CHAPTER 1

Overview

Our goal is to describe a stream of geometric group theory connecting many of the classically considered groups arising in combinatorial group theory with right-angled Artin groups. The nexus here are the "special cube complexes" whose fundamental groups embed naturally in right-angled Artin groups.

Nonpositively curved cube complexes, which Gromov introduced merely as a convenient source of examples [**28**], have come to take an increasingly central status within parts of geometric group theory – especially among groups with a comparatively small number of relations. Their ubiquity is explained by Sageev's construction [**65**] which associates a dual cube complex to a group that has splittings or even 'semi-splittings' i.e. codimension-1 subgroups.

Right-angled Artin groups, which at first appear to be a synthetic class of particularly simple groups, have taken their place as a natural target – possibly even a "universal receiver" for groups that are well-behaved and that have good residual properties and many splittings or at least "semi-splittings".

We begin by reviewing nonpositively curved cube complexes and a disk diagram approach to them – first entertained by Casson (above dimension 2). These disk diagrams are used to understand their hyperplanes and convex subcomplexes. While many of the essential properties of CAT(0) cube complexes can be explained using the CAT(0) triangle comparison metric, we have not adopted this viewpoint. It seems that the most important characteristic properties of CAT(0) cube complexes arise from their hyperplanes, and these are exposed very well through disk diagrams – and the view will serve us further when we take up small-cancellation theory.

Special cube complexes are introduced as cube complexes whose immersed hyperplanes behave in an organized way and avoid various forms of self-intersections. CAT(0) cube complexes are high-dimensional generalizations of trees, and likewise, from a certain viewpoint, special cube complexes play a role as high-dimensional "generalized graphs". In particular they allow us to build (finite) covering spaces quite freely, and admit natural virtual retractions onto appropriate "generalized immersed subgraphs" just like ordinary 1-dimensional graphs. The fundamental groups of special cube complexes embed in right-angled Artin groups – because of a local isometry to the cube complex of a naturally associated raag. Since right-angled Artin groups embed in (and are closely allied with) right-angled Coxeter groups, this means that one can obtain linearity and residual finiteness by verifying virtual specialness.

We describe some criteria for verifying that a nonpositively curved cube complex is virtually special – the most fundamental is the condition that double hyperplane cosets are separable. A deeper criterion [**36**] arises from a nonpositively curved cube complex that splits along an embedded 2-sided hyperplane into one or two smaller nonpositively curved cube complexes. Under good enough conditions the resulting cube complex is virtually special:

THEOREM 1.1 (Specializing Amalgams). *Let Q be a compact nonpositively curved cube complex with $\pi_1 Q$ hyperbolic. Let P be a 2-sided embedded hyperplane in Q such that $\pi_1 P \subset \pi_1 Q$ is malnormal and each component of $Q - N_o(P)$ is virtually special. Then Q is virtually special.*

A subgroup H of G is "codimension-1" if H splits G into two or more "deep components" – like an infinite cyclic subgroup of a surface group. In his PhD thesis, Sageev understood that when G acts minimally on a CAT(0) cube complex $\widetilde{X}$, the stabilizers of hyperplanes are virtually codimension-1 subgroups of G. He contributed an important converse to this:

CONSTRUCTION 1.2 (dual CAT(0) cube complex). Given a group G and a collection $H_1, \ldots, H_r$ of codimension-1 subgroups, one obtains an action of G on a *dual CAT(0) cube complex* – whose hyperplane stabilizers are commensurable with conjugates of the H_i.

We review Construction 1.2 in the context of Haglund-Paulin wallspaces, and describe some results on the finiteness properties of the action of G on the CAT(0) cube complex $\widetilde{X}$. The main point is that if we can produce sufficiently many quasiconvex codimension-1 subgroups in the hyperbolic group G, then we can apply Construction 1.2 to obtain a proper cocompact action of G on a CAT(0) cube complex. This is how we prove the following result [**43**]:

THEOREM 1.3 (Cubulating Amalgams). *Let G be a hyperbolic group that splits as $A *_C B$ or $A *_{C^t = D}$ where C is malnormal and quasiconvex. Suppose A, B are each fundamental groups of compact cube complexes. And suppose that some technical conditions hold (and these hold when A, B are virtually special). Then G is the fundamental group of a compact nonpositively curved cube complex.*

A *hierarchy* for a group G is a way to repeatedly build it starting with trivial groups (but sometimes other basic pieces) by repeatedly taking amalgams $A *_C B$ and $A *_{C^t = D}$ whose vertex groups have shorter length hierarchies. The hierarchy is *quasiconvex* if at each stage the amalgamated subgroup C is a finitely generated that embeds by a quasi-isometric embedding, and similarly, the hierarchy is *malnormal* if C is malnormal in $A *_C B$ or $A *_{C^t = D}$.

Taken together, Theorem 1.1 and Theorem 1.3 inductively provide the following target for virtual specialness – a malnormal variant of our main result in Theorem 1.7.

THEOREM 1.4 (Malnormal Quasiconvex Hierarchy). *Suppose G has a malnormal quasiconvex hierarchy. Then G is virtually compact special.*

Cubical Small-cancellation Theory: A presentation $\langle a, b, \dots \mid W_1, \dots, W_r \rangle$ is $C'(\frac{1}{n})$ if for any "piece" P (i.e. a subword that occurs in two or more ways among the relators) in a relator W_i we have $|P| < \frac{1}{n}|W_i|$. For $n \geq 6$ the group G of the presentation is hyperbolic and disk diagram methods provide a very explicit understanding of many properties of G.

The presentation above can be reinterpreted as $\langle X \mid Y_1, \dots, Y_r \rangle$ where X is a bouquet of loops and each $Y_i \to X$ is an immersed circle corresponding to W_i, and the group G of the presentation is then $\pi_1 X / \langle\langle \pi_1 Y_1 \dots, \pi_1 Y_r \rangle\rangle$. We generalize this to a setting where X is a nonpositively curved cube complex and each $Y_i \to X$ is a local isometry. We also offer a notion of $C'(\frac{1}{n})$ small-cancellation theory for such "cubical presentations". The main results of classical small-cancellation theory – Greendlinger's lemma and the ladder theorem (and other results involving annular diagrams) have quite explicit generalizations. In particular, we obtain the following result which generalizes the classification of finite trees: T is either a single vertex, is a subdivided arc, or has three or more leaves:

THEOREM 1.5. *If D is a reduced diagram in a cubical $C'(\frac{1}{24})$ presentation then either D is a single 0-cell or cone-cell, or D is a "ladder" consisting of a sequence of cone-cells, or D has at least three spurs and/or cornsquares and/or shells.*

The undefined terms in Theorem 1.5 are described in Chapter 9, but the reader might wish to take a glimpse at Figure 9.5.

One motivation for introducing a cubical small-cancellation theory is that when the "relators" Y_i also have given wallspace structures, then there are natural walls – and hence usually codimension-1 subgroups – in the group G, generalizing the same phenomenon for $C'(\frac{1}{6})$ groups.

This cubical small-cancellation theory helps to coordinate the proof of the following result:

THEOREM 1.6 (Malnormal Special Quotient Theorem). *Let G be hyperbolic and virtually compact special. Let $\{H_1, \dots, H_r\}$ be an almost malnormal collection of subgroups. There exist finite index subgroups $H_1', \dots, H_r'$ such that $G/\langle\langle H_1', \dots, H_r' \rangle\rangle$ is virtually compact special and hyperbolic.*

Most of our exposition circulates around the proof of Theorem 1.6. Assuming that $G = \pi_1 X$, we first choose a collection of local isometries $Y_i \to X$ with $\pi_1 Y_i = H_i$. We then choose appropriate finite covers $\widehat{Y}_i$ (the H_i' will be $\pi_1 \widehat{Y}_i$) such that the group $\bar{G}$ of $\langle X \mid \widehat{Y}_1, \dots, \widehat{Y}_r \rangle$ has a finite index subgroup $\bar{G}'$ with a malnormal quasiconvex hierarchy (we have hidden a few steps here) that can be obtained by cutting along hyperplanes in a finite cover $\widehat{X}$. Thus $\bar{G}'$ is virtually special by Theorem 1.4.

THEOREM 1.7 (Quasiconvex Hierarchy). *Suppose G is hyperbolic and has a quasiconvex hierarchy. Then G is virtually compact special.*

Proving Theorem 1.7 depends on proving virtual specialness of the amalgamated free products and HNN extensions that arise at each stage of the hierarchy. Given a splitting, say $G = A*_{C^t=D}$, the plan is to find a finite index subgroup G' with an almost malnormal quasiconvex hierarchy and conclude by applying Theorem 1.4. To do this, we verify separability of C by applying Theorem 1.6 to quotient subgroups of C using an argument inducting on $\mathrm{Height}(G, C)$. This idea of repeatedly quotienting with an induction on height was independently discovered by Agol-Groves-Manning.

Generalizations of Theorem 1.7 hold in many (and conjecturally all) cases when G is hyperbolic relative to abelian subgroups. We describe how to deduce these generalizations from Theorem 1.7. The proof of separability essentially involves a generalization of Theorem 1.7 to provide virtually special parabolic quotienting. However cubulating requires some additional work.

1.1. Applications

We describe three notable classes of groups with quasiconvex hierarchies in Chapter 16:

Limit groups have hierarchies given by Kharlamopovich-Miasnikov and by Sela, and are thus virtually special.

Every one-relator group has a Magnus-Moldavanskii hierarchy – and for one relator groups with torsion this hierarchy is a quasiconvex hierarchy. (Though technically one must pass to a torsion-free finite index subgroup.) This resolves Baumslag's conjecture that every one-relator group with torsion is residually finite – indeed they are virtually special and thus linear and have separable quasiconvex subgroups.

For a hyperbolic 3-manifold M with an incompressible surface S the Haken hierarchy of M yields a quasiconvex hierarchy for $\pi_1 M$ provided $\pi_1 S$ is geometrically finite, and so $\pi_1 M$ is virtually special. When the hyperbolic 3-manifold has a geometrically finite incompressible surface, we thus find that $\pi_1 M$ is subgroup separable: Indeed, the geometrically finite subgroups are quasiconvex and hence separable using virtual specialness, and the virtual fiber subgroups are easily seen to be separable, and there are no other subgroups by the Tameness Theorem [**1, 13**]. A second corollary is that when the hyperbolic 3-manifold M is Haken, in the sense that it has an incompressible surface S, then M is virtually fibered. Indeed, either S is a virtual fiber, or it is geometrically finite, and in the latter case $\pi_1 M$ is virtually in a raag and thus virtually RFRS and so Agol's fibering criterion applies [**2**].

1.2. A Scheme for Understanding Groups

The above discussions are instances of partial success in implementing the following "grand plan" for understanding many groups:

- Find codimension-1 subgroups in a group G.
- Produce the dual CAT(0) cube complex $\widetilde{C}$ upon which G acts.
- Verify that G acts properly and relatively cocompactly on $\widetilde{C}$ by examining the extrinsic nature of the codimension-1 subgroups.
- Consequently G is the fundamental group of a nonpositively curved cube complex $C = G\backslash\widetilde{C}$. (Or C is an orbihedron if G has torsion.)
- Find a finite covering space $\widehat{C}$ of C, such that $\widehat{C}$ is special.
- The specialness reveals many structural secrets of G. For instance, G is linear since it embeds in $SL_n(\mathbb{Z})$, and the geometrically well-behaved subgroups of G are separable.

There is much work to be done to determine exactly when the above plan can be applied, and there are certainly groups where the plan is impossible to implement – e.g. any nonlinear group. However, when the plan is successful, it provides a very useful viewpoint. We conclude that in many cases, especially when G has comparatively few relators, we see that:

Though G might arise as the fundamental group of a small 2-complex or 3-manifold, in many cases one should sacrifice this small initial presentation in favor of a much larger and higher-dimensional object that is a nonpositively curved special cube complex, and has the advantage of being far more organized, thus revealing important structural aspects of G.

CHAPTER 2

Nonpositively Curved Cube Complexes

This chapter serves as a quick introduction to nonpositively curved cube complexes. We quickly review the basic definitions in Section 2.1. A variety of examples are provided in Section 2.2. The fundamental examples, which arise from right-angled Artin groups, are described in Section 2.3. While Gromov introduced nonpositively curved cube complexes as a source of examples with metric nonpositive curvature in the sense that the thin triangle comparison condition holds, we have adopted a combinatorial viewpoint here. Metric arguments are not usually critical here, and though a traditional geometric group theory attitude can sometimes coordinate or motivate a proof, it seems to make things messier to work simultaneously in two categories. The key feature that gives nonpositively curved cube complexes their characteristic properties are the hyperplanes which we describe in Section 2.4 and will continually revisit in Chapter 3 and subsequently.

2.1. Definitions

An *n-cube* is a copy of $[-1,1]^n$. Its *faces* are restrictions of some coordinates to ± 1. We regard the faces as cubes in their own right.

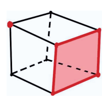

FIGURE 2.1. Two faces in a 3-cube

A *cube complex* X is a cell complex obtained by gluing cubes together along faces. The identification maps of faces are modeled by isometries – so this is entirely combinatorial.

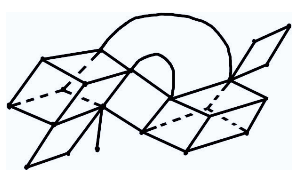

FIGURE 2.2. A cube complex

The *link* of a 0-cube v of X is the simplex-complex whose n-simplices are corners of $(n+1)$-cubes adjacent with v. So link(v) is the "ϵ-sphere" about v in X.

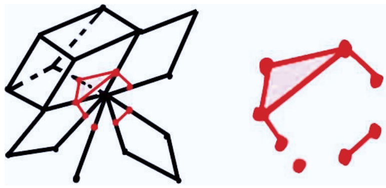

FIGURE 2.3. Link within a cube complex

A *flag complex* is a simplicial complex with the property that $n+1$ vertices span an n-simplex if and only if they are pairwise adjacent. Thus a flag complex is determined completely by a simplicial graph. Note that a graph Γ is flag if and only if girth$(\Gamma) \geq 4$.

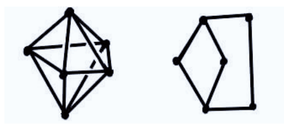

FIGURE 2.4. A 2-dimensional and 1-dimensional flag complex

The cube complex X is *nonpositively curved* if link(v) is a flag complex for each $v \in X^0$. We say $\widetilde{X}$ is a *CAT(0) cube complex* if it is nonpositively curved and simply-connected.

REMARK 2.1. A CAT(0) cube complex also has a genuine CAT(0) triangle comparison metric where each n-cube is isometric to $[-1,1]^n$ however this is not usually the best viewpoint here.

We use the notation $\widetilde{X}$ since a CAT(0) cube complex $\widetilde{X}$ arises as the universal cover of a nonpositively curved cube complex X.

EXAMPLE 2.2 (Not nonpositively curved). The cube complex homeomorphic to an n-sphere obtained by identifying two n-cubes identified along their boundaries is not nonpositively curved for $n \geq 2$. Indeed, the link of each 0-cube is not simplicial.

The cube complex obtained by removing one of the eight 3-cubes around the origin in $\mathbb{R}^3$ is not nonpositively curved. The link of the central 0-cube is isomorphic to the simplicial complex obtained from an octahedron by removing a single open 2-simplex.

2.2. Some Favorite 2-Dimensional Examples

EXAMPLE 2.3 (Dehn Complexes). A link projection is *alternating* if the curves travel alternately above and below at crossings. The projection is *prime* if each embedded circle σ in the plane that intersects the projection P transversely in exactly two noncrossing points has the property that the part of P on either the inside or on the outside of σ consists of a single arc. A projection that is not prime is illustrated on the right of Figure 2.8. (Any link that is both prime and alternating in the usual sense has a projection that is both prime and alternating.)

Let L be a link in S^3. The *Dehn complex* X of L is a square complex that embeds in $S^3 - L$. The 2-dimensional cube complex X has exactly two 0-cubes v_b, v_t which are positioned at the "bottom" and at the "top" of the projection plane. We give the projection the checkerboard coloring. There is a 1-cube for each region of the projection – the 1-cubes associated with black regions are oriented from v_t to v_b and the 1-cubes associated with white regions are oriented from v_b to v_t. There is a 2-cube for each crossing of P, its attaching map is a length 4 path that travels up and down around the crossing (following the boundary of a saddle). We refer the reader to Figure 2.6.

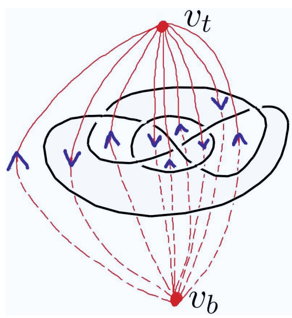

FIGURE 2.5. The 1-skeleton of the Dehn complex

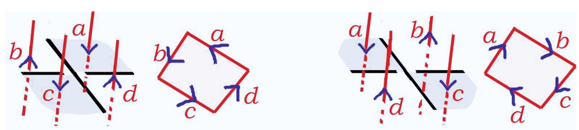

FIGURE 2.6. Squares correspond to crossings

Weinbaum discovered that the link projection is prime and alternating if and only if the Dehn complex X is nonpositively curved, but he formulated this in terms of C(4)-T(4) small-cancellation complexes [**72, 52**]. The explanation we give here is from [**81**].

We first observe that link(v_t) embeds in the projection diagram (with the traditional omitted parts indicating the nature of the crossing points).

To see that X is nonpositively curved exactly when P is prime and alternating, we note that the checkerboard coloring shows that link(v_t) is bipartite so it suffices to verify that there are no 2-cycles. But the two

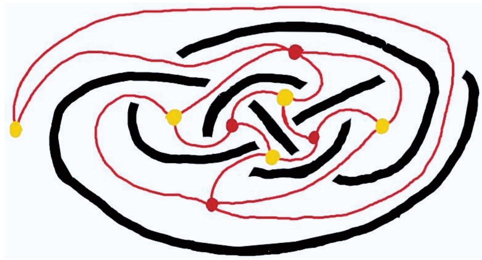

FIGURE 2.7. link(v_t) embeds in the projection diagram

different types of 2-cycles would show that P is either not alternating or not prime.

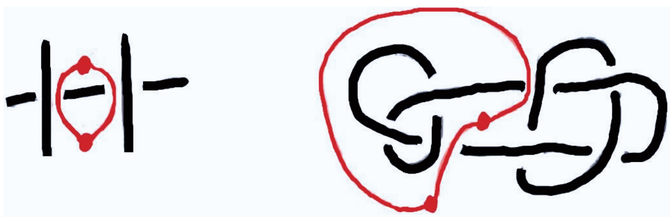

FIGURE 2.8. 2-cycles that indicate that P is not alternating or not prime.

EXAMPLE 2.4 (Graphs of graphs). Let X be a topological space that decomposes as a graph Γ of spaces where each vertex space X_v is a graph and each edge space $X_e \times [-1, 1]$ is the product of a graph and an interval. Suppose the attaching maps $\phi_{e-} : X_e \times \{-1\} \to X_{\iota(e)}$ and $\phi_{e+} : X_e \times \{+1\} \to X_{\tau(e)}$ are combinatorial immersions. Then X is a nonpositively curved square complex. Recall that map is an *immersion* if it is locally injective.

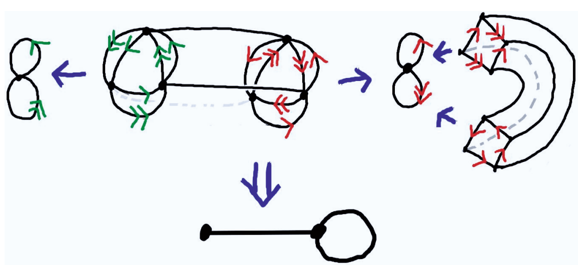

FIGURE 2.9. A graph of spaces with two vertex spaces and two edge spaces.

EXAMPLE 2.5 (Amalgams along $\mathbb{Z}$). Any group of the form $F_2 *_{\mathbb{Z}} F_2$ is isomorphic to the fundamental group of a graph of spaces yielding a nonpositively curved square complex as above. Indeed, after conjugating, we can assume the group is of the form $\langle a,b\rangle *_{U=V} \langle c,d\rangle$ where U is a cyclically reduced word in $a^{\pm1}, b^{\pm1}$, and V is a cyclically reduced word in $c^{\pm1}, d^{\pm1}$. We can think of the words U, V as immersed combinatorial paths in the corresponding bouquets of circles. The group thus arises as a fundamental group of a graph of spaces where each vertex space is a bouquet of two circles, and the edge space is a cylinder attached along the cycles U and V. We then subdivide the vertex spaces to ensure that $|U| = |V|$, and so the cylinder can be compatibly subdivided into squares. This subdivision rarely works for an HNN extension $F_2 *_{\mathbb{Z}^t=\mathbb{Z}}$. Understanding how to deal with such HNN examples was one of the motivations for this research.

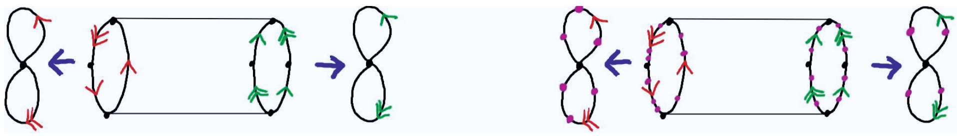

FIGURE 2.10. Subdivide the left and right vertex spaces so that the cylinder edge space has bounding circles of the same length as on the right.

A $\mathcal{VH}$-*complex* is a square complex whose 1-cells are divided into two classes: *vertical* and *horizontal*, and where attaching maps of 2-cells are length 4 paths that alternate between vertical and horizontal edges.

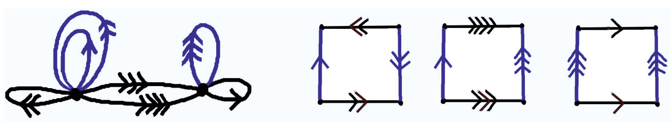

FIGURE 2.11. A $\mathcal{VH}$-complex

When X arises as a graph of spaces as in Example 2.4, then link(v) is a bipartite graph, and X is a nonpositively curved $\mathcal{VH}$-complex. The vertical 1-cells are in the vertex spaces, and the horizontal 1-cells project to edges of the underlying graph. Note that the nonpositive curvature holds because the attaching maps are immersions. We now draw attention to two classes of graphs of spaces that arise from restrictions on the nature of the attaching maps of edges spaces:

EXAMPLE 2.6 (Complete Square Complexes). When all attaching maps $\phi_{e\pm}$ are covering maps of graphs, then X is a *complete square complex* which means that each link(v) is a complete bipartite graph. In this case, the universal cover $\widetilde{X}$ is isomorphic to the cartesian product of two trees. These can be surprisingly complicated and we refer to [**82**], and to [**47**] for some simple such examples. Most notably, Burger-Mozes gave such examples where $\pi_1 X$ is infinite simple [**12**].

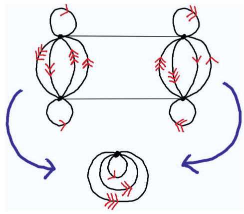

FIGURE 2.12. A complete square complex with 6 squares.

EXAMPLE 2.7 (Clean $\mathcal{VH}$-complexes). When all attaching maps $\phi_{e\pm}$ are combinatorial embeddings, then X is *clean*. Sometimes X might not be clean but has a finite covering space that is clean.

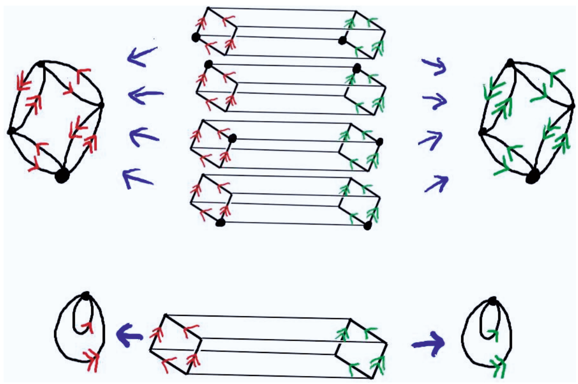

FIGURE 2.13. A $\mathcal{VH}$-complex for a genus 2 surface, together with a degree 4 clean cover

REMARK 2.8. When X is a compact nonpositively curved $\mathcal{VH}$-complex, then X has a (single or) double cover $\widehat{X}$ that splits as a graph of spaces as in Example 2.4 and Figure 2.9 and so $\pi_1\widehat{X}$ splits as a corresponding graph of groups. In this case X has a clean finite cover precisely when the edge groups of $\pi_1\widehat{X}$ are separable [**81**].

2.3. Right-Angled Artin Groups

The *right-angled Artin group* (*raag*) or *graph group* $G(\Gamma)$ associated to the simplicial graph Γ has the following presentation:

$$\Big\langle\ g_v : v \in \text{Vertices}(\Gamma)\ \Big|\ [g_u, g_v] : (u,v) \in \text{Edges}(\Gamma)\ \Big\rangle \tag{2.1}$$

EXAMPLE 2.9.

(1) $G($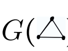$) \cong \mathbb{Z}^3$
(2) $G($$) \cong F_3$
(3) $G($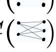$) \cong F_2 \times F_3$
(4) $G($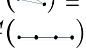$) \cong \pi_1(S^3 -$ 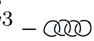$)$

When girth$(\Gamma) \geq 4$, the standard 2-complex R of Presentation (2.1) is already a nonpositively curved square complex. In general, we must also add higher dimensional cubes. We define $R(\Gamma)$ to be the cube complex obtained from the standard 2-complex of Presentation (2.1) by adding an n-cube for each collection of n pairwise commuting generators – i.e. for each n-clique in Γ. $R(\Gamma)$ is often called the *Salvetti complex* [**16**].

We record the above construction as follows:

THEOREM 2.10. *For each* [*finite*] *simplicial graph* Γ*, there is a* [*compact*] *nonpositively curved cube complex* $R(\Gamma)$ *such that the associated raag* $G(\Gamma)$ *can be identified with* $\pi_1 R(\Gamma)$*. In particular, the standard 2-complex of the defining presentation for* $G(\Gamma)$ *equals the 2-skeleton of* $R(\Gamma)$*.*

EXAMPLE 2.11. A 3-dimensional example that is worth thinking through to illustrate the construction proving Theorem 2.10 is indicated in Figure 2.14. Notice that $R(\Gamma)$ has only one 0-cell v, and that link(v) contains two copies of Γ: one "ascending" and one "descending" and has additional simplices corresponding to corners of cubes that are mixed.

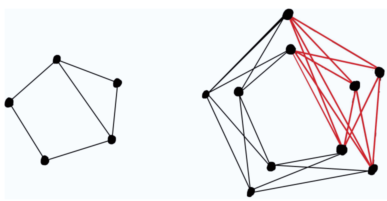

FIGURE 2.14. Γ is on the left, and link(v) is on the right. Note that link(v) contains an octahedron whose eight 2-simplices correspond to the eight corners of the added 3-cube.

PROPOSITION 2.12. *Raags have the following properties:*

(1) *They are residually torsion-free nilpotent* [**22**].
(2) *They are residually finite rational solvable* [**2**].
(3) *They are linear* [**46**].
(4) *They embed in right-angled Coxeter groups and hence in* $SL_n(\mathbb{Z})$ [**44, 20**].

We refer to Charney's survey paper for more about right-angled and other Artin groups [**15**].

2.4. Hyperplanes

A *midcube* is a subspace of a cube obtained by restricting one coordinate to 0. A *hyperplane* is a connected subspace of a CAT(0) cube complex that intersects each cube in a single midcube or in $\emptyset$. The hyperplane H is said to be *dual* to each 1-cube that it intersects.

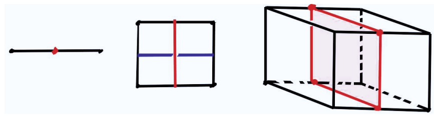

FIGURE 2.15. Midcubes in a 1-cube, 2-cube, and 3-cube

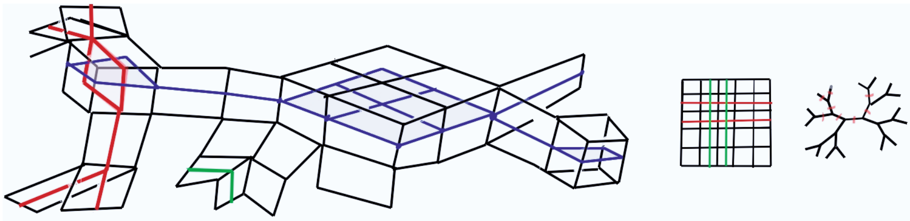

FIGURE 2.16. Some hyperplanes in a 3-dimensional CAT(0) cube complex, and in the plane, and in a tree.

We record some fundamental properties of hyperplanes given in [**65**]. We will describe proofs in Chapter 3.

THEOREM 2.13. *Let $\widetilde{X}$ be a CAT(0) cube complex*

(1) *Each midcube lies in a unique hyperplane of $\widetilde{X}$.*
(2) *A hyperplane H of $\widetilde{X}$ is itself a CAT(0) cube complex (regard midcubes as cubes).*
(3) *The cubical neighborhood $N(H) \cong H\times[-1,1]$ is a convex subcomplex of $\widetilde{X}$ called the* carrier *of H.*
(4) *$\widetilde{X} - H$ consists of two components.*

We adopt a highly combinatorial view here: A *geodesic* in $\widetilde{X}$ is a combinatorial edge path in the 1-skeleton of $\widetilde{X}$ that is a geodesic in $\widetilde{X}^1$ with respect to the graph metric. A subcomplex $\widetilde{Y}$ of $\widetilde{X}$ is *convex* if for any geodesic γ in $\widetilde{X}$, if the endpoints of γ are in $\widetilde{Y}$ then $\gamma \subset \widetilde{Y}$.

REMARK 2.14. It is natural to use the L^1 metric on a CAT(0) cube complex $\widetilde{X}$, where distance equals the length of the shortest path that is piecewise parallel to axes. The inclusion $\widetilde{X}^1 \subset \widetilde{X}$ is then an isometric embedding on $\widetilde{X}^0$, where we use the graph metric on $\widetilde{X}^1$. We note that $\mathsf{d}(p,q) = \#(p,q)$ for $p,q \in \widetilde{X}^0$, where $\#(p,q)$ denotes the number of hyperplanes separating p and q.