



AN ALTERNATE PROOF OF WISE'S MALNORMAL SPECIAL QUOTIENT THEOREM

IAN AGOL¹, DANIEL GROVES² and JASON FOX MANNING³

¹ University of California, Berkeley, 970 Evans Hall #3840, Berkeley, CA 94720-3840, USA;
email: ianagol@math.berkeley.edu

² Department of Mathematics, Statistics, and Computer Science,
University of Illinois at Chicago, 322 Science and Engineering Offices (M/C 249),
851 S. Morgan St., Chicago, IL 60607-7045, USA;
email: groves@math.uic.edu

³ Department of Mathematics, 310 Malott Hall, Cornell University, Ithaca, NY 14853, USA;
email: jfmanning@math.cornell.edu

Received 11 September 2014; accepted 12 November 2015

Abstract

We give an alternate proof of Wise's malnormal special quotient theorem (MSQT), avoiding cubical small cancellation theory. We also show how to deduce Wise's Quasiconvex Hierarchy Theorem from the MSQT and theorems of Hsu and Wise and Haglund and Wise.

2010 Mathematics Subject Classification: 20F67 (20E08)

1. Introduction

The role of *subgroup separability* is a central one in geometric group theory and topology. In particular, as witnessed by Scott's Criterion [Sco78], it is intimately linked to the problem of promoting immersions to embeddings in finite covers. After the work of Kahn and Markovic [KM12], the Virtual Haken Conjecture was reduced to proving separability of certain surface subgroups of 3-manifold groups. The key to the first author's proof of this separability in [Ago13] was the notion of *virtually special cube complexes*, as developed by Wise and his collaborators.

In [HW08], Haglund and Wise proved that if G is a hyperbolic group which acts properly and cocompactly on a CAT(0) cube complex then the action is

© The Author(s) 2016. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

virtually special if and only if all quasiconvex subgroups of G are separable. This leads to the desirability of finding criteria which show that a group acts in a virtually special (and proper cocompact) way on a CAT(0) cube complex.

The results of Hsu and Wise [HW15] and Haglund and Wise [HW12] together show that when a hyperbolic group G has a *malnormal quasiconvex hierarchy* then it has the desired type of action. See Section 3 for the definition of hierarchies and Section 6 for more details, in particular Theorem 6.1 which combines the results of [HW15, HW12]. In the unpublished work [Wis], Wise proves that hyperbolic groups with *quasiconvex hierarchies* also admit such actions, and hence all their quasiconvex subgroups are separable. Wise's result is much more broadly applicable than Theorem 6.1 because subgroups appearing in the hierarchy in many natural applications will not be malnormal. See Section 10 for more details and a proof of Wise's result, which appears as Theorem 10.2 in this paper.

Subsequently, the first author [Ago13] proved that *any* hyperbolic group which acts properly and cocompactly on a CAT(0) cube complex acts in a virtually special way, with no further hypotheses. One of the main purposes of this paper is to explain some of the ingredients (from [Wis]) that go into this proof, so we do not assume the results from [Ago13] here. In fact, this paper is intended to provide an alternative account of the results needed for [Ago13] without relying on any of the results from [Wis]. It is important to note that Wise proves various results in greater generality than is required for [Ago13], and in this paper we do not recover the full strength of the results in [Wis]. In particular, we work throughout with hyperbolic groups, whereas certain of Wise's results are proved for more general relatively hyperbolic groups.

As we explain in Section 10 (and is made clear in [Wis]), the key result for proving Wise's Theorem 10.2 is the *Malnormal Special Quotient Theorem (MSQT)*, and the main purpose of the current paper is to provide a new proof of this theorem. In fact, we prove Theorem 2.7 below, which is a generalization of Wise's MSQT.

Our key innovation in the proof is a new kind of hierarchy for virtually special hyperbolic groups (described in Section 5). It is a consequence of work of Haglund and Wise [HW08] that any virtually special hyperbolic group has a finite-index subgroup with a malnormal quasiconvex hierarchy terminating in the trivial group. However, the MSQT requires us to 'kill' certain subgroups of a malnormal family \mathcal{P} of quasiconvex 'peripheral' subgroups. We construct (in Theorem 2.11) a new hierarchy (of a different finite-index subgroup) which is quasiconvex but does not terminate in the trivial group. The peripheral subgroups of the finite-index subgroup are conjugate to finite-index subgroups of the elements of \mathcal{P} . The virtue of this new hierarchy is that its peripheral subgroups

are elliptic at every stage of the hierarchy, and furthermore when they intersect an edge group of a splitting in an infinite set they are entirely contained in the edge group. (See Section 3.5 and particularly Definition 3.31 for more details.) Therefore (using certain results about Relatively Hyperbolic Dehn Filling—see Section 7), the hierarchy and the filling are operations which ‘commute,’ and we find a hierarchy of the filled group. Using the results of Haglund and Wise and Hsu and Wise described above, we deduce that the filled group is virtually special, as required. In Section 4, we prove a Combination Theorem which (as with most combination theorems) states that certain subgroups are amalgamated free products or HNN extensions, and are also quasiconvex. This allows us to see that the hierarchy which we build topologically has the expected algebraic structure, which in turn allows us to deduce that the hierarchy is quasiconvex. We now provide a more detailed and technical outline of the main work in this paper.

2. Outline

Special cube complexes are defined in [HW08, Definition 3.2]. Note that [HW08, Proposition 3.10] states that a special complex has a finite A -special cover, and many people use the term ‘special’ to mean A -special. Since we are only interested in this property virtually, the distinction is not important. In fact, we do not need to use the definition of special cube complexes in this paper. Rather, we use the notion as a black box by quoting results, particularly those from [HW08, HW15, HW12].

DEFINITION 2.1. We say that a group G is *special* if it is the fundamental group of a compact special cube complex. (Other authors use the term *compact special*.)

A group is *virtually special* if it has a finite-index subgroup which is special.

REMARK 2.2. If G is a hyperbolic group which acts properly and cocompactly on a CAT(0) cube complex \tilde{X} then G is virtually special if and only if all quasiconvex subgroups of G are separable [HW08, Theorems 1.3 and 1.4]. Therefore, being virtually special is a property of the (hyperbolic) group, not the particular cube complex upon which it acts properly and cocompactly.

As mentioned in the introduction, the Main Theorem from [Ago13] shows that any hyperbolic group which acts properly and cocompactly on a CAT(0) cube complex is virtually special. However, in this paper we are giving an alternative account of some of the ingredients of the proof of this theorem, and so we do not assume this result from [Ago13].

NOTATION 2.3. Let A be a group. The notation $B \dot{<} A$ indicates that $B < A$ and $|A : B| < \infty$. Similarly, the notation $B \dot{\triangleleft} A$ indicates $B \triangleleft A$ and $|A : B| < \infty$.

DEFINITION 2.4. Let $\mathcal{P} = \{P_1, \dots, P_m\}$ be a finite collection of subgroups of a group G (called a *peripheral structure* for G). A choice of subgroups (called *filling kernels*) $\{N_i \triangleleft P_i\}_{i=1}^m$ gives rise to a (Dehn) filling of (G, \mathcal{P}) :

$$G \rightarrow G(N_1, \dots, N_m)$$

with kernel $\langle\langle \cup_i N_i \rangle\rangle_G$. A Dehn filling $G \rightarrow G(N_1, \dots, N_m)$ is called *peripherally finite* if $N_i \triangleleft P_i$ for all i .

We want to apply Dehn filling to relatively hyperbolic pairs (G, \mathcal{P}) where G is itself hyperbolic. To see when we can do this, recall the definition of a malnormal collection of subgroups.

DEFINITION 2.5. A collection \mathcal{P} of subgroups of G is said to be *malnormal* (respectively *almost malnormal*) if for any $P, P' \in \mathcal{P}$, and $g \in G$, either $P^g \cap P'$ is trivial (respectively finite), or $P = P'$ and $g \in P$.

The following well-known characterization of relatively hyperbolic pairs (G, \mathcal{P}) where G is hyperbolic follows from results of Bowditch and Osin.

THEOREM 2.6 [Bow12, Osi06]. *Suppose G is hyperbolic and \mathcal{P} is a finite collection of subgroups of G . The pair (G, \mathcal{P}) is relatively hyperbolic if and only if \mathcal{P} is an almost malnormal collection of quasiconvex subgroups.*

Proof. The forward direction follows from some much more general facts about parabolic subgroups of relatively hyperbolic groups: If (G, \mathcal{P}) is any relatively hyperbolic pair, then the collection \mathcal{P} is almost malnormal by [Osi06, Proposition 2.36]. Moreover, the elements of \mathcal{P} are undistorted in G [Osi06, Lemma 5.4]. Undistorted subgroups of a hyperbolic group are quasiconvex.

The other direction is proved by Bowditch [Bow12, Theorem 7.11]. \square

Here is our main result:

THEOREM 2.7 (Main theorem). *Let G be hyperbolic and virtually special, and suppose (G, \mathcal{P}) is relatively hyperbolic. There are subgroups $\{\dot{P}_i \dot{\triangleleft} P_i\}$ so that if $\bar{G} = G(N_1, \dots, N_m)$ is any filling with $N_i < \dot{P}_i$ and P_i/N_i virtually special and hyperbolic for all i , then \bar{G} is hyperbolic and virtually special.*

Since finite groups are virtually special and hyperbolic, we recover the following theorem of Wise.

COROLLARY 2.8 (The malnormal special quotient theorem, [Wis]). *Let G be hyperbolic and virtually special, and suppose (G, \mathcal{P}) is relatively hyperbolic. There are subgroups $\{P_i \triangleleft P_i\}$ so that if $\bar{G} = G(N_1, \dots, N_m)$ is any peripherally finite filling with $N_i < P_i$ for all i , then \bar{G} is hyperbolic and virtually special.*

The proof of Theorem 2.7 relies on the following theorem, which is [Wis, Theorem 11.2] in Wise's manuscript. In Section 3, we explain the terminology, and in Section 6, we recall how the result follows almost immediately from the work of Hsu and Wise [HW15] and Haglund and Wise [HW12].

THEOREM 6.1 (Malnormal quasiconvex hierarchy). *If G is a hyperbolic group with a malnormal quasiconvex hierarchy terminating in a collection of virtually special groups, then G is virtually special.*

Thus, in order to prove Theorem 2.7, it suffices to prove that \bar{G} (under suitable hypotheses on the filling) is a hyperbolic group which (virtually) has a malnormal quasiconvex hierarchy terminating in a collection of virtually special groups. The key results for proving that this is the case are Theorems 2.11 and 2.12 below.

NOTATION 2.9 (Peripheral structure on a finite-index subgroup). Given $G' \triangleleft G$, a peripheral structure \mathcal{P} on G induces a peripheral structure \mathcal{P}' on G' : For each $P \in \mathcal{P}$, let $\mathcal{E}_0(P) = \{gPg^{-1} \cap G' \mid g \in G\}$, and let $\mathcal{E}(P)$ be obtained from $\mathcal{E}_0(P)$ by choosing one element of each G' -conjugacy class of subgroup. Note that $\mathcal{E}(P)$ is finite, and each element is conjugate to a finite-index subgroup of P . Now let

$$\mathcal{P}' = \bigsqcup_{P \in \mathcal{P}} \mathcal{E}(P).$$

We summarize all this with the notation $(G', \mathcal{P}') \triangleleft (G, \mathcal{P})$.

Note that (G', \mathcal{P}') is relatively hyperbolic if and only if (G, \mathcal{P}) is.

REMARK 2.10. The collection \mathcal{P}' just defined can also be recovered from topology, if we keep in mind the fundamental group of a subcomplex $Y \subset X$ is only defined up to conjugacy in $\pi_1 X$. Suppose $G = \pi_1 X$ and \mathcal{Y} is a collection of subcomplexes of X whose fundamental groups correspond to \mathcal{P} . Let $X' \rightarrow X$ be a regular finite-degree cover, so $\pi_1 X' = G' \triangleleft G$. Let \mathcal{Y}' be the collection of

preimages of elements of \mathcal{Y} in X' . Then \mathcal{P}' is the set of fundamental groups of elements of \mathcal{Y}' .

Our main result is obtained from the following two statements. (See Section 3 below for definitions of the terminology relating to hierarchies.)

THEOREM 2.11 (Virtual relative hierarchy). *Given (G, \mathcal{P}) relatively hyperbolic, with G hyperbolic and virtually special, there is some $(G', \mathcal{P}') \triangleleft (G, \mathcal{P})$ which has a malnormal quasiconvex fully \mathcal{P}' -elliptic hierarchy terminating in \mathcal{P}' .*

Theorem 2.11 is proved in Section 5.6.

THEOREM 2.12 (Hierarchy filling). *Suppose (G, \mathcal{P}) is relatively hyperbolic with G hyperbolic and having a malnormal quasiconvex fully \mathcal{P} -elliptic hierarchy terminating in \mathcal{P} . For all sufficiently long fillings,*

$$(G, \mathcal{P}) \longrightarrow (\bar{G}, \bar{\mathcal{P}}),$$

with all elements of $\bar{\mathcal{P}}$ hyperbolic, the group \bar{G} is hyperbolic and has a malnormal quasiconvex hierarchy terminating in $\bar{\mathcal{P}}$.

Theorem 2.12 is proved in Section 8.2.

In Section 9, we prove the main theorem from the two theorems above. Briefly one argues as follows: Given (G, \mathcal{P}) as in the theorem, one passes to a finite-index subgroup (G_0, \mathcal{P}_0) satisfying the conclusion of Theorem 2.11. This subgroup has a malnormal quasiconvex hierarchy terminating in peripheral subgroups. Theorem 2.12 says that any sufficiently long filling of (G_0, \mathcal{P}_0) still has such a hierarchy. Now, assume a filling (in addition to being sufficiently long) has the property that P/N is always virtually special, for $P \in \mathcal{P}_0$ and N the corresponding filling kernel. Using the Relatively Hyperbolic Dehn Filling Theorem 7.2, these P/N are precisely the peripheral subgroups of the filling. Using the Malnormal Quasiconvex Hierarchy Theorem 6.1 the filling is virtually special. If the filling was chosen appropriately, it gives a filling of (G, \mathcal{P}) which is therefore also virtually special.

2.1. Conventions. Let Δ be a geodesic triangle with vertices $\{x, y, z\}$ in a geodesic space X . The triangle determines a metric tripod T_Δ , and a comparison map $\phi : \Delta \rightarrow T_\Delta$, isometric on each side of Δ . The leg of T_Δ adjacent to $\phi(x)$ has length equal to the *Gromov product*

$$(y, z)_x := \frac{1}{2}(d(x, y) + d(x, z) - d(y, z)).$$

The triangle Δ is said to be δ -thin if $\phi^{-1}(t)$ has diameter at most δ , for all $t \in T_\Delta$. For us a space is δ -hyperbolic if all the triangles are δ -thin. (For the relationship between this and other definitions of hyperbolicity, see [BH99, Ch. III.H].)

3. Hierarchies

3.1. Hierarchies of groups. We refer to Serre [Ser80] for more detail on the basic theory of graphs of groups, their fundamental groups, and the relationship with actions on trees. All the underlying graphs of a graph of groups will be finite. In contrast to Serre, we do not require that the maps from edge groups to vertex groups be injective. Rather, in our applications a graph of groups is found from a graph of spaces via van Kampen's Theorem, and much of our work is to prove that the maps from edge spaces to vertex spaces are π_1 -injective.

Recall that for Serre, a graph is a pair of sets (V, E) , with a fixed-point free involution $e \mapsto \bar{e}$ on E (exchanging an oriented edge with its orientation reverse), and a 'terminus' map $t : E \rightarrow V$ mapping an edge to the vertex it is oriented toward. The 'initial' point of an edge is $i(e) := t(\bar{e})$.

DEFINITION 3.1. A *graph of groups* (Γ, \mathcal{G}) is:

- (1) a graph $\Gamma = (V, E)$;
- (2) an assignment $\mathcal{G} : V \sqcup E \rightarrow \mathbf{Grp}$ of a group to each edge and vertex of Γ , satisfying $\mathcal{G}(\bar{e}) = \mathcal{G}(e)$; and
- (3) for each edge $e \in E$, a homomorphism $\phi_e : \mathcal{G}(e) \rightarrow \mathcal{G}(t(e))$.

The graph of groups is *faithful* if all homomorphisms ϕ_e are injective.

REMARK 3.2. For any graph of (pointed) spaces (see Section 3.2 below), there is a corresponding graph of groups obtained by applying the π_1 functor to the spaces and maps, and using the same underlying graph. The groupoid of homotopy classes of paths between basepoints of vertex spaces is isomorphic to the fundamental groupoid of the corresponding graph of groups. Similarly, the fundamental group of the graph of spaces is isomorphic to the fundamental group of the graph of groups.

DEFINITION 3.3. A *graph of groups structure for G* is a graph of groups (Γ, \mathcal{G}) together with an isomorphism between G and $\pi_1(\Gamma, \mathcal{G}, v)$, the fundamental group of the graph of groups based at some vertex v . The structure is *degenerate* if the graph Γ consists of a single point v , with $\mathcal{G}(v) = G$, and the isomorphism is the identity map. We will say the structure is *faithful* if (Γ, \mathcal{G}) is faithful.

DEFINITION 3.4. A *hierarchy of length 0* for G is the degenerate graph of groups structure for G . Let $n > 0$. A *hierarchy of length n* for G is a (finite) graph of groups structure for G , together with a hierarchy of length $(n - 1)$ for each vertex group.

Note that a hierarchy of length 1 for G is the same thing as a graph of groups structure for G . Longer hierarchies can be thought of as ‘multilevel graphs of groups.’

DEFINITION 3.5 (Levels of a hierarchy, terminal groups). Let G_v be a vertex group of this graph of groups structure. We say that G_v is at *level 1* of the hierarchy. The groups at level 1 of the hierarchy for G_v are at *level 2* of the hierarchy for G , and so on. If \mathcal{H} is a hierarchy, we sometimes write ‘ $H \in \mathcal{H}$ ’ to mean that H is at level i of \mathcal{H} for some i .

If a hierarchy has length n , the groups at level n are called the *terminal groups* of the hierarchy.

It is convenient in this definition to allow any or all of the graphs of groups to be degenerate, though of course we are mostly interested in the case when at least some of the graphs of groups at each level are nondegenerate.

DEFINITION 3.6. A *faithful* hierarchy is one for which every graph of groups occurring in the hierarchy is faithful.

In a faithful hierarchy, each vertex or edge group at any level embeds in G , and so can be thought of as a subgroup of G . In general, $H \in \mathcal{H}$ only gives a subgroup up to conjugacy in G , but we can pin it down if necessary by choosing maximal trees in all the underlying graphs of the hierarchy.

DEFINITION 3.7. Let \mathcal{H} be a faithful hierarchy of G , and let \mathcal{P} be a collection of subgroups of G . If every terminal group of \mathcal{H} is conjugate to an element of \mathcal{P} , we say that the hierarchy *terminates in \mathcal{P}* .

REMARK 3.8 (Warning). Even in the faithful case, our terminology differs from Wise’s in [Wis12, Wis] in two ways. First, Wise’s hierarchies consist of a sequence of *one-edge* splittings, and not general graphs of groups. Our hierarchies could be converted to ones satisfying this property, at the expense of lengthening them. Second, Wise usually requires the terminal groups to be trivial (sometimes virtually special), and we do not have any such requirement in general.

Sometimes hierarchies have some extra properties which we like. Here are some of them.

DEFINITION 3.9. Let \mathcal{H} be a faithful hierarchy for the group G . Each $H \in \mathcal{H}$ comes equipped with a graph of groups structure. Write $\mathcal{E}(H)$ for the set of edge groups of this structure. We say the hierarchy \mathcal{H} is:

- (1) *quasiconvex* if G is finitely generated and, for all $H \in \mathcal{H}$, $K \in \mathcal{E}(H)$, K is quasiisometrically embedded in H ; and
- (2) *(almost) malnormal* if, for all $H \in \mathcal{H}$, $K \in \mathcal{E}(H)$, K is (almost) malnormal in G .

REMARK 3.10 (Different kinds of malnormal hierarchy). There are a number of reasonable choices for the definition of a malnormal hierarchy. A stronger version might require $\mathcal{E}(H)$ to be a malnormal collection in G (or just in H). A weaker (and more common) version just requires the elements of $\mathcal{E}(H)$ to be malnormal in H and not necessarily in G .

3.2. Homotopy hierarchies of spaces. Hierarchies of spaces are multilevel graphs of spaces, so we start by recalling what is meant by a graph of spaces.

DEFINITION 3.11. A *graph of spaces* (Γ, \mathcal{X}) consists of the following data:

- (1) A graph $\Gamma = (V, E)$;
- (2) an assignment of path-connected spaces $\mathcal{X} : V \sqcup E \rightarrow \mathbf{Top}$ satisfying $\mathcal{X}(\bar{e}) = \mathcal{X}(e)$;
- (3) for each edge e a continuous map $\psi_e : \mathcal{X}(e) \rightarrow \mathcal{X}(t(e))$.

For each $v \in V$, let $E_v = \{e \in E \mid t(e) = v\}$ be the collection of edges incident on v . We combine all the maps ψ_e for edges incident on v to get a map

$$\Psi_v : \bigsqcup_{e \in E_v} \mathcal{X}(e) \rightarrow \mathcal{X}(v).$$

The *semistar* $S(v)$ is defined to be the mapping cylinder of Ψ_v . The *realization* $R(\Gamma, \mathcal{X})$ is then defined to be the union of the semistars, glued together using the identifications $\mathcal{X}(e) = \mathcal{X}(\bar{e})$. We will sometimes abuse language and refer to this realization as a graph of spaces.

DEFINITION 3.12. A *graph of spaces structure* on X is a homotopy equivalence $h : R(\Gamma, \mathcal{X}) \rightarrow X$ for some graph of spaces (Γ, \mathcal{X}) . The structure is *degenerate* if Γ consists of a single point labeled by X , and the homotopy equivalence is the identity.

REMARK 3.13. A graph of spaces (Γ, \mathcal{X}) gives rise to a graph of groups (Γ, \mathcal{G}) in a straightforward way. Indeed, for each vertex or edge space $\mathcal{X}(x)$, one may choose a basepoint b_x . The edge or vertex group $\mathcal{G}(x)$ is defined to be $\pi_1(\mathcal{X}(x), b_x)$. The maps ψ_e need not be basepoint preserving, so we also make a choice, for each edge e , of a path σ_e from $\psi_e(b_e)$ to $b_{t(e)}$. Such a path gives an identification of $\mathcal{G}(v)$ with $\pi_1(\mathcal{X}(v), \psi_e(b_e))$. Using this identification, we can define $\phi_e : \mathcal{G}(e) \rightarrow \mathcal{G}(t(e))$ to be $(\psi_e)_*$.

The fundamental groupoid of this graph of groups [Hig76] is isomorphic to the groupoid of homotopy classes of paths in $R(\Gamma, \mathcal{X})$ with endpoints in $\{b_v \mid v \text{ a vertex of } \Gamma\}$.

Conversely, given any graph of groups we may build a corresponding graph of basepointed spaces in a straightforward way, realizing each edge and vertex space by a $K(\pi, 1)$, and setting the maps ψ_e to be continuous pointed maps inducing the homomorphisms ϕ_e .

DEFINITION 3.14. A *hierarchy of length 0* for a space X is the degenerate graph of spaces structure on X .

A *hierarchy of length n* for X is a graph of spaces structure on X , together with a hierarchy of length $(n - 1)$ on each vertex space.

The *levels* of the structure are the collections of vertex spaces at each stage. That is, the vertex spaces at level one are the vertex spaces of the graph of spaces structure on X ; those at level 2 are the vertex spaces of the spaces at level 1, and so on.

The vertex spaces at level n of a length n hierarchy are called the *terminal spaces*.

At each level of a hierarchy, there is a collection of homotopy equivalences connecting the graphs of spaces at that level with the vertex spaces of the previous level. In the following two examples, these homotopy equivalences are actually homeomorphisms. However, note that in our proof of the Virtual Hierarchy Theorem 2.11, we need to consider hierarchies in which the homotopy equivalences are *not* homeomorphisms.

EXAMPLE 3.15 (Special cube complex). The hyperplanes of a (compact) special cube complex X induce a hierarchy of spaces which is particularly simple. Order

the hyperplanes H_1, \dots, H_n . Each hyperplane H_i is two-sided, so the (open) cubes which intersect it give a product neighborhood $N(H_i) \cong H_i \times (-1, 1)$. Each hyperplane H_i is the edge space for a graph of spaces structure with underlying graph either an edge (if H_i is separating) or a loop (otherwise). The (one or two) components of $X \setminus N(H_i)$ are the vertex spaces. The order allows us to use these decompositions to define a hierarchy.

Level 0 is the cube complex X ; level 1 consists of the components of $X \setminus N(H_1)$; level k consists of the components of $X \setminus (\bigcup_{i=1}^k N(H_i))$. The edge spaces at level k are the components of $H_k \setminus (\bigcup_{i=1}^{k-1} N(H_i))$. The terminal spaces are points.

EXAMPLE 3.16 (Haken 3-manifold). Each surface in the Haken hierarchy has a product neighborhood, so we can regard this hierarchy as a hierarchy of spaces in an obvious way. The terminal spaces are balls.

REMARK 3.17. Using Remark 3.13, each graph of spaces gives rise to a graph of groups, via some choices of basepoints and connecting arcs. A hierarchy of spaces therefore gives rise to a hierarchy of groups in a straightforward way.

There is also a general construction of a hierarchy of spaces from any hierarchy of groups, following the last paragraph of Remark 3.13. In this construction the ability to use homotopy equivalences rather than homeomorphisms is very convenient.

3.3. Induced hierarchies on subgroups. The following result follows easily by considering the action of $H < G$ on the Bass–Serre tree corresponding to a graph of groups structure (Γ, \mathcal{G}) for G .

PROPOSITION 3.18 [Bas93]. *Suppose that G has a graph of groups structure (Γ, \mathcal{G}) , and that $H \triangleleft G$. Then H has an induced graph of groups structure $(\tilde{\Gamma}, \mathcal{H})$ so that*

- (1) *Each vertex group of $(\tilde{\Gamma}, \mathcal{H})$ is equal to $(K^g \cap H) \triangleleft K^g$ for some vertex group K of (Γ, \mathcal{G}) and some $g \in G$.*
- (2) *Each edge group of $(\tilde{\Gamma}, \mathcal{H})$ is equal to $(K^g \cap H) \triangleleft K^g$ for some edge group K of (Γ, \mathcal{G}) and some $g \in G$.*

Proposition 3.18 has the following corollary, by induction on n :

COROLLARY 3.19. *Suppose that G has a hierarchy \mathcal{H} of length n , and that $H \triangleleft G$. Then the hierarchy on G induces a hierarchy \mathcal{H}' of length n on H . Moreover, for each $i \in \{0, \dots, n\}$,*

- (1) Each vertex group at level i of \mathcal{H}' is equal to $(K^g \cap H) \triangleleft K^g$ for some vertex group K at level i of \mathcal{H} and some $g \in G$.
- (2) Each edge group at level i of \mathcal{H}' is equal to $(K^g \cap H) \triangleleft K^g$ for some edge group K at level i of \mathcal{H} and some $g \in G$.

3.4. Quasiconvex hierarchies of special groups.

3.4.1. *Quasiconvexity.* The following lemma is easy, thinking of the Cayley graph of G as coarsely having the structure of a tree of spaces.

LEMMA 3.20. *Suppose that G has a graph of groups structure (Γ, \mathcal{G}) . If the edge groups of (Γ, \mathcal{G}) are quasiisometrically embedded in G , then so are the vertex groups.*

COROLLARY 3.21. *If \mathcal{H} is a quasiconvex hierarchy of G , then all the edge and vertex groups of \mathcal{H} are quasiisometrically embedded in G .*

The following lemma is straightforward.

LEMMA 3.22. *If \mathcal{H} is a quasiconvex hierarchy of G , and $G_0 \triangleleft G$, then the induced hierarchy \mathcal{H}_0 on G_0 is a quasiconvex hierarchy.*

3.4.2. *Malnormality.* Recall that a subgroup $H < G$ is *malnormal* if $H \cap gHg^{-1}$ is trivial for all $g \notin H$, and is *almost malnormal* if $H \cap gHg^{-1}$ is finite for all $g \notin H$. The failure of (almost) malnormality of a subgroup is measured by the height.

DEFINITION 3.23 (Height). Let $H < G$. The *height* of H in G is the largest number $n \geq 0$ so that there are n distinct cosets $\{g_1H, \dots, g_nH\}$, so that the intersection of conjugates $\bigcap_i g_iHg_i^{-1}$ is infinite. Thus finite groups have height 0, infinite almost malnormal subgroups have height 1, and so on.

The following lemma is straightforward and left to the reader.

LEMMA 3.24. *Suppose $H < G$, and $G_0 \triangleleft G$. Let $H_0 = H \cap G_0$. If H is malnormal in G , then H_0 is malnormal in G_0 . More generally, the height of H_0 in G_0 is at most the height of H in G .*

DEFINITION 3.25. Say that $H < G$ is *virtually malnormal* if there is a finite-index subgroup $G_0 < G$ so that if $H_0 = H \cap G_0$, then H_0 is malnormal in G_0 .

DEFINITION 3.26. The subgroup $H < G$ is *separable in G* (or just *separable* when G is understood) if

$$H = \bigcap \{G_0 \mid H < G_0 \triangleleft G\}.$$

PROPOSITION 3.27 [HW08, HW09]. *If G is hyperbolic and virtually torsion-free, and $H < G$ is quasiconvex and separable, then H is virtually malnormal.*

Proof. Hruska and Wise [HW09, Theorem 9.3] implies that a separable quasiconvex subgroup of a hyperbolic group is virtually *almost* malnormal. Indeed they show that there is some $G_0 \triangleleft G$ containing H so that H is almost malnormal in G_0 . By hypothesis, G is virtually torsion-free. Let $G_1 \triangleleft G$ be a torsion-free subgroup, and let $G_2 = G_0 \cap G_1$. We then have $H_2 = H \cap G_2$ malnormal in G_2 , using Lemma 3.24. \square

THEOREM 3.28 [HW08, Theorem 1.3]. *Let G be hyperbolic and virtually special, and let $H < G$ be quasiconvex. Then H is separable in G .*

COROLLARY 3.29. *If G is hyperbolic and virtually special, and $H < G$ is quasiconvex, then H is virtually malnormal.*

Proof. Since G is virtually special it is linear, in particular virtually torsion-free. The subgroup H is separable by Theorem 3.28. We can therefore apply Proposition 3.27. \square

THEOREM 3.30. *If \mathcal{H} is a quasiconvex hierarchy of a hyperbolic virtually special group G , then for some $G_0 \triangleleft G$, the induced hierarchy \mathcal{H}_0 of G_0 is a malnormal hierarchy.*

Proof. The finitely many edge groups of \mathcal{H} are all virtually malnormal by Proposition 3.29. Lemma 3.24 implies malnormality is preserved by passing to further finite-index subgroups. We may therefore find $G_0 \triangleleft G$ so that for every edge group H of \mathcal{H} , the intersection $H \cap G_0$ is malnormal in G_0 . By Corollary 3.19, every edge group of the induced hierarchy is of the form $H^g \cap G_0$ for some such edge group H , and some $g \in G$. Conjugation by $g \in G$ gives an automorphism of G_0 , so all these edge groups are malnormal in G_0 . In particular, any edge group at level i is malnormal in the vertex group at level $i - 1$ which contains it, so the hierarchy \mathcal{H}_0 is a malnormal hierarchy. \square

3.5. Relative hierarchies. In this paper, we are particularly interested in hierarchies relative to a family of peripheral subgroups.

DEFINITION 3.31 (\mathcal{P} -elliptic hierarchy). Let G be a group with a hierarchy \mathcal{H} , and let \mathcal{P} be a family of subgroups of G . Recall that any vertex group $H \in \mathcal{H}$ comes equipped with some graph of groups structure $H \cong \pi_1(\Gamma, \mathcal{G}, v)$. We say that \mathcal{H} is \mathcal{P} -elliptic if it satisfies the following condition:

Suppose that $H \in \mathcal{H}$, and suppose that P' is any G -conjugate of $P \in \mathcal{P}$ satisfying $P' \subseteq H$. Then P' is H -conjugate into a vertex group of (Γ, \mathcal{G}) .

A \mathcal{P} -elliptic hierarchy \mathcal{H} is *fully \mathcal{P} -elliptic* if whenever P' is any G -conjugate of $P \in \mathcal{P}$ and whenever E is an edge group in \mathcal{H} then either $P' \cap E$ is finite or else $P' < E$.

In other words, a hierarchy is fully \mathcal{P} -elliptic if the elements of \mathcal{P} are never cut up by the splittings in the hierarchy. (We think of ‘cutting’ along edge groups, which is the algebraic consequence of cutting along subspaces of a topological space.)

Recall the notation $(G_0, \mathcal{P}_0) \triangleleft (G, \mathcal{P})$, which means that $G_0 \triangleleft G$, and that the elements of \mathcal{P}_0 are representatives of all the G_0 -conjugacy classes of subgroup of the form $P^g \cap G_0$ where $P \in \mathcal{P}$ and $g \in G$. The following result is straightforward.

LEMMA 3.32. *Suppose that $(G_0, \mathcal{P}_0) \triangleleft (G, \mathcal{P})$ and that G has a fully \mathcal{P} -elliptic hierarchy \mathcal{H} . Then the induced hierarchy \mathcal{H}_0 of G_0 is fully \mathcal{P}_0 -elliptic.*

The results in this section can be easily assembled to give:

THEOREM 3.33. *Suppose that G is hyperbolic and virtually special and that (G, \mathcal{P}) is relatively hyperbolic. Suppose G is given a quasiconvex fully \mathcal{P} -elliptic hierarchy \mathcal{H} **terminating** in \mathcal{P} . Then for some $(G_0, \mathcal{P}_0) \triangleleft (G, \mathcal{P})$, and any $(G', \mathcal{P}') \triangleleft (G, \mathcal{P})$ so that $G' \leq G_0$, the induced hierarchy \mathcal{H}' is a quasiconvex malnormal fully \mathcal{P}' -elliptic hierarchy **terminating** in \mathcal{P}' .*

Proof. By Theorem 3.30, there is a $(G_0, \mathcal{P}_0) \triangleleft (G, \mathcal{P})$ so that the induced hierarchy \mathcal{H}_0 is malnormal. The induced hierarchy \mathcal{H}' on (G', \mathcal{P}') remains quasiconvex and malnormal by Lemmas 3.22 and 3.24. Lemma 3.32 shows that the hierarchy is fully \mathcal{P}' -elliptic. Finally, the ‘moreover’ part of Corollary 3.19 shows that the terminal groups of \mathcal{H}' are a subset of conjugates of \mathcal{P}' . However, since the elements of \mathcal{P}' are always elliptic in all graphs of groups decompositions, all elements of \mathcal{P}' must appear as conjugates of terminal vertex groups of \mathcal{H}' . \square

The results in this section (particularly Theorem 3.33 above) are used in the proof of Theorem 2.11 in Section 5 below.

4. A metric combination theorem

In this section, we prove a combination theorem in the setting of δ -hyperbolic CAT(0) spaces which is analogous to the Baker–Cooper combination theorem in hyperbolic manifolds [BC08]. This combination theorem is Theorem 4.5 below. It is worth noting that this theorem is about CAT(0) spaces in general, not just about cube complexes. Our application of Theorem 4.5 is to show that the hierarchy of spaces that we build in the proof of Theorem 2.11 in Section 5 is faithful and quasiconvex.

DEFINITION 4.1. Let $R \geq 0$. Suppose that $A \subseteq Y$ is connected. Let $\pi_A : Y^A \rightarrow Y$ be the cover corresponding to the image of $\pi_1(A)$ in $\pi_1(Y)$. There is a canonical lift $L : A \rightarrow Y^A$. Let $\tilde{N}_R(A)$ be the R -neighborhood of $L(A)$ in Y^A .

We say that A is R -embedded in Y if π_A is injective on $\tilde{N}_R(A)$, and call the image $\pi_A(\tilde{N}_R(A))$ a *tubular neighborhood* of A .

If A is not connected, but the components of A are R -embedded with disjoint tubular neighborhoods, we say that A is R -embedded in Y .

DEFINITION 4.2 (Elevation). Let W be connected, and let $\phi : W \rightarrow Z$ be some map. If $\pi : \hat{Z} \rightarrow Z$ is a cover, then ϕ may not lift to \hat{Z} but there is some (minimal) cover $\pi_W : \hat{W} \rightarrow W$ so that the composition $\phi \circ \pi_W$ lifts to a map $\hat{\phi} : \hat{W} \rightarrow \hat{Z}$. Such a lift (or its image) is called an *elevation* of W to \hat{Z} . If ϕ is an inclusion map, then the elevations of W are just the connected components of $\pi^{-1}(W)$.

Most of the time, we are interested in the images of elevations, and not the precise maps. In those cases, we say two elevations are *distinct* if they have different images.

DEFINITION 4.3. Suppose that $f : \mathbb{N} \rightarrow \mathbb{N}$ is an affine function $f(n) = Kn + C$. We say that $\mathcal{B} = \{B_i\}$ forms an f -separated family of subcomplexes of Y if for any two distinct elevations U_1 and U_2 of elements of \mathcal{B} to \tilde{Y} (the universal cover) and any $D \geq 0$ we have

$$\text{Diam}(N_D(U_1) \cap N_D(U_2)) \leq f(D).$$

Note that if there are at least two distinct elevations of elements of \mathcal{B} , then any function f as in the definition must have $K \geq 1$.

The key example of an f -separated family is given by the following proposition:

PROPOSITION 4.4. Let (G, \mathcal{H}) be a relatively hyperbolic pair with G hyperbolic. Suppose that Y is a space with $\pi_1 Y = G$ and that for each $H \in \mathcal{H}$

there is a path-connected subspace $B_H \subseteq Y$ with $\pi_1(B_H)$ conjugate to H . Then $\mathcal{B} = \{B_H\}$ is f -separated for some affine function f .

Proof. Fix a Cayley graph Γ for G . This graph is δ -hyperbolic for some $\delta \geq 0$, and the cosets gH for $H \in \mathcal{H}$ are uniformly λ -quasiconvex in Γ for some $\lambda > 0$. We may assume δ and λ are both integers.

We first prove the analogous statement for cosets of elements of \mathcal{H} in G . Namely, we find a number K so that for any H_1, H_2 in \mathcal{H} (not necessarily distinct), any two distinct cosets g_1H_1 and g_2H_2 , and any $t \in \mathbb{N}$,

$$\text{diam}(N_t(g_1H_1) \cap N_t(g_2H_2)) \leq Kt + K. \tag{4.1}$$

There are finitely many pairs H_1, H_2 in \mathcal{H} so it suffices to fix (not necessarily distinct) H_1 and H_2 and find a K which works for that pair. Also note that it suffices to verify the inequality (4.1) for $g_1 = 1$.

To find K , we first observe that there are only finitely many double cosets H_1gH_2 which intersect the ball of radius $R := 2\lambda + 2\delta$ about 1. Let $\mathcal{C} = \{c_1, \dots, c_n\} \subset B_R(1)$ be a set of representatives of these cosets. In case $H_1 = H_2$, we omit the representative for the double coset $H_1 = H_1H_1$ from \mathcal{C} .

CLAIM 4.4.1. For each $c \in \mathcal{C}$, the diameter $\text{diam}(N_R(H_1) \cap N_R(cH_2))$ is finite.

Proof of Claim 4.4.1. If the diameter is infinite, then the limit sets of H_1 and H_2^c would intersect. By [GMRS98, Lemmas 2.6 and 2.7], there would be an infinite order element in $H_1 \cap H_2^c$, contradicting malnormality of \mathcal{H} . \square

Let

$$M = \max_{c \in \mathcal{C}} \text{diam}(N_R(H_1) \cap N_R(cH_2)) + 2\lambda + 1.$$

CLAIM 4.4.2. Unless $H_1 = H_2$ and $g \in H_1$, we have $\text{diam}(N_t(H_1) \cap N_t(gH_2)) \leq 6t + M$.

Proof of Claim 4.4.2. Suppose not. Then there are elements $x, y \in N_t(H_1) \cap N_t(gH_2)$ with $d(x, y) > 6t + M$. There are therefore elements $a_x, a_y \in H_1$ and $b_x, b_y \in H_2$ with

$$\max\{d(a_x, x), d(gb_x, x), d(a_y, y), d(gb_y, y)\} \leq t.$$

Consider a pair of δ -thin triangles with vertex sets $\{a_x, gb_y, gb_x\}$ and $\{a_x, gb_y, a_y\}$. (See Figure 1.) Note that $d(a_x, gb_y) \geq d(x, y) - 2t > M + 4t$. Moreover, the Gromov products $(gb_x, gb_y)_{a_x}$ and $(a_x, a_y)_{gb_y}$ are each at most $2t$. It follows




Figure 1. The cosets H_1 and gH_2 must come close.

that there are subsegments $[p_1, q_1] \subseteq [a_x, a_y]$ and $[p_2, q_2] \subseteq [gb_x, gb_y]$ of length at least M , and whose endpoints satisfy $d(p_1, p_2) \leq 2\delta$ and $d(q_1, q_2) \leq 2\delta$. The λ -quasiconvexity of H_1 and H_2 then implies that there are $a'_x, a'_y \in H_1$ and $b'_x, b'_y \in H_2$ so that

$$\min\{d(a'_x, a'_y), d(gb'_x, gb'_y)\} \geq M - 2\lambda, \tag{4.2}$$

but

$$\max\{d(a'_x, gb'_x), d(a'_y, gb'_y)\} \leq 2\lambda + 2\delta = R. \tag{4.3}$$

This gives a contradiction: the inequality (4.3) implies that H_1gH_2 intersects $B_R(1)$ nontrivially, so we must have $\text{diam}(N_R(H_1) \cap N_R(gH_2)) \leq M - 2\lambda - 1$. On the other hand, (4.2) implies that $\text{diam}(N_R(H_1) \cap N_R(gH_2)) \geq M - 2\lambda$. \square

We can therefore take $K(H_1, H_2) = \max\{6, M\}$, and let $K = \max\{K(H, H') \mid H, H' \in \mathcal{H}\}$.

We now derive the Proposition from the group-theoretic statement. Let $\tilde{Y} \rightarrow Y$ be the universal cover, and for each i let E_i be the elevation of B_i to \tilde{Y} which is preserved by H_i . For some $\mu \geq 1, \epsilon \geq 0$, we can find equivariant maps $\phi : \tilde{Y} \rightarrow \Gamma$ and $\psi : \Gamma \rightarrow \tilde{Y}$, so that:

- (1) ϕ and ψ are (μ, ϵ) -quasiisometries;
- (2) ϕ and ψ are ϵ -quasiinverses; and
- (3) for each $H_i \in \mathcal{H}$ and each $g \in G$, the Hausdorff distances $d_{\text{Haus}}(\phi(gB_i), gH_i)$ and $d_{\text{Haus}}(\psi(gH_i), gB_i)$ are bounded above by ϵ .

Now let $t > 0$, and choose elevations U_1 and U_2 of elements B_{i_1} and $B_{i_2} \in \mathcal{B}$. We have $U_1 = g_1E_{i_1}$ and $U_2 = g_2E_{i_2}$. We have $\phi(N_t(U_j)) \subseteq N_{\mu t + 2\epsilon}(g_jH_{i_j})$ for $j = 1, 2$. Thus, $\phi(N_t(U_1) \cap N_t(U_2)) \subseteq N_{\mu t + 2\epsilon}(g_1H_1) \cap N_{\mu t + 2\epsilon}(g_2H_2)$ has diameter at most $K(\mu t + 2\epsilon) + K$. Thus $\psi \circ \phi(N_t(U_1) \cap N_t(U_2))$ has diameter

at most $\mu(K(\mu t + 2\epsilon) + K) + \epsilon$. Since ϕ and ψ are ϵ -quasiinverses, we get the affine bound

$$\text{diam}(N_r(U_1) \cap N_r(U_2)) \leq \mu(K(\mu t + 2\epsilon) + K) + 2\epsilon,$$

as required. □

THEOREM 4.5 (Combination theorem). *Let $\delta > 0$ and let $f : \mathbb{N} \rightarrow \mathbb{N}$ be a nondecreasing affine function. Suppose that Y is a compact NPC space with δ -hyperbolic universal cover \tilde{Y} . For ϵ_B large enough (in terms of δ), and ϵ_A large enough (in terms of δ, f , and ϵ_B), the following combination theorem holds.*

Suppose:

- (1) $\mathcal{A} = \bigsqcup A_i$ and $\mathcal{B} = \bigsqcup B_j$ are embedded locally convex subsets (where the A_i are the connected components of \mathcal{A} and the B_j are the connected components of \mathcal{B});
- (2) \mathcal{B} forms an f -separated family;
- (3) \mathcal{A} is ϵ_A -embedded in Y ;
- (4) Each $B_j \subset \mathcal{B}$ is ϵ_B -embedded in Y ; and
- (5) Θ is a connected component of $\mathcal{A} \cup \mathcal{B}$.

Then Θ is π_1 -injective with λ -quasiconvex elevation to \tilde{Y} , where λ is a function of δ only.

In applications, the components of \mathcal{B} are thought of as ‘peripheral’ subcomplexes. With slightly stronger assumptions on how \mathcal{A} and \mathcal{B} interact, we are able to rule out ‘accidental parabolics.’

DEFINITION 4.6. Let \mathcal{B} be an embedded locally convex subset of the nonpositively curved space Y , and let $\mathcal{E} \subset Y$. We say that \mathcal{E} has an *accidental \mathcal{B} -loop* if there is a homotopically essential loop $\xi \subset \mathcal{E}$ so that:

- (1) ξ is (freely) homotopic to a geodesic loop in \mathcal{B} ; and
- (2) no positive power of ξ is homotopic in \mathcal{E} to a geodesic loop in \mathcal{B} .

This notion is stable under finite covers:

LEMMA 4.7. *Let $\mathcal{B}, \mathcal{E}, Y$ be as in Definition 4.6, and let $\pi : \tilde{Y} \rightarrow Y$ be a finite cover. If $\tilde{\mathcal{B}} = \pi^{-1}(\mathcal{B})$ and \mathcal{E}' is an elevation of \mathcal{E} to \tilde{Y} then \mathcal{E} has an accidental \mathcal{B} -loop if and only if \mathcal{E}' has an accidental $\tilde{\mathcal{B}}$ -loop.*

Proof. Let γ be an accidental \mathcal{B} -loop in \mathcal{E} , and let γ' be an elevation to \mathcal{E}' . By condition (1) of Definition 4.6, γ is freely homotopic to a geodesic loop in \mathcal{B} . We can lift the homotopy to get that γ' is freely homotopic to a geodesic loop in $\tilde{\mathcal{B}}$. If γ' were freely homotopic to a loop in $\tilde{\mathcal{B}} \cap \mathcal{E}'$, then we could project the homotopy to get a homotopy of γ'' into $\mathcal{B} \cap \mathcal{E}$ for some $n \geq 1$, contradicting condition (2) of Definition 4.6.

Conversely, if γ is an accidental $\tilde{\mathcal{B}}$ -loop in \mathcal{E}' , we can show that $\bar{\gamma} = \pi(\gamma)$ is an accidental \mathcal{B} -loop in \mathcal{E} . Projecting the homotopy of γ into $\tilde{\mathcal{B}}$ gives a homotopy of $\bar{\gamma}$ into \mathcal{B} . Moreover, if some positive power of $\bar{\gamma}$ were homotopic in \mathcal{E} to a loop in \mathcal{B} we could lift that homotopy to show some positive power of γ was homotopic in \mathcal{E}' to a loop in $\tilde{\mathcal{B}}$. □

We'll also prove:

PROPOSITION 4.8. *Let Y , \mathcal{A} , and \mathcal{B} satisfy the assumptions of the Combination Theorem 4.5. Suppose moreover that \mathcal{A} has no accidental \mathcal{B} -loops. If Θ is a component of $\mathcal{A} \cup \mathcal{B}$, then Θ has no accidental \mathcal{B} -loops.*

The remainder of this section is devoted to proving Theorem 4.5 and Proposition 4.8.

4.1. Broken geodesics. The following Lemma can be proven using essentially the same argument as [BH99, III.H.1.13].

LEMMA 4.9. *Let $l \geq 0$. Let $c = c_1 \cdots c_n$ be a path from p to q in a δ -hyperbolic geodesic space, satisfying:*

- (1) *for each i , the subpath $c_i = [p_i, p_{i+1}]$ is geodesic;*
- (2) *for each $i \neq 1$, the Gromov product $(p_{i-1}, p_{i+1})_{p_i} \leq l$;*
- (3) *for each $i \notin \{1, n\}$, the length of c_i is strictly bigger than $2l + 8\delta$.*

Then if γ is any geodesic from p to q , the Hausdorff distance between γ and c is at most $l + 5\delta$.

4.2. Proof of Theorem 4.5. We will show that the following assumptions on ϵ_A, ϵ_B are sufficient.

$$\epsilon_B > 50\delta \tag{4.4}$$

$$\epsilon_A > 2 \max\{f(\epsilon_B + 2\delta) + \delta + \epsilon_B, 2\epsilon_B\} \tag{4.5}$$

Recall that given $S \subseteq Y$ connected, and $\pi : Y' \rightarrow Y$ any cover, an *elevation* of S to Y' is a component of $\pi^{-1}(S)$. We are particularly interested in elevations to the universal cover \tilde{Y} of components of \mathcal{A} and \mathcal{B} .

DEFINITION 4.10. Any elevation of a connected component of \mathcal{A} to \tilde{Y} is called an *\mathcal{A} -elevation*. Any geodesic in an \mathcal{A} -elevation is called an *\mathcal{A} -arc*. We define *\mathcal{B} -elevations* and *\mathcal{B} -arcs* similarly.

4.2.1. π_1 -injectivity. In this subsection, we show that under the assumptions (4.4) and (4.5), the space Θ described in the statement of Theorem 4.5 is π_1 -injective. Notice that Θ is itself an NPC space, since it is obtained from NPC spaces by isometrically gluing together convex subspaces (see [BH99, Proposition II.11.6]). However, it is worth remarking that the inclusion $\Theta \hookrightarrow Y$ is (typically) not a local isometry, and in fact Theorem 4.5 is obvious in the case that it is.

We argue by contradiction, so suppose that $k \neq 1$ is in $K := \ker(\pi_1 \Theta \rightarrow \pi_1 Y)$. Let $\tilde{\Theta}$ be the universal cover of Θ . By possibly moving basepoints, we can suppose k is represented by a loop γ whose elevation to $\tilde{\Theta}$ is a biinfinite geodesic. Since k is in K , the loop γ lifts to the universal cover \tilde{Y} of Y . The loop is a concatenation of alternating \mathcal{A} -arcs and \mathcal{B} -arcs, for example

$$\gamma = a_1 b_2 \cdots a_{n-1} b_n, \quad (4.6)$$

where each a_i is an \mathcal{A} -arc, and each b_i is a \mathcal{B} -arc. We regard the expression (4.6) as a ‘cyclic word’ whose syllables are arcs. To simplify notation below, the indices $1, \dots, n$ should be taken to be elements of $\mathbb{Z}/n\mathbb{Z}$.

Since \mathcal{A} -elevations and \mathcal{B} -elevations intersect, there may be some choice in the expression (4.6); we always assume the expression is chosen so that the *syllable length* n is minimized. Note that γ cannot consist of a single \mathcal{A} or \mathcal{B} -arc, since otherwise some \mathcal{A} or \mathcal{B} -elevation contains a nontrivial geodesic loop. Moreover, by moving the basepoint, we can assume that the first arc is an \mathcal{A} -arc, and the last is a \mathcal{B} -arc, as in (4.6). In other words, the syllable length is even.

LEMMA 4.11. *The syllable length is not 2.*

Proof. We have already noted that the syllable length is not 1. If the syllable length were 2, we would have a geodesic bigon in \tilde{Y} , with one side in some \mathcal{A} -elevation and the other side in some \mathcal{B} -elevation. This contradicts convexity of the \mathcal{A} and \mathcal{B} -elevations. \square

Since the syllable length is an even number bigger than 2, there are at least two \mathcal{B} -arcs (and the same number of \mathcal{A} -arcs) in γ . We describe a way to ‘shortcut’




Figure 2. Shortcutting around the \mathcal{A} -arcs.

the \mathcal{A} -arcs and obtain an ϵ_B -local $(1, 5\delta)$ -quasigeodesic loop from γ . This will contradict Lemma 4.9.

Each \mathcal{A} -arc a_j lies in some \mathcal{A} -elevation E_j . Let a_j and a_{j+2} be two consecutive \mathcal{A} -syllables. The intervening \mathcal{B} -syllable b_{j+1} must leave E_j , or otherwise the expression (4.6) could be shortened. Since E_j is convex it follows that $E_{j+2} \neq E_j$. The ϵ_A -embeddedness of \mathcal{A} implies that $d(E_j, E_{j+2}) \geq 2\epsilon_A$. (A similar argument shows that successive \mathcal{B} -syllables must lie in separate \mathcal{B} -elevations, though there is no useful lower bound on the distance between these elevations.)

Since the elevations E_j and E_{j+2} are of distance at least $2\epsilon_A$ from one another, there must be a point p_{j+1} on b_{j+1} satisfying $d(p_{j+1}, E_j) = \epsilon_A/2$ and another point q_{j+1} on b_{j+1} satisfying $d(q_{j+1}, E_{j+2}) = \epsilon_A/2$. (See Figure 2.) For each odd j we choose a geodesic arc α_j in \tilde{Y} from q_{j-1} to p_{j+1} . For even j , we let β_j be the subarc of b_j between p_j and q_j . Concatenating all these, we obtain a piecewise geodesic loop

$$\sigma = \alpha_1\beta_2 \cdots \alpha_{n-1}\beta_n.$$

Observe $|\beta_j| \geq \epsilon_A$ for each even j .

LEMMA 4.12. *Let j be odd. Then $|\alpha_j| > 2\epsilon_B$.*




Figure 3. α -arcs are long.

Proof. As argued above, the arc α_j goes between points $q := q_{j-1}$ and $p := p_{j+1}$ in distinct \mathcal{B} -elevations E_{j-1} and E_{j+1} . Let m be the midpoint of α_j . By way of contradiction, suppose that $d(m, p) = d(m, q) \leq \epsilon_B$.

Consider a geodesic triangle with base a_j and apex m (see Figure 3). Let $x \in E_j \cap E_{j-1}$ and $y \in E_j \cap E_{j+1}$ be the endpoints of a_j , and let $t = (x, y)_m$ be the Gromov product of the endpoints of a_j with respect to m . Let η be the initial subsegment of $[m, x]$ of length t . The geodesic segment η lies in an $\epsilon_B + \delta$ -neighborhood of E_{j-1} and in an $\epsilon_B + 2\delta$ -neighborhood of E_{j+1} . It follows that the length of η is at most $f(\epsilon_B + 2\delta)$. Hence $d(m, E_j) \leq f(\epsilon_B + 2\delta) + \delta$, and

$$\epsilon_A/2 = d(p, E_j) \leq f(\epsilon_B + 2\delta) + \delta + \epsilon_B,$$

a contradiction to (4.5). □

So σ is a piecewise geodesic, and each geodesic piece has length at least $2\epsilon_B$. We must now bound the Gromov products at the corners.

LEMMA 4.13. *Let j be even. The Gromov product of q_{j-2} and q_j at p_j is at most δ . The Gromov product of p_j and p_{j+2} at q_j is at most δ .*

Proof. We prove the first assertion. The proof of the second assertion is identical up to shuffling labels.




Figure 4. Showing Gromov products are small.

We use similar notation to that in Lemma 4.12: $p = p_j, q = q_{j-2}$, and $q' = q_j$. Let x be the endpoint of a_{j-1} meeting b_{j-2} , and let y be the endpoint meeting b_j (see Figure 4).

Because the points $\{y, p, q'\}$ lie on a geodesic, at most one of the quantities $(y, q)_p, (q', q)_p$ is greater than δ . Indeed, consider the two δ -thin triangles with vertex sets $\{p, q, y\}$ and $\{p, q, q'\}$. Letting $M = \min\{(y, q)_p, (q', q)_p\}$, we see that there are points $p_1 \in [p, y]$ and $p_2 \in [p, q']$ with $d(p_i, p) = M$ for each i , but so that $d(p_1, p_2) \leq 2\delta$. Since $\{y, p, q'\}$ lie on a geodesic, we have $d(p_1, p_2) = 2M$, so $M \leq \delta$.

We will show that in fact $(y, q)_p > \delta$, from which we deduce $(q, q')_p \leq \delta$ as required. Consider the pair of triangles with vertex sets $\{p, q, y\}$ and $\{x, y, q\}$. There are two cases, depending on the relative sizes of the Gromov products at y . Let $D = (y, q)_p$ be the quantity we are trying to bound from below.

Case 1: $(p, q)_y \leq (x, q)_y$

In this case, (see the left-hand side of Figure 5), there is a path from p to a point on $a_{j-1} \subseteq E_{j-1}$ of length $D + 2\delta$. Since $d(p, E_{j-1}) = \epsilon_A/2$, this implies $D \geq \epsilon_A/2 - 2\delta > \delta$, as required.

Case 2: $(p, q)_y > (x, q)_y$

(See the right-hand side of Figure 5.) Let $\Delta = (p, q)_y - (x, q)_y$ be the difference in the Gromov products at y , and note that there is a path from p to a point on $a_j \subseteq E_{j-1}$ of length $D + 2\delta + \Delta$. Moreover, there is a length Δ subsegment of $[y, q]$ which is of distance $\leq \delta$ from both $[y, p] \subseteq E_j$




Figure 5. Two possibilities for the quadrilateral with vertices (y, p, q, x) . Case 1 is on the left, Case 2 on the right. The sides of the internal dashed triangles have length at most δ .

and $[x, q] \subseteq E_{j-2}$. But this implies that $\Delta \leq f(\delta)$. Thus

$$\epsilon_A/2 \leq D + 2\delta + \Delta \leq D + 2\delta + f(\delta),$$

which implies $D \geq \epsilon_A/2 - (2\delta + f(\delta)) > \delta$ as required. □

We now have that σ is a broken geodesic with Gromov products in the corners of at most δ , and made of pieces of length at least $2\epsilon_B$. Since σ is a loop, applying Lemma 4.9 with $l = \delta$ we see that σ must lie a Hausdorff distance at most 6δ from any point of σ . Since σ contains a geodesic subsegment of length at least $2\epsilon_B > 6\delta$, this is a contradiction. □

4.2.2. *Quasiconvexity.* The argument for quasiconvexity involves similar ideas. Let p and q be points in the elevation $\tilde{\Theta}$ to \tilde{Y} . These are connected by a (unique) $\tilde{\Theta}$ -geodesic γ which can be written as a concatenation of \mathcal{A} -arcs and \mathcal{B} -arcs

$$\gamma = s_1 \cdots s_n \tag{4.7}$$

where n is minimal subject to the constraint that each s_j is a geodesic arc either in a single \mathcal{A} -elevation or a single \mathcal{B} -elevation. We call n the *syllable length* of γ .

Consider the first interior \mathcal{B} syllable of γ , which will be s_2 or s_3 , and the last interior \mathcal{B} syllable which is s_{n-2} or s_{n-1} . Then we may write $\gamma = \alpha\theta\omega$, where α and ω consist of at most 2 syllables, and θ has at most $n - 2$ syllables which begin and end with \mathcal{B} syllables. Let p' and q' be the initial and terminal points of θ . As in the proof of π_1 -injectivity, we modify θ to a path σ avoiding the \mathcal{A} -syllables. This path will not be in $\tilde{\Theta}$, but will be a quasigeodesic which is a controlled distance away from $\tilde{\Theta}$, and we then apply quasigeodesic stability to find that $[p', q']$ is not far from σ .




Figure 6. Modifying a subpath of γ to a quasigeodesic σ between p' and q' . Note that the blue segments at the beginning and end of γ may be absent, possibly changing the numbering.

Let $[p', q']$ be the geodesic joining p' and q' in \tilde{Y} . If we can show that $[p', q']$ is uniformly bounded distance from $\tilde{\Theta}$, then it follows that the geodesic $[p, q] \subset \tilde{Y}$ will be a bounded distance from $\tilde{\Theta}$, since the polygonal arc $\alpha[p', q']\omega$ is made of at most 5 segments, so that $[p, q] \subset \mathcal{N}_{3\delta}(\alpha[p', q']\omega)$, and therefore a bounded distance from $\tilde{\Theta}$.

Let $J = \{j_0, j_0 + 2, \dots, j_{\max} - 2, j_{\max}\}$ be the set of indices of \mathcal{B} syllables s_j of θ . Thus $j_0 \in \{2, 3\}$, and $j_{\max} \in \{n - 2, n - 1\}$. For each $j \in J$, the \mathcal{B} -arc s_j travels between distinct \mathcal{A} -elevations E_{j-1} and E_{j+1} , which are of distance at least $2\epsilon_A$ away from one another. There are therefore points p_j and q_j on s_j so that $d(p_j, E_{j-1}) = \epsilon_A/2$ and $d(q_j, E_{j+1}) = \epsilon_A/2$, and the subsegment $[p_j, q_j]$ has length at least ϵ_A .

We now let σ be a broken geodesic with ordered vertices $(p', q_{j_0}, \dots, p_{j_{\max}}, q')$. (See Figure 6.)

LEMMA 4.14. *The path σ is contained in a 2δ -neighborhood of γ (and hence of $\tilde{\Theta}$).*

Proof. See Figure 6. The path σ is built from two kinds of subarcs. The first are paths already in γ . The second, pictured as dotted green arcs in the Figure, are geodesics. These geodesics naturally lie in geodesic triangles or quadrilaterals, the remaining edges of which are in γ . Using δ -hyperbolicity, the result is immediate. \square

By the same argument used to prove Lemma 4.13, the path σ has Gromov products in the corners bounded by δ .

We may apply Lemma 4.9 with $l = \delta$ to see that σ has Hausdorff distance at most 6δ from $[p', q']$. Lemma 4.14 then implies that $[p', q']$ lies in an 8δ -neighborhood of $\tilde{\Theta}$. Then $[p, q]$ lies in an 11δ neighborhood of $\tilde{\Theta}$ as observed above, since it lies in a 3δ neighborhood of $\alpha[p', q']\omega$, which is a

polygonal arc made of at most 5 geodesic segments, all of which but $[p', q']$ lie in $\tilde{\Theta}$.

We conclude that $\tilde{\Theta}$ is 11δ -quasiconvex. □

4.3. Proof of Proposition 4.8. As in the proof of Theorem 4.5, we assume

$$\epsilon_B > 50\delta, \quad \text{and} \\ \epsilon_A > 2 \max\{f(\epsilon_B + 2\delta) + \delta + \epsilon_B, 2\epsilon_B\}.$$

We argue by contradiction, supposing that there is an essential Θ -loop γ which is freely homotopic in Y to an essential loop η in some component B of \mathcal{B} , but that no positive power of γ is homotopic in Θ to a loop in \mathcal{B} . Adjusting the loops by a homotopy, we may assume that η is geodesic and that γ is Θ -geodesic, and can be written as a sequence of \mathcal{A} -syllables a_i and \mathcal{B} -syllables b_i ; for example

$$\gamma = a_1 b_2 \cdots a_{n-1} b_n,$$

where n is the syllable length of γ , and the expression is chosen to minimize this length in the Θ -homotopy class of γ .

LEMMA 4.15. *The syllable length of γ is not 1.*

Proof. If the syllable length of γ is 1, then there are two possibilities, $\gamma \subseteq \mathcal{A}$ or $\gamma \subseteq \mathcal{B}$. The first case is ruled out by the hypothesis that \mathcal{A} has no accidental \mathcal{B} -loops, the second by assumption that γ is not homotopic in Θ to a loop in \mathcal{B} . □

We suppose therefore that the syllable length of γ is greater than 1, so there is at least one \mathcal{B} -syllable, and at least one \mathcal{A} -syllable. We can lift the homotopy between γ and η to the universal cover \tilde{Y} , obtaining elevations $\tilde{\gamma}$ and $\tilde{\eta}$ which fellow-travel one another (see Figure 7).

LEMMA 4.16. *Successive \mathcal{B} -arcs of $\tilde{\gamma}$ are contained in distinct \mathcal{B} -elevations.*

Proof. Suppose that two successive \mathcal{B} -arcs are contained in the same \mathcal{B} -elevation E . Since the elevation is convex, the intervening \mathcal{A} -arc is also contained in E . It follows that γ is homotopic in Θ to a loop with one fewer \mathcal{A} -syllable, contradicting the minimality of the syllable length. □

We can perform the same ‘shortcutting’ operation as in the proof of Theorem 4.5, modifying $\tilde{\gamma}$ to a piecewise geodesic σ which contains long ($\geq \epsilon_A$) subsegments of the \mathcal{B} -arcs of γ , and which has Gromov products bounded by δ . Using Lemma 4.9 and δ -hyperbolicity we can see that the broken geodesic σ lies




Figure 7. Fellow-traveling lifts of η and γ . \mathcal{A} -arcs are red; \mathcal{B} -arcs are blue.

within Hausdorff distance 8δ of $\tilde{\eta}$. The broken geodesic σ contains a subsegment of each \mathcal{B} -arc of $\tilde{\gamma}$ of length at least ϵ_A . The geodesic $\tilde{\eta}$ is contained in a single \mathcal{B} -elevation. By Lemma 4.16, not all the \mathcal{B} -arcs of $\tilde{\gamma}$ are contained in the same \mathcal{B} -elevation as $\tilde{\eta}$. We deduce $\epsilon_A \leq f(8\delta)$, which contradicts the hypothesis on ϵ_A . This completes the proof of Proposition 4.8. \square

5. Fully \mathcal{P} -elliptic hierarchies

In this section, we prove Theorem 2.11. First recall the statement.

THEOREM 2.11. *Given (G, \mathcal{P}) relatively hyperbolic, with G hyperbolic and virtually special, there is some $(G', \mathcal{P}') \triangleleft (G, \mathcal{P})$ which has a malnormal quasiconvex fully \mathcal{P}' -elliptic hierarchy terminating in \mathcal{P}' .*

We describe a fairly general procedure for generating relative hierarchies of spaces from cube complex pairs. Then we talk about how to ensure such a hierarchy is faithful, quasiconvex, and malnormal, by passing to carefully chosen finite covers.

5.1. Augmented cube complexes. We first establish some terminology useful for dealing with cube complex pairs.

Let X be an NPC cube complex, and $\mathcal{Z} = \bigsqcup_{i=1}^n Z_i$ a union of connected NPC cube complexes which admits a locally isometric immersion $\Phi = \bigsqcup_{i=1}^n \phi_i : \mathcal{Z} \rightarrow X$. We will call (X, \mathcal{Z}) a *cube complex pair*. Then the mapping cylinder of Φ ,

$$C_\Phi = X \sqcup (\mathcal{Z} \times [0, 1]) / \{(z, 1) \sim \Phi(z)\} \tag{5.1}$$

naturally has the structure of an NPC cube complex. Call such a complex the *augmented cube complex* based on the pair (X, \mathcal{Z}) .

There are canonical inclusions $X \hookrightarrow C_\Phi$ and $\mathcal{Z} \hookrightarrow C_\Phi$ (this second via $z \mapsto (z, 0)$). The subset $X \subseteq C_\Phi$ is a deformation retract of C_Φ . The components Z_i of $\mathcal{Z} \subseteq C_\Phi$ are referred to as *peripheral subcomplexes*. The hyperplanes $Z_i \times \{\frac{1}{2}\}$

are called *peripheral hyperplanes*, and all other hyperplanes of C_ϕ are called *nonperipheral hyperplanes*.

Any cover $\tilde{X} \rightarrow X$ also gives rise to a cover $C_{\tilde{\phi}} \rightarrow C_\phi$, where $\tilde{\Phi} : \tilde{\mathcal{Z}} \rightarrow \tilde{X}$ is made up of all the elevations of the maps ϕ_i .

DEFINITION 5.1 (Augmented hyperplanes). Let $\Phi : \mathcal{Z} \rightarrow X$ be a locally isometric immersion of NPC cube complexes, and let $C = C_\phi$ be the augmented complex. Let W be a nonperipheral hyperplane of C . The *augmented hyperplane* $A(W)$ is the component of $W \cup \mathcal{Z}$ containing W .

In other words, the augmented hyperplane $A(W)$ is the union of W with any components of \mathcal{Z} which meet W . In general, augmented hyperplanes are not particularly well behaved. They are not convex subsets of the augmented complex, are not locally separating, and may not be π_1 -injective. We will nonetheless be able to find situations in which they give a faithful hierarchy.

5.2. Double-dot hierarchy. In this subsection, we start with an NPC cube complex X , together with some collection of subcomplexes \mathcal{Z} , and construct a particular hierarchy on a space homotopic to a ‘generalized double cover’ of X . The stages of the hierarchy (though not its terminal spaces) depend on an ordering of the hyperplanes of X . In the rest of the paper, the hierarchy is referred to as the *double-dot hierarchy for the pair* (X, \mathcal{Z}) .

DEFINITION 5.2 (cf. [Wis, Construction 9.1]). Let X be any cube complex. Any hyperplane $W \subset X$ gives rise to a map from $i_W : \pi_1 X \rightarrow \mathbb{Z}/2$, measuring the mod-2 intersection between a loop and the hyperplane. Let \mathcal{W} be the set of embedded, 2-sided, nonseparating hyperplanes in X . The *double-dot cover* $\tilde{X} \rightarrow X$ is the cover corresponding to the kernel of the map

$$h_{\mathcal{W}} = \bigoplus_{W \in \mathcal{W}} i_W : \pi_1 X \rightarrow \bigoplus_{W \in \mathcal{W}} \mathbb{Z}/2.$$

Given any augmented complex C_ϕ , we now describe a ‘relative’ hierarchy of spaces structure on the double-dot cover using the augmented hyperplanes of C_ϕ . An example of an augmented cube complex is shown in Figure 8, together with its double-dot cover. Parts of the double-dot hierarchy for this complex are shown in Figures 9 and 10.

Let $C = C_\phi$ be an augmented cube complex, where $\Phi : \mathcal{Z} \rightarrow X$ is a locally isometric immersion of cube complexes. Let \mathcal{W} be the set of embedded, 2-sided, nonseparating hyperplanes of C . Choose an order (W_1, \dots, W_n) for these hyperplanes. For $i \in \{1, \dots, n\}$, let A_i be the augmented hyperplane $A(W_i)$.




Figure 8. An example of an augmented complex (at right) and its double-dot cover. X is a figure-eight, and \mathcal{Z} consists of two circles. There are two nonseparating hyperplanes, red and orange.

Let $c : \ddot{C} \rightarrow C$ be the double-dot cover, and let $\ddot{Z} = c^{-1}(\mathcal{Z})$. Fix a basepoint $\tilde{p} \in \ddot{C}$ so that $p = c(\tilde{p})$ lies in the complement of $\bigcup \mathcal{W}$. The complementary components of $c^{-1}(\bigcup \mathcal{W})$ are labeled by elements of $\bigoplus_{\mathcal{W}}(\mathbb{Z}/2)$. For $\mathbf{t} \in \bigoplus_{\mathcal{W}}(\mathbb{Z}/2)$, let $V_{\mathbf{t}}$ be the closure of the union of components labeled by \mathbf{t} . Note that $V_{\mathbf{t}}$ need not be connected. The terminal vertex spaces of our hierarchy are components of $V_{\mathbf{t}} \cup \ddot{Z}$, where \mathbf{t} ranges over $\bigoplus_{\mathcal{W}}(\mathbb{Z}/2)$.

The order on \mathcal{W} determines a hierarchy of spaces \mathcal{H} as follows: For each i , let $\mathcal{W}_i = \{W_1, \dots, W_i\}$, and let $M_i = \bigoplus_{\mathcal{W}_i} \mathbb{Z}/2$. The complementary components of $\bigcup \mathcal{W}_i$ are labeled by elements of M_i . For each $\mathbf{t} \in M_i$, let $K_{\mathbf{t}}$ be the closure of the part labeled \mathbf{t} . We define the \mathbf{t} -vertex spaces to be those components of $K_{\mathbf{t}} \cup \ddot{Z}$ which intersect $K_{\mathbf{t}}$. The vertex spaces at level i are the \mathbf{t} -vertex spaces, for \mathbf{t} ranging over M_i .

The edge spaces at level i are (some of the) components of pairwise intersections of vertex spaces. We now describe these intersections.

DEFINITION 5.3. Let A be the closure of a component of $c^{-1}(W_i) \setminus \bigcup_{j < i} c^{-1}(W_j)$ in \ddot{C} . We call A a *partly cut up elevation of W_i* .

The following lemma is immediate from the construction.

LEMMA 5.4. Any two vertex spaces at level i are either disjoint or intersect in a union of peripheral complexes and disjoint, partly cut up elevations of W_i .

We now describe the graph of spaces (Γ, \mathcal{X}) at level i , associated to a vertex space V at level $i - 1$. The space V is a \mathbf{t} -vertex space, for some $\mathbf{t} \in M_{i-1}$. Under

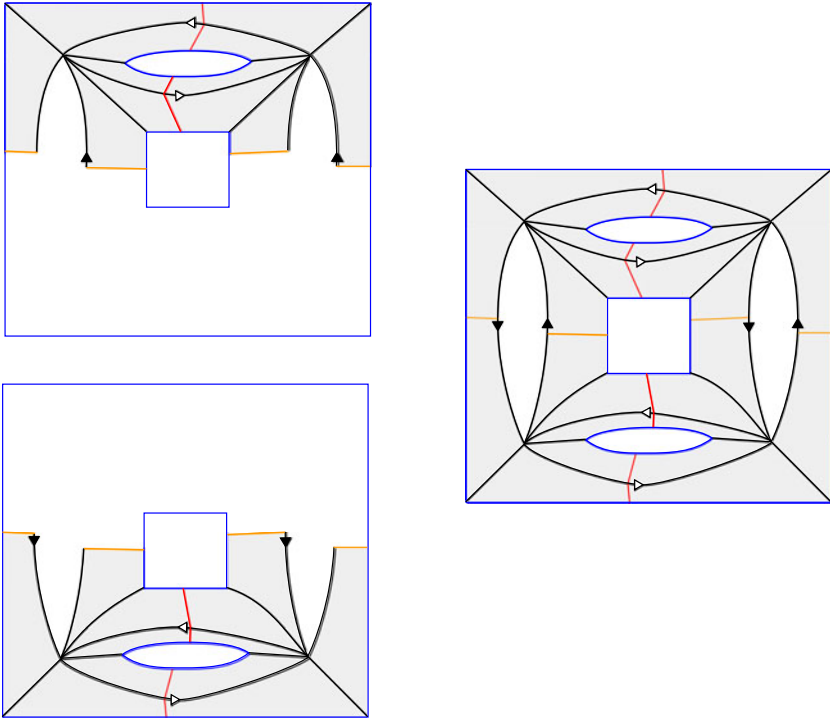


Figure 9. Level 1 of a double-dot hierarchy, cutting along the elevated orange augmented hyperplane. There are two edge spaces, each consisting of two elevations of the orange hyperplane together with an elevation of a blue peripheral circle.

the canonical projection $M_i \rightarrow M_{i-1}$, the fiber above \mathbf{t} consists of two elements $\{\mathbf{t}^+, \mathbf{t}^-\}$, so V is a union of \mathbf{t}^+ -vertex spaces $\mathcal{V}^+ = \{V_1^+, \dots, V_p^+\}$ and \mathbf{t}^- -vertex spaces $\mathcal{V}^- = \{V_1^-, \dots, V_m^-\}$. These form the vertex spaces of the graph of spaces associated to V . No two \mathbf{t}^+ -vertex spaces intersect nontrivially, and neither do any \mathbf{t}^- -vertex spaces. Thus, the incidence graph of the set of vertex spaces is bipartite. We modify the incidence graph slightly to a graph Γ by repeating edges representing multiple components of intersection. Thus the unoriented edges of Γ are in bijective correspondence with the components of

$$\left(\bigcup \mathcal{V}^+\right) \cap \left(\bigcup \mathcal{V}^-\right).$$

The edge-space-to-vertex-space maps are the inclusion maps, and the homotopy equivalence $R(\Gamma, \mathcal{X}) \rightarrow V$ is given in the obvious way, by inclusion on the vertex spaces, and projection of the mapping cylinders to their images.

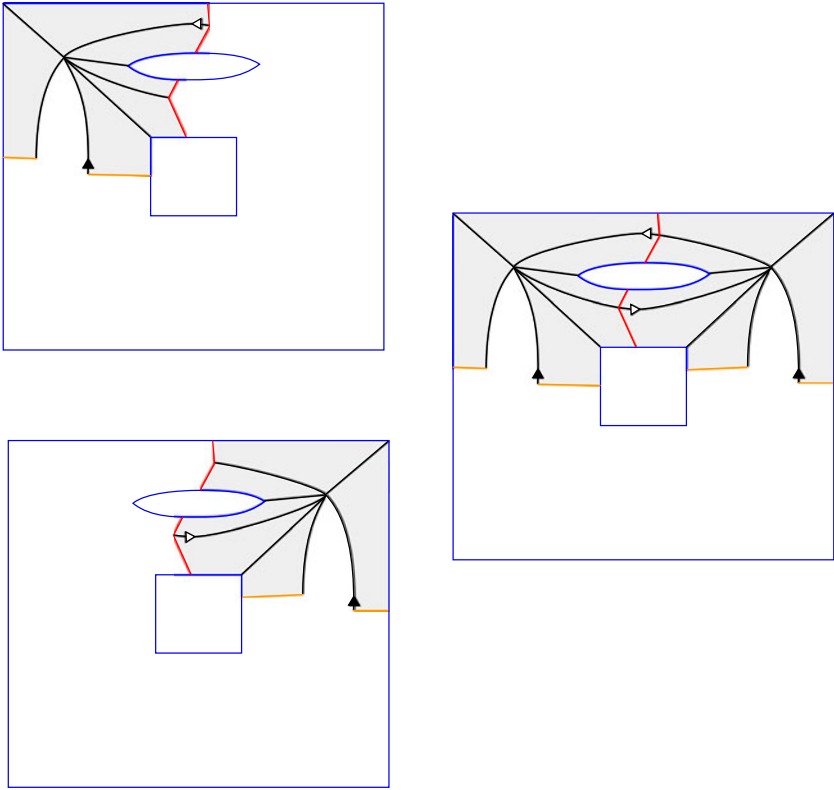


Figure 10. Level 2 of a double-dot hierarchy, cutting one of the vertex spaces along the elevated red augmented hyperplane. There is a single edge space.

5.3. Terminal spaces. In general, the hierarchy described in Section 5.2 above may not be faithful, and even if it is, may not be quasiconvex, malnormal, or have reasonable terminal spaces. However, if every hyperplane of X is embedded, nonseparating and two-sided, the terminal spaces are particularly nice:

LEMMA 5.5. *Suppose that the elements of \mathcal{Z} are embedded and locally isometrically embedded in X . Let C be the augmented cube complex coming from the pair (X, \mathcal{Z}) , and let \ddot{C} be the double-dot cover. Let \ddot{Z} be the union of the peripheral subcomplexes of \ddot{C} . Suppose that every hyperplane of X is embedded, 2-sided, and nonseparating.*

Let Y be a terminal space of the double-dot hierarchy. Then Y has a graph of spaces structure (Γ, \mathcal{Y}) , where:

- (1) Γ is bipartite, with red and black vertices;
- (2) if v is a black vertex, $\mathcal{Y}(v)$ is contractible;
- (3) if v is a red vertex, $\mathcal{Y}(v)$ is a component of \check{Z} ; and
- (4) every edge space is contractible.

Proof. The hypotheses imply that the set of hyperplanes \mathcal{W} from Definition 5.2 is the entire collection of nonperipheral hyperplanes of C .

After cutting along each of the W_i in turn (for varying i), what remains are: parts of X cut along all nonperipheral hyperplanes, joined (via parts of the cylinders in the trivial mapping cylinder of the augmentation) to components of \check{Z} . The graph of spaces decomposition of Y is obtained by cutting along the (remaining parts of the) peripheral hyperplanes. Since we have now cut along *all* the hyperplanes of \check{C} , the remaining parts which do not intersect the peripheral subcomplexes are cubical polyhedra (in the language of [Ago13]) and are contractible. The other parts are homotopy equivalent to components of \check{Z} . Finally, the edge spaces are cubical polyhedra in the peripheral hyperplanes, and hence are also contractible. \square

COROLLARY 5.6. *With the same assumptions as Lemma 5.5, the fundamental group of a terminal space is a free product $(\ast_{i=1}^p G_i) \ast F$ where F is finitely generated free, and each G_i is the fundamental group of some component of \check{Z} .*

5.4. Finding a faithful hierarchy. In order to turn the topological statements of the last subsection into group-theoretic statements, we must find conditions under which the double-dot hierarchy is faithful. This uses the combination theorem from Section 4.

5.4.1. Lemmas about R -embedded subsets. The following lemma is straightforward.

LEMMA 5.7. *Suppose that A is R -embedded in B . Let $\pi : \tilde{B} \rightarrow B$ be a finite cover, and let $\tilde{A} = \pi^{-1}(A)$. Then \tilde{A} is R -embedded in \tilde{B} .*

We will need to know that certain intersections of R -embedded sets are R -embedded (Lemma 5.9). First we need to understand the fundamental groups of those intersections.


LEMMA 5.8. *Let A and B be locally convex subsets of the NPC space C . Let I be a component of $A \cap B$, and let $p \in I$. Then $\pi_1(I, p) = \pi_1(A, p) \cap \pi_1(B, p)$.*

Proof. If $\sigma \in (\pi_1(A, p) \cap \pi_1(B, p)) \setminus \pi_1(I, p)$, then represent σ as a based geodesic loop γ_A in A and as another such loop γ_B in B . Unless they coincide, the homotopy between them lifts to a nondegenerate geodesic bigon in the CAT(0) universal cover of C . Therefore, $\gamma_A = \gamma_B \subseteq A \cap B$. \square

LEMMA 5.9. *Let $R > 0$, and suppose A and B are R -embedded locally convex subsets of the NPC space C . Then $A \cap B$ is R -embedded.*

Proof. We first show that components of $A \cap B$ are R -embedded, and then show they are further than $2R$ from one another.

Let I be a component of $A \cap B$. Then I is contained in $A_0 \cap B_0$, where A_0 is a connected component of A and B_0 is a connected component of B . Choosing $p \in I$ determines a diagram of covers



with canonical inclusions $N_R(A_0) \subseteq C^{A_0}$, $N_R(B_0) \subseteq C^{B_0}$, and $N_R(I) \subseteq C^I$. We suppose by contradiction that there are two points $x \neq y$ in $N_R(I)$ so that $\pi_I(x) = \pi_I(y)$. Since $\phi_{A_0}(N_R(I)) \subseteq N_R(A_0)$ and A is R -embedded, it must be the case that $\phi_{A_0}(x) = \phi_{A_0}(y)$. Similarly $\phi_{B_0}(x) = \phi_{B_0}(y)$. Let σ_x be an arc joining x to the lifted basepoint \tilde{p} in I , and let σ_y join y to \tilde{p} . Then $\gamma = \pi_I(\sigma_x \sigma_y)$ is a loop in C based at p which lifts to both C^{A_0} and C^{B_0} , but not to C^I , contradicting Lemma 5.8.

Now suppose that I_1 and I_2 are two different components of $A \cap B$. By way of contradiction, suppose σ is an arc joining I_1 to I_2 of length strictly less than $2R$. For $i = 1, 2$, let A_i be the connected component of A containing I_i , and let B_i be the connected component of B containing I_i . Since A is R -embedded and σ has length less than $2R$, we must have $A_2 = A_1$; similarly $B_2 = B_1$. Again using R -embeddedness, $\sigma \cup A_1$ lifts to C^{A_1} . It follows that σ is homotopic real endpoints to an arc τ_A in A . Since A is locally convex, we may choose τ_A to be geodesic. Similarly, σ is a homotopic real endpoint to a geodesic τ_B in B . But this gives rise to a nondegenerate geodesic bigon in the universal cover \tilde{C} , contradicting the CAT(0) inequality. \square

5.4.2. *Criterion for the double-dot hierarchy to be faithful.* We are particularly interested in subspaces coming from cutting a cube complex up along walls.

Let Y be an NPC cube complex, and let $\tilde{Y} \rightarrow Y$ be the double-dot cover. Let W be a nonseparating 2-sided embedded hyperplane of Y , and let \tilde{W} be the preimage of W under the double-dot cover. Then \tilde{W} cuts Y into two (not necessarily connected) parts, V^+ and V^- .

LEMMA 5.10. *If W is R -embedded then V^+ is R -embedded, as is V^- .*

Proof. Each component of V^+ is bounded by a collection of elevations of W . Lemma 5.7 implies that the union of these elevations is R -embedded, so V^+ is as well. \square

Combining Lemmas 5.10 and 5.9 we obtain:

COROLLARY 5.11. *Let \mathcal{W} be a collection of 2-sided, embedded, nonseparating, R -embedded hyperplanes in a cube complex Y , and let $\mathbf{t} \in \bigoplus_{\mathcal{W}} \mathbb{Z}/2$. If $V_{\mathbf{t}}$ is that part of the double-dot cover labeled by \mathbf{t} then $V_{\mathbf{t}}$ is R -embedded.*

DEFINITION 5.12. Let $\phi : Z \rightarrow X$ be a locally isometric immersion. We say that ϕ (or sometimes Z) is *superconvex* if whenever \tilde{Z} is an elevation of Z to the universal cover of X , then \tilde{Z} contains every biinfinite geodesic which lies in a bounded neighborhood of \tilde{Z} .

LEMMA 5.13 ([Hag08] (cf. [SW11])). *Let G be hyperbolic and the fundamental group of an NPC cube complex X . Let $H < G$ be a quasiconvex subgroup. Then there is a superconvex immersion $\phi : Z \rightarrow X$ with $\phi_*(\pi_1(Z)) = H$.*

Proof. Let Λ_H be the limit set of H in $\partial G = \partial X$. By quasiconvexity and δ -hyperbolicity, any geodesic joining elements of ∂H stays uniformly close to H . Thus, there is a compact set K in \tilde{X} so that $H \cdot K$ contains all such geodesics. It follows by δ -hyperbolicity of \tilde{X} that quasiconvexity and combinatorial quasiconvexity (in the sense of [Hag08, Definition 2.24]) are the same notion. Therefore, by [Hag08, Theorem H] the combinatorial convex hull Y of $H \cdot K$ is H -cocompact. We take $Z = Y/H$. \square

DEFINITION 5.14. Let (X, \mathcal{Z}) be a cube complex pair, so that the immersion $\Phi : \mathcal{Z} \rightarrow X$ is superconvex on each component of \mathcal{Z} . Then (X, \mathcal{Z}) is a *superconvex pair*.

In order to ensure our edge spaces do not have any accidental parabolics, we need the following lemma:

LEMMA 5.15. *Let (X, \mathcal{Z}) be a superconvex pair, with each component of \mathcal{Z} embedded, and let C be the corresponding augmented cube complex. For $n \geq 1$ let $\{W_1, \dots, W_n\}$ be a collection of embedded, 2-sided, nonseparating hyperplanes of C . Let P be a component of $W_n \setminus \bigcup_{i < n} W_i$. Then P has no accidental \mathcal{Z} -loops.*

Proof. Recall that the augmented space C is the mapping cylinder of a locally isometric immersion Φ from $\mathcal{Z} = \bigsqcup_{j=1}^k Z_j$ to X , where each Z_j is a connected NPC cube complex. Let $\phi_j = \Phi|_{Z_j}$. Consider a cube τ of C which intersects, but is not contained in X . This cube can be identified with $\sigma \times [0, 1]$ where $\sigma = \sigma \times \{0\}$ is a cube of Z , and $\sigma \times \{1\} = \Phi(\sigma)$. We observe that under this identification, $P \cap \tau = (P \cap \sigma) \times [0, 1]$. Since Z_j is embedded, we get a copy of $Z_j \times [0, 1] \subseteq C$, and the foregoing argument shows that the intersection of P with this product is exactly $(P \cap Z_j \times \{0\}) \times [0, 1]$.

Now, let $\alpha \subseteq P$ be a geodesic loop which is homotopic (in C) to a geodesic loop α' in $Z_j \subseteq \mathcal{Z}$.

Suppose first that α is contained in $C \setminus X$. It is then possible to homotope α along the product structure into \mathcal{Z} . Since the intersection of P with the product structure is itself a product, this homotopy takes place inside P , so α is not an accidental \mathcal{Z} -loop.

We may therefore assume that α intersects X . We claim that α must be contained in X . Suppose not. By reparameterizing the geodesic loop $\alpha : [0, L] \rightarrow C$, we may suppose that $\alpha(0) = \alpha(L) \in X$, but that $\alpha(t) \notin X$ for small positive t . Let R be the smallest positive number so that $\alpha(R) \in X$, and let $\beta = \alpha|_{[0, R]}$. Then β has image in some $Z_j \times [0, 1]$, so we may write $\beta(t) = (z(t), s(t))$ for $z(t) \in Z_j, s(t) \in [0, 1]$. The projection $s(t)$ must also be locally geodesic, so the quantity $s'(t)$ is constant. Since $s(0) = s(R) = 1$, we have $s'(t) = 0$ for all $t \in [0, R]$, contradicting the supposition that $\alpha(t) \notin X$ for small positive t .

Lifting to the universal cover, we obtain a biinfinite geodesic $\tilde{\alpha} \subseteq \tilde{X}$ which lies in a bounded neighborhood of the image of some elevation of $\phi_j : Z_j \rightarrow X$. Since ϕ_j is superconvex, $\tilde{\alpha}$ actually lies in this elevation. But this means that α is in $P \cap \phi_j(Z_j) = P \cap (Z_j \times \{1\})$. Since P intersects $Z_j \times [0, 1]$ in a product, α is homotopic in P into Z_j , and therefore into \mathcal{Z} . Thus α is not an accidental \mathcal{Z} -loop. □

THEOREM 5.16. *Suppose that (X, \mathcal{Z}) is an NPC cube complex pair, and suppose that C is the associated augmented complex. Suppose that the universal cover of C is δ -hyperbolic, and that there exists some function f so that the family \mathcal{Z} is f -separated. Let ϵ_A and ϵ_B be constants sufficient for the Combination Theorem 4.5.*

- (1) Suppose that each nonseparating embedded 2-sided hyperplane of C is ϵ_A -embedded, and that each component Z of \mathcal{Z} is ϵ_B -embedded. Then the double-dot hierarchy of \ddot{C} is faithful, and the corresponding hierarchy of fundamental groups is quasiconvex.
- (2) Suppose further that the pair (X, \mathcal{Z}) is superconvex. Then the edge spaces of the double-dot hierarchy have no accidental $\ddot{\mathcal{Z}}$ -loops.

Proof. For the first conclusion, it suffices to show all the edge spaces are π_1 -injective, with quasiconvex elevations to the universal cover \tilde{C} of C .

Let $\{W_1, \dots, W_n\}$ be the ordered list of embedded, nonseparating, 2-sided hyperplanes in C . Let V be a vertex space of the hierarchy at level $i - 1$. Then the nonperipheral part of V is marked by some $\mathbf{t} \in M_{i-1} = \bigoplus_{\{W_1, \dots, W_{i-1}\}} \mathbb{Z}/2$. Let $V_{\mathbf{t}}$ be the part of \ddot{C} marked by \mathbf{t} . By Corollary 5.11, the set $V_{\mathbf{t}}$ is ϵ_A -embedded.

The vertex space V is a component of $V_{\mathbf{t}} \cup \ddot{\mathcal{Z}}$. Let $\{\mathbf{t}^+, \mathbf{t}^-\}$ be the fiber above \mathbf{t} under the projection $M_i \rightarrow M_{i-1}$.

Let E be an edge space at level i of the hierarchy. Then E is a component of the intersection of two vertex spaces V^+ and V^- , marked by \mathbf{t}^+ and \mathbf{t}^- , respectively. Let \ddot{W}_i be the preimage of W_i in the double-dot cover \ddot{C} . By Lemma 5.7, \ddot{W}_i is ϵ_A -embedded.

Let $\mathcal{A} = V_{\mathbf{t}} \cap \ddot{W}_i$. By Lemma 5.9, \mathcal{A} is ϵ_A -embedded.

Let $\mathcal{B} = \ddot{\mathcal{Z}} \subseteq \ddot{C}$. Then E is a component of $\mathcal{A} \cup \mathcal{B}$. The first conclusion of the theorem follows from the Combination Theorem 4.5 for \mathcal{A} , \mathcal{B} , and setting $\Theta = E$.

For the second conclusion, we argue as follows. At level i of the hierarchy the edge spaces are components of $\mathcal{A} \cup \mathcal{B}$ where the components of \mathcal{A} are elevations of components of $W_i \setminus \bigcup_{j < i} W_j$. Lemmas 5.15 and 4.7 imply that the components of \mathcal{A} have no accidental \mathcal{B} -loops. We can therefore apply Proposition 4.8 to deduce that E has no accidental $\ddot{\mathcal{Z}}$ -loops. □

5.5. Scott’s Criterion and separability. We briefly recall Scott’s Criterion [Sco78] about separability as it is used multiple times in the proof of Theorem 2.11 as an application of separability. We ignore basepoints in what follows since doing this does not cause any ambiguity or confusion.

Suppose that X is a connected complex and that H is a subgroup of $\pi_1(X)$. Let X^H be the cover of X corresponding to H . Scott’s Criterion states that H is separable in $\pi_1(X)$ if and only if the following condition holds:

For every finite subcomplex Δ of X^H there exists an intermediate covering

$$X^H \rightarrow X_{\Delta} \rightarrow X$$

so that $X_{\Delta} \rightarrow X$ is a finite-degree covering and Δ embeds in X_{Δ} .

We use this (often implicitly) to promote certain immersions to embeddings in finite covers, and for other purposes that are made clear by this criterion.

5.6. Proof of Theorem 2.11. In the following, we make repeated use of Scott's criterion, together with the fact (referred to as *QCERF*) that quasiconvex subgroups of virtually special groups are separable (Theorem 3.28).

Step 1. We start with a relatively hyperbolic pair (G, \mathcal{P}) so that G is hyperbolic and virtually special. Since G is virtually special, there is some $(G_0, \mathcal{P}_0) \triangleleft (G, \mathcal{P})$ (see Notation 2.9) so that $G_0 = \pi_1 X_0$, where X_0 is a special cube complex.

LEMMA 5.17. *If X_0 is a compact special cube complex, then some isometrically embedded subcomplex Y of X_0 is homotopy equivalent to X_0 and has the property that every hyperplane of Y gives a nontrivial splitting of $G_0 = \pi_1 X_0$.*

Proof. Since X_0 is special, every hyperplane is embedded and two-sided. Thus each hyperplane gives some one-edge splitting of G_0 . Suppose W is a hyperplane giving a trivial splitting of G_0 , with open cubical neighborhood $N(W)$. Since the group-theoretic splitting is trivial, $X_0 \setminus N(W)$ must have two components, at least one of which is homotopy equivalent to X_0 . Let X' be a component of $X_0 \setminus N(W)$ homotopy equivalent to X_0 . The total number of cubes (of all dimensions) in X' is less than the total number of cubes in X_0 . Thus we can only perform this procedure finitely many times, eventually arriving at a subcomplex which every hyperplane splits nontrivially. \square

We may therefore suppose (possibly replacing X_0 by a subcomplex) every hyperplane of X_0 gives a nontrivial splitting of G_0 . Any finite cover of X_0 therefore has the same property.

Let $H < G_0$ be the fundamental group of a hyperplane W . Since W is convex in the (Gromov hyperbolic and CAT(0)) universal cover of X_0 , H is a quasiconvex subgroup. Using QCERF we can pass to a finite cover in which every elevation of W is nonseparating. (For the idea, see [Lub96, Lemma 2.2].) Therefore, there is a finite regular cover $X_1 \rightarrow X_0$ in which every hyperplane is nonseparating.

On the level of fundamental groups we have $G_1 = \pi_1 X_1 \triangleleft G_0$ with an induced peripheral structure $(G_1, \mathcal{P}_1) \triangleleft (G_0, \mathcal{P}_0)$ (again as in Notation 2.9). We can suppose further that $(G_1, \mathcal{P}_1) \triangleleft (G, \mathcal{P})$ by passing to a further finite-index subgroup if necessary.

Step 2. The subgroups \mathcal{P}_1 are quasiconvex in G_1 so we can represent them by superconvex locally isometric immersions $\{Z(P) \xrightarrow{\phi_P} X_1 \mid P \in \mathcal{P}_1\}$

(Lemma 5.13). Setting $\mathcal{Z}_1 = \bigcup_{P \in \mathcal{P}_1} Z(P)$ we have a superconvex cube complex pair (X_1, \mathcal{Z}_1) , and we can form the associated augmented complex C_1 . The following easy lemma ensures that finite covers of C_1 are of the same form.

LEMMA 5.18. *Any regular finite cover C' of C_1 has the properties:*

- (1) *C' is the augmented complex for a superconvex pair (X', \mathcal{Z}') with*

$$(G', \mathcal{P}') = (\pi_1(X'), \{\pi_1(Z) \mid Z \text{ a component of } \mathcal{Z}'\}) \triangleleft (G_1, \mathcal{P}_1)$$

and

- (2) *Every nonperipheral hyperplane of C' is nonseparating.*

Step 3. We fix some parameters. Since G is a hyperbolic group, the universal cover \tilde{C}_1 of C_1 is δ -hyperbolic for some δ . Since (G, \mathcal{P}) is relatively hyperbolic, so is (G_1, \mathcal{P}_1) . In particular, \mathcal{P}_1 is a malnormal collection. Proposition 4.4 then implies that the family of elevations of components of \mathcal{Z}_1 to \tilde{C}_1 is f -separated for some affine f . We let ϵ_A and ϵ_B be constants which are sufficiently large for the Combination Theorem 4.5 to work. (For example we may take $\epsilon_B = 50\delta + 1$ and $\epsilon_A = 2 \max\{f(\epsilon_B + 2\delta) + \delta + \epsilon_B, 2\epsilon_B\} + 1$.)

Fundamental groups of hyperplanes and of peripheral complexes are quasiconvex in G . Using QCERF we may pass to a finite regular cover $C_2 \rightarrow C_1$ associated to a pair (X_2, \mathcal{Z}_2) so that:

- (1) every nonperipheral hyperplane is ϵ_A -embedded in C_2 ; and
- (2) every peripheral subcomplex is ϵ_B -embedded in C_2 .

Step 4. Let $\ddot{G}_2 = \pi_1(\ddot{C}_2)$, and let

$$\ddot{\mathcal{P}}_2 = \{\pi_1(\tilde{Z} < \pi_1(\ddot{C}_2)) \mid \tilde{Z} \text{ an elevation of a component } Z \subseteq \mathcal{Z}_2 \text{ to } \ddot{C}_2\}.$$

We now can show that the double-dot hierarchy \mathcal{H} on \ddot{C}_2 is faithful and quasiconvex using Theorem 5.16(1). Since the peripheral subcomplexes are never cut up, this hierarchy is $\ddot{\mathcal{P}}_2$ -elliptic. Since the pair (X_2, \mathcal{Z}_2) is superconvex, Theorem 5.16(2) implies that the edge spaces of the hierarchy have no accidental $\ddot{\mathcal{Z}}_2$ -loops. This implies that the hierarchy is fully $\ddot{\mathcal{P}}_2$ -elliptic.

The cube complex C_2 is the augmented complex associated to a pair (X_2, \mathcal{Z}_2) so that every hyperplane of X_2 is embedded, 2-sided, and nonseparating (from Step 1), so Corollary 5.6 applies. We deduce that the hierarchy of groups on \ddot{G}_2 terminates in free products of free groups and elements of $\ddot{\mathcal{P}}_2$.

Step 5. Pass to a further finite-index subgroup $(G_3, \mathcal{P}_3) \triangleleft (\ddot{G}_2, \ddot{P}_2)$ so the hierarchy induced by \mathcal{H} is a malnormal hierarchy. This can be done by Theorem 3.33.

The terminal groups of the induced hierarchy are finite index in the terminal groups of \mathcal{H} . In particular, they are free products of free groups and elements of \mathcal{P}_3 . We can then continue the hierarchy (quasiconvexly, malnormally, and fully \mathcal{P}_3 -elliptically) to one which terminates in \mathcal{P}_3 . \square

6. The (malnormal) special combination theorem

The malnormal quasiconvex hierarchy theorem of Wise is essential to our proof of Theorem 2.12. We explain how the result follows from the work in [HW15, HW12].

THEOREM 6.1 (Malnormal quasiconvex hierarchy [Wis, Theorem 11.2]). *If G is a hyperbolic group with a malnormal quasiconvex hierarchy terminating in a collection of virtually special groups, then G is virtually special.*

Theorem 6.1 follows from the following (malnormal version of the) Special Combination Theorem, by induction on the length of the hierarchy.

THEOREM 6.2 (Malnormal special combination [HW15] + [HW12]). *Suppose that G is a hyperbolic group and that $G = \pi_1(\Gamma, \mathcal{G})$ for some faithful graph of groups (Γ, \mathcal{G}) where:*

- (1) *the edge groups of (Γ, \mathcal{G}) are malnormal and quasiconvex in G ; and*
- (2) *the vertex groups of (Γ, \mathcal{G}) are virtually special.*

Then G is virtually special.

REMARK 6.3. In Section 10, below we remark how the Malnormal Special Quotient Theorem and Dehn filling can be used to remove the condition of malnormality from (1) above in order to obtain another, much more powerful result of Wise from [Wis].

Proof of Theorem 6.2. The result follows from the case that Γ has a single edge by induction on the number of edges in Γ . We therefore assume that Γ is a one-edge splitting.

[HW15, Main Theorem 8.1] implies that G acts properly and cocompactly on a CAT(0) cube complex X .

Case 1: $G = A *_C B$.

Since C is quasiconvex, there is a C -cocompact convex subcomplex $X(C) \subseteq X$ [Hag08, Theorem 2.28]. Denote the quotient by the C -action $Y(C) = C \backslash X(C)$. Since A and B are quasiconvex, there are similarly an A -cocompact subcomplex $X(A)$ and a B -cocompact subcomplex $X(B)$ with quotients $Y(A) = A \backslash X(A)$ and $Y(B) = B \backslash X(B)$. By enlarging $X(A)$ and $X(B)$ if necessary, we may assume their intersection contains $X(C)$.

The inclusions of $X(C)$ into $X(A)$ and $X(B)$ induce local isometries $f : Y(C) \rightarrow Y(A)$ and $g : Y(C) \rightarrow Y(B)$.

These give the necessary data to build a graph of spaces Z with underlying graph a single edge. It is easily verified that Z is a nonpositively curved cube complex, and the edge space $Y(C)$ is a separating hyperplane. Since A and B are assumed to be virtually special groups, the vertex spaces $Y(A)$ and $Y(B)$ must be virtually special cube complexes, as explained in Remark 2.2. Since $\pi_1(Y(C)) = C$ is malnormal in G , we may apply [HW12, Theorem 8.5] to conclude that Z is virtually special.

*Case 2: $G = A *_C$.*

This case is similar, but the graph of spaces construction is slightly different. We suppose that $C < A$, and choose $t \in G$, $\phi : C \rightarrow A$ so that G has the presentation:

$$\langle A, t \mid c = t\phi(c)t^{-1}, c \in C \rangle.$$

Let $C' = \phi(C)$. Choose some C -cocompact convex subcomplex $X(C) \subseteq X$. Note that $X(C') := t^{-1}X(C)$ is C' -cocompact. Since C and C' are both subgroups of the quasiconvex subgroup A , we can choose an A -cocompact convex subcomplex $X(A) \subseteq X$ containing both $X(C)$ and $X(C')$.

We now form the quotient spaces $Y(A) = A \backslash X(A)$, $Y(C) = C \backslash X(C)$, and $Y(C') = C' \backslash X(C')$. Let $\iota : Y(C) \rightarrow Y(A)$ and $\iota' : Y(C') \rightarrow Y(A)$ be the maps induced by the inclusions in X . The translation $x \mapsto t^{-1}x$ on X restricts to a homeomorphism $X(C) \rightarrow X(C')$, which descends to a homeomorphism $\tau : Y(C) \rightarrow Y(C')$. We now have the data to build a graph of spaces with a single vertex space $Y(A)$ and a single edge space $Y(C)$. The two maps of the edge space to $Y(A)$ are ι and $\iota' \circ \tau$, both local isometries.

The space Z thus constructed is a nonpositively curved cube complex whose fundamental group is isomorphic to G . Moreover, it contains a hyperplane H with fundamental group equal to C , and $Z \setminus N(H)$ is virtually special. We may apply [HW12, Theorem 8.5] to conclude that Z is virtually special. \square

REMARK 6.4. The preceding argument does not use the assumption that the edge groups of the hierarchy are malnormal in G , but only the ‘local’ information that they are malnormal in the next level up.

7. Relative hyperbolicity and Dehn filling

In this section, we state the group-theoretic Dehn filling results needed to prove Theorem 2.12.

7.1. Group-theoretic Dehn filling. We first recall the definition of group-theoretic Dehn filling (Definition 2.4). From a group pair (G, \mathcal{P}) and a collection $\{N_i \triangleleft P_i \mid P_i \in \mathcal{P}\}$, we obtain a quotient of G and a filling map

$$\pi : G \rightarrow G(N_1, \dots, N_m) = G / \left\langle\left\langle \bigcup_i N_i \right\rangle\right\rangle. \quad (7.1)$$

DEFINITION 7.1. Let (G, \mathcal{P}) be relatively hyperbolic. A statement is true for all sufficiently long fillings if there is a finite set $B \subset G \setminus \{1\}$ so that the statement holds for all fillings as in (7.1) with $B \cap (\bigcup_i N_i) = \emptyset$.

The basic group-theoretic Dehn filling result is the following (but see [DGO11] for an interesting generalization).

THEOREM 7.2 (Relatively hyperbolic Dehn filling [Osi07], cf. [GM08]). *Let G be a group and $\mathcal{P} = \{P_1, \dots, P_m\}$ a collection of subgroups so that (G, \mathcal{P}) is relatively hyperbolic. Let $F \subset G$ be finite. For all sufficiently long fillings*

$$\phi : G \rightarrow \tilde{G} := G(N_1, \dots, N_m);$$

- (1) $\ker(\phi|_{P_i}) = N_i$ for $i = 1, \dots, m$;
- (2) $(\tilde{G}, \{\phi(P_1), \dots, \phi(P_m)\})$ is relatively hyperbolic; and
- (3) $\phi|_F$ is injective.

7.2. Quasiconvex Dehn filling results. In this subsection, we recall the definition of relatively quasiconvex subgroup and the statements relating quasiconvexity to Dehn fillings.

DEFINITION 7.3. Let (G, \mathcal{P}) be a relatively hyperbolic group (relatively) generated by a finite subset $S \subseteq G$. The relative Cayley graph $\hat{\Gamma} = \Gamma(G, S \sqcup (\bigcup \mathcal{P}))$ is then a δ -hyperbolic, fine graph [Bow12, Osi06]. A subgroup $H < G$ is said to be relatively quasiconvex in (G, \mathcal{P}) if there is a constant λ so that whenever $v \in G$ lies on a $\hat{\Gamma}$ -geodesic with endpoints in H , then $d_S(v, H) \leq \lambda$.

REMARK 7.4. There are many equivalent notions of relatively quasiconvex subgroups—[Hru10] (cf. [Hru]), [AGM09, MMP10].

The following result follows immediately from [Hru10, Theorem 1.5] or [MP12, Theorem 1.1].

PROPOSITION 7.5. *Suppose that G is hyperbolic, that H is quasiconvex in G and that (G, \mathcal{P}) is relatively hyperbolic. Then H is a relatively quasiconvex subgroup of G .*

THEOREM 7.6 [Hru10]. *A relatively quasiconvex subgroup H of a relatively hyperbolic group (G, \mathcal{P}) has a (finite) collection of peripheral subgroups \mathcal{D} with respect to which it is relatively hyperbolic. Moreover, the peripheral subgroups \mathcal{D} of H are conjugate into elements of \mathcal{P} .*

REMARK 7.7. By modifying \mathcal{D} if necessary, we can always suppose:

- (1) each element of \mathcal{D} is infinite; and
- (2) every infinite intersection of H with a conjugate of some $P \in \mathcal{P}$ is conjugate in H to an element of \mathcal{D} .

DEFINITION 7.8. Let (G, \mathcal{P}) be relatively hyperbolic. A relatively quasiconvex subgroup H of G is *fully quasiconvex* if for any $P \in \mathcal{P}$ and any $g \in G$ the subgroup $H \cap P^g$ is either finite or else has finite index in P^g .

DEFINITION 7.9 (*H*-fillings). Let G be hyperbolic relative to $\mathcal{P} = \{P_1, \dots, P_m\}$, and let $H < G$. We say that a filling $G \rightarrow G(N_1, \dots, N_m)$ is an *H-filling* if, whenever $gP_i g^{-1} \cap H$ is infinite, for $P_i \in \mathcal{P}$ and $g \in G$, it follows that $gN_i g^{-1} \subseteq H$.

DEFINITION 7.10 (Induced filling). If H is relatively quasiconvex in (G, \mathcal{P}) then Theorem 7.6 gives H a peripheral structure \mathcal{D} so that (H, \mathcal{D}) is relatively hyperbolic. We suppose \mathcal{D} satisfies the conditions of Remark 7.7. Consider an *H*-filling

$$\pi : G \rightarrow G(N_1, \dots, N_m).$$

Let $D_j \in \mathcal{D}$. There is some $P_i \in \mathcal{P}$ and some $g \in G$ with $g^{-1}D_j g \subseteq P_i$. We define $K_j = gN_i g^{-1}$. (Although D_j determines P_i it does not quite determine g . However, note that since N_i is normal in P_i , the group K_j is independent of the conjugating element g .) Because π is an *H*-filling, we have $K_j \triangleleft D_j$, so these

groups determine a filling of H

$$\pi_H : H \rightarrow H(K_1, \dots, K_n),$$

called the *induced filling of H* .

The next theorem summarizes results from [AGM09], where G is assumed to be torsion-free and a slightly different definition of H -filling is used. As explained in [MMP10, Appendix B] (cf. [Ago13, Appendix]), the torsion-free assumption is unnecessary, so long as we use Definition 7.9.

THEOREM 7.11. *Let (G, \mathcal{P}) be relatively hyperbolic, let $H < G$ be fully relatively quasiconvex, and let $F \subset G$ be finite. For all sufficiently long H -fillings $\phi : G \rightarrow G(N_1, \dots, N_m)$ of G :*

- (1) [AGM09, Proposition 4.3] $\phi(H)$ is fully relatively quasiconvex.
- (2) [AGM09, Proposition 4.4] $\phi(H)$ is isomorphic to the induced filling (see Definition 7.10). More precisely if $\phi_H : H \rightarrow H(K_1, \dots, K_n)$ is the induced filling map, then $(\ker \phi) \cap H = \ker \phi_H$.
- (3) [AGM09, Proposition 4.5] $\phi(F) \cap \phi(H) = \phi(F \cap H)$.

The next statement we need is immediate from the discussion preceding Corollary A.46 in [Ago13]. Note that this discussion does not rely in any way on the Malnormal Special Quotient Theorem, even though the main result of that appendix does. Also note that in the statements of the Appendix to [Ago13], the assumption is made that the peripheral structure \mathcal{P} is the structure induced by the quasiconvex group H . This assumption is not used until we argue that height strictly decreases after filling. For the rest (including the following statement) it is only important that H is fully quasiconvex in (G, \mathcal{P}) .

THEOREM 7.12 (Height reduction). *Suppose that (G, \mathcal{P}) is relatively hyperbolic. Suppose that $H \leq G$ is fully quasiconvex of height k . For all sufficiently long H -fillings $G \rightarrow \bar{G}$, the image \bar{H} of H in \bar{G} has height at most k .*

8. Induced splittings of quotients

In this section, we prove Theorem 2.12, which shows that (with appropriate assumptions) nice hierarchies of (G, \mathcal{P}) descend to nice hierarchies of Dehn fillings of (G, \mathcal{P}) .

8.1. Definition of the filled hierarchy. Let (G, \mathcal{P}) be relatively hyperbolic. Let \mathcal{H} be a quasiconvex fully \mathcal{P} -elliptic hierarchy of G . In this setting, every

vertex or edge group of \mathcal{H} is fully relatively quasiconvex. Let $\pi : (G, \mathcal{P}) \rightarrow (\bar{G}, \bar{\mathcal{P}})$ be a Dehn filling of (G, \mathcal{P}) . In this subsection, we describe a (not necessarily faithful) hierarchy on $(\bar{G}, \bar{\mathcal{P}})$. We call this hierarchy $\bar{\mathcal{H}}$ the *filled hierarchy*.

Each of the groups in the filled hierarchy is the induced filling (in the sense of Definition 7.9) of some group in the hierarchy \mathcal{H} .

At level 0, the filled hierarchy $\bar{\mathcal{H}}$ consists of the degenerate graph of groups structure on \bar{G} . Suppose we have defined the filled hierarchy down to level i less than the length of the hierarchy \mathcal{H} , and let $\bar{A} \in \bar{\mathcal{H}}$ be a vertex group at level i . The group \bar{A} is the induced filling of a vertex group A at level i of \mathcal{H} . This group A comes equipped with a graph of groups structure $\alpha : \pi_1(\Gamma, \mathcal{A}, a_0) \rightarrow A$, from which we must build a graph of groups structure on \bar{A} . This graph of groups structure has the exact same underlying graph Γ and base vertex $a_0 \in \Gamma$, and differs only in the assignment $\bar{\mathcal{A}}$ of groups and homomorphisms, which is defined as follows: If x is a vertex or edge of Γ , then $\bar{\mathcal{A}}_x$ is defined to be the induced filling of \mathcal{A}_x . Write $\pi_x : \mathcal{A}_x \rightarrow \bar{\mathcal{A}}_x$ for the induced filling map. Let e be an edge of Γ with $t(e) = v$, and let $\phi_e : \mathcal{A}_e \rightarrow \mathcal{A}_v$ be the homomorphism coming from the graph of groups. We have to fill in the square:

$$\begin{array}{ccc} \mathcal{A}_e & \xrightarrow{\pi_e} & \bar{\mathcal{A}}_e \\ \phi_e \downarrow & & \downarrow \\ \mathcal{A}_v & \xrightarrow{\pi_v} & \bar{\mathcal{A}}_v \end{array}$$

The necessary condition $\ker \pi_e \subseteq \ker(\pi_v \circ \phi_e)$ holds, so we get an induced map $\bar{\phi}_e : \bar{\mathcal{A}}_e \rightarrow \bar{\mathcal{A}}_v$. These induced maps complete the definition of the graph of groups $(\Gamma, \bar{\mathcal{A}})$.

We emphasize that the edge maps for this graph of groups may not be injective. In our applications (for example Corollary 8.4), we must establish injectivity.

To complete the construction of $\bar{\mathcal{H}}$ we need an isomorphism $\bar{\alpha} : \pi_1(\Gamma, \bar{\mathcal{A}}, a_0) \rightarrow \bar{A}$. Coming from \mathcal{H} we have an isomorphism $\alpha : \pi_1(\Gamma, \mathcal{A}, a_0) \rightarrow A$. Also, the induced fillings $\pi_x : \mathcal{A}_x \rightarrow \bar{\mathcal{A}}_x$ commute with the edge maps, so they suffice to define a surjection $\pi_\Gamma : \pi_1(\Gamma, \mathcal{A}, a_0) \rightarrow \pi_1(\Gamma, \bar{\mathcal{A}}, a_0)$. To define $\bar{\alpha}$ we need to complete the square

$$\begin{array}{ccc} \pi_1(\Gamma, \mathcal{A}, a_0) & \xrightarrow{\pi_\Gamma} & \pi_1(\Gamma, \bar{\mathcal{A}}, a_0) \\ \alpha \downarrow & & \downarrow \\ A & \xrightarrow{\pi_A} & \bar{A}, \end{array} \tag{8.1}$$

where the map at the bottom is the induced filling map of A .

LEMMA 8.1. *The diagram (8.1) can be filled in with an isomorphism $\bar{\alpha}$.*

Proof. To complete the diagram with an isomorphism $\bar{\alpha}$ it is necessary only to show that $\ker \pi_\Gamma = \ker(\pi_A \circ \alpha)$.

We first show $\ker \pi_\Gamma \subseteq \ker(\pi_A \circ \alpha)$. Let $k \in \ker \pi_\Gamma$. Then there is an expression

$$k = \prod_i k_i^{g_i}$$

where each $g_i \in \pi_1(\Gamma, \mathcal{A}, a_0)$, and each k_i lies in the kernel of some π_{v_i} where v_i is a vertex of Γ . (We assume we have made some choice of maximal tree in Γ so we can regard \mathcal{A}_v as a subgroup of $\pi_1(\Gamma, \mathcal{A}, a_0)$.) To show $\ker \pi_\Gamma \subseteq \ker(\pi_A \circ \alpha)$ it therefore suffices to show $\ker \pi_v \subseteq \ker \pi_A \circ \alpha$ for each v . We therefore assume without loss of generality that $k \in \ker \pi_v$ for some vertex v .

The hierarchy is assumed quasiconvex and fully \mathcal{P} -elliptic. Thus, \mathcal{A}_v is fully relatively quasiconvex in (G, \mathcal{P}) . Even better, there is an induced peripheral structure $\mathcal{D}_v = \{D_1, \dots, D_n\}$ on \mathcal{A}_v so that for each i , there is an infinite $P_{j_i} \in \mathcal{P}$ and a $g_i \in G$ so that $D_i = (P_{j_i})^{g_i}$. Since $\pi_v : \mathcal{A}_v \rightarrow \bar{\mathcal{A}}_v$ is the filling induced by π , the element k has an expression

$$k = \prod_\beta n_\beta^{a_\beta}$$

with each n_β in some induced filling kernel $N_\beta \subseteq D_\beta$ and each $a_\beta \in \mathcal{A}_v$. It suffices then to show each of these n_β is in $\ker \pi_A \circ \alpha$. But this is clear, since the conjugate $(P_{j_\beta})^{g_\beta}$ must be conjugate in A to some element of the induced peripheral structure on A , and so n_β is conjugate into some filling kernel of the induced filling π_A on A .

Conversely, if $k \in \ker(\pi_A \circ \alpha)$, then $k = \prod_{i=1}^l \alpha^{-1}(k_i^{g_i})$ where each k_i is in some filling kernel $K_{j_i} \triangleleft D_{j_i}$ for the induced filling π_A . But since the hierarchy is fully \mathcal{P} -elliptic, the subgroup $\alpha^{-1}(D_{j_i})$ is conjugate into some vertex group \mathcal{A}_v of (Γ, \mathcal{A}) . It follows that $\alpha^{-1}(k_i)$ is conjugate into $\ker \pi_v$ for some $v \in \Gamma$, so $\alpha^{-1}(k_i)$ is in $\ker \pi_\Gamma$. □

8.2. Proof of Theorem 2.12. The first lemma is a straightforward application of the definitions.

LEMMA 8.2. *Let G be hyperbolic, and \mathcal{P} be a family of subgroups so that (G, \mathcal{P}) is relatively hyperbolic. Let \mathcal{H} be a quasiconvex fully \mathcal{P} -elliptic hierarchy of G , and let $A < G$ be an edge or vertex group of \mathcal{H} . Then A is fully relatively quasiconvex in (G, \mathcal{P}) , and every filling of (G, \mathcal{P}) is an A -filling.*

LEMMA 8.3. *Let G be hyperbolic, and \mathcal{P} be a family of subgroups so that (G, \mathcal{P}) is relatively hyperbolic. Let \mathcal{H} be a quasiconvex fully \mathcal{P} -elliptic hierarchy of G , and let A be an edge or vertex group of \mathcal{H} . Then for all sufficiently long fillings*

$$\phi : (G, \mathcal{P}) \rightarrow (\bar{G}, \bar{\mathcal{P}}),$$

the following hold whenever A is a vertex or edge group of \mathcal{H} :

- (1) *The subgroup $\bar{A} = \phi(A)$ is fully relatively quasiconvex in $(\bar{G}, \bar{\mathcal{P}})$.*
- (2) *If \bar{G} is hyperbolic, then \bar{A} is quasiconvex in \bar{G} .*
- (3) *The subgroup \bar{A} is isomorphic to the induced filling of A .*
- (4) *The height of \bar{A} in \bar{G} is at most the height of A in G .*

Proof. As there are only finitely many edge and vertex groups occurring in \mathcal{H} , statements (1), (3), and (4) follow from 7.11(1), 7.11(2) (with $F = \emptyset$), and 7.12, respectively. The statement (2) follows from (1) by [Hru10, Theorem 10.5] (which implies that in this case \bar{A} is undistorted in \bar{G}) and [BH99, Corollary III.Γ.3.6]. □

In particular, Lemma 8.3(3) has the immediate consequence that the filled hierarchy is faithful:

COROLLARY 8.4. *Let G be hyperbolic, and \mathcal{P} be a family of subgroups so that (G, \mathcal{P}) is relatively hyperbolic. Let \mathcal{H} be a quasiconvex fully \mathcal{P} -elliptic hierarchy of G . Then for all sufficiently long fillings*

$$\pi : (G, \mathcal{P}) \rightarrow (\bar{G}, \bar{\mathcal{P}}),$$

the filled hierarchy $\bar{\mathcal{H}}$ on \bar{G} is faithful.

Proof of Theorem 2.12. As in the statement, suppose that (G, \mathcal{P}) is a relatively hyperbolic pair with G hyperbolic, and suppose G is equipped with a malnormal quasiconvex fully \mathcal{P} -elliptic hierarchy \mathcal{H} terminating in \mathcal{P} . Let

$$\pi : (G, \mathcal{P}) \rightarrow (\bar{G}, \bar{\mathcal{P}})$$

be a filling which is sufficiently long with respect to Lemma 8.3 and Corollary 8.4. By Corollary 8.4, the quotient \bar{G} inherits a hierarchy $\bar{\mathcal{H}}$ with the same underlying graphs and with each of the edge or vertex groups equal to the image under π of a corresponding edge or vertex group of \mathcal{H} . In particular, all these edge and vertex groups are quasiconvex, by Lemma 8.3(2), so $\bar{\mathcal{H}}$ is a

quasiconvex hierarchy. By 8.3(4), the edge groups of $\bar{\mathcal{H}}$ are all malnormal in \bar{G} , so $\bar{\mathcal{H}}$ is a malnormal hierarchy. The construction of the filled hierarchy ensures the terminal groups are the induced fillings of the terminal groups of \mathcal{H} . Since the terminal groups of \mathcal{H} are in \mathcal{P} , the terminal groups of $\bar{\mathcal{H}}$ are in $\bar{\mathcal{P}}$. \square

9. Proof of the main theorem

Before we prove the main theorem, we introduce some terminology.

DEFINITION 9.1. Suppose that $(G', \mathcal{P}') \triangleleft (G, \mathcal{P})$, and that $\{N'_j \triangleleft P'_j \mid P'_j \in \mathcal{P}'\}$ is a collection of filling kernels.

We say that $\{N'_j\}$ is *equivariantly chosen* if (i) whenever gP'_jg^{-1} and $hP'_k h^{-1}$ both lie in P_i , then $gN'_jg^{-1} = hN'_kh^{-1}$; and (ii) Each such gN'_jg^{-1} is normal in the P_i which contains it.

An *equivariant filling* of (G', \mathcal{P}') is a filling with an equivariantly chosen set of filling kernels.

The following observation is the reason for considering equivariant fillings.

PROPOSITION 9.2. An equivariant filling $(G', \mathcal{P}') \rightarrow (\bar{G}', \bar{\mathcal{P}}')$ determines a filling $(G, \mathcal{P}) \rightarrow (\bar{G}, \bar{\mathcal{P}})$, so that $(\bar{G}', \bar{\mathcal{P}}') \triangleleft (\bar{G}, \bar{\mathcal{P}})$.

We now apply Wise’s Malnormal Quasiconvex Hierarchy Theorem (Theorem 6.1), together with our Theorems 2.11 and 2.12, to establish the main result.

Proof of Theorem 2.7. Let (G, \mathcal{P}) satisfy the hypotheses of Theorem 2.7. Theorem 2.11 implies that there is a subgroup $(G', \mathcal{P}') \triangleleft (G, \mathcal{P})$ with a malnormal quasiconvex fully \mathcal{P}' -elliptic hierarchy terminating in \mathcal{P}' .

According to Theorem 2.12, any sufficiently long filling $(\bar{G}', \bar{\mathcal{P}}')$ of (G', \mathcal{P}') has a quasiconvex malnormal fully $\bar{\mathcal{P}}'$ -elliptic hierarchy terminating in $\bar{\mathcal{P}}'$. Since G is virtually special, it is residually finite, so there exist peripherally finite fillings of (G', \mathcal{P}') which are sufficiently long for the conclusions of Theorem 2.12 to hold. Let $G'(K_1, \dots, K_M)$ be some such peripherally finite filling, which we may assume is also sufficiently long for the conclusions of Theorem 7.2 to hold. We then modify it to be equivariant with respect to G , by setting

$$K'_i = \bigcap \{K_j^g \mid g \in G, \#(K_j^g \cap P_i) = \infty\}.$$

We note that these modified filling kernels are contained in the old ones, so the modified filling $G'(K'_1, \dots, K'_M)$ will still be sufficiently long for the conclusions of Theorems 2.12 and 7.2 to hold.

The filling $G'(K'_1, \dots, K'_M)$ is still peripherally finite, and by Proposition 9.2 it moreover determines a filling of G . For $P_j \in \mathcal{P}$, define \dot{P}_j to be equal to $(K'_i)^g$ for some (any) i, g so that $(K'_i)^g \subset P_j$. The equivariance ensures there is no ambiguity in this definition.

Now consider any filling $G \rightarrow G(N_1, \dots, N_m)$ so that for each i we have (i) $N_i \triangleleft P_i$; (ii) $N_i < \dot{P}_i$; and (iii) P_i/N_i is virtually special and hyperbolic. There is an induced equivariant filling

$$G' \longrightarrow G'(N'_1, \dots, N'_M)$$

so that $N'_j < K'_j$ and $N'_j \triangleleft P'_j$ for each j . This filling is sufficiently long that the conclusion of Theorem 2.12 holds. In particular, the filling $G'(N'_1, \dots, N'_M)$ has a malnormal quasiconvex hierarchy terminating in $\bar{\mathcal{P}}' = \{P'_j/N'_j\}$. By Theorem 7.2, the pair $(G'(N'_1, \dots, N'_M), \bar{\mathcal{P}}')$ is relatively hyperbolic; since the groups in $\bar{\mathcal{P}}'$ are hyperbolic, so is the filling. Thus $G'(N'_1, \dots, N'_M)$ is a hyperbolic group with a malnormal quasiconvex hierarchy terminating in a collection of hyperbolic, virtually special groups. By Theorem 6.1, $G'(N'_1, \dots, N'_M)$ is virtually special.

By Proposition 9.2, $G'(N'_1, \dots, N'_M) \triangleleft G(N_1, \dots, N_m)$, so the filling $G(N_1, \dots, N_m)$ of the original group is also virtually special. □

10. Quasiconvex virtual hierarchies

In [Wis], Wise proves a theorem about groups with *virtual* hierarchies, which are slight generalizations of hierarchies allowing finite-index subgroups along the way. We recall the definition and statement of Wise’s Quasiconvex Hierarchy Theorem.

DEFINITION 10.1 [Wis, Definition 11.5]. Let \mathcal{QVH} denote the smallest class of hyperbolic groups that is closed under the following operations.

- (1) $\{1\} \in \mathcal{QVH}$.
- (2) If $G = A *_B C$ and $A, C \in \mathcal{QVH}$ and B is finitely generated and quasiisometrically embedded in G then $G \in \mathcal{QVH}$.
- (3) If $G = A *_B$ and $A \in \mathcal{QVH}$ and B is finitely generated and quasiisometrically embedded then $G \in \mathcal{QVH}$.
- (4) If $H < G$ with $|G : H| < \infty$ and $H \in \mathcal{QVH}$ then $G \in \mathcal{QVH}$.

THEOREM 10.2 [Wis, Theorem 13.3]. *A hyperbolic group is in \mathcal{QVH} if and only if it is virtually (compact) special.*

This is stated as [Ago13, Theorem 2.9] (with the extra assumption that the group is torsion-free). Other than the Malnormal Special Quotient Theorem, it is the result from [Wis] that is used in [Ago13]. It is used in [Ago13, Theorem 3.1].

In this section, we finish by showing how to deduce Wise’s Quasiconvex Hierarchy Theorem from the Malnormal Special Quotient Theorem and some Dehn filling results from the [Ago13, Appendix].

10.1. Induced peripheral structure from a quasiconvex subgroup. We first briefly recall a construction from [AGM09, Ago13]. Let G be a hyperbolic group, and let H be a quasiconvex subgroup. We describe peripheral structures \mathcal{D} on H and \mathcal{P} on G so that:

- (1) (H, \mathcal{D}) and (G, \mathcal{P}) are relatively hyperbolic;
- (2) H is fully relatively quasiconvex in (G, \mathcal{P}) and \mathcal{D} satisfies the conditions given in Remark 7.7; and
- (3) every element of \mathcal{P} is commensurable in G with at least one element of \mathcal{D} .

We refer to [AGM09, Ago13] for proofs of the above facts.

By [AGM09, Corollary 3.5], there are only finitely many H -conjugacy classes of minimal infinite subgroups of the form

$$H \cap g_2 H g_2^{-1} \cap \dots \cap g_j H g_j^{-1}. \tag{10.1}$$

The *malnormal core* \mathcal{D} of H is formed as follows:

- (1) Let \mathcal{D}'' be the set of minimal infinite subgroups of the form (10.1).
- (2) Form \mathcal{D}' from \mathcal{D}'' by replacing each $D \in \mathcal{D}''$ with its H -commensurator.
- (3) Form \mathcal{D} from \mathcal{D}' by choosing one representative of each H -conjugacy class.

The *induced peripheral structure* \mathcal{P} in G is then formed from \mathcal{D} via two similar steps:

- (1) Form \mathcal{P}' from \mathcal{D} by replacing each $D \in \mathcal{D}$ with its G -commensurator.
- (2) Form \mathcal{P} from \mathcal{P}' by choosing one representative of each G -conjugacy class.

In this section, we use the following refinement of Theorem 7.12 from [Ago13]. Note that the proof of this result is elementary and does not use the Malnormal Special Quotient Theorem or indeed any results about cube complexes.

THEOREM 10.3 [Ago13, Theorem A.47]. *Let H be a quasiconvex subgroup of the hyperbolic group G , let \mathcal{D} be the malnormal core of H in G and let \mathcal{P} be the peripheral structure induced on G by H . For all sufficiently long peripherally finite fillings $\phi : G \rightarrow \tilde{G} = G(N_1, \dots, N_m)$, the height of $\phi(H)$ in \tilde{G} is strictly less than that of H in G .*

10.2. Induced peripheral structure from a one-edge splitting. In this subsection, we suppose that G is a hyperbolic group, and that G has a one-edge splitting with edge group $H < G$. That is, either $G = A *_H B$ or $G = A *_\phi$, where $\phi : H \rightarrow H'$ is an isomorphism of two subgroups of G . We suppose further that the edge group H is quasiconvex.

To the splitting corresponds a Bass–Serre tree on which G acts, the edges of which are in one-to-one correspondence with left cosets of H in G . Let e be the edge corresponding to the trivial coset. Let \mathcal{T} be the collection of maximal subtrees of T with infinite stabilizer, and let $\mathcal{T}_e = \{T \in \mathcal{T} \mid e \subseteq T\}$. Because the height of H in G is finite, each tree in \mathcal{T} is finite.

OBSERVATION 10.4. *With the above notation:*

- (1) *Each $D \in \mathcal{D}$ preserves some tree in \mathcal{T}_e ; this gives a one-to-one correspondence of \mathcal{D} with the orbit space $H \backslash \mathcal{T}_e$.*
- (2) *Each $P \in \mathcal{P}$ also preserves some tree in $\mathcal{T}_e \subset \mathcal{T}$, giving a one-to-one correspondence of \mathcal{P} with the orbit space $G \backslash \mathcal{T}$.*
- (3) *Since each tree $T \in \mathcal{T}$ is finite, each conjugate gPg^{-1} with $P \in \mathcal{P}$ has a finite-index subgroup which is equal to the stabilizer of a tree in \mathcal{T} .*

LEMMA 10.5. *Let G be hyperbolic, with a one-edge splitting over a quasiconvex subgroup H , and with \mathcal{D}, \mathcal{P} as above. The vertex groups of the splitting are fully relatively quasiconvex in (G, \mathcal{P}) .*

Proof. Let P be a conjugate of an element of \mathcal{P} which intersects a vertex group A in an infinite subgroup. A finite-index subgroup $P' < P$ is the stabilizer of some tree $T \in \mathcal{T}$. The vertex group A is the stabilizer of some vertex v in the Bass–Serre tree. Since $A \cap P$ is infinite, so also is $A \cap P'$. This infinite subgroup $A \cap P'$ therefore stabilizes all of $T \cup \{v\}$. It follows that $v \in T$, since T is a maximal tree with infinite stabilizer. Thus P' fixes v , and so $P' < A$. \square

10.3. Virtual torsion-freeness. In this subsection, we show that a group with a quasiconvex one-edge splitting and virtually special vertex groups is virtually torsion-free.

LEMMA 10.6. *Suppose that (Γ, \mathcal{G}) is a (faithful) finite graph of groups with finite edge groups and virtually torsion-free vertex groups. Then $\pi_1(\Gamma, \mathcal{G})$ is virtually torsion-free.*

Proof. Each vertex group \mathcal{G}_v contains some finite-index normal N_v which is torsion-free. We can form a new graph of groups (Γ, \mathcal{G}') with the same underlying graph of groups Γ and the same edge groups, but with $\mathcal{G}'_v = \mathcal{G}_v/N_v$. Since the kernels of the maps $\mathcal{G}_v \rightarrow \mathcal{G}'_v$ are torsion-free, each edge-inclusion map $\phi_e : \mathcal{G}_e \rightarrow \mathcal{G}_{t(e)}$ induces a well-defined monomorphism $\phi'_e : \mathcal{G}'_e \rightarrow \mathcal{G}'_{t(e)}$.

Because the vertex groups of (Γ, \mathcal{G}) are quotients of the vertex groups of (Γ, \mathcal{G}') and the edge-inclusion maps of (Γ, \mathcal{G}') are induced by those of (Γ, \mathcal{G}) , there is an induced surjection

$$q : \pi_1(\Gamma, \mathcal{G}) \rightarrow \pi_1(\Gamma, \mathcal{G}').$$

Since the maps $\mathcal{G}_v \rightarrow \mathcal{G}'_v$ have torsion-free kernel, so does q . But $\pi_1(\Gamma, \mathcal{G}')$ is a graph of finite groups, so it is virtually free. Let $F \triangleleft \pi_1(\Gamma, \mathcal{G}')$ be free. Then $q^{-1}(F) \triangleleft \pi_1(\Gamma, \mathcal{G})$ is torsion-free. □

PROPOSITION 10.7. *Suppose the hyperbolic group G admits a one-edge splitting with edge group equal to the quasiconvex subgroup H . Suppose the vertex groups of this splitting are virtually special. Then G is virtually torsion-free.*

Proof. We argue by induction on the height n of H in G . If $n = 0$, then the result follows from Lemma 10.6, since virtually special groups are virtually torsion-free.

If $n > 0$, then let \mathcal{D} be the malnormal core of H , and let \mathcal{P} be the peripheral structure on G induced by H . By Lemma 10.5, the vertex groups of the splitting over H are fully relatively quasiconvex. The elements of \mathcal{P} are virtually quasiconvex subgroups of vertex groups, so they are virtually special. Therefore, there exist arbitrarily long peripherally finite fillings of (G, \mathcal{P}) . Since the vertex groups are fully relatively quasiconvex, there are even arbitrarily long peripherally finite fillings which are both H -fillings and A -fillings for every vertex group A .

By the Malnormal Special Quotient Theorem (Corollary 2.8), there are also arbitrarily long peripherally finite such fillings so that the induced fillings of the vertex groups are themselves virtually special. Choose some such peripherally finite filling $\phi : G \rightarrow \bar{G} = G(N_1, \dots, N_m)$, long enough so that:

- (1) ϕ has no nontrivial torsion elements in its kernel (this is possible because there are only finitely many conjugacy classes of torsion elements in a hyperbolic group);

- (2) ϕ restricts to the induced filling on H and on each vertex group (see 7.11(2) with $F = \emptyset$);
- (3) \tilde{G} is hyperbolic and $\tilde{H} = \phi(H)$ is quasiconvex in \tilde{G} (see 7.11(1) with $F = \emptyset$);
- (4) the height of $\phi(H)$ in \tilde{G} is strictly less than the height of H in G (Theorem 10.3).

Because the filling ϕ is an H -filling as well as an A -filling for each vertex group A , Lemma 8.1 implies that the (hyperbolic) quotient \tilde{G} has a 1-edge splitting over \tilde{H} . Also, \tilde{H} is quasiconvex, and has height smaller than that of H . Induction implies that \tilde{G} is virtually torsion-free. Let $\tilde{G}_0 < \tilde{G}$ be torsion-free, and let $G_0 = \phi^{-1}(\tilde{G}_0)$. The group G_0 is a torsion-free finite-index subgroup of G . \square

10.4. Proof of Theorem 10.2.

We now prove Theorem 10.2.

Proof. We have already noted in Example 3.15 that if G is hyperbolic and special then it admits a particular nice hierarchy and it is easy to see that this means any hyperbolic virtually special group is in $Q\mathcal{VH}$. Thus, it is the converse that we must prove—any hyperbolic group in $Q\mathcal{VH}$ is virtually special.

Certainly we know that $\{1\}$ is (virtually) special. Also, if H is virtually special and $H < G$ then G is clearly also virtually special. Therefore, it suffices to prove that if G admits a one-edge splitting with virtually special vertex group(s) and quasiconvex edge group H then G is virtually special.

We induct on the height of H in G . If the height is 0, then H is finite. Proposition 10.7 implies that there is some $G_0 < G$ which is torsion-free. But then G_0 is a graph of virtually special groups with trivial edge groups. In particular G_0 has a malnormal quasiconvex hierarchy terminating in virtually special groups. By Theorem 6.1, G_0 (and hence G) is virtually special.

Now suppose the height of H is positive. The following is the key claim:

CLAIM. H is separable in G .

Suppose the claim has been proved. By Proposition 10.7, G is virtually torsion-free. By Proposition 3.27, separable quasiconvex subgroups of G are virtually malnormal. Therefore, there is a finite-index (torsion-free) normal subgroup $G_0 < G$ so that for any G -conjugate H' of H , the subgroup $H' \cap G_0$ is malnormal in G_0 . The induced graph of groups for the hyperbolic group G_0 has malnormal and quasiconvex edge groups, and so G_0 (and hence G) is virtually special by Theorem 6.1.

Proof of claim. This argument follows the same structure as that for Proposition 10.7, as we induct on the height of H . If H is finite (height 0) then we have already shown that G is virtually special, so in particular H is separable.

Now suppose that the height of H is $k > 0$. We wish to show that H is separable, so choose $g \notin H$. We again consider the malnormal core \mathcal{D} of H and the induced peripheral structure \mathcal{P} on G , and choose a peripherally finite filling $\phi : G \rightarrow \bar{G} = G(N_1, \dots, N_m)$ whose induced fillings on the vertex groups are virtually special (by Corollary 2.8), and so that furthermore:

- (1) ϕ restricts to the induced filling on H and on the vertex groups (7.11(2));
- (2) \bar{G} is hyperbolic and $\bar{H} = \phi(H)$ is quasiconvex in \bar{G} (7.11(1));
- (3) the height of \bar{H} in \bar{G} is strictly less than k (Theorem 10.3); and
- (4) $\phi(g) \notin \bar{H}$ (7.11(3) with $F = \{g\}$).

As before (again using the results of Section 8 and in particular Lemma 8.1), the quotient \bar{G} has a 1-edge splitting over \bar{H} , and \bar{H} is quasiconvex and has smaller height than H . The vertex groups are virtually special. By induction there is a finite quotient $\pi : \bar{G} \rightarrow Q$ so that $\pi(\phi(g)) \notin \pi(\bar{H})$. But then $\pi \circ \phi$ gives a map to a finite quotient separating g from H . □

Acknowledgments

We thank the referee for helpful comments and corrections. The first author is supported by NSF grant DMS-1105738, the Miller Institute, and the Simons Foundation. The second author is supported by NSF Grant DMS-0953794. The third author is supported by NSF Grant DMS-1104703.

References

[Ago13] I. Agol, ‘The virtual Haken conjecture’, *Doc. Math.* **18** (2013), 1045–1087. With an appendix by Agol, Daniel Groves, and Jason Manning.

[AGM09] I. Agol, D. Groves and J. F. Manning, ‘Residual finiteness, QCERF and fillings of hyperbolic groups’, *Geom. Topol.* **13**(2) (2009), 1043–1073.

[BC08] M. Baker and D. Cooper, ‘A combination theorem for convex hyperbolic manifolds, with applications to surfaces in 3-manifolds’, *J. Topol.* **1**(3) (2008), 603–642.

[Bas93] H. Bass, ‘Covering theory for graphs of groups’, *J. Pure Appl. Algebra* **89**(1–2) (1993), 3–47.

[Bow12] B. H. Bowditch, ‘Relatively hyperbolic groups’, *Internat. J. Algebra Comput.* **22**(3) (2012), 1250016, 66.

[BH99] M. R. Bridson and A. Haefliger, ‘Metric spaces of non-positive curvature’, in *Grundlehren der mathematischen Wissenschaften*, Vol. 319 (Springer, Berlin, 1999).

- [DGO11] F. Dahmani, V. Guirardel and D. Osin, ‘Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces’, Preprint, 2011, <http://front.math.ucdavis.edu/1111.7048>.
- [GMRS98] R. Gitik, M. Mitra, E. Rips and M. Sageev, ‘Widths of subgroups’, *Trans. Amer. Math. Soc.* **350**(1) (1998), 321–329.
- [GM08] D. Groves and J. F. Manning, ‘Dehn filling in relatively hyperbolic groups’, *Israel J. Math.* **168** (2008), 317–429.
- [Hag08] F. Haglund, ‘Finite index subgroups of graph products’, *Geom. Dedicata* **135** (2008), 167–209.
- [HW08] F. Haglund and D. T. Wise, ‘Special cube complexes’, *Geom. Funct. Anal.* **17**(5) (2008), 1551–1620.
- [HW12] F. Haglund and D. T. Wise, ‘A combination theorem for special cube complexes’, *Ann. of Math. (2)* **176**(3) (2012), 1427–1482.
- [Hig76] P. J. Higgins, ‘The fundamental groupoid of a graph of groups’, *J. Lond. Math. Soc. (2)* **13**(1) (1976), 145–149.
- [Hru10] G. C. Hruska, ‘Relative hyperbolicity and relative quasiconvexity for countable groups’, *Algebr. Geom. Topol.* **10**(3) (2010), 1807–1856.
- [Hru] G. C. Hruska, Erratum to: ‘Relative hyperbolicity and relative quasiconvexity for countable groups’, In preparation.
- [HW09] G. C. Hruska and D. T. Wise, ‘Packing subgroups in relatively hyperbolic groups’, *Geom. Topol.* **13**(4) (2009), 1945–1988.
- [HW15] T. Hsu and D. T. Wise, ‘Cubulating malnormal amalgams’, *Invent. Math.* **199**(2) (2015), 293–331.
- [KM12] J. Kahn and V. Markovic, ‘Immersing almost geodesic surfaces in a closed hyperbolic three manifold’, *Ann. of Math. (2)* **175**(3) (2012), 1127–1190.
- [Lub96] A. Lubotzky, ‘Free quotients and the first Betti number of some hyperbolic manifolds’, *Transform. Groups* **1**(1–2) (1996), 71–82.
- [MMP10] J. F. Manning and E. Martínez-Pedroza, ‘Separation of relatively quasiconvex subgroups’, *Pacific J. Math.* **244**(2) (2010), 309–334.
- [MP12] E. Martínez-Pedroza, ‘On quasiconvexity and relatively hyperbolic structures on groups’, *Geom. Dedicata* **157** (2012), 269–290.
- [Osi06] D. V. Osin, ‘Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems’, *Mem. Amer. Math. Soc.* **179**(843) (2006), vi+100.
- [Osi07] D. V. Osin, ‘Peripheral fillings of relatively hyperbolic groups’, *Invent. Math.* **167**(2) (2007), 295–326.
- [SW11] M. Sageev and D. T. Wise, ‘Periodic flats in CAT(0) cube complexes’, *Algebr. Geom. Topol.* **11**(3) (2011), 1793–1820.
- [Sco78] P. Scott, ‘Subgroups of surface groups are almost geometric’, *J. Lond. Math. Soc. (2)* **17**(3) (1978), 555–565.
- [Ser80] J.-P. Serre, *Trees*, (Springer, Berlin, 1980), Translated from the French by John Stillwell.
- [Wis12] D. T. Wise, ‘From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry’, CBMS Regional Conference Series in Mathematics, vol. 117 (Published for the Conference Board of the Mathematical Sciences, Washington, DC, 2012).
- [Wis] D. T. Wise, The structure of groups with a quasiconvex hierarchy, <http://comet.lehman.cuny.edu/behstock/cbms/program.html>, unpublished manuscript.