \leq D$, the wall projection $\mathsf{WProj}_{\bar{X}}(\bar{A}' \to \bar{A})$ is trivial.

(5) Let $(\bar{X}, \bar{A}) \to (X_0, A_0)$ be any further cover. Then for any other elevation $\bar{A}'$ of A, if the hyperplanes between $\bar{A}, \bar{A}'$ are M-close, then $\mathsf{WProj}_{\bar{X}}(\bar{A}' \to \bar{A})$ is trivial.

These properties are stable under covers, so we obtain them consecutively. First we choose a finite special cover $X_1 \to X$. Let A_1 be an elevation of $A \to X$ to X_1. We form the canonical completion $\mathsf{C}(A_1, X_1) \to X$ and note that its based elevation is injective. Let $X_2 \to X$ be a finite regular cover factoring through $\mathsf{C}(A_1, X_1) \to X$, and note that all elevations of A to X_2 are injective. Then we apply Lemma 4.9 and Corollary 4.10 to get a finite cover $X_3 \to X_2$ with arbitrarily high embedding radii of the desired subcomplexes. To get the two last properties we consecutively apply Lemmas 5.7 and 5.5.

Let $\bar{A} \to A_0$ be the further finite cover in the statement of Theorem 5.3. Using separability we complete $\bar{A} \to X_0$ to a finite cover $\dot{X} \to X_0$. We alert the reader to the following simplification of notation: in any further cover $\check{X} \to \dot{X}$, we will always consider a preferred elevation of A that is isomorphic to $\bar{A}$, and we will still denote this elevation by $\bar{A}$. By Lemma 4.8, the subspace $\bar{A} \subset \check{X}$ is R_1-embedded, and so for any radius $0 \leq R' \leq R_1$ we have $(\bar{A})^{+R'} \subset \check{X}$. Moreover, if $\check{X} \to \dot{X}$ factors through $\hat{X} \to \dot{X}$, then the covering map $\check{X} \to \hat{X}$ induces an isomorphism between the R'-thickenings of $\bar{A}$.

Step 2: *Connected intersection of thickened elevations and hyperplanes.* We claim that there is a further finite cover $\ddot{X} \to \dot{X}$ with an isomorphic elevation of $\bar{A}$ such that the following holds: for any hyperplane $\ddot{H}$ of $\ddot{X}$ dual to an edge of $\bar{A}$ and for any pair of connected locally convex subcomplexes $\ddot{Y}, \ddot{Z} \subset \ddot{X}$ satisfying $\bar{A} \subset \ddot{Y} \subset (\bar{A})^{+R_1}$, $N(\ddot{H}) \subset \ddot{Z} \subset N(\ddot{H})^{+R_1}$, the intersection $\ddot{Y} \cap \ddot{Z}$ is connected.

Indeed for any hyperplane $\dot{H}$ of $\dot{X}$ dual to an edge e of $\bar{A}$ and for any pair of connected locally convex subcomplexes $\dot{Y}, \dot{Z} \subset \dot{X}$ satisfying $\bar{A} \subset \dot{Y} \subset (\bar{A})^{+R_1}$, $N(\dot{H}) \subset \dot{Z} \subset N(\dot{H})^{+R_1}$, let $C_{\dot{Y}\dot{Z}}$ denote the component of $\dot{Y} \cap \dot{Z}$ containing e. By Lemma 4.20, there is a finite cover $\dot{X}_{\dot{H},\dot{Y},\dot{Z}} \to \dot{X}$ with an isomorphic elevation of $\dot{Y}$ (and hence of $\bar{A}$) such that the elevation of $\dot{Z}$ at e intersects $\dot{Y}$ connectedly. The fiber-product of the various covers $\dot{X}_{\dot{H},\dot{Y},\dot{Z}} \to \dot{X}$ contains a natural isomorphic elevation of $\bar{A}$ that is contained in a connected component $\ddot{X}$ of the fiber-product.

We claim that $\ddot{X}$ has the required connectedness property. This is clear for $\ddot{Y}$ arbitrary and $\ddot{Z} = N(\ddot{H})$ or $\ddot{Z} = N(\ddot{H})^{+R_1}$, because in that case $\ddot{Y}, \ddot{Z}$ cover subcomplexes $\dot{Y}, \dot{Z}$. The result follows for an intermediate $N(\ddot{H}) \subset \ddot{Z} \subset N(\ddot{H})^{+R_1}$ using Lemma 4.18.

Step 3: *Geometric properties of the union of $\bar{A}$ and a hyperplane.* We now consider the collection of spaces $\ddot{U}_1, \ldots, \ddot{U}_k, \ddot{V}_1, \ldots, \ddot{V}_\ell$ that are subspaces of $\ddot{X}$ that arise in the following two ways:

- the union $\ddot{U}_i$ of $\bar{A}$ and a hyperplane $\ddot{H}$ passing through it;
- the union $\ddot{V}_j$ of $\bar{A}$, a distant elevation $\ddot{A}'$, and a hyperplane H passing through both.

Here we say that an elevation $\ddot{A}'$ of A is *distant* if $d(\ddot{A}', \bar{A}) > D$.

Each $\ddot{U}_i, \ddot{V}_j$ is quasiconvex, and we consider their locally convex thickenings $\ddot{B}_i \to \ddot{X}, \ddot{C}_j \to \ddot{X}$. Below we describe precisely these constructions.

Observe that $\bar{A}$ is wall-injective in $\ddot{X}$ and a similar property holds more generally for any intermediate locally convex subcomplex $\bar{A} \subset \ddot{Y} \subset (\bar{A})^{+R_1}$. Indeed let e_1, e_2 be 1-cells of $\ddot{Y}$ dual to a hyperplane $\ddot{H}$ that is also dual to some edge of $\bar{A}$. Since $\ddot{Y} \cap N(\ddot{H})$ is connected by Step 2, there is a path p in $\ddot{Y} \cap N(\ddot{H})$ from e_1 to e_2. By local convexity, e_1pe_2 travels on a sequence of squares of $\ddot{Y}$ dual to $\ddot{H}$, and we deduce that e_1, e_2 are parallel inside $\ddot{Y}$. In particular, any hyperplane $\ddot{H}$ that intersects $\bar{A}$ actually meets $\bar{A}$ along a single hyperplane of $\bar{A}$, which we denote by $\ddot{H}_{\bar{A}}$.

We then form the space $\ddot{U}$ by gluing $\bar{A}$ and $N(\ddot{H})$ along their connected intersection. By construction, the embedding radii of both $\bar{A}$ and $\ddot{H}$ in $\ddot{X}$ are $> R_1$. Recall that R, K denote the constants of Corollary 4.15 and at the very beginning of Step 1 we chose $R_1 \geq R$. By Corollary 4.15, the map $\ddot{U} \to \ddot{X}$ factors through a local isometry $\ddot{B} \to \ddot{X}$ so that $\ddot{U} \to \ddot{B}$ is an injective π_1-isomorphism. Furthermore, there are connected locally convex subcomplexes $\mathcal{A}_{\ddot{H}}, \ddot{\mathcal{H}}$ of $\ddot{B}$ such that $\bar{A} \subset \mathcal{A}_{\ddot{H}} \subset (\bar{A})^{+K}$, $N(\ddot{H}) \subset \ddot{\mathcal{H}} \subset (N(\ddot{H}))^{+K}$, $\ddot{B} = \mathcal{A}_{\ddot{H}} \cup \ddot{\mathcal{H}}$; the intersection $\mathcal{A}_{\ddot{H}} \cap \ddot{\mathcal{H}}$ is connected and contained inside $\big(\bar{A} \cap N(\ddot{H})\big)^{+K}$.

Since $R_1 \geq R \geq K$, the map $\ddot{B} \to \ddot{X}$ is injective on both $\mathcal{A}_{\ddot{H}}, \ddot{\mathcal{H}}$. We claim that, in fact, $\ddot{B} \to \ddot{X}$ is injective. By local injectivity of $\ddot{B} \to \ddot{X}$, it suffices to check that $\mathcal{A}_{\ddot{H}} \cap \ddot{\mathcal{H}}$ is a connected subspace of $\ddot{X}$. This follows from Step 2, where we established the connectedness of $\mathcal{A}_{\ddot{H}} \cap \ddot{\mathcal{H}}$ provided $\bar{A} \subset \mathcal{A}_{\ddot{H}} \subset (\bar{A})^{+R_1}$, $N(\ddot{H}) \subset \ddot{\mathcal{H}} \subset N(\ddot{H})^{+R_1}$. (Recall that $K \leq R_1$.)

Since $\mathcal{A}_{\ddot{H}} \subset \bar{A}^{+K}$, we see that the embedding radius of $\mathcal{A}_{\ddot{H}}$ is $\geq R_1 - K \geq R$, and likewise $\ddot{\mathcal{H}}$ has embedding radius $\geq R$. Since $\bar{A} \subset (\mathcal{A}_{\ddot{H}})^{+R} \subset (\bar{A})^{+R_1}$ and $N(\ddot{H}) \subset \ddot{\mathcal{H}}^{+R} \subset (N(\ddot{H}))^{+R_1}$, Step 2 implies that $(\mathcal{A}_{\ddot{H}})^{+R} \cap \ddot{\mathcal{H}}^{+R}$ is connected. We can now apply Lemma 4.11 to see that the embedding radius of $\ddot{B}$ is $\geq R$.

Assume now that $\ddot{H}$ cuts an elevation $\ddot{A}'$ with $d(\ddot{A}', \bar{A}) > D$. Let γ be a connected component of $\ddot{B} \cap \ddot{A}'$ that contains an edge dual to $\ddot{H}$. Let $\ddot{W}$ be the space obtained by attaching $\ddot{B}$ to $\ddot{A}'$ along γ. We denote by $\ddot{V} \subset \ddot{W}$ the π_1-surjective subspace $\ddot{U} \cup \ddot{A}'$. By Corollary 4.15, the map $\ddot{W} \to \ddot{X}$ factors through a local isometry $\ddot{C} \to \ddot{X}$ so that $\ddot{W} \to \ddot{C}$ is an injective π_1-isomorphism. Furthermore, there are connected locally convex subcomplexes $\mathcal{B}, \mathcal{A}'_\gamma$ of $\ddot{C}$ such that $\ddot{B} \subset \mathcal{B} \subset \ddot{B}^{+K}$, $\ddot{A}' \subset \mathcal{A}'_\gamma \subset (\ddot{A}')^{+K}$, $\ddot{C} = \mathcal{B} \cup \mathcal{A}'_\gamma$; the intersection $\mathcal{B} \cap \mathcal{A}'_\gamma$ is connected and contained inside γ^{+K}.

Step 4: *Construction of* $\bar{X} \to X$. The inclusion $\bar{A} \to \ddot{X}$ factors through the various local isometries $\ddot{B}_i, \ddot{C}_j \to \ddot{X}$. We apply Corollary 4.26 to the collection of local isometries $\{\ddot{B}_i \to \ddot{X}\}, \{\ddot{C}_j \to \ddot{X}\}$ relative to $\bar{A}$. This yields a cover $(\bar{X}, \bar{A})$ in which their elevations $\{\bar{B}_i, \bar{C}_j\}$ are injective and have pairwise connected intersection.

We note that $\bar{A}$ is wall-injective in $\bar{X}$. Indeed assume two edges $\bar{a}_1, \bar{a}_2$ of $\bar{A}$ are dual to a hyperplane $\bar{H}$ of $\bar{X}$. By projecting $\bar{H}$ to the hyperplane $\ddot{H}$ of $\ddot{X}$, we note that $\bar{a}_1, \bar{a}_2$ are also parallel in $\ddot{X}$. Since $\bar{A}$ is wall-injective in $\ddot{X}$, there is a hyperplane of $\bar{A}$ dual to both $\bar{a}_1, \bar{a}_2$. The conclusion follows since $\bar{X} \to \ddot{X}$ induces an isomorphism $\bar{A} \to \bar{A}$.

Step 5: *Verifying that wall projections are trivial.* Let $\bar{A}' \neq \bar{A}$ be an elevation of A, and let $\ddot{A}'$ denote the image of $\bar{A}'$ inside $\ddot{X}$. If $d(\ddot{A}', \bar{A}) \leq D$ and $\ddot{A}' \neq \bar{A}$, then $\mathsf{WProj}_{\ddot{X}}(\ddot{A}' \to \bar{A})$ is already trivial by Step 1 and we are done by Lemma 5.2. Otherwise either $\ddot{A}' = \bar{A}$ or $\ddot{A}'$ is distant from $\bar{A}$ (in the sense that $d(\ddot{A}', \bar{A}) > D$). In each of these cases we will deduce the triviality of the wall-projection by showing M-closeness of the hyperplanes cutting through both $\bar{A}$ and $\bar{A}'$.

In the first case consider any two hyperplanes $\bar{H}_1, \bar{H}_2$ that pass through both $\bar{A}$ and $\bar{A}'$. Their images in $\ddot{X}$ are $\ddot{H}_1, \ddot{H}_2$. For each i, let $\ddot{U}_i$ and $\ddot{B}_i$ be the spaces from Step 3 associated to $(\bar{A}, \ddot{H}_i)$. Since $\bar{A}'$ maps to $\bar{A}$ and $\bar{H}_1, \bar{H}_2$ map to $\ddot{H}_1, \ddot{H}_2$, we have $\bar{A}' \subset \bar{U}_1 \cap \bar{U}_2 \subset \bar{B}_1 \cap \bar{B}_2$ where $\bar{U}_i \subset \bar{B}_i$ denotes the elevation of $\ddot{U}_i$ that contains $\bar{A}$.

There is an edge-path $\bar{\sigma}$ inside the connected subcomplex $\bar{B}_1 \cap \bar{B}_2$ that starts at $\bar{A}$ and ends in $\bar{A}'$. Since $\ddot{B}_1 = (\mathcal{A}_{\ddot{H}_1} \cup \ddot{\mathcal{H}}_1) \subset (\bar{A}^{+K} \cup \ddot{H}_1^{+K})$, the locally convex subcomplex $\bar{B}_1$ is contained in a union of elevations of $\bar{A}^{+K}$ and $\ddot{H}_1^{+K}$. Since $R_1 \geq K+1$, we have $\bar{A}^{+K} \cap \bar{A}' = \emptyset$. Thus there is a first vertex $\bar{p}$ on $\bar{\sigma}$ that is not in $\bar{A}^{+K}$. Note that $\bar{p} \in \bar{A}^{+(K+1)}$. Since $R_1 \geq 2K+1$, this vertex $\bar{p}$ does not belong to any elevation of $\bar{A}^{+K} \subset \ddot{X}$ to $\bar{X}$. Thus $\bar{p} \in (\bar{e}_1')^{+K}$, where $\bar{e}_1'$ is an edge dual to a hyperplane $\bar{H}_1'$ mapping to the hyperplane $\ddot{H}_1$ and contained inside $\bar{B}_1$. We now check that $\bar{H}_1' = \bar{H}_1$. Given an edge $\bar{e}_1$ dual to $\bar{H}_1$ and contained in $\bar{A}$, we have $\bar{e}_1', \bar{e}_1 \subset \bar{A}^{+(2K+2)}$. The images $\ddot{e}_1', \ddot{e}_1$ of these edges under $\bar{X} \to \ddot{X}$ are both dual to $\ddot{H}_1$. Since $R_1 \geq 2K+2$, we have seen in Step 3 that $\ddot{e}_1', \ddot{e}_1$ are parallel inside $\bar{A}^{+(2K+2)}$. Since $\bar{X} \to \ddot{X}$ induces an isomorphism $\bar{A}^{+(2K+2)} \to \bar{A}^{+(2K+2)}$, it follows that $\bar{e}_1', \bar{e}_1$ are dual to the same hyperplane of $\bar{A}^{+(2K+2)}$, and so $\bar{H}_1' = \bar{H}_1$. Similarly $\bar{p} \in \bar{H}_2^{+K}$. It follows that there is a path of length $\leq 2K\dim(X)$ from $N(\bar{H}_1)$ to $N(\bar{H}_2)$ and so $\bar{H}_1, \bar{H}_2$ are M-close. See Remark 4.1 for the relationship between combinatorial neighborhoods and cubical thickenings.

The second case (when $d(\ddot{A}', \bar{A}) > D$) is similar except that we use locally convex thickenings $\ddot{C}_1, \ddot{C}_2$ of the spaces built from $\bar{A}$, $\ddot{A}'$, together with $\ddot{H}_1, \ddot{H}_2$ respectively. There is no difference from the previous explanation at the vertex of $\bar{\sigma}$ leaving the K-thickening of $\bar{A}$: it comes within a uniform distance of both $\bar{H}_1$ and $\bar{H}_2$. And such a vertex exists because D is large enough. We shall now provide the details.

Consider any two hyperplanes $\bar{H}_1, \bar{H}_2$ cutting both $\bar{A}$ and $\bar{A}'$. Choose an edge $\bar{a}_i'$ of $\bar{A}'$ dual to $\bar{H}_i$, and denote by a_i' the image of $\bar{a}_i'$ inside $\ddot{X}$. The images of $\bar{H}_1, \bar{H}_2$ in $\ddot{X}$ are $\ddot{H}_1, \ddot{H}_2$. For each i, let $\ddot{B}_i$ be the locally convex thickening of $\bar{A} \cup N(\ddot{H}_i)$. We let γ_i denote the connected component of a_i' inside $\ddot{B}_i \cap \ddot{A}'$. Let $\ddot{V}_i, \ddot{W}_i, \ddot{C}_i$ be the immersed spaces associated to $\ddot{B}_i, \ddot{A}', \gamma_i$ that we constructed in Step 3. Since $\bar{A}'$ maps to $\ddot{A}'$ and $\bar{H}_1, \bar{H}_2$ map to $\ddot{H}_1, \ddot{H}_2$, we have $\bar{A}' \subset \bar{W}_1 \cap \bar{W}_2 \subset \bar{C}_1 \cap \bar{C}_2$.

There is an edge-path $\bar{\sigma}$ inside the connected subcomplex $\bar{C}_1 \cap \bar{C}_2$ that starts at $\bar{A}$ and ends in $\bar{A}'$. The following inclusion shows that the locally convex thickening $\bar{C}_1$ is contained in the union of elevations of $\ddot{A}'^{+K}$, elevations of $\bar{A}^{+2K} \subset \ddot{X}$, and elevations of $\ddot{H}_1^{+2K}$:

$$\ddot{C}_1 = \mathcal{B}_1 \cup \mathcal{A}_1' \subset \left(\mathcal{A}_{\ddot{H}_1} \cup \ddot{\mathcal{H}}\right)^{+K} \cup \ddot{A}'^{+K} \subset \bar{A}^{+2K} \cup \ddot{H}_1^{+2K} \cup \ddot{A}'^{+K}.$$

Since $D > (2K)\dim(X)$, we have $\bar{A}^{+2K} \cap \bar{A}' = \emptyset$, thus there is a first vertex $\bar{p}$ on $\bar{\sigma}$ that is not in $\bar{A}^{+2K}$. Since $R_1 \geq 4K+1$, this vertex $\bar{p}$ does not belong to the $2K$-thickening of any elevation of $\bar{A} \subset \ddot{X}$. Since $D > (3K+1)\dim(X)$, the point $\bar{p}$ does not belong to the K-thickening of any elevation of $\ddot{A}'$ contained in $\bar{C}_1$. It follows that $\bar{p}$ belongs to the $2K$-thickening of an elevation of $N(\ddot{H}_1)$. Equivalently, $\bar{p} \in (\bar{e}_1')^{+2K}$ for some edge $\bar{e}_1'$ dual to a hyperplane $\bar{H}_1'$ mapping to the hyperplane $\ddot{H}_1$ and contained inside $\bar{B}_1$. The argument ends in the same manner as in the first case: since $R_1 \geq 4K+2$, we have $\bar{H}_1' = \bar{H}_1$, and so any two hyperplanes between $\bar{A}$ and $\bar{A}'$ are within a distance $4K\dim(X) \leq M$. □

5.D. *A variation on the theme.* Later, we will need the following result, which is a consequence of Theorem 5.3.

COROLLARY 5.8. *Let X be a compact virtually special cube complex, and let $A, B \to X$ be compact local isometries with $\pi_1 A \subset \pi_1 X$ malnormal, $\pi_1 B \subset \pi_1 X$ malnormal. Each conjugate of $\pi_1 A$ has trivial intersection with $\pi_1 B$. More precisely, if $a \to A$ and $b \to B$ are immersed circles that are homotopic to each other in X, then they are null-homotopic.*

Assume $\pi_1 X$ is word hyperbolic. Then there exists a finite cover $A_0 \to A$ such that any further finite cover $\bar{A} \to A_0$ can be completed to a finite special cover $\bar{X} \to X$ with the following properties:

(1) *All elevations of $A \to X, B \to X$ to $\bar{X}$ are injective.*
(2) *$\bar{A}$ is wall-injective.*
(3) *Every elevation of A distinct from $\bar{A}$ has trivial wall-projection onto $\bar{A}$.*
(4) *Every elevation of B has trivial wall-projection onto $\bar{A}$.*

Proof. Step 1: *The auxiliary pair $C \to Y$.* We choose base points $\bar{a}, \bar{b}$ in A, B, and let a, b be their images inside X. We then consider the space Y obtained by adding a single 1-cube e to X with origin at a and with endpoint at b. We also form a connected cube complex C by setting $C = A \sqcup [0,1] \sqcup B/_{\bar{a}=0,\bar{b}=1}$. We denote by $\bar{e}$ the image of $[0,1]$ inside C. Mapping $\bar{e}$ to e we get a natural map $C \to Y$.

Step 2: *Geometric properties of Y.* Observe that $\widetilde{Y}$ is a hyperbolic CAT(0) cube complex and $C \to Y$ is a local isometry. Indeed the universal cover $\widetilde{Y}$ is a tree-like space, where the vertex spaces are disjoint copies of $\widetilde{X}$, connected by the edges mapping to e.

Step 3: *$\pi_1 C \subset \pi_1 Y$ is malnormal.* This can be proven by either simple disc diagram arguments, combinatorial group theory arguments involving normal forms, or geometric considerations in the universal cover. We leave the details to the reader.

Step 4: *Y is virtually special.* Let $\widehat{X} \to X$ be a special cover of finite degree d. The preimage of a consists of d points, and the preimage of b consists of d points. We then choose a one-to-one correspondence between preimages of a and preimages of b, and glue d edges accordingly. The resulting cube complex $\widehat{Y}$ covers Y, and it is special. Indeed the union of special cube complexes meeting along vertices is itself special since it straightforwardly satisfies the condition of Lemma 2.6.

Step 5: *Constructing* A_0. We can now apply Theorem 5.3 to the local isometry $C \to Y$. We thus obtain a finite cover $C_0 \to C$ such that any further finite cover $\bar{C} \to C_0$ extends to a finite special cover $\bar{Y} \to Y$ where all elevations of $C \to Y$ are injective and have trivial wall-projection onto $\bar{C}$, provided they are distinct from $\bar{C}$, and moreover $\bar{C} \subset \bar{Y}$ is wall-injective. We choose $A_0 \subset C_0$ to be a fixed elevation of $A \subset C$ to the covering space $C_0 \to C$.

Step 6: *Conclusion.* We now verify that $A_0 \to A$ has the desired property. Let $\bar{A} \to A_0$ be any finite cover. C_0 is special since it is a locally convex subcomplex of the special complex Y_0, and thus $A \to A_0 \subset C_0$ extends to a finite cover $\bar{C} \to C_0$.

We next further complete $\bar{C} \to C_0$ to a finite special cover $\bar{Y} \to Y$ with the properties of Theorem 5.3. We denote by $\bar{X}$ the elevation of $X \subset Y$ that contains $\bar{A}$, and we claim that $\bar{X}$ has the desired properties.

$\bar{X}$ is special since it is a locally convex subcomplex of the special cube complex $\bar{Y}$. Note that $\bar{A}$ is wall-injective in $\bar{C}$, thus also in $\bar{Y}$ and *a fortiori* in $\bar{X}$. Each elevation of $A \to X$ or $B \to X$ to $\bar{X}$ extends to an elevation of $C \to Y$, and is thus injective.

Consider an elevation $E \neq \bar{A}$ of either A or B to $\bar{X}$. Let $\bar{C}'$ be the elevation of C containing E. We first treat the case that $\bar{C}' = \bar{C}$. Observe that E and $\bar{A}$ have no common hyperplane in $\bar{C}$ since $\bar{C}$ is the disjoint union of covers of A and B attached together along isolated 1-cells. The wall injectivity of $\bar{C} \subset \bar{X}$ implies that E and $\bar{A}$ have no common hyperplane in $\bar{X}$ either. Thus $\mathsf{WProj}_{\bar{X}}(E \to \bar{A}) = \bar{A}^0$ and is thus trivial.

In the other case where $\bar{C}' \neq \bar{C}$, we see that $\mathsf{WProj}_{\bar{X}}(E \to \bar{A})$ is contained in $\mathsf{WProj}_{\bar{Y}}(\bar{C}' \to \bar{C})$ and is thus trivial. □

6. The main technical result: A symmetric covering property

Let P be an embedded 2-sided hyperplane in the cube complex Q. Let $N_o(P)$ be the *open cubical neighborhood* of P consisting of all open cubes of X intersecting P. Note that $N_o(P) \cong P \times (-1, 1) \subset P \times [-1, 1]$. The map $N_o(P) \hookrightarrow Q$ extends to a map $\phi : P \times [-1, 1] \to Q$.

Let A and B denote $P \times \{-1\}$ and $P \times \{+1\}$. We refer to A, B as the *sides* of P. Let $X = Q - N_o(P)$. Let $\alpha : A \to X$ and $\beta : B \to X$ denote the restrictions of ϕ. Let p denote a basepoint of P, and let $a, b = (p, \pm 1)$ be the corresponding basepoints in A, B. We regard the images of a, b to be corresponding basepoints of X.

THEOREM 6.1. *Let Q be a compact connected nonpositively curved cube complex, and let P be a hyperplane in Q such that the following hold*:

(1) $\pi_1 Q$ *is word-hyperbolic.*
(2) P *is an embedded, nonseparating,* 2-*sided hyperplane in* Q.

(3) *$\pi_1 P$ is malnormal in $\pi_1 Q$.*
(4) *$X = Q - N_o(P)$ is virtually special.*

For any finite cover $\widehat{X}$ of X, there is a finite regular cover $\overset{\bowtie}{X}$ factoring through $\widehat{X}$ such that $\overset{\bowtie}{X} \to X$ induces the same cover on each side A, B of P.

Proof. We will show that for each finite cover $\widehat{X}$, there is a finite regular cover $\overset{\bowtie}{X}$ factoring through $\widehat{X}$ such that the isomorphism $\gamma = \beta\alpha^{-1}$ from A to B lifts to an isomorphism $\overset{\bowtie}{\gamma}$ from some elevation $\overset{\bowtie}{A}$ to some elevation $\overset{\bowtie}{B}$.

Step 1: (Covers with trivial wall projection onto $\widehat{A}_a$ and $\widehat{B}_b$). There exist special, connected, finite based covers $\widehat{X}_a$ and $\widehat{X}_b$ factoring through $\widehat{X}$ such that

(1) The based elevation $\widehat{A}_a$ of A to $\widehat{X}_a$ is injective and wall-injective in $\widehat{X}_a$.
(2) The based elevation $\widehat{B}_b$ of B to $\widehat{X}_b$ is injective and wall-injective in $\widehat{X}_b$.
(3) The isomorphism $\gamma : A \to B$ lifts to an isomorphism $\widehat{\gamma} : \widehat{A}_a \to \widehat{B}_b$.
(4) $\mathsf{WProj}_{\widehat{X}_a}(\breve{B} \to \widehat{A}_a)$ is trivial for each elevation $\breve{B}$ of B to $\widehat{X}_a$.
(5) $\mathsf{WProj}_{\widehat{X}_a}(\breve{A} \to \widehat{A}_a)$ is trivial for each elevation $\breve{A}$ of A to $\widehat{X}_a$ with $\breve{A} \neq \widehat{A}_a$.
(6) $\mathsf{WProj}_{\widehat{X}_b}(\breve{A} \to \widehat{B}_b)$ is trivial for each elevation $\breve{A}$ of A to $\widehat{X}_b$.
(7) $\mathsf{WProj}_{\widehat{X}_b}(\breve{B} \to \widehat{B}_b)$ is trivial for each elevation $\breve{B}$ of B to $\widehat{X}_b$ with $\breve{B} \neq \widehat{B}_b$.

This follows by Corollary 5.8.

Step 2: We form the following canonical completions and note that the inclusion maps hold by Lemma 3.12 and the isomorphism is obtained from the isomorphism $\widehat{A}_a \cong \widehat{B}_b$:

$$\mathsf{C}(\widehat{A}_a, \widehat{X}_a) \hookleftarrow \mathsf{C}(\widehat{A}_a, \widehat{A}_a) \cong \mathsf{C}(\widehat{B}_b, \widehat{B}_b) \hookrightarrow \mathsf{C}(\widehat{B}_b, \widehat{X}_b).$$

Step 3: There exist based covers $\bar{A} \to A$ and $\bar{B} \to B$ such that

(1) The isomorphism $\gamma : A \to B$ lifts to an isomorphism $\bar{\gamma} : \bar{A} \to \bar{B}$ so we have the following commutative diagram:

$$\begin{array}{ccc} \bar{A} & \to & \bar{B} \\ \downarrow & & \downarrow \\ \widehat{A}_a & \to & \widehat{B}_b. \end{array}$$

(2) $\bar{A}$ factors through each elevation of A to $\mathsf{C}(\widehat{A}_a, \widehat{X}_a)$ and to $\mathsf{C}(\widehat{B}_b, \widehat{X}_b)$.
(3) $\bar{B}$ factors through each elevation of B to $\mathsf{C}(\widehat{A}_a, \widehat{X}_a)$ and to $\mathsf{C}(\widehat{B}_b, \widehat{X}_b)$.

Indeed we simply choose based covers of A and B that factor through all elevations, and then a common cover using the isomorphism $\gamma : A \to B$.

Step 4: The canonical retraction map $\mathsf{C}(\widehat{A}_a, \widehat{X}_a)$ together with the cover $\bar{A} \to \widehat{A}_a$ induces the covers $\widehat{\mathsf{C}(\widehat{A}_a, \widehat{X}_a)}$ and $\widehat{\mathsf{C}(\widehat{A}_a, \widehat{A}_a)}$. Similarly, we obtain $\widehat{\mathsf{C}(\widehat{B}_b, \widehat{X}_b)}$ and $\widehat{\mathsf{C}(\widehat{B}_b, \widehat{B}_b)}$ so we have the following commutative diagrams:

$$
\begin{array}{ccccccc}
\widehat{\mathsf{C}(\widehat{A}_a, \widehat{X}_a)} & \to & \bar{A} & \qquad & \widehat{\mathsf{C}(\widehat{A}_a, \widehat{A}_a)} & \to & \bar{A} \\
\downarrow & & \downarrow & & \downarrow & & \downarrow \\
\mathsf{C}(\widehat{A}_a, \widehat{X}_a) & \to & \widehat{A}_a, & & \mathsf{C}(\widehat{A}_a, \widehat{A}_a) & \to & \widehat{A}_a, \\
\widehat{\mathsf{C}(\widehat{B}_b, \widehat{B}_b)} & \to & \bar{B} & & \widehat{\mathsf{C}(\widehat{B}_b, \widehat{X}_b)} & \to & \bar{B} \\
\downarrow & & \downarrow & & \downarrow & & \downarrow \\
\mathsf{C}(\widehat{B}_b, \widehat{B}_b) & \to & \widehat{B}_b, & & \mathsf{C}(\widehat{B}_b, \widehat{X}_b) & \to & \widehat{B}_b.
\end{array}
$$

The isomorphism between $\bar{A} \to \widehat{A}_a$ and $\bar{B} \to \widehat{B}_b$, and the isomorphism between the $\mathsf{C}(\widehat{A}_a, \widehat{A}_a) \to \widehat{A}_a$ and $\mathsf{C}(\widehat{B}_b, \widehat{B}_b) \to \widehat{B}_b$ induce the following commutative diagram:

$$
\begin{array}{ccccccc}
\widehat{\mathsf{C}(\widehat{A}_a, \widehat{X}_a)} & \hookleftarrow & \widehat{\mathsf{C}(\widehat{A}_a, \widehat{A}_a)} & \cong & \widehat{\mathsf{C}(\widehat{B}_b, \widehat{B}_b)} & \hookrightarrow & \widehat{\mathsf{C}(\widehat{B}_b, \widehat{X}_b)} \\
\downarrow & & \downarrow & & \downarrow & & \downarrow \\
\mathsf{C}(\widehat{A}_a, \widehat{X}_a) & \hookleftarrow & \mathsf{C}(\widehat{A}_a, \widehat{A}_a) & \cong & \mathsf{C}(\widehat{B}_b, \widehat{B}_b) & \hookrightarrow & \mathsf{C}(\widehat{B}_b, \widehat{X}_b).
\end{array}
$$

Step 5:

— Let $\overset{\bowtie}{X}_a$ be the smallest regular cover factoring through each component of $\widehat{\mathsf{C}(\widehat{A}_a, \widehat{X}_a)}$.
— Let $\overset{\bowtie}{X}_b$ be the smallest regular cover factoring through each component of $\widehat{\mathsf{C}(\widehat{B}_b, \widehat{X}_b)}$.
— Let $\overset{\bowtie}{X}$ be the smallest regular cover factoring through $\overset{\bowtie}{X}_a$ and $\overset{\bowtie}{X}_b$.
— Let $\overset{\bowtie}{A}$ be the smallest regular cover factoring through each component of $\widehat{\mathsf{C}(\widehat{A}_a, \widehat{A}_a)}$.
— Let $\overset{\bowtie}{B}$ be the smallest regular cover factoring through each component of $\widehat{\mathsf{C}(\widehat{B}_b, \widehat{B}_b)}$.
— Let $\overset{\bowtie}{A}_a$ and $\overset{\bowtie}{A}_b$ denote the elevations of A to $\overset{\bowtie}{X}_a$ and $\overset{\bowtie}{X}_b$.
— Let $\overset{\bowtie}{B}_a$ and $\overset{\bowtie}{B}_b$ denote the elevations of B to $\overset{\bowtie}{X}_a$ and $\overset{\bowtie}{X}_b$.

It is clear that the isomorphism $\gamma : A \to B$ lifts to an isomorphism $\overset{\bowtie}{\gamma} : \overset{\bowtie}{A} \to \overset{\bowtie}{B}$. We will show that $\overset{\bowtie}{A} \cong \overset{\bowtie}{A}_a$ since they factor through each other and that $\overset{\bowtie}{A}$ factors through $\overset{\bowtie}{A}_b$. It will follow that each elevation of A to $\overset{\bowtie}{X}$ is isomorphic to $\overset{\bowtie}{A}$. An analogous argument shows that $\overset{\bowtie}{B} \cong \overset{\bowtie}{B}_b$ and that $\overset{\bowtie}{B}$ factors through $\overset{\bowtie}{B}_a$. Consequently, each elevation of B to $\overset{\bowtie}{X}$ is isomorphic to $\overset{\bowtie}{B}$.

Since $\widehat{\mathsf{C}(\widehat{A}_a, \widehat{A}_a)} \subset \widehat{\mathsf{C}(\widehat{A}_a, \widehat{X}_a)}$, it is obvious that $\overset{\bowtie}{A}_a$ factors through $\overset{\bowtie}{A}$. Since $\overset{\bowtie}{X}_a$ is the smallest regular cover induced by $\widehat{\mathsf{C}(\widehat{A}_a, \widehat{X}_a)}$, to see that $\overset{\bowtie}{A}$ factors through $\overset{\bowtie}{A}_a$, it suffices to check that $\overset{\bowtie}{A}$ factors through each elevation of A to $\widehat{\mathsf{C}(\widehat{A}_a, \widehat{X}_a)}$.

There are two cases to consider, according to the history of A, as it follows a sequence of elevations indicated below:

$$\begin{array}{ccccccc}
A_3 & \to & A_2 & \to & A_1 & \to & A \\
\downarrow & & \downarrow & & \downarrow & & \downarrow \\
\widehat{\mathsf{C}(\widehat{A}_a, \widehat{X}_a)} & \to & \mathsf{C}(\widehat{A}_a, \widehat{X}_a) & \to & \widehat{X}_a & \to & X.
\end{array}$$

If $A_1 = \widehat{A}_a$ is the base elevation of A, then by Lemma 3.13, $A_2 \subset \mathsf{C}(\widehat{A}_a, \widehat{A}_a) \subset \mathsf{C}(\widehat{A}_a, \widehat{X}_a)$. Since A_3 is contained in $\widehat{\mathsf{C}(\widehat{A}_a, \widehat{A}_a)}$, we see that $\overset{\bowtie}{A}$ factors through A_3.

If $A_1 \neq \widehat{A}_a$, then $\mathsf{WProj}_{\widehat{X}_a}(A_2 \to \widehat{A}_a)$ is trivial, and so by Lemma 3.16, A_2 is nullhomotopic in the retraction map $\mathsf{C}(\widehat{A}_a, \widehat{X}_a) \to \widehat{A}_a$. Thus $A_3 \cong A_2$. But $\overset{\bowtie}{A}$ factors through $\bar{A}$, which factors through A_2, which is isomorphic to A_3.

To see that $\overset{\bowtie}{A}$ factors through $\overset{\bowtie}{A}_b$, we show that $\bar{A}$ factors through $\overset{\bowtie}{A}_b$ by showing $\bar{A}$ factors through each elevation of A to $\widehat{\mathsf{C}(\widehat{B}_b, \widehat{X}_b)}$. Again consider the history of elevations of A to $\widehat{\mathsf{C}(\widehat{B}_b, \widehat{X}_b)}$:

$$\begin{array}{ccccccc}
A_3 & \to & A_2 & \to & A_1 & \to & A \\
\downarrow & & \downarrow & & \downarrow & & \downarrow \\
\widehat{\mathsf{C}(\widehat{B}_b, \widehat{X}_b)} & \to & \mathsf{C}(\widehat{B}_b, \widehat{X}_b) & \to & \widehat{X}_b & \to & X.
\end{array}$$

Observe that A_1 has trivial wall projection onto $\widehat{B}_b$, and so by Lemma 3.16, A_2 is nullhomotopic in the retraction map $\mathsf{C}(\widehat{B}_b, \widehat{X}_b) \to \widehat{B}_b$. Thus $A_3 \cong A_2$. But $\bar{A}$ factors through A_2, and hence $\overset{\bowtie}{A}$ factors through $\bar{A}$, which factors through $A_3 \cong A_2$. □

7. Subgroup separability of quasiconvex subgroups

Without hyperbolicity, the conclusion of Theorem 6.1 permits a proof of residual finiteness of $\pi_1 Q$ using the normal form theorem for graphs of groups. In this section we prove the following stronger property from this conclusion together with hyperbolicity.

THEOREM 7.1. *Let Q be a compact nonpositively curved cube complex with $\pi_1 Q$ word-hyperbolic. Let P be an embedded nonseparating 2-sided hyperplane in Q. Let $X = Q - N_o(P)$ be the cube complex obtained from Q by deleting the open cubical neighborhood of P.*

Assume that $\pi_1 X$ has separable double cosets of quasiconvex subgroups. Suppose that for each finite cover $\widehat{X} \to X$, there is finite regular cover $\overset{\bowtie}{X} \to X$ factoring through $\widehat{X}$ such that $\overset{\bowtie}{X}$ induces the same cover on each side A, B of P. Then every quasiconvex subgroup of $\pi_1 Q$ is separable.

Definition 7.2. The space Y has the *cocompact convex core property* with respect to a subgroup H of $\pi_1 Y$ if for each compact subset $C \subset \widetilde{Y}$, there is a convex subcomplex S of $\widetilde{Y}$ such that $C \subset S$ and S is H-stable and H-cocompact.

LEMMA 7.3. *Let X be a compact nonpositively curved cube complex X with $\pi_1 X$ word-hyperbolic. Then X has the cocompact convex core property with respect to each quasiconvex subgroup.*

Proof. This follows from Theorem 4.2. $\square$

The proof of Theorem 7.1 follows the scheme for proving subgroup separability given in [Wis00], and as was done there, it is possible to relax the initial hypotheses a bit. For instance, we actually show that $\pi_1 Q$ is separable with respect to subgroups H whenever Q has the convex core property with respect to H, under the assumption that $\widetilde{X}$ has separable double cosets of subgroups with the cubical convex core property.

Proof of Theorem 7.1. We regard Q as a graph of spaces, whose open edge space is $N_o(P)$, and whose vertex space is X. Note that the attaching maps of the edge space are $A \to X$ and $B \to X$.

Choose a basepoint in Q that also lies in X. Let $\dot{Q}$ denote the based cover corresponding to a quasiconvex subgroup of $\pi_1 Q$. Let $\sigma \in \pi_1 Q - \pi_1 \dot{Q}$, and let $\dot{\sigma}$ be the based lift of σ. The space $\dot{Q}$ has an induced graph of space structure whose vertex spaces are components of the preimage of X and whose open edge spaces are components of the preimage of $N_o(P)$. Since $\pi_1 \dot{Q}$ is finitely generated, we can let $\dot{Q}'$ be a π_1-isomorphic locally-convex subspace corresponding to a finite subgraph of spaces such that $\dot{Q}'$ contains $\dot{\sigma}$. Note that the vertex and edge spaces of $\dot{Q}'$ are isomorphic to vertex and edge spaces of $\dot{Q}$. By the cocompact convex core property for Q, let Y be a compact core of $\dot{Q}$ containing $\dot{\sigma}$ and also containing a preimage of the basepoint within each vertex space of $\dot{Q}'$. Note that we can assume that $Y \subset \dot{Q}'$ since $\dot{Q}'$ is itself locally convex.

Since vertex and edge spaces of $\dot{Q}$ are connected and locally convex, the intersection of Y with each vertex and edge space of $\dot{Q}$ is a connected and locally convex subcomplex by Lemma 4.18.1. These intersections correspond precisely to the vertex and edge spaces of a graph of space structure for Y. Let $Y_i = Y \cap \dot{X}_i$ denote the vertex spaces of Y where $\dot{X}_1, \ldots, \dot{X}_n$ denote the vertex spaces of $\dot{Q}$ that have nonempty intersection with Y. We emphasize that $Y_i \hookrightarrow \dot{X}_i$ induces a π_1-isomorphism.

For each i, we shall now use the separability of certain double cosets to produce a finite cover $\widehat{X}_i \to X$ with the following injectivity properties:

(1) $\dot{X}_i \to X$ factors as $\dot{X}_i \to \widehat{X}_i \to X$ and $Y_i \subset \dot{X}_i$ embeds in $\widehat{X}_i$ under this map.
(2) The distinct elevations $\dot{A}_{ij} \to \dot{X}_i$ of $A \to X$ whose images intersect $Y_i \subset \dot{X}_i$ factor through distinct elevations in $\widehat{X}_i$.
(3) The distinct elevations $\dot{B}_{ik} \to \dot{X}_i$ of $B \to X$ whose images intersect $Y_i \subset \dot{X}_i$ factor through distinct elevations in $\widehat{X}_i$.

We first determine the double cosets related to Properties (2) and (3). There are finitely many elevations of A whose images intersect Y_i since A and Y_i are compact (and similarly for B). For each i, these distinct elevations correspond precisely to finitely many double cosets of the form $\pi_1\dot{X}_i\alpha_{ij}\pi_1 A$ and $\pi_1\dot{X}_i\beta_{ik}\pi_1 B$. More accurately, we assume that the basepoint of Q and X is the image of the basepoint of A and that δ is a path from this A-basepoint to the B-basepoint. The second collection of double cosets are of the form $\pi_1\dot{X}_i\beta_{ik}\delta\pi_1 B\delta^{-1}$.

We now determine the cosets related to Property (1). Enumerate the finitely many pairs p_ℓ, q_ℓ of distinct vertices of Y_i that map to the same vertex of X, and for each ℓ, let ω_ℓ be a path in Y_i from the basepoint to p_ℓ and let σ_ℓ be a path from p_ℓ to q_ℓ. Note that the projections of the paths $\omega_\ell, \sigma_\ell, \omega_\ell^{-1}$ to X are concatenatable to a closed path $\gamma_\ell = \omega_\ell\sigma_\ell\omega_\ell^{-1}$. A based lift of Y_i embeds in a based cover of $\ddot{X}$ of X precisely if each path $\gamma_\ell \notin \pi_1\ddot{X}$.

By double coset separability (note that these compact local isometries have quasiconvex fundamental groups), for each i, we are able to choose a finite index normal subgroup $N_i < \pi_1 X$ such that the $N_i\pi_1\dot{X}_i\alpha_{ij}\pi_1 A$ are all disjoint from each other, the $N_i\pi_1\dot{X}_i\beta_{ik}\delta\pi_1 B\delta^{-1}$ are all disjoint from each other, and finally each $N_i\pi_1\dot{X}_i\gamma_{i\ell}$ is disjoint from $N_i\pi_1\dot{X}_i$. It follows that the based covering space $\widehat{X}_i$ with $\pi_1\widehat{X}_i = N_i\pi_1\dot{X}_i$ has the properties enumerated above.

Let Z denote the space obtained from Y by extending each Y_i to $\widehat{X}_i$, so Z is the quotient of the disjoint union $Y \sqcup (\sqcup_i \widehat{X}_i)$ obtained by identifying each Y_i with its embedded image in $\widehat{X}_i$. Note that there is an induced map $Z \to Q$. If all elevations of A, B to the $\widehat{X}_i$ were isomorphic, then we could "close" Z to a finite cover $\bar{Q} \to Q$. We would then have $Y \subset \bar{Q}$, thus $\sigma \notin \pi_1\bar{Q}$, and also $H \subset \pi_1\bar{Q}$, and we would be done.

Let $\widehat{X}$ denote a finite cover factoring through each $\widehat{X}_i$. By hypothesis, let $\overset{\bowtie}{X}\to \widehat{X}$ be a finite regular cover such that $\overset{\bowtie}{X}$ induces the same cover $\overset{\bowtie}{A}, \overset{\bowtie}{B}$ on each side A, B of P. Choose a one-to-one correspondence between the elevations of A and B to $\overset{\bowtie}{X}$, and attach copies of $\overset{\bowtie}{P} \times I$ to form a finite based cover $\overset{\bowtie}{Q}\to Q$ whose vertex space is $\overset{\bowtie}{X}$.

Let $\overset{\bowtie}{Z}= Z\otimes_Q \overset{\bowtie}{Q}$ be the based component of the fiber-product of $\overset{\bowtie}{Q}\to Q$ and $Z \to Q$. Note that each vertex space of $\overset{\bowtie}{Z}$ is isomorphic to $\overset{\bowtie}{X}$ since it is a component of the fiber-product of $\widehat{X}_i$ and $\overset{\bowtie}{X}$. Moreover, note that for each edge space of $\overset{\bowtie}{Z}$, its two ends are contained in copies of $\overset{\bowtie}{A}$ and $\overset{\bowtie}{B}$. We extend this edge space to a copy of $\overset{\bowtie}{P} \times I$.

The remaining copies of $\overset{\bowtie}{A}$ and $\overset{\bowtie}{B}$ do not have incident edge spaces, and a quick count shows that there are the same number of each. We attach edge spaces between them using an arbitrary one-to-one correspondence.

The result is a finite based cover $\bar{Q}$ of Q. Indeed the construction gives a natural combinatorial map $\bar{Q} \to Q$ that is a local isomorphism. The key point to understanding that this is a covering map is that there is exactly one incoming and outgoing edge space (at a vertex space) for each elevation of A, B. If we had not been careful to maintain the partial one-to-one correspondence between elevations when we constructed Z from Y, then extending edge spaces of $\overset{\bowtie}{Z}$ to finite covers of edge spaces of Q could result in multiple incoming or outgoing edge spaces attached along the same elevation of $\overset{\bowtie}{A}$ or $\overset{\bowtie}{B}$.

Finally, observe that the element σ is separated from $\pi_1\dot{Q}$ in the right representation on cosets of $\pi_1\bar{Q}$. Indeed the endpoint of $\bar{\sigma}$ is not in the preimage of the basepoint of Y in $\overset{\bowtie}{Z}\subset \bar{Q}$. Thus $\pi_1 Y$ is separated from σ in the right coset representation since they act differently on the base coset. $\square$

COROLLARY 7.4. *Let Q be a compact connected nonpositively curved cube complex, and let P be a hyperplane in Q such that the following hold*:

(1) *$\pi_1 Q$ is word-hyperbolic.*
(2) *P is an embedded 2-sided hyperplane in Q.*
(3) *$\pi_1 P$ is malnormal in $\pi_1 Q$.*
(4) *Each component of $Q - N_o(P)$ is virtually special.*

Then every quasiconvex subgroup of $\pi_1 Q$ is separable.

Proof. The Corollary follows when P is not separating by combining Theorems 7.1 and 6.1. The double coset separability hypothesis holds as in Remark 4.19.

Assume now that P is separating. Let Q' be obtained from Q by adding a new 1-cube that connects the components of X. Observe that $Q \to Q'$ is a local isometry. Obviously $\pi_1 Q'$ is word-hyperbolic. It is easy to verify that $\pi_1 P$ is still malnormal in $\pi_1 Q'$. And $Q' - N_o(P)$ is virtually special: the disjoint union of connected special finite covers of the two connected components attached together by a collection of 1-cubes provides a special cover. We conclude using the nonseparating case since any quasiconvex subgroup of $\pi_1 Q$ maps to a quasiconvex subgroup of $\pi_1 Q'$. $\square$

8. Main theorem: Virtual specialness of malnormal cubical amalgams

In [HaW08] we proved the following criterion.

PROPOSITION 8.1. *Let C be a compact nonpositively curved cube complex, and suppose that $\pi_1 C$ is word-hyperbolic. Then C is virtually special if and only if every quasiconvex subgroup of $\pi_1 C$ is separable.*

Combining Corollary 7.4 and Proposition 8.1 we obtain the following result.

THEOREM 8.2. *Let Q be a compact connected nonpositively curved cube complex, and let P be a hyperplane in Q such that the following hold*:

(1) *$\pi_1 Q$ is word-hyperbolic.*
(2) *P is an embedded 2-sided hyperplane in Q.*
(3) *$\pi_1 P$ is malnormal in $\pi_1 Q$.*
(4) *Each component of $Q - N_o(P)$ is virtually special.*

Then Q is virtually special.

Let us reconsider the statement of Theorem 1.2. The cube complex $A \cup_M B$ equals the quotient of $A \sqcup \big(M \times [-1, 1]\big) \sqcup B$ by using the maps $A \leftarrow M$ and $M \to B$ to identify $M \times \{\pm 1\}$ with the images of M in A, B. Note that the nonpositive curvature of $A \cup_M B$ holds because the attaching maps are local isometries. Thus the statement of Theorem 1.2 corresponds to the separating case of Theorem 8.2 by letting $P = M \times \{0\}$ be the hyperplane of $A \cup_M B$ so that $N_o(P) = M \times (-1, 1)$.

THEOREM 8.3. *A compact nonpositively curved cube complex C is virtually special provided the following hold*:

(1) *$\pi_1 C$ is word-hyperbolic*,
(2) *each hyperplane of C is 2-sided and embeds*,
(3) *$\pi_1 D$ is malnormal in $\pi_1 C$ for each hyperplane D of C.*

Proof. We repeatedly cut along hyperplanes and apply Theorem 8.2. Here it is convenient to remove open regular neighborhoods of hyperplanes when cutting. The base case where Q consists of a single vertex is reached after finitely many cuts. $\square$

Remark 8.4. We note that a "malnormal hyperplane hierarchy" gives a more general formulation using Theorem 8.2.

We now present a generalization of Theorem 8.2 that works in the presence of torsion. While the statement of Theorem 8.2 could be recast elegantly in terms of "cubical orbihedra," we will instead work in the universal cover to avoid extra definitions. We emphasize that the hypothesis in the statement below is equivalent to that of Theorem 8.2 when G is torsion-free.

THEOREM 8.5. *Let G act properly and cocompactly on a* CAT(0) *cube complex $\widetilde{Q}$. Let $\widetilde{P}$ be a hyperplane, and let $H = \text{Stabilizer}(\widetilde{P})$. Suppose the following hold*:

(1) *G is word-hyperbolic.*
(2) *H is an almost malnormal subgroup of G.*
(3) *$g\widetilde{P} \cap \widetilde{P} = \emptyset$ unless $g \in H$.*
(3′) *$H\backslash\widetilde{P}$ embeds in $H\backslash N(\widetilde{P})$ as a 2-sided hyperplane.*
(4) *For each component $\widetilde{X}$ in $\widetilde{Q} - GN_o(\widetilde{P})$ the group* Stabilizer$(\widetilde{X})$ *has a torsion-free finite index subgroup J such that $J\backslash\widetilde{X}$ is special.*

Then G has a torsion-free finite index subgroup G' such that $G'\backslash\widetilde{Q}$ is special.

All steps in the proof of Theorem 8.2 and its supporting results generalize to this torsion permitting context. The reader may wish to read through the proofs with the $\widetilde{X}$ or orbihedra viewpoint in mind, as it leads to a logically simpler proof. We will instead hitch the proof to a torsion-free scenario by using the proof of Theorem 6.1.

Proof. Without loss of generality, we will assume that G acts transitively on the components of $\widetilde{Q} - GN_o(\widetilde{P})$. Let $\widetilde{A}$ and $\widetilde{B}$ be two sides of $\widetilde{P}$ in $\widetilde{Q}$, and choose $g \in G$ such that $\widetilde{A}$ and $g\widetilde{B}$ both lie on the same complementary component $\widetilde{X}$. Let $J_o = \text{Stabilizer}(\widetilde{X})$, let $J_1 \subset J$ be a finite index subgroup of the hypothesized torsion-free special group J, and let $X_1 = J_1\backslash(\widetilde{X})$. Let $A_1 = (J_1 \cap H)\backslash\widetilde{A}$, and let $B_1 = (J_1 \cap H^g)\backslash\widetilde{B}$. We concede that A_1 and B_1 are unlikely to be isomorphic. However, we then pass to a finite cover $\widehat{X}_2$ with elevations A_2, B_2 of A_1, B_1 such that A_2 and B_2 are isomorphic with an isomorphism "induced by $\widetilde{P}$." The various translates of $\widetilde{A}$ in $\widetilde{X}$ project to a family of subspaces $A_2 = A_{20}, A_{21}, \ldots, A_{2m}$ in X_2, and there is likewise a family $B_2 = B_{20}, B_{21}, \ldots, B_{2n}$. We emphasize that these subspaces form a malnormal family. We now apply the proof of Theorem 6.1 to obtain a finite cover $\overset{\bowtie}{X} \to X_2$ corresponding to a normal subgroup of J_o and such that the elevations $\overset{\bowtie}{A}, \overset{\bowtie}{B}$ of A_2, B_2 are isomorphic (as induced by $\widetilde{P}$). The finite index normal subgroup $\pi_1 \overset{\bowtie}{X}$ of J_o induces a virtually free quotient $G \to \bar{G}$, and all torsion in G survives in the quotient $\bar{G}$. Let $\bar{G}^*$ be a torsion-free finite index subgroup of $\bar{G}$, and let G^* denote its preimage in G. Let $X = G^*\backslash\widetilde{X}$. Let $P_1, \dots, P_\ell$ denote the finitely many 2-sided embedded disjoint hyperplanes that are images of translates of $\widetilde{P}$. Note that each $\pi_1 P_i$ is malnormal in $\pi_1 X$. The result now follows using ℓ applications of Theorem 8.2. □

9. Virtually special ⇔ separable hyperplanes

A subgroup H of K is *almost malnormal* if $H^k \cap H$ is finite for each $k \in K - H$.

LEMMA 9.1. *Let H be a separable quasiconvex subgroup of a word-hyperbolic group G. Then G has a finite index subgroup K that contains H as an almost malnormal subgroup.*

Proof. As proven in [GMRS98] (see also [HrW09]), there are finitely many double cosets Hg_iH such that $g_iHg_i^{-1} \cap H$ is infinite. By separability, we can choose K containing H but not containing any g_i. Thus H is almost malnormal in K. □

We are now able to obtain the following characterization of virtual specialness in the word-hyperbolic case. It remains an open problem whether such a characterization holds in general. See [HaW08] for a characterization using double hyperplane cosets.

THEOREM 9.2. *Let C be a compact nonpositively curved cube complex such that $\pi_1 C$ is word-hyperbolic. Then C is virtually special if and only $\pi_1 D$ is separable in $\pi_1 C$ for each immersed hyperplane D of C.*

Proof. For each hyperplane D_i, apply Lemma 9.1 to obtain a finite cover $C_i \to C$ such that $\pi_1 D_i$ is malnormal in $\pi_1 C_i$. Then let C' be a regular cover factoring through all the C_i's, and observe that each hyperplane of C_i has malnormal fundamental group. We can then pass to a finite cover $\bar{C}$ such that each hyperplane is embedded.

Indeed for each immersed hyperplane $D \to C$, let $N \to C$ denote its immersed cubical regular neighborhood. By convexity of $\widetilde{N} \subset \widetilde{C}$, we see that N embeds in the cover $\widehat{C}_N$ with $\pi_1 \widehat{C} \cong \pi_1 N$. By separability, we see that N embeds in a finite cover $\bar{C}_N$ of C.

We now let $\bar{C}$ be a finite regular cover factoring through $\bar{C}_N$ as N varies over all regular neighborhoods of immersed hyperplanes. At this stage producing a further cover in which hyperplanes are 2-sided is easily done [HaW08]. We conclude by applying Theorem 8.3. □

10. Uniform arithmetic hyperbolic manifolds of simple type

The main goal of this section is to show that certain arithmetic lattices in $\mathcal{H}^n = \mathrm{Isom}(\mathbb{H}^n)$ are virtually special. In the first subsection we describe conditions on a hyperbolic lattice that imply virtual specialness. Assuming that there are sufficiently many codimension-1 immersed closed geodesic submanifolds, we can cubulate the group and then apply Theorem 9.2 to obtain virtual specialness. In the second subsection we verify that uniform arithmetic hyperbolic lattices of simple type satisfy this geodesic submanifold criterion and are thus virtually special.

Our application to subgroup separability of quasiconvex subgroups generalizes earlier results in [ALR01] as well as more recent work of Agol [Ago06] that remarkably pushed Scott's original reflection group idea to handle many arithmetic examples up to dimension 11.

10.A. *Criterion for virtual specialness of closed hyperbolic manifolds.*

Theorem 10.1. *Let G be a uniform lattice in $\mathcal{H}^n$. Let $H_1, \ldots, H_k$ be isometric copies of $\mathbb{H}^{n-1}$ in $\mathbb{H}^n$. Suppose that* $\text{Stabilizer}_G(H_i)$ *acts cocompactly on H_i for each i. Suppose there exists D such that any length D geodesic intersects gH_i for some $g \in G$ and $1 \leq i \leq k$. Then G acts properly and cocompactly on a* CAT(0) *cube complex C. Moreover, G contains a finite index subgroup F such that $F\backslash C$ is a special cube complex.*

The cubulation utilizes Sageev's construction, and can be deduced from the following formulation which we quote from [HrW10].

Proposition 10.2. *Let G act cocompactly on a δ-hyperbolic* CAT(0) *space X. Let $H_1, \ldots, H_k$ in X be a set of convex hyperplanes in X. Suppose that the union of their translates $T = \{gH_i : g \in G, 1 \leq i \leq k\}$ is locally finite in X. Suppose there exists D such that any geodesic segment of length D crosses some hyperplane gH_i. Then G acts properly and cocompactly on a* CAT(0) *cube complex C. Moreover, the distinct hyperplanes Y of C are in one-to-one correspondence with distinct hyperplanes gH_i, and* $\text{Stabilizer}(Y) = \text{Stabilizer}(gH_i)$.

Lemma 10.3. *Let G be a finitely generated subgroup of $\mathcal{H}^n$. Let H be an isometric copy of $\mathbb{H}^{n-1}$ in $\mathbb{H}^n$. Then* $\text{Stabilizer}(H)$ *is a separable subgroup of G.*

Proof. Let r be the reflection along H. Let $G' = \langle G, r\rangle$, and observe that G' is finitely generated and hence residually finite since it is linear. Observe that the centralizer $\text{Cent}_{G'}(r)$ is a separable subgroup of G'. Indeed if $k \notin \text{Cent}_{G'}(r)$, then $[r, k] \neq 1$. Let $G' \to \bar{G}'$ be a finite quotient in which $[\bar{r}, \bar{k}] \neq \bar{1}$. Then the preimage of $\text{Cent}_{\bar{G}'}(\bar{r})$ in G' separates k from $\text{Cent}_{G'}(r)$. Finally, observe that $\text{Stabilizer}_G(H) = \text{Cent}_{G'}(r) \cap G$ is separable in G. Indeed, $\text{Stabilizer}_{\mathcal{H}^n}(H) = \text{Cent}_{\mathcal{H}^n}(r)$. $\square$

Proof of Theorem 10.1. The proper and cocompact action of G on a CAT(0) cube complex C follows from Proposition 10.2. The stabilizer of each hyperplane of C equals the stabilizer of a hyperplane gH_i in $\mathbb{H}^n$ and is therefore separable by Lemma 10.3. Since G is residually finite, and there are finitely many torsion elements, we can pass to a finite index subgroup G' that is torsion-free. We can then apply Theorem 9.2 to $G'\backslash C$ to obtain a finite index subgroup F of G' such that $F\backslash C$ is special. □

10.B. *Uniform arithmetic hyperbolic lattices of "simple type"*. The results in this subsection were motivated by the possibility of cubulating arithmetic hyperbolic manifolds of simple type using the plethora of totally geodesic submanifolds. We are grateful to Nicolas Bergeron for pointing out that the density of the commensurator was an easy way to see that there are sufficiently many such submanifolds to apply our criterion. After we developed this viewpoint on the virtual specialness of these lattices, an alternate treatment was proposed in [BHW11], which uses the double coset separability criterion along the lines of the proof of virtual specialness of Coxeter groups [HaW10].

THEOREM 10.4. *Let G be a uniform arithmetic lattice in $\mathcal{H}^n$ of simple type. Then*

(1) *G acts properly and cocompactly on a* CAT(0) *cube complex C.*
(2) *G contains a finite index subgroup F such that $F\backslash C$ is special.*

Proof. This follows from Lemma 10.10, where we verify the criterion of Theorem 10.1. □

Combining with [HaW08], we obtain the following consequence.

COROLLARY 10.5. *Every quasiconvex subgroup of G is a virtual retract and is hence separable.*

To prove Theorem 10.4, we will show that G contains sufficiently many subgroups acting on codimension-1 hyperplanes. Before embarking on the proof, it will be helpful to state an explicit characterization of a simple type arithmetic lattice and to note two of their elementary properties.

Remark 10.6 (simple type arithmetic lattices in $\mathrm{SO}(1,n)$). Let $\mathbb{F}$ be a totally real algebraic number field. Let $\mathcal{O}$ be the ring of integers in $\mathbb{F}$, and let $a_1,\dots,a_n \in \mathcal{O}$ be such that

(1) each a_j is positive,
(2) each $\sigma(a_j)$ is negative for every place $\sigma \neq 1$.

Let $G = \mathrm{SO}(a_1x_1^2+\cdots+a_nx_n^2-x_{n+1}^2;\mathbb{R}) \cong \mathrm{SO}(n,1)$, and let $G_{\mathcal{O}} \subset G$ be the subgroup corresponding to matrices with entries in $\mathcal{O}$ and with determinant 1. Then $G_{\mathcal{O}}$ is an *arithmetic lattice of simple type in G*. Moreover, $G_{\mathcal{O}}$ is uniform (i.e., cocompact) if and only if $(0,\dots,0)$ is the only solution in $\mathcal{O}^{n+1}$ to the equation $a_1x_1^2+\cdots+a_nx_n^2-x_{n+1}^2=0$.

Up to commensurability and conjugation, the "simple type lattices" are given precisely by the previous construction, which we quote from [Mor12]. Aside from some exceptional families that appear for $n = 3, 7$, these simple type lattices are the only arithmetic lattices in $\mathcal{H}^n$ for odd n. For even n, there is an additional family of uniform lattices arising from quaternion algebras.

By imposing the restriction $x_j = 0$ for some $1 \leq j \leq n$, and noting that the corresponding equation $x_{n+1}^2 = \sum_{i \neq j} a_i x_i^2$ still admits only the trivial solution, we obtain the following immediate corollary of the result specified in Remark 10.6.

COROLLARY 10.7. *Let $G_{\mathcal{O}}$ be a uniform arithmetic lattice of simple type in* $\mathrm{SO}(1, n)$. *Then $G_{\mathcal{O}}$ contains a subgroup K' that stabilizes an isometric copy of $\mathbb{H}^{n-1}$ in $\mathbb{H}^n$ and is itself a uniform arithmetic lattice* (*of simple type*).

Definition 10.8. The *commensurator* of a subgroup $\mathcal{H}$ in $\mathcal{G}$ is the following subgroup of $\mathcal{G}$:

$$\mathrm{Comm}(\mathcal{H}, \mathcal{G}) = \{g \in \mathcal{G} : [\mathcal{H} : \mathcal{H} \cap g\mathcal{H}g^{-1}] < \infty \text{ and } [\mathrm{g}\mathcal{H}\mathrm{g}^{-1} : \mathcal{H} \cap \mathrm{g}\mathcal{H}\mathrm{g}^{-1}] < \infty\}.$$

Since $\mathrm{Comm}(G_{\mathcal{O}}, \mathcal{H}^n)$ obviously contains $\mathrm{SO}(1, n; \mathbb{Q})$, the following result is readily verified for arithmetic lattices of simple type.

PROPOSITION 10.9. *Let G be an arithmetic lattice in $\mathcal{H}^n$. Then the group* $\mathrm{Comm}(G, \mathcal{H}^n)$ *is dense in $\mathcal{H}^n$.*

Dense means that $\mathcal{H}^n$ is the closure of the subspace $\mathrm{Comm}(G, \mathcal{H}^n)$, where we view $\mathcal{H}^n$ as a topological space in the ordinary way as a Lie group. In fact, the converse to Proposition 10.9 holds and is a deeper result of Margulis, which we do not need.

LEMMA 10.10. *Let G be a uniform arithmetic hyperbolic lattice in $\mathcal{H}^n$ of simple type. There are totally geodesic codimension*-1 *submanifolds $H_1, \ldots, H_k$* (*which we call hyperplanes*) *such that*

(1) $\mathrm{Stabilizer}_G(H_i)$ *acts cocompactly on H_i.*
(2) *The set of hyperplanes $\{gH_i : g \in G, 1 \leq i \leq k\}$ is locally finite.*
(3) *There exists D such that any length D geodesic crosses some gH_i.*

Proof. By Corollary 10.7, there is an $(n-1)$-dimensional hyperplane $H \subset \mathbb{H}^n$ such that $\mathrm{Stabilizer}_G(H)$ acts properly and cocompactly on H. Let $c \in \mathrm{Comm}(G)$. We claim that $\mathrm{Stabilizer}_G(cH)$ acts cocompactly on cH. Indeed let S denote the stabilizer of H in $\mathcal{H}^n$. Since $[G : G \cap G^c] < \infty$, we have $[\mathrm{Stabilizer}_G(cH) : \mathrm{Stabilizer}_G(cH) \cap G^c] < \infty$, and it is enough to prove that $\mathrm{Stabilizer}_G(cH) \cap G^c$ is cocompact on cH. Now $\mathrm{Stabilizer}_G(cH) \cap G^c = S^c \cap G \cap G^c$. By assumption $G \cap G^c$ is of finite index in G^c, thus $S^c \cap G \cap G^c$ is of finite index in S^c. Since S is cocompact on H, it follows that S^c — the stabilizer of cH in G^c — is also cocompact on cH, which ends the argument.

By choosing various elements $c_i \in \mathrm{Comm}(G)$, we are thus able to produce a collection of hyperplanes $H_i = c_iH$, each of which has cocompact stabilizer in G.

Let us check that the cocompactness of $\mathrm{Stabilizer}_G(H)$ on H implies the local finiteness of the family of subsets $\{gH\}_{g\in G}$. By cocompactness, there is a ball A centered at some point of H such that for any $p \in H$, there exists $h \in \mathrm{Stabilizer}_G(H)$ with $h^{-1}p \in A$. For any ball B of $\mathbb{H}^n$ and any g such that $gH \cap B \neq \emptyset$, there exists an $h \in \mathrm{Stabilizer}_G(H)$ such that $ghA \cap B = \emptyset$. Since G acts properly, it follows that the set $\{g \in G, gH \cap B \neq \emptyset\}$ is the union of finitely many cosets $g_1 \mathrm{Stabilizer}_G(H) \cup \cdots \cup g_k \mathrm{Stabilizer}_G(H)$. Thus only finitely many translates of H meet B.

It thus remains to verify that the third property holds for an appropriate choice of $c_1, \ldots, c_k$. Since G is uniform, we can choose a closed radius r ball A such that $GA = \mathbb{H}^n$. Let B, C denote the balls with same center as A and with radius $r + 1, r + 2$.

As any closed convex subset of $\mathbb{H}^n$, B is the intersection of the closed half-spaces containing B. In this intersection we may restrict to the family of half-spaces whose complement meets the sphere ∂C in a nonempty open subset. By compactness of ∂C, there is a finite collection of half-spaces $K_1, \ldots, K_m$ containing B and such that the union of the complements of K_j covers ∂C. In other words, the polytope $\Pi_0 = \cap_i K_i$ is contained in the interior of C.

We now approximate the polyhedron Π (containing B) by a polyhedron Π' whose faces span hyperplanes that are translates of H by elements of the commensuratorand such that $A \subset \Pi' \subset C$. The complements of the K_i provide an open covering of the sphere ∂C. By density of the commensurator, there exists $c_1, \ldots, c_m \in \mathrm{Comm}(G)$ such that $H_i := c_iH$ is so near to ∂K_i that each H_i is disjoint of A, and the complements of the half-spaces K_i' of $\mathbb{H}^n$ bounded by H_i and containing A provide a covering of ∂C. This exactly means that we have $A \subset \Pi' \subset C$.

We now show that any geodesic γ of length $D = 2(r + 2)$ intersects some gH_i. Since D is the diameter of C, any geodesic with initial point in $A \subset \Pi' \subset C$ and length $> D$ has its terminal point outside C, thus crosses the boundary of Π', which means it intersects some bounding hyperplane H_i. Now by the choice of A, the initial point of any geodesic may be translated into A by an element of G. Thus the geodesic intersects some translate of some H_i. $\square$

Acknowledgement. We are grateful to Ian Agol, Nathan Dunfield, Piotr Przytycki, and especially to the referee for many helpful corrections that improved the exposition.

References

[Ago06] I. Agol, untitled, 2006, in preparation.

[ALR01] I. Agol, D. D. Long, and A. W. Reid, The Bianchi groups are separable on geometrically finite subgroups, *Ann. of Math.* **153** (2001), 599–621. MR 1836283. Zbl 1067.20067. http://dx.doi.org/10.2307/2661363.

[BŚ99] W. Ballmann and J. Świątkowski, On groups acting on nonpositively curved cubical complexes, *Enseign. Math.* **45** (1999), 51–81. MR 1703363. Zbl 0989.20029. http://dx.doi.org/10.5169/seals-64441.

[BHW11] N. Bergeron, F. Haglund, and D. T. Wise, Hyperplane sections in arithmetic hyperbolic manifolds, *J. Lond. Math. Soc.* **83** (2011), 431–448. MR 2776645. Zbl 1236.57021. http://dx.doi.org/10.1112/jlms/jdq082.

[BS00] M. Bonk and O. Schramm, Embeddings of Gromov hyperbolic spaces, *Geom. Funct. Anal.* **10** (2000), 266–306. MR 1771428. Zbl 0972.53021. http://dx.doi.org/10.1007/s000390050009.

[BS05] S. Buyalo and V. Schroeder, Embedding of hyperbolic spaces in the product of trees, *Geom. Dedicata* **113** (2005), 75–93. MR 2171299. Zbl 1090.54029. http://dx.doi.org/10.1007/s10711-005-3124-9.

[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., Novosibirsk, 1997, pp. 91–109, 190. MR 1624115. Zbl 0906.20025.

[Git99] R. Gitik, Ping-pong on negatively curved groups, *J. Algebra* **217** (1999), 65–72. MR 1700476. Zbl 0936.20019. http://dx.doi.org/10.1006/jabr.1998.7789.

[GMRS98] R. Gitik, M. Mitra, E. Rips, and M. Sageev, Widths of subgroups, *Trans. Amer. Math. Soc.* **350** (1998), 321–329. MR 1389776. Zbl 0897.20030. http://dx.doi.org/10.1090/S0002-9947-98-01792-9.

[Gro87] M. Gromov, Hyperbolic groups, in *Essays in Group Theory*, *Math. Sci. Res. Inst. Publ.* **8**, Springer-Verlag, New York, 1987, pp. 75–263. MR 0919829. Zbl 0634.20015.

[Hag] F. Haglund, Aspects combinatoires de la théorie géométrique des groupes, Université Paris Sud, Orsay, 2007, Habilitation.

[Hag08] ———, Finite index subgroups of graph products, *Geom. Dedicata* **135** (2008), 167–209. MR 2413337. Zbl 1155.53025. http://dx.doi.org/10.1007/s10711-008-9270-0.

[HaW08] F. Haglund and D. T. Wise, Special cube complexes, *Geom. Funct. Anal.* **17** (2008), 1551–1620. MR 2377497. Zbl 1155.53025. http://dx.doi.org/10.1007/s00039-007-0629-4.

[HaW10] ———, Coxeter groups are virtually special, *Adv. Math.* **224** (2010), 1890–1903. MR 2646113. Zbl 1195.53055. http://dx.doi.org/10.1016/j.aim.2010.01.011.

[HrW09] G. C. Hruska and D. T. Wise, Packing subgroups in relatively hyperbolic groups, *Geom. Topol.* **13** (2009), 1945–1988. MR 2497315. Zbl 1188.20042. http://dx.doi.org/10.2140/gt.2009.13.1945.

[HrW10] G. C. Hruska and D. T. Wise, Finiteness properties of cubulated groups, 2010, preprint. Available at https://docs.google.com/file/d/0B45cNx80t5-2U05ST2liTEpEU3c/edit?pli=1.

[Min06] A. Minasyan, Separable subsets of GFERF negatively curved groups, *J. Algebra* **304** (2006), 1090–1100. MR 2264291. Zbl 1175.20034. http://dx.doi.org/10.1016/j.jalgebra.2006.03.050.

[Mor12] D. W. Morris, Introduction to arithmetic groups, 2012. arXiv math/0106063.

[Rol] M. A. Roller, Poc-sets, median algebras and group actions. an extended study of Dunwoody's construction and Sageev's theorem. Available at www.personal.soton.ac.uk/gan/Roller.pdf.

[Sag95] M. Sageev, Ends of group pairs and non-positively curved cube complexes, *Proc. London Math. Soc.* **71** (1995), 585–617. MR 1347406. Zbl 0861.20041. http://dx.doi.org/10.1112/plms/s3-71.3.585.

[Sco78] P. Scott, Subgroups of surface groups are almost geometric, *J. London Math. Soc.* **17** (1978), 555–565. MR 0494062. Zbl 0412.57006. http://dx.doi.org/10.1112/jlms/s2-17.3.555.

[Wis00] D. T. Wise, Subgroup separability of graphs of free groups with cyclic edge groups, *Q. J. Math.* **51** (2000), 107–129. MR 1760573. Zbl 0991.05056. http://dx.doi.org/10.1093/qmathj/50.1.107.

[Wis02] ———, The residual finiteness of negatively curved polygons of finite groups, *Invent. Math.* **149** (2002), 579–617. MR 1923477. Zbl 1040.20024. http://dx.doi.org/10.1007/s002220200224.

(Received: February 25, 2009)
(Revised: February 25, 2012)

Université de Paris XI (Paris-Sud), Orsay, France
E-mail: frederic.haglund@math.u-psud.fr

McGill University, Montreal, Quebec, Canada
E-mail: wise@math.mcgill.ca