Combinatorial Morse Theory

> Groups, as men, will be known by their actions.

In this essay we aim to give an account of Bestvina and Brady's landmark paper [BB], where they develop a version of Morse theory without calculus. Their main theorem is a triumph for geometric and topological methods in group theory because it sheds light on the geometric nature of the a priori algebraic finiteness properties and gives a very different and tractable approach to determining said finiteness properties: for example in homology the existence of various long exact sequences makes calculating homology much easier than trying to find strange resolutions or prove that none exist. As the epigraph suggests, all of these benefits are reaped by considering suitable actions on geometric objects.

The exposition begins with a review of metric geometry taken from [BH] before moving on to the proof of the Main Theorem of [BB] and a discussion of some immediate further applications, though the paper has been cited over 500 times in the literature so naturally the discussion is far from exhaustive. The paper is also known for proving that at most one of the Whitehead conjecture and Eilenberg-Ganea conjecture is true, though there is now an alternative approach in [Ho99].

The author would like to express his heartfelt thanks to Professor Henry Wilton for a great many insightful comments and helpful discussions, which have been indispensible in the study of the material and the writing of this essay.

Contents

1 The CAT(κ) conditions

The exposition in sections 1, 2, and 5 follows [Wi].
A wide variety of conditions on metric spaces that try to generalise the notion of curvature on manifolds have been given by mathematicians, most of them being equivalent to each other in some sense. One of the most common, and the one that will feature in this essay, is the $CAT(\kappa)$ condition, based on thinness of the triangles: a standard exercise for students who have first come across the hyperbolic plane is to prove that for any triangle, any side lies in some neighbourhood of the other two sides, with the size of the neighbourhood independent of the triangle. In the Euclidean plane, this is incredibly false. To make this idea precise will require some background. [1] Sections 1 and 2 will develop some general theory of $CAT(\kappa)$ spaces that will be used later.
Denote by M_κ the unique connected, complete, 2-dimensional Riemannian manifold of constant curvature κ, and D_κ its diameter. Note

- $M_1 \equiv S^2, D_1 = \pi$
- $M_0 \equiv \mathbb{R}^2, D_0 = \infty$
- $M_{-1} \equiv \mathbb{H}^2, D_{-1} = \infty$

and the other M_κ are just scaled copies of these.
Let (X, d) be a complete, geodesic, proper metric space. Given two points p and q, denote a geodesic between them by $[p, q]$, and define a triangle with vertices p, q, r to be the union of the geodesics $[p, q] \cup [q, r] \cup [r, p]$. Note that this is an egregious abuse of notation since the geodesics, and hence the triangle, may not be unique.
Let $\Delta = \Delta(x_1, x_2, x_3)$ be a geodesic triangle in X, and suppose that it has perimeter at most $2D_\kappa$ (so that the triangle inequality can hold). Then there is, up to isometry, a unique triangle $\overline{\Delta} = \overline{\Delta}(\overline{x_1}, \overline{x_2}, \overline{x_3}) \subseteq M_\kappa$ with $d_X(x_i, x_j) = d_{M_\kappa}(\overline{x_i}, \overline{x_j})$, called the *comparison triangle* for Δ. There is a natural way to define a surjection $\overline{\Delta} \to \Delta$ such that restricting to an edge gives an isometry, so given $y \in [x_i, x_j]$ there is a well-defined *comparison point* $\overline{y} \in [\overline{x_i}, \overline{x_j}]$. Since geodesics in X may intersect each other (the map is an isometry only when restricted to an edge), the map is not necessarily an injection, which means that a point $y \in \Delta$ may have up to three comparison points (but this is still well-defined because we specify the edge as well as the point).

Definition. A complete, geodesic metric space (X, d) is CAT(κ) if, for any geodesic triangle Δ of perimeter at most $2D_\kappa$ and any $p, q \in \Delta$, the comparison points $\overline{p}, \overline{q} \in \overline{\Delta}$ satisfy $d_X(p, q) \leq d_{M_\kappa}(\overline{p}, \overline{q})$
If X is locally CAT(κ), it is said to be of curvature at most κ. A locally CAT(0) space is called non-positively curved.

Imagine for example a pair of geodesics emanating from a point $x \in X$. If, locally, the points on the geodesics near x are closer than their comparison points in the plane, that suggests that the space is somehow pushing them closer. Also note that some authors, e.g. [BH], do not require X to be complete, and refer to spaces with the additional requirement that they are complete as Hadamard spaces. It is also proved in [BH] that a space is CAT(κ) iff it is CAT(κ') for all

[1] Technically the condition described is known as the hyperbolic plane being δ-slim.

$\kappa' > \kappa$.
Examples of CAT(0) spaces include

- (real) inner product spaces
- trees, which are in fact CAT(κ) for any κ
- products $X \times Y$ of CAT(0) spaces X, Y when the product is given the l^2 norm.

The next lemma proves convexity of the metric, an important property of CAT(0) spaces in general.

Lemma 1. *Let X be a CAT(0) space, $\gamma, \delta : [0,1] \to X$ be geodesics. Then*

$$d(\gamma(t), \delta(t)) \leq (1-t)d(\gamma(0), \delta(0)) + td(\gamma(1), \delta(1))$$

Proof. If $\gamma(0) = \delta(0)$ then apply the CAT(0) inequality followed by the fact that in Euclidean space, $d(\overline{\gamma(t)}, \overline{\delta(t)}) = td(\overline{\gamma(1)}, \overline{\delta(1)})$.
For the general case, divide the quadrilateral formed by $\gamma(0), \delta(0), \gamma(1), \delta(1)$ into two triangles and apply the previous case. □

Corollary 2. *If X is a CAT(0) space it is uniquely geodesic, i.e. there is a unique geodesic joining each pair of points*

From this we deduce:

Lemma 3. *Let X be a proper, uniquely geodesic space. Then the geodesics vary continuously with their endpoints.*

Proof. Suppose $x_n \to x$ and $y_n \to y$. Let $\gamma_n = [x_n, y_n]$, $\gamma = [x, y]$. WLOG take the domains of all geodesics to be $[0,1]$.
Claim: $\gamma_n \to \gamma$ pointwise

Proof. If not, then there is a t_0 and an $\epsilon > 0$ such that $d(\gamma(t_0), \gamma_{n_i}(t_0)) > \epsilon$ for some subsequence n_i. Convexity of the metric implies that all the γ_n are contained in a closed, hence compact, ball B of radius R, so $d(\gamma_n(s), \gamma_n(t)) < 2R|s-t|$ for all s, t. The γ_n are an equicontinuous family of maps $[0,1] \to B$, and hence there is a subsequence of the γ_{n_i} that converges uniformly to a geodesic from x to y (on every interval apply continuity to show it is an isometric embedding, hence geodesic), which by uniqueness is γ. □

The convergence is uniform: let $\epsilon > 0$. If $d(\gamma_n(t_0), \gamma_(t_0)) < \epsilon/3$ then $d(\gamma_n(t), \gamma_(t)) < \epsilon$ whenever $|t - t_0| < \frac{\epsilon}{6R}$. Apply compactness of $[0,1]$. □

Proposition 4. *Any CAT(0) space X is contractible.*

Proof. For each $y \in X$ let $\gamma(\cdot, y)$ be the unique geodesic from x to y. Then $F : X \times [0,1] \to X, (y,t) \mapsto \gamma(1-t, y)$ is a homotopy equivalence from X to $\{x\}$. □

A sign that $CAT(0)$ is a suitable generalisation to metric spaces of the notion of non-positive curvature is that classical theorems of Riemannian geometry have appropriate generalisations:

Theorem 5. *(Hopf-Rinow) Let X be a length space. If X is complete and locally compact then it is proper and geodesic.*

Theorem 6. *(Cartan-Hadamard) Let X be a complete, locally compact, connected length space of non-positive curvature. Then the universal cover $\tilde{X}$, with the induced length metric, is $CAT(0)$.*

For proofs of these, see [Ba90].

2 Gromov's Link Condition

The link condition is a first indication that cube complexes can be nice to work with. It will reduce the question of existence of a non-positively curved metric on a cube complex, a priori a non-trivial task potentially involving lots of intricate geometric arguments, to a purely combinatorial check that even a computer can do.

Definition. A locally finite cell complex X is *Euclidean* if every cell is isometric to a convex polyhedron in Euclidean space and the attaching maps are isometries from the lower-dimensional cell to a face of the new cell.

Any such X inherits a length metric which is proper and geodesic by Hopf-Rinow. Note that the torus, as the quotient of a square, and the 2-sphere (as a single cube), are cube complexes. The torus has a flat embedding into $\mathbb{R}^4$ so can be given a non-positively curved metric, but the sphere cannot. Apart from the moral obstruction that spherical geometry should be very far away from anything remotely resembling hyperbolic geometry, the sphere is simply connected so by Cartan-Hadamard, if it were non-positively curved it would be CAT(0) and hence contractible.

Definition. Let X be a geodesic space. The *link* of a point $x_0 \in X$, denoted $Lk(x_0)$, is the space of unit-speed geodesics $\gamma : [0, a] \to X$ with $\gamma(0) = x_0$, modulo the equivalence relation that $\gamma_1 \sim \gamma_2$ iff they coincidence on some interval $[0, \epsilon)$ with $\epsilon > 0$. [2]

The link of a point is really just a neighbourhood of the point, capturing the behaviour very near the point. $Lk(x_0)$ is a cell complex as well: the intersection of $Lk(x_0)$ with a cell of X of dimension n is a cell of dimension $n - 1$. For example, in the torus constructed as above the link of the unique vertex looks like

[2]An important theorem in mathematics is that mathematicians can't agree on definitions. An alternative definition of the link for vertices of the complex, which may be more intuitive, is as follows: Let X be a Euclidean complex and v be a vertex of X. Let $\epsilon > 0$ be much smaller than the length of the shortest 1-cell attached to v (which exists by local finiteness). Then the *link* of v is

$$Lk(v) = S_v(\epsilon) = \{x \in X : d(x, v) = \epsilon\}$$

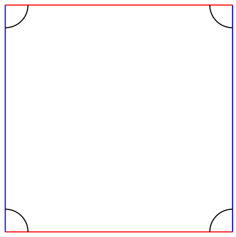

so is S^1. At the corner of a cube, the link is also (homeomorphic to) S^1 but looks a bit different:

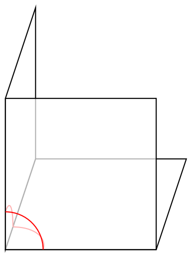

If one were to flatten the cube however, it wouldn't be possible to do so isometrically while keeping the S^1 intact. This is because the total length of the link is too short, or phrased differently, the angle is too small. This is a sign of positive curvature, and in fact the angle will be a useful metric. This can be generalised to any metric space, but since the complex is Euclidean, there is a cheat.
On each cell, the link is part of a sphere, so has a natural spherical metric, which is a length metric. These glue together to a length metric on $Lk(x_0)$. This metric is denoted by $\angle_{x_0}$, which is of interest due to the following theorem:

Theorem 7. *(Gromov's link condition) Any Euclidean complex X is non-positively curved iff $Lk(x_0)$ is $CAT(1)$ for every $x_0 \in X$.*

For a proof of this, see Chapter II.5 of [BH]

2.1 Injectivity radius and systole

We will need a technical lemma establishing when local to global lifting holds for general CAT(κ) spaces.

Definition. Let X be a geodesic metric space. The *injectivity radius* of X is the smallest $r \geq 0$ such that there are distinct geodesics in X with common endpoints of length r. The *systole* of X is the length of the smallest isometrically embedded circle in X.

An isometrically embedded circle gives distinct geodesics with common endpoints, so the systole is at least twice the injectivity radius.
In the more general setting of curvature at most κ, some of the previous results hold but crucially convexity of the metric doesn't hold when $\kappa > 0$. Not all is lost, since short geodesics are still unique.

Proposition 8. *Let X be a compact geodesic metric space of curvature at most κ. Then X is not $CAT(\kappa)$ if and only if it contains an isometrically embedded* *circle of length less than $2D_{(\kappa)}$. If it does, then it contains a circle of length equal to twice the injectivity radius of X; in particular, twice the injectivity radius is equal to the systole.*

2.2 Cube complexes

Definition. A Euclidean complex is a *cube complex* if every cell is isometric to a cube.

As mentioned before, the link of the vertex of a complex is again a (simplicial) complex. Denote by $\square^n$ the regular $n-$cube in $\mathbb{R}^n$ with vertex at the origin, and edges defined by the unit basis vectors. In the case of a regular $n-$cube, the link is a subset of the unit sphere S^{n-1} which is homeomorphic to an $(n-1)-$simplex, so every edge in the link of a cube complex has length $\frac{\pi}{2}$ in the inherited spherical metric. These are known as *all-right spherical* simplicial complexes,

Definition. A simplicial complex L is *flag* if, for every $k \geq 2$, whenever $K \subseteq L^{(1)}$ is a subcomplex of the 1-skeleton that is isomorphic to the 1-skeleton of an $n-$simplex, there is an $n-$simplex Σ in L whose 1-skeleton $\Sigma^{(1)} = K$.

Informally, everything that can be filled in has been filled in. We are finally in a position to reduce geometry to a computer check.

Theorem 9. *An all-right spherical simplicial complex L is $CAT(1)$ iff it is flag.*

Note that the barycentric subdivision of any simplicial complex is flag, so there is no topological obstruction.

Proof. Since the link of a vertex in an all-right spherical simplicial complex is again an all-right spherical simplicial complex, it is natural to try to induct on the largest dimension of any cube in the cube complex[3]. We will need the fact that L is locally CAT(1) iff the link of every vertex is CAT(1). The proof of Gromov's link condition goes through to show this fact. The base case of a 0-dimensional cube complex, a discrete set of points, is trivial.

Suppose L is CAT(1). Links are also CAT(1) and therefore, by induction, are flag. Suppose $K \subseteq L^{(1)}$ is isomorphic to the $(n-1)$-skeleton of an $n-$simplex and let v be a vertex of L contained in K. The link $Lk(v)$ is flag, and $Lk(v) \cap K$ is an $(n-2)-$simplex, which bounds an $(n-1)-$simplex in $Lk(v)$. Therefore, K bounds an $n-$simplex in L.

For the converse, suppose that L is flag. Links of vertices are also flag and so, by induction, are CAT(1). Therefore L is of curvature at most one by the Link Condition. By Proposition 8, it remains to show that L has no isometrically embedded, locally geodesic circle of length less than 2π. Suppose therefore that γ is such a locally geodesic circle.

Suppose that $x \in L$ and that γ intersects $B_x(\frac{\pi}{2})$. As before, fix $\overline{x}$ in S^2 and let $\overline{\gamma}$ be the development of γ into S^2. Then $\overline{\gamma} \cap B_{\overline{x}}(\frac{\pi}{2})$ has length π, and it follows that this is also the length of the intersection of γ with $B_x(\frac{\pi}{2})$.

Let u, v be vertices of L such that γ intersects $B_u(\frac{\pi}{2})$ and $B_v(\frac{\pi}{2})$. Because γ is of length less than 2π, it follows from the previous paragraph that some

[3]One might worry whether this exists. By local finiteness, combined with the CAT(1) condition meaning that only small triangles need to be considered, this will exist in some neighbourhood of a given vertex, which is sufficient.

point of γ is contained $B_u(\frac{\pi}{2}) \cap B_v(\frac{\pi}{2})$, and so u and v are distance less than π apart. Therefore $d(u, v) = \frac{\pi}{2}$. So the set of vertices of every open simplex that γ touches span the 1-skeleton of a simplex and hence span a simplex, because L is flag. So γ is contained in a simplex, which is absurd. $\square$

A cube complex X whose links are all $CAT(1)$ spaces is non-positively curved by the above and Gromov's link condition, so its universal cover also has a $CAT(0)$ metric and is therefore contractible, meaning that this cube complex is in fact a $K(G, 1)$ for $G = \pi_1(X)$. Note that in the last paragraphs of the above proof, the following lemma has been proved:

Lemma 10. *Let v be a vertex in an all-right spherical complex and let B be the ball of radius $\frac{\pi}{2}$ about v. Let $x, y \in \partial B$, and let γ be a geodesic from x to y which intersects the interior of B. Then γ has length at least π.*

Definition. A subcomplex M of an all-right piecewise spherical complex N is said to be convex if whenever $a, b \in M$ satisfy $d_N(a, b) < \pi$, the geodesic $[a, b]$ is contained in M.

Definition. A subcomplex $M \subset N$ is said to be *full* if whenever a set of vertices of M spans a simplex $\tau \subset N$, then $\tau \subset M$.

Notice how similar this is to the condition of being flag.

Lemma 11. *Let $M \subset N$ be a full subcomplex of an all-right spherical simplicial complex N. Then M is convex in N*

This lemma will be used later to determine the homotopy type of certain complexes in the proof of the Main Theorem.

Proof. Suppose that $M \subset N$ is full, and that $a, b \in M$ satisfy $d_N(a, b) < \pi$. We have to show that the geodesic $[a, b]$ is contained in M. Let σ_a (respectively σ_b) denote the minimal simplex of N which contains a (respectively b). Note that σ_a and σ_b are contained in M.
If $a = b$ the result is trivial. If $a \neq b$ the set of all simplices σ which intersect $[a, b]$ nontrivially in their interiors is non-empty.

Claim. The vertices of such σ are contained in the union of the set of vertices of σ_a and the set of vertices of σ_b.

Proof of claim. Suppose σ is a simplex of N whose interior intersects $[a, b]$ nontrivially, and let $v \in \sigma$ be a vertex. The path $[a, b]$ intersects the open star about v (= the open ball in N of radius $\pi/2$ about v). Now if both $d_N(a, v)$ and $d_N(v, b)$ are at least $\pi/2$ the previous lemma implies that the length of $[a, b]$ is at least π contradicting the assumption $d_N(a, b) < \pi$. Thus one of $d_N(a, v)$ and $d_N(v, b)$ is strictly less than $\pi/2$. Suppose $d_N(a, v) < \pi/2$. Then either $a = v$ or a lies in the interior of a simplex with vertex v. In either case, v is a vertex of σ_a, and we are done. $\square$

Hence, $\sigma \subset M$ by fullness. But $[a, b]$ is contained in the union of such σ, and so $[a, b] \subset M$ as required. $\square$

Observe that the claim is a natural way to use the condition of fullness, much like checking the existence of a simplex in a flag complex is reduced to checking the existence of the $1-$skeleton.

Definition. Let L be a simplicial complex, and σ a simplex of L. Denote by $St(\sigma, L)$ the closed star of σ, i.e. the subcomplex of L consisting of all simplices which contain σ. Denote by $St'(\sigma, L)$ the subcomplex of L consisting of all simplices which contain a face of σ.

Proposition 12. *Let L be a flag complex equipped with the all right $CAT(1)$ metric and let σ be a simplex of L*

1. *$St'(\sigma, L)$ is contractible*
2. *Let τ be another simplex of L such that $d(a, b) < \frac{\pi}{2}$ for any $a \in \tau$ and $b \in \sigma$. Then $\sigma \cap \tau \neq \varnothing$.*

Proof. 1: For each simplex ρ in L, the closed star $St(\rho, L)$ is contractible as it is a union of simplices, any subcollection of which intersects in a common (contractible) face containing ρ.
Given a set $\{v_1 \dots v_n\}$ of vertices of σ let $< v_1 \dots v_n >$ denote the face of σ that they span. Since L is flag,

$$St(v_1, L) \cap \dots \cap St(v_n, L) \subset St(< v_1 \dots v_n >, L)$$

and the reverse inclusion trivially holds, so $St(v_1, L) \cap \dots \cap St(v_n, L) = (< v_1 \dots v_n >, L)$ is contractible. Hence

$$St'(\sigma, L) = \bigcup \{St(v, L) : v \text{ a vertex of } \sigma\}$$

is also contractible.
2: Consider the simplicial map from $St'(\sigma, L)$ to a $1-$simplex (of length $\frac{\pi}{2}$ since L is all right) mapping σ to a vertex and the frontier [4] $Fr(St'(\sigma, L))$ to the other vertex. This is a distance non-increasing simplicial map on each simplex of $St'(\sigma, L)$ so is nonincreasing in the spherical metric.
Now suppose that $a \in \tau$ and $b \in \sigma$ are such that $d(a, b) < \frac{\pi}{2}$, but $\sigma \cap \tau = \varnothing$. Then a geodesic γ from a to b intersects $Fr(St'(\sigma, L))$, and so $\gamma \cap St'(\sigma, L)$ maps onto the 1-simplex. But this implies that $d(a, b)$ is at least the distance in the $1-$simplex, which is $\frac{\pi}{2}$, a contradiction. □

3 Finiteness properties of groups

In this section H is a discrete group, and R is a unital commutative ring with $1 \neq 0$, which will be viewed as a trivial module over the group ring RH. Finiteness represents a degree of control over a mathematical object: when reasoning about finite sets, for example, one can appeal to induction. When generalising to well-ordered sets, again what makes induction work is the fact that any decreasing sequence is finite. In geometric group theory, the groups studied are usually not finite, but finitely generated or even finitely presentable. This section will describe further hierarchies of finiteness conditions algebraically, and the main theorem will show that there is in fact associated geometry.

Definition. A module P is *projective* if and only if for every surjective module homomorphism $f : N \twoheadrightarrow M$ and every module homomorphism $g : P \to M$,

[4] the closure minus the interior, also known as the boundary

there exists a module homomorphism $h : P \to N$ such that $f \circ h = g$
Equivalently, P is projective if and only if every short exact sequence of the form

$$0 \to A \to B \to P \to 0$$

is split exact.
Equivalently, P is projective if and only if it is a direct summand of a free module

Definition. A group H is said to be of type $FP_n(R)$ (also written $H \in FP_n(R)$) if there exists a resolution (exact sequence)

$$P_n \to P_{n-1} \cdots \to P_0 \to R \to 0$$

of the trivial RH module R by finitely generated projective $RH-$modules P_i. H is said to be of type $FP(R)$ if the resolution can be taken to be of the form

$$0 \to P_n \to P_{n-1} \cdots \to P_0 \to R \to 0$$

for some n, i.e. a finite resolution.

Definition. A group is of type $FL_n(R)$, (respectively $FL(R)$), if it satisfies the conditions of the previous definition with all instances of the word 'projective' replaced by the word 'free'.

Since projectives are precisely direct summands of free modules, FP_n and FH_n are equivalent by just taking direct sums with complements, and $FL \implies FP$ since free modules are projective, but the converse is open.
Note that any module has resolutions by free, projective, or injective modules, but showing existence involves taking, say, a free module on the generators bijecting with the generators of the previous module, which may be infinite at some point in the tower, so the crux of the definition is in requiring finiteness of the resolution.
Resolutions of this form are useful because of the following fact: Projective resolution of a module M is unique up to a chain homotopy, i.e., given two projective resolutions $P_0 \to M$ and $P_1 \to M$ of M there exists a chain homotopy between them. This means properties calculated from a specifically chosen resolution, such as the homology, are actually invariants of M and can be used to compare different modules M.

Definition. A group H is said to be of type F_n if there exists an Eilenberg-MacLane space $K(H, 1)$ with finite $n-$skeleton, and of type F if the $K(H, 1)$ is finite.

Lemma 13 (Schanuel)**.** *If $0 \to K \to P \to M \to 0$ and $0 \to K' \to P' \to M \to 0$ are short exact sequences of $R-$modules and P and P' are projective, then $K \oplus P'$ is isomorphic to $K' \oplus P$.*

Proof. Define the following submodule of $P \oplus P'$, where $\phi : P \to M \text{and} \phi' : P' \to M$:

$X = \{(p, q) \in P \oplus P' : \phi(p) = \phi'(q)\}$. The map $\pi : X \to P$, where π is defined as the projection of the first coordinate of X into P, is surjective. Since ϕ' is surjective, for any p in P, one may find a $q \in P'$ such that $\phi(p) = \phi'(q)$.

This gives $(p, q) \in X$ with $\pi(p, q) = p$. Now examine the kernel of the map:
$\ker \pi = \{(0, q) : (0, q) \in X\}$
$= \{(0, q) : \phi'(q) = 0\}$
$\cong \ker \phi' \cong K'$.

We may conclude that there is a short exact sequence $0 \to K' \to X \to P \to 0$. Since P is projective this sequence splits, so $X \cong K' \oplus P$. Similarly, we can write another map $\pi : X \to P'$, and the same argument as above shows that there is another short exact sequence

$$0 \to K \to X \to P' \to 0$$

and so $X \cong P' \oplus K$. Combining the two equivalences for X gives the desired result. □

This is used to prove the following lemma, showing that the finiteness conditions $FP_n(R)$ are generally different:

Proposition 14. *If there exists a resolution*

$$0 \to Z_n \to P_n \to P_{n-1} \cdots \to P_0 \to R \to 0$$

where the P_i are all finitely generated projective RH modules, and Z_n isn't finitely generated over RH, then H is of type $FP_n(R)$ but not $FP_{n+1}(R)$.

Proof. See [Br]. □

Motivated by homological considerations from algebraic topology, there is another family of finiteness conditions:

Definition. A group H is said to be of type $FH_n(R)$ if it acts freely, faithfully, properly discontinuously, cellularly, and cocompactly on a cell complex X with $\tilde{H}_i(X, R) = 0$ for all $i \leq n - 1$.
It is said to be of type $FH(R)$ if the complex X is $R-$acyclic.

It is natural to ask how these various finiteness properties relate to each other.

Lemma 15.

- $H \in FH_n(R) \implies H \in FP_n(R)$
- $H \in FH(R) \implies H \in FP(R)$
- *If H acts freely, faithfully, properly discontinuously, cellularly, and cocompactly on an $(n-1)$ connected cell complex X, then H is of type F_n*

Proof. For the first two, H acts on a cell complex X whose reduced chain complex with coefficients in R is the desired resolution. Note that vanishing of the reduced homology groups is a statement of exactness at the relevant modules in the resolution, and since the RH module structure comes from the action of H on X, so compactness of the quotient X/H gives finite generation of the chain groups when viewed as $RH-$modules
For the last one, X/H has no homotopy groups in dimensions 2 to $n-1$, so it suffices to kill the higher homotopy groups from dimension n up by adding cells of dimension $n+1$ and up, which is a standard way of producing a $K(H, 1)$. Note that the $n-$skeleton is finite throughout the process. □

As can be seen in the proof, one of the most common methods of producing resolutions is by algebraic topology.
Proposition 14 gives a further (topological) method of distinguishing the various finiteness properties:

Lemma 16. *Suppose a group H acts freely, faithfully, properly discontinuously, cellularly, and cocompactly on a cell complex X which satisfies*

- $\tilde{H}_i(X, R) = 0$ *for all* $i \leq n-1$.
- $\tilde{H}_n(X, R)$ *isn't finitely generated as an* $RH-$*module*

Then H is of type $FH_n(R)$ but not $FP_{n+1}(R)$.

3.1 Summary of relations

$$F_n \implies FH_n \iff FL_n \implies FP_n$$
$$F \implies FH \implies FL \implies FP$$

It is unknown whether $FP \implies FL$ or $FL \implies FH$ hold.
It will be shown later that $F \implies FH$ and $F_n \implies FH_n$ for $n \geq 2$ are not reversible, but note that $FH_1 \implies F_1$, e.g. by the previous lemma.
A group of type FL_1 is finitely generated and seen to be of type FH_1 by considering its Cayley complex with respect to a presentation for this generating set. If the group is of type FL_2 then the first homology of the quotient of this Cayley complex is finitely generated. Add 2-cells equivariantly to the Cayley complex kills the first homology and shows the group is of type FH_2. Hence FL_n is equivalent to FH_n for $n = 1, 2$ but for larger n this remains open.

4 Morse functions

In this section Morse functions and their basic properties are introduced. The results on extending homotopy type are analogous to those in the smooth case.

Definition. Let X be a CW-complex. X is said to be an *affine cell-complex* if it has the following extra structure. For every cell e there is a convex polyhedral cell C_e in some fixed $\mathbb{R}^m$ and a *characteristic function* $\chi_e : C_e \to e$ such that the restriction of χ_e to any face of C_e is a characteristic function of another cell, possibly composed by a partial affine homeomorphism of $\mathbb{R}^m$. An admissible characteristic function for e is any function obtained from χ_e by precomposing with a partial affine homeomorphism of $\mathbb{R}^m$.

Definition. Let X be an affine cell complex. A map $f : X \to \mathbb{R}$ is a *Morse function* if it satisfies

- for every cell e of $X f\chi_e : C_e \to \mathbb{R}$ extends to an affine map $\mathbb{R}^m \to \mathbb{R}$ and $f\chi_e$ is constant only when $dim(e) = 0$
- the image of the 0-skeleton is discrete in $\mathbb{R}$

These are discrete analogues of critical points being non-degenerate in the smooth case. For a non-empty closed $J \subset \mathbb{R}$ denote by X_J the set $f^{-1}(J)$ and let $X_t = X_{\{t\}}$, called the *level-sets* of the Morse function f.
The set $X_{(-\infty,t]}$ is called the *sublevel set* corresponding to X_t

Lemma 17. *If $J \subset J' \subset \mathbb{R}$ are connected and $X_{J'} \backslash X_J$ contains no vertices of X, then $X_J \hookrightarrow X_{J'}$ is a homotopy equivalence*

Proof. For each cell e of X and each admissible characteristic function $\chi_e : C_e \to X$ we construct a deformation retraction [5] $H_t^{\chi_e}$ of $C_e \cap (f\chi_e)^{-1}(J')$ to $C_e \cap (f\chi_e)^{-1}(J)$ such that

- if χ_e is precomposed with a partial affine homeomorphism h then $H_t^{\chi_e h} = h^{-1} H_t^{\chi_e} h$
- the restriction of $H_t^{\chi_e h}$ to a face of C_e is the deformation retraction associated to that face.

Observe that if C is a convex cell in Euclidean space and F, G are two disjoint faces with F top dimensional and G either top dimensional or a vertex, then any deformation retraction from $\overline{\partial C \backslash F}$ to G extends to a deformation retraction from C to G. This then allows us to define the deformation retraction inductively on $\dim((e)$. ☐

4.1 Ascending and descending links

Definition. A piecewise linear or PL manifold is a manifold homeomorphic to a simplicial complex that is also equipped with a maximal atlas of charts with transition functions given by piecewise linear functions.

The PL structure on an affine cell complex can be given inductively over skeleta in such a way that the characteristic functions are automatically PL. Let $X^{(i)}$ denote the $i-$skeleton and note $X^{(0)}$ is trivially PL. All attaching maps $\partial C \to X^{(i)}$ are PL since the restrictions to faces are PL by induction, and each cell is a convex Euclidean polyhedron so there is a natural way to give the $X^{(i+1)}$ skeleton a PL structure, which works.
When $\alpha : A \to B$ is a PL map between polyhedra and $a \in A$ is an isolated point of $\alpha^{-1}\alpha(a)$, then α induces a map $\alpha_* : Lk(a, A) \to Lk(\alpha(a), B)$. In particular, when $\chi_e : C_e \to X$ is an admissible characteristic function of a cell e of an affine cell complex X and w is a vertex of C_e, this allows us to identify the link $Lk(v, X)$ with $\cup\{\chi_{e*}(Lk(w, C_e)) : \chi_e(w) = v\}$

Definition. The ascending link, also denoted $\uparrow -$ link, is

$$Lk_\uparrow(v, X) = \cup\{\chi_{e*}(Lk(w, C_e)) : \chi_e(w) = v \text{and} f\chi_e \text{has a minimum at} w\}$$

The descending link, also denoted $\downarrow -$ link, is

$$Lk_\downarrow(v, X) = \cup\{\chi_{e*}(Lk(w, C_e)) : \chi_e(w) = v \text{and} f\chi_e \text{has a maximum at} w\}$$

These are trivially subcomplexes of $Lk(v, V)$.

Lemma 18. *(Morse) Let $f : X \to \mathbb{R}$ be a Morse function on an affine cell complex as above. Suppose $J \subset J' \subset \mathbb{R}$ are closed and connected, $inf J = inf J'$, and $J' \backslash J$ contains only one point r of f(0-cells). Then $X_{J'}$ is homotopy equivalent to X_J with the copies of $Lk_\downarrow(v, X)$ (v any vertex with $f(v) = r$) coned off. (Note: some vertices in X may have trivial descending links so define the cone of an empty set to be a vertex.)*

[5]This notion of deformation retraction is called a *strong* deformation retraction by some people.

An analogous statement holds when $inf J = inf J'$ is replaced by $sup J = sup J'$ and $Lk_{\downarrow}(v, X)$by $Lk_{\uparrow}(v, X)$, and the same proof works after noticing that f is a Morse function if and only if $-f$ is too, and the ascending links of f are the descending links of $-f$ and vice versa. Henceforth, for all statements which are about ascending and descending links, only one case will be treated since the proof works for the other case as well after using this observation.

Proof. The argument is similar to that of the previous lemma. Since $X_{J' \cap (-\infty, r]} \hookrightarrow X_{J'}$ is a homotopy equivalence by Lemma 17 we may assume that sup $J' = r$. Let $r - \epsilon = sup J$. If a cell e of X has the property that min $f|_e > r$ then e is disjoint from $X_{J'}$. For any admissible characteristic function $\chi_e : Ce \to X$ of any other cell e we construct, inductively on dim e, a deformation retraction of $(f\chi_e)^{-1}(-\infty, r])$ onto the subset

$$(f\chi_e)^{-1}(-\infty, r - \epsilon]) \cup \bigcup \{F : F \text{is a face of} C_e \text{with} f\chi_e(F) \subset (-\infty, r]\}$$

satisfying the naturality properties stated in the proof of Lemma 17.
These deformation retractions induce a deformation retraction of $X_{J'}$ onto X_J with the cones attached as stated in the Lemma. □

Corollary 19. *Let $f : X \to \mathbb{R}$ be a Morse function on an affine cell complex as above. Suppose that $J \subset J' \subset \mathbb{R}$ are nonempty and connected.*

1. *If each ascending and descending link is homologically $n-$ connected, then $X_J \hookrightarrow X_{J'}$ induces isomorphisms in $\tilde{H}_i$ for $i \leq n$ and an epimorphism on $\tilde{H}_{n+1}$*
2. *If each ascending and descending link is simply connected, then $X_J \hookrightarrow X_{J'}$ induces an isomorphism on π_1.*
3. *If each ascending and descending link is connected, then $X_J \hookrightarrow X_{J'}$ induces an epimorphism on π_1.*

Proof. Since the image of the zero skeleton is discrete, one can show this inductively by applying Lemma 17 and Lemma 18 to get the 'shape' and finish off using the Mayer Vietoris sequence for the first part, and the Seifert-van Kampen theorem for the second and third parts. □

5 Right-angled Artin groups

In this section, following both [BB] and [Wi], we introduce an important class of groups called right-angled Artin groups, or RAAGs, and some of their basic properties. The Main Theorem will be about these groups. The appendix will describe related results that show other reason why RAAGs are of interest.

Definition (right-angled Artin group)**.** Let N be a simplicial graph, i.e. a graph as a graph theorist (or maths Olympiad contestant) would consider. Then

$$A_N = \langle V(N) \mid [u, v] = 1 \text{ for all } (u, v) \in E(N) \rangle$$

is the *right-angled Artin group*, or *graph group* of N.

Let L be the unique flag complex with $L^{(1)} = N$. A_N is also denoted by A_L. Note that this group has a natural surjective map ϕ to $\mathbb{Z}$ by mapping each generator to $1 \in \mathbb{Z}$ and we will be interested in finiteness properties of the kernel, denoted H_L.

Example. If N is the discrete graph on n vertices, then $A_N = F_n$. In this case $\Sigma = N$.

Example. If N is the complete graph on n vertices, then $A_N = \mathbb{Z}^n$. In this case Σ is the $n-1$ dimensional tetrahedron.

Definition. Associated to a right-angled Artin group A_Γ is a cube complex S_Γ constructed as follows. Begin with a wedge of circles attached to a point x_0 and labeled by the generators $s_1, \ldots, s_n$. For each edge, say from s_i to s_j in Γ, attach a 2-torus with boundary labeled by the relator $s_i s_j s_i^{-1} s_j^{-1}$. For each triangle in Γ connecting three vertices s_i , s_j , s_k, attach a 3-torus with faces corresponding to the tori for the three edges of the triangle. Continue this process, attaching a $k-$torus for each set of k mutually commuting generators (i.e., generators spanning a complete subgraph in Γ). The resulting space, S_Γ, is called the Salvetti complex for A_Γ. It is clear, by construction, that the fundamental group of S_Γ is A_Γ.

Alternatively, given a simplicial group N, the *Salvetti complex* $\mathcal{S}_N$ is the cube complex defined as follows:

- Set $\mathcal{S}_N^{(2)}$ is the presentation complex for A_N.
- For any immersion of the 2-skeleton of a d-dimensional cube, we glue in a d-dimensional cube to $\mathcal{S}_N^{(2)}$.

Alternatively, we have a natural inclusion $\mathcal{S}_N^{(2)} \subseteq (S^1)^{|V(N)|}$, and $\mathcal{S}_N$ is the largest subcomplex whose 2-skeleton coincides with $\mathcal{S}_N^{(2)}$.
In fact, there is a recipe for getting the link out of N.

Definition. The *double* $D(L)$ of a simplicial complex L is defined as follows:

- The vertices are $\{v_1^+, \ldots, v_n^+, v_1^-, \ldots, v_n^-\}$, where $\{v_1, \ldots, v_n\}$ are the vertices of L.
- The simplices are those of the form $\langle v_{i_0}^{\pm}, \ldots, v_{i_k}^{\pm} \rangle$, where $\langle v_{i_0}, \ldots, v_{i_k} \rangle \in L$.

In [BB] this is called the spherical complex $S(L)$. Note that

- $D(L)$ contains many copies of L, especially L^+, which is spanned by the v_i^+, and L^-, which is spanned by the v_i^-.
- L^+ (and also L^-) is a retract of $D(L)$, using the map that sends $v_i^{\pm}$ to v_i. Note also that $D(L)$ is flag iff L is flag.

Lemma 20. *The link of the unique vertex x_0 of S_L is isomorphic to $D(L)$*

Proof. For $v \in L^{(0)}$, by construction there are precisely two corresponding vertices in $Lk(x_0)$, which are denoted $v^{\pm}$ according to the orientation of the 1-cell labelled by v. A set of vertices $\{v_0 \ldots v_n\}$ spans a simplex in L iff the only face of $[0,1]^{L^{(0)}}$ spanned by the corresponding directions is a cube (using the second definition of the Salvetti complex). This contributes $2^{n+1} n-$cells to $Lk(x_0)$, one for every possible choice of $\pm$ signs for the $n+1$ vertices $\{v_0 \ldots v_n\}$. $\square$

The double of a flag complex is flag so by the link condition S_L is non-positively curved. Due to the observation after Theorem 9 the Salvetti complex is a $K(A_L, 1)$. We will study the kernel H_L by considering a map from S_L to S^1 inducing the map ϕ on π_1 that lifts to a $\phi-$equivariant Morse function.

Theorem 21. *Construct a map $l : S_L \to S^1$ by isometrically mapping each circle in the construction of S_L, representing an oriented generator, to the loop representing $1 \in \mathbb{Z} = \pi_1(S^1)$. Each higher dimensional cube in the construction of S_L is glued on isometrically, so if it has coordinates $(x_1, \ldots, x_n)$ mapping it to $(x_1 + \cdots + x_n) \mod 1 \in \mathbb{R}/\mathbb{Z} = S^1$ defines l on all of S_L. By construction this induces the homomorphism ϕ, and*

1. *The lift of l to universal covers is a ϕ equivariant Morse function $f : X \to \mathbb{R}$.*

2. *All $\uparrow -$ and $\downarrow -$links of X with respect to f are isomorphic to L.*

Proof. By inspection of the definition, l is non-constant on every cube, so its lift f is too. The image of the zero skeleton is just the lifts of the basepoint of the image S^1, i.e. $\mathbb{Z} \subset \mathbb{R}$. Hence f is a Morse function.
At any point in the pre-image of the $0-$skeleton, the link is isomorphic to $D(L)$ by Lemma 20, and moreover the vertices $v^{\pm}$ indicate whether the vertex lies in the ascending link or the descending link. Thus the ascending (resp. descending) link consists of precisely those simplices which are spanned by the v^+ (resp. v^-), and we have already remarked that these are isomorphic to L. $\square$

6 The Main Theorem

In this section we follow [BB] and describe the homotopy type of (sub-)level sets to prove their main theorem.

Theorem 22. *Let L be a finite flag complex, A_L be the associated right angled Artin group, and $\phi : A_L \to \mathbb{Z}$ be the associated homomorphism with kernel H_L.*

1. *$H \in FP_{n+1}(R)$ if and only if L is homologically $n-$connected.*
2. *$H \in FP(R)$ if and only if L is acyclic.*
3. *H is finitely presented if and only if L is simply connected.*

The first two are statements about homology, so it is unsurprising that it is possible to prove them by computing the homology of sub-level sets, as is done in section 7 of [BB]. That calculation is omitted here as all implications will follow from the description of the homotopy type of (sub-)level sets. However, the right to left implications are general facts about Morse functions, unrelated to the structure of a right angled Artin group, so we record the proof here:

Theorem 23. *Suppose a group G acts freely, faithfully, properly discontinuously, cellularly, and cocompactly on a contractible affine cell complex X such that if e is a cell of X with characteristic function χ_e and $g \in G$, then $g\chi_e$ is an admissible characteristic function for $g(e)$. Suppose $\phi : G \to \mathbb{Z}$ is an epimorphism with kernel H, and that $f : X \to \mathbb{R}$ is a $\phi-$equivariant Morse function (when $\mathbb{Z}$ acts on $\mathbb{R}$ by translations).*

1. *If all $\uparrow -$links and $\downarrow -$links are homologically $n-$connected, then $H \in FH_{n+1}(R)$.*
2. *If all $\uparrow -$links and $\downarrow -$links are $R-$acyclic, then $H \in FH(R)$.*
3. *If all $\uparrow -$links and $\downarrow -$links are simply connected, then H is finitely presented.*

Proof. 1: By Corollary 19, for $t < s$ the inclusion $X_{(-\infty,t]} \hookrightarrow X_{(-\infty,s]}$ induces isomorphisms on $\tilde{H}_i$ for $i \leq n$ and an epimorphism in $\tilde{H}_{n+1}$. Since $X = \cup_{r\in\mathbb{Z}} X_{(-\infty,r]}$ is contractible, it is acyclic, and homology commutes with direct limits so $\tilde{H}_i(X_{(-\infty,t]}) = 0$ for $i \leq n$ and all t. Similarly, $\tilde{H}_i(X_{[t,\infty)}) = 0$ for $i \leq n$ and all t. Applying Mayer-Vietoris to $X = X_{(-\infty,t]} \cup X_{[t,\infty)})$ shows that the intersection X_t is homologically $n-$connected. But since H is the kernel, it acts freely and cocompactly on the level set, so is of type $FH_{n+1}(R)$.
2: The proof of 1 shows that X_t is acyclic, so the action of H on the level set is again the desired action. 3: From 1, $\tilde{H}_0(X_t) = 0$ so X_t is connected. By Corollary 19 $X_t \hookrightarrow X$ induces an isomorphism on π_1. Hence X_t is simply connected and H is finitely presented by item 3 of Lemma 15. □

Some immediate applications to distinguishing finiteness conditions:

1. Let $L = S^n$, triangulated as the $(n+1)-$fold join of $0-$spheres. The associated group is F_2^{n+1}, and from the homology of spheres one sees H_L is of type FP_n but not FP_{n+1}. This example first appeared in full generality in [Bi76]

2. So far the ring of coefficients has played no real role. Here is an example where it does: Let p be a prime. Consider the (appropriately triangulated) Moore space L obtained by attaching an $(n+1)-$cell to an $n-$sphere $(n \geq 1)$ via an attaching map of degree p. Writing down the corresponding long exact sequence for a pair in homology shows that L is acyclic over any field of characteristic $\neq p$, but over fields of characteristic p it is only $(n-1)-$connected.

3. Let L be an acyclic nonsimply-connected finite flag complex of dimension 2. Then H_L is FP (over all rings) but it is not finitely presented. Furthermore, the cohomological dimension of H_L is 2 since the level sets are 2-dimensional and acyclic. It is natural to ask if such groups H_L have 2-dimensional Eilenberg-Mac Lane spaces. Thus one obtains a family of potential counterexamples to the Eilenberg-Ganea conjecture. We will explore this idea further after the proof of the main theorem.

6.1 Sheets

Definition. A subspace F of a complete geodesic metric space is called a flat of dimension k (more briefly, a $k-$flat, or just a flat if the dimension is unspecified) if it is isometric to Euclidean space $\mathbb{E}^k$. When X is the universal cover of a Salvetti complex as considered above, define a *sheet* in X to be a flat which is a subcomplex of the pre-image in X of one of the tori in Q_L.

Sheets will be viewed as building blocks for X and will be instrumental in determining the homotopy type of (sub-)level sets in the next section. In general, there may be many more flats than sheets: when L consists of two points, the associated RAAG is F_2 and X is the tree with every vertex of degree 4 with edges directed and labelled a or b. Then the flats will correspond to paths through edges all of the same label, of which there are countably many, but there are uncountably many biinfinite geodesics.
Given a vertex $v \in X$ and $A \subset X$ a union of sheets, we define $C(v, A)$ to be the cone on $Lk(v, A)$, and consider it as a small neighborhood of v in A. Similarly, $C_{\downarrow}(v, A)$ is just the cone on $Lk_{\downarrow}(v, A)$.
There is a natural retraction $r_v : Lk(v, X) \to Lk_{\downarrow}(v, X)$ mapping the vertices of the form v^+ to v^- and fixing the v^-, which extends simplicially to give r_v. This further extends to a map on cones $C(v, X) \to C_{\downarrow}(v, A)$, also denoted by r_v. Some further properties:

Proposition 24. *1. X is covered by sheets.*

2. *The intersection of any collection of sheets is either empty, a vertex, or a sheet.*

3. *All (sub-)level sets of f restricted to a sheet are contractible. Moreover, all $\uparrow -$ and $\downarrow -$links of the restriction are single simplices.*

4. *The retraction r_v preserves sheets through v, i.e. if $A \subset X$ is a union of sheets and $v \in A$, then the restriction of r_v induces retractions $Lk(v, A) \to Lk_{\downarrow}(v, A)$ and $C(v, X) \to C_{\downarrow}(v, A)$, also denoted by r_v. Moreover, $r_v^{-1}(Lk_{\downarrow}(v, A)) = Lk(v, A)$.*

Proof. 1: Q_L is a union of tori.
2: This holds for the intersection of two sheets: if the sheets project to tori which have intersection strictly containing the unique point in Q_L, the projections intersect in some torus T, so the intersection of the sheets lies in the pre-image of T. Now induct on the number of sheets.
3: The restriction of f to a $k-$dimensional sheet is given by the linear map $(x_1, \dots, x_k) \mapsto x_1 + \dots + x_k$ where $\mathbb{R}^k$ is cubed by $\mathbb{Z}^k$. Hence the sub-level sets are half-spaces cut out by a hyperplane, which are contractible.
4: Immediate from definitions. □

6.2 The homotopy type of sublevel sets

In this section X will be built up out of sheets.

Lemma 25. *Let w be a vertex of X and let K be the union of a collection of sheets containing w. Let J be a closed connected interval in $\mathbb{R}$. Then*

1. *K is contractible.*
2. *All $\uparrow$ and $\downarrow$ links of K at vertices apart from w are contractible, and at w they are naturally isomorphic.*
3. *$K_J = K \cap X_J$ is homotopy equivalent to $Lk_\downarrow(w, K)$ when $w \notin X_J$ and is contractible if $w \in X_J$*

Proof. Since the intersection of sheets containing w is again a sheet containing w, K is a cone with conepoint w, hence is contractible.
As before, it suffices to show this for descending links. Let $x \neq w$ be a vertex of K. There is a sheet S containing w and x which is minimal with this property with respect to inclusion (given by intersecting all sheets containing both). Then $Lk_\downarrow(x, K)$ is the union of simplices of the form $Lk_\downarrow(x, S')$ where S' ranges over the sheets in K that pass through x and w, all of which contain $Lk_\downarrow(x, S)$, so the union is contractible.
There is a natural central symmetry defined on each sheet S through w mapping ascending links to descending links of each sheet bijectively, and in a compatible way with respect to inclusions of sheets, so these combine to give an isomorphism between $Lk_\downarrow(x, K)$ and $Lk_\uparrow(x, K)$.
Suppose $w \in X_J$. By Lemma 18 $K_J \hookrightarrow K$ is a homotopy equivalence, so K_J is contractible. If $w \notin X_J$, WLOG assume that w is above X_J, i.e. $f(w) > t$ for all $t \in J$. The proof in the other case is exactly the same. Choose $\epsilon \in (0, 1)$ such that $f(w) - t > \epsilon$ for all $t \in J$, which is possible since J is closed. It follows from Lemma 17 that

$$K_J \hookrightarrow K_{(-\infty, f(w)-\epsilon]} \stackrel{K}{\hookleftarrow}_{f(w)-\epsilon}$$

are homotopy equivalences, so it suffices to show $K_{f(w)-\epsilon}$ is homotopy equivalent to $Lk_\downarrow(w, K)$. Now observe that since K is a union of sheets, $\{Lk_\downarrow(w, S)\}$ and $\{S \cap X_{f(w)-\epsilon}\}$ are closed contractible covers of $Lk_\downarrow(w, K)$ and $K_{f(w)-\epsilon}$, which are both closed under intersections. But there is a 1-1 order preserving correspondence between $\{Lk_\downarrow(w, S)\}$ and $\{S \cap X_{f(w)-\epsilon}\}$, so the sets that they cover, i.e. $Lk_\downarrow(w, K)$ and $K_{f(w)-\epsilon}$, are homotopy equivalent. □

Lemma 26. *Let X be a CAT(0) cube complex, $Q \subset X$ a cube, and $c \in X$ a vertex. Then the minimal distance from c to Q is attained at a unique vertex of Q.*

Proof. Uniqueness: suppose that there are two points $x, y \in Q$ realizing the minimum distance. Cubes are convex in X so the geodesic $[x, y]$ lies in Q. Applying the CAT(0) inequality to the geodesic triangle xyc shows that interior points of $[x, y]$ are closer to c than the endpoints, contradicting the choice of x and y.
Attained at a vertex: Let γ be a distance minimizing geodesic from c to Q, and suppose that γ doesn't end at a vertex of Q. Then γ meets Q in the interior of some face. Let e denote an edge of this face.
For each point $p \in \gamma$ choose a minimal cube Q_p containing p.

Claim. There is a cellular map from the union of all Q_p over $p \in \gamma$ to a $1-$simplex taking γ to an interior point of the simplex.

Proof. Let Q_0 be the minimal cube containing the edge e and the initial segment of γ. Here γ is parametrized so that it starts at Q and ends at c. This cube contains the face of Q at which c originates, and so has a metric product structure, defined as the (l^2) product of the metric on the codimension one face and metric on the edge e. Since the cubes are Euclidean, this gives rise to a notion of dot product. Since c minimizes the distance from c to Q we see that $Q_0 \cap c$ is perpendicular to the $e-$coordinate. The metric product structure induces a codimension-1 metric product structure on the face of intersection of Q_0 and the next cube Q_1 of the collection along the geodesic c. This extends uniquely to a codimension-1 product structure of Q_1 and so on. Finally, we map the $e-$coordinate onto a 1-simplex, collapsing the other coordinates to a point. The geodesic c remains perpendicular to the $e-$coordinate throughout, and so maps to an interior point of the 1-simplex.

□

Since c lies on the geodesic γ, it also gets mapped to an interior point, so can't have been a vertex, contradicting the initial assumption. □

For the next result we need the following more general definition of a link, which can be found in [BH]:

Definition. Let S be a simplex in an abstract simplicial complex K. The link of S in K, denoted $Lk(s, K)$, is the subcomplex of K consisting of those simplices T such that $T \cap S = \varnothing$ and $T \cup S$ is a simplex of K.

One can see that this agrees with the previous definition when S is a single vertex and that this is downward-closed so is indeed a subcomplex.
Now we build up X as a union of sheets and study the change in homotopy type inductively. Choose a base vertex v in X and an ordering $v = v_1, v_2, v_3 \ldots$ of all the vertices of X such that $d(v, v_i) \leq d(v, v_j)$ whenever $i \leq j$. Let K_i denote the union of all sheets through v_i.

Lemma 27. *K_i is a convex subset of X.*

Proof. It suffices to show that geodesics in K_i are also geodesics in X. This proof breaks down into a series of reduction steps.
Step 1: Geodesics which lie inside a cube in K_i are locally geodesics in X, so the places we have to check are where geodesics pass through vertices or pass between cubes through a common face.
Suppose a geodesic in K_i passes through a vertex y. The two directions of the geodesics define points in $Lk(y, X)$. To be a local geodesic, the distance between these two points should be at least π (consider two rays passing through the centre of a sphere in Euclidean space: if the spherical distance between the points on the sphere were less than π one could shorten the path by joining points on the sphere by a line). This holds in K_i, so in this case it suffices to show

$$Lk(y, K_i) \subset Lk(y, X) \text{ is convex (A)}$$

Similarly, if a geodesic in K_i intersects a cube Q (the common face) in a single interior point,it will be a local geodesic in X if the link of Q in K_i is a convex subcomplex of the link of Q in X. Let y be any vertex of Q and let σ_Q be the simplex of $Lk(y, X)$ determined by Q.
There is a natural identification of $Lk(Q, K_i)$ [6] with $Lk(\sigma_Q, Lk(y, K_i))$ as follows: if $T \cup Q$ determines a simplex of K_i (and hence a $|T|$ dimensional face of $Lk(Q, K_i)$), then each subset of vertices of T determines a simplex that contains y so it determines a subset of vertices in $Lk(\sigma_Q, Lk(y, K_i))$ that spans a face.
$Lk(Q, X)$ can similarly be identified with $Lk(\sigma_Q, Lk(y, X))$ so the convexity condition becomes

$$Lk(\sigma_Q, Lk(y, K_i)) \subset Lk(\sigma_Q, Lk(y, X)) \text{ is convex (B)}$$

Step 2: By Lemma 11 it suffices to show

$$Lk(y, K_i) \text{ is a full subcomplex of } Lk(y, X) \text{ (A')}$$

for all vertices $y \in K_i$, and

$$Lk(\sigma_Q, Lk(y, K_i)) \text{ is a full subcomplex of } Lk(\sigma_Q, Lk(y, X)) \text{ (B')}$$

for all cubes $Q \subset K_i$ and vertices $y \in Q$. B' follows from A' by the following observation: Let M be a full subcomplex of a simplicial complex N and $\sigma \subset M$ be a simplex. Then $Lk(\sigma, M)$ is a full subcomplex of $Lk(\sigma, N)$.
Step 3: Observe that if M is a full subcomplex of a simplicial complex N, then the associated double $D(M)$ is a full subcomplex of the $D(N)$, and that $D(Lk_\downarrow) = Lk$ by Lemma 20. Hence it suffices to show

$$Lk_\downarrow(y, K_i) \text{ is a full subcomplex of } Lk_\downarrow(y, X) = L \text{ (A")}$$

Step 4: Here we show property (A") holds. Since X is a union of sheets, $Lk_\downarrow(v_i, K_i) = Lk_\downarrow(v_i, X)$, so A" clearly holds. Let $y \in K_i$ be a vertex other than v_i. Then the geodesic $[v_i, y]$ defines a unique point of $Lk_\downarrow(y, X)$, which is in fact a point in $Lk_\downarrow(y, K_i)$ since K_i is a geodesic cone on v_i. The image under the retraction r_y gives a point $p \in Lk_\downarrow(y, K_i)$. Again because K_i is a geodesic

[6] Technically the definition of links is for simplicial complexes, but since we are working with cube complexes, it is easy to obtain a triangulation and the link is independent of choice of triangulation.

cone on v_i, any simplex of $Lk_\downarrow(y, K_i)$ comes from a sheet through y and v_i, so the simplex must contain p. Thus one sees that $Lk_\downarrow(y, K_i)$ is just the union U of all simplices $\sigma \subset L$ which contain p, and τ is a simplex of U if and only if there is a simplex σ containing both τ and p.
This is full in L: let σ_p denote the minimal simplex of L which contains p. Then σ_p is a face of each simplex containing p. In particular, if $v \in U$ is a vertex, then either $v \in \sigma_p$ or $\{v\} \cup \sigma_p$ spans a simplex. Hence, by induction on j, if $\{w_1 \dots w_j\} \subset U$ is a collection of vertices which span a simplex $\sigma \subset L$ then $\{w_1 \dots w_j\} \cup \sigma_p$ spans a simplex of U. Thus $\sigma \subset U$, so U is a full subcomplex of L. □

Now fix $n > 1$, and let $K = K_n \cap (K_1 \cup \dots K_{n-1})$.

Lemma 28. *K is the union of sheets containing w and at least one v_j with $j < n$.*

Proof. The union is contained in K so it suffices to show that given $x \in K$ there exists a sheet $S \subset X$ through v_n which contains both x and some v_j with $j < n$. Suppose $v \in K$ is a vertex. $v \in K$ implies $v \in K_n \cap K_i$ for some $i < n$. Let K_v denote the union of all sheets through v. Then $v_i, v_n \in K_v$, which is convex by Lemma 27, so contains the geodesic $[v_i, v_n]$.
Let Q denote the minimal cube containing v_n and the end segment of the geodesic $[v_i, v_n]$. By minimality, $Q \subset K_v$, so Q lies in a sheet S through v. Now it suffices to show that one of the vertices of Q is in the set $\{v_1 \dots v_{n-1}\}$.
By construction of the ordering on the v_i, the geodesic $[v_1, v_i]$ is not longer than $[v_1, v_n]$, and so the CAT(0) inequality implies that the distance from v_1 to any interior point of $[v_i, v_n]$ is strictly less than $d(v_1, v_n)$. In particular this is true for points of Q which lie in the interior of $[v_i, v_n]$. Thus v_n isn't the closest point of Q to v_1. By Lemma 26 the closest point is some vertex, and the ordering implies that this vertex is in the set $\{v_1 \dots v_{n-1}\}$.
Now suppose $x \in K$ isn't a vertex. Then $x \in S' \subset K_n \cap K_i$ for some sheet S' and some $i < n$. By the above argument each vertex of S' is contained in a sheet through v_n and some $v_i, i < n$. There are only finitely many sheets through v_n so some sheet S must contain a maximal general position subset of vertices of S'. This must then contain S', otherwise we would be able to find a larger sheet. □

Lemma 29. *$Lk_\downarrow(v_n, K) \cong Lk_\uparrow(v_n, K)$ is contractible.*

Proof. Let a be the point in $Lk(v_n, X)$ determined by the geodesic $[v_n, v]$ and let $b = r(a) \in Lk_\downarrow(v_n, X)$ be the image of a under the retraction r_{v_n}. Let σ denote the smallest simplex of $Lk_\downarrow(v_n, X)$ that contains b.
Let $\mathbb{S}$ denote the collection of simplices τ in $L = Lk_\downarrow(v_n, X)$ such that the sheet through v_n corresponding to τ contains one of the v_i.

Claim. Every face of σ is in $\mathbb{S}$.

Proof of claim. Let S_0 be the sheet through v_n such that $Lk_\downarrow(v_n, S_0) = \sigma$. By Proposition 24 (4) $Lk(v_n, S_0)$ contains a. Let Q be the smallest cube that contains v_n and such that $Lk(v_n, Q)$ contains a. Thus $Q \subset S_0$ and $\sigma = r_{v_n}(Lk(v_n, Q))$. Now we show each vertex of Q which is distance one from v_n is in $\{v_1 \dots v_{n-1}\}$, so σ and all its faces are in $\mathbb{S}$.

Let v' be a vertex of Q which is distance one from v_n. Then the angle at v_n defined by $[v_n, v]$ and $[v_n, v']$ is less than $\pi/2$. Thus the distance from v to a point on the edge $[v_n, v']$ which is close to v_n is less than $d(v, v_n)$. Lemma 26 implies that $d(v, v') = d(v, [v_n, v']) < d(v, v_n)$, so $v' \in \{v_1 \dots v_{n-1}\}$.

□

Claim. $\mathbb{S} = \{\tau \text{a simplex of } L : \tau \cap \sigma \neq \varnothing\}$

Proof of claim. For $\subset$, let τ be an element of $\mathbb{S}$ and let S be the corresponding sheet in K. Then S contains v_i for some $i < n$. In the triangle $\delta(v_n, v, v_i)$ we have $d(v, v_n) \geq d(v, v_i)$ and thus the angle at v_n is $< \pi/2$. Hence $Lk(v_n, S)$ contains a point at distance $< \pi/2$ from a defined above, so $\tau = Lk_{\downarrow}(v_n, S) = r_{v_n}(Lk(v_n, S))$ contains a point at distance $< \pi/2$ from $b = r_{v_n}(a)$. Thus it must intersect σ by Proposition 12 (2).
For the reverse inclusion, if $S_1 \subset S_2$ are sheets and $Lk_{\downarrow}(v_n, S_1) \in \mathbb{S}$ then $Lk_{\downarrow}(v_n, S_2) \in \mathbb{S}$ by definition of $\mathbb{S}$. In particular, if a vertex of a simplex τ of L is in $\mathbb{S}$, then so is τ. If $\tau \cap \sigma \neq \varnothing$ then τ contains a vertex of σ. But all vertices of σ are in $\mathbb{S}$ by the first claim, so $\tau \in \mathbb{S}$. □

By the previous lemma $Lk_{\downarrow}(v_n, K)$and $Lk_{\uparrow}(v_n, K)$ are isomorphic to the simplicial complex determined by $\mathbb{S}$, and the second claim implies that this simplicial complex is $St'(\sigma, L)$, which is contractible by Proposition 12 (1). □

Theorem 30. *X_J is homotopy equivalent to a wedge of L's, one for every vertex of X not in X_J.*

Proof. Write X_J as the increasing union of $X(n) := X_J \cap (K_1 \cup \dots \cup K_n)$ and note $X(n) = X(n-1) \cup (X_J \cap K_n)$. Furthermore,

$$X(n-1) \cap (X_J \cap K_n) = X_J \cap (K_n \cap (K_1 \cup \dots \cup K_n - 1))$$

is homotopic to $Lk_{\downarrow}(v_n, K = K_n \cap (K_1 \cup \dots \cup K_n - 1))$ (by Lemma 25 3), which is contractible by the previous lemma. Lemma 25 3 also implies that $X_J \cap K_n$ is either contractible or homotopy equivalent to L depending on whether $v_n \in X_J$. It follows inductively that $X(n)$ is homotopy equivalent to a wedge of L's, one for every v_i, $i \leq n$, not contained in X_J. Since inclusions $X_(n-1) \hookrightarrow X(n)$ respect the wedge structure, the theorem follows. □

The 'if' direction of all three has been done previously, so now we use this theorem to show the 'only if' direction. We need one more result before proving the Main Theorem.

Proposition 31. *Let H be a finitely presented group. Suppose a group H acts freely, faithfully, properly discontinuously, cellularly, and cocompactly on a connected cell complex Y. Then it is possible to attach to Y finitely many $H-$orbits of $2-$cells so that the resulting complex is simply connected.*

Proof. Since H is finitely presented, its presentation complex P is compact, and H acts freely, faithfully, properly discontinuously, cellularly, and cocompactly on its simply connected Cayley 2-complex $\tilde{P}$. Construct an $H-$equivariant $\alpha : Y^{(2)} \to \tilde{P}$ as follows: fix a vertex v in Y and map it to a vertex w in $\tilde{P}$. Pick representatives for the $H-$orbits of the $1-$cells attached to v and map them to $1-$cells attached to w. This leaves only one possibility for mapping the $0-$cells

attached to these 1−cells, i.e. the neighbours of v. Repeat until all $H-$orbits of 0−cells are mapped. This terminates in finite time since the quotient $H\backslash Y$ is compact so there are only finitely many orbits of 0− and 1−cells, and all orbits are reached since Y is connected. Extend this $H-$equivariantly. Now there is a natural extension of the map to 2−cells, which is well-defined since $\tilde{P}$ is simply connected. Similarly, there is an $H-$equivariant cellular map $\beta : \tilde{P}^{(1)} \to Y^{(1)} \subset Y$.
Attach one orbit of 2−cells to Y for every orbit of 2−cells in $\tilde{P}$, of which there are only finitely many since H is finitely presented. Then β extends to an equivariant cellular map $\tilde{beta} : \tilde{P}^{(2)} \to Y^{(2)}$. The map $Y^{(1)} \times \{0,1\} \to Y$ defined by $(y,0) \mapsto y, (y,1) \mapsto \beta\alpha(y)$ extends to an equivariant cellular map $F : Y^{(1)} \times \{0,1\} \to Y \cup Y^{(0)} \times [0,1] \to Y$ since Y is connected. Attach an orbit of 2−cells to Y for every orbit of 1−cells in Y to extend F to an equivariant cellular map $\tilde{F} : Y^{(1)} \times [0,1] \to Y$. For example, consider when $Y^(1)$ consists of a single edge. Then $Y^{(1)} \times \{0,1\} \to Y \cup Y^{(0)} \times [0,1]$ is the boundary of a square, and the 2−cell added is just the interior of the square.
Let l be a loop in $Y^(1)$. Then $\tilde{F}$ is a homotopy between l and $\beta\alpha(l)$. $\alpha(l)$ is nullhomotopic in $\tilde{P}^{(2)}$, thus so is $\tilde{\beta}\alpha(l) = \beta\alpha(l)$. □

Proof of Main Theorem. 1, 2: Let n be the smallest integer such that $\tilde{H_n}(L) \neq 0$. Then $\tilde{H_n}(X_t)$ is not a finitely generated $RH-$module as there are infinitely many vertices not in the level set X_t, so by Lemma 16 H is not of type FP_{n+1}.
3: If L is not connected, then by 1 H is not finitely generated, let alone finitely presented. If L is connected but not simply-connected, then $\pi_1(X_t)$ is the free product of $\pi_1(L)$'s, one for every vertex not in X_t. Suppose H were finitely presented. By the previous proposition $\pi_1(X_t)$ is generated by the H translates of finitely many loops, and since X is contractible each of these finitely many loops is null homotopic in X. Since a homotopy is a map from a compact space, its image is also compact, hence all (finitely many) will be contained in some $X_{[t-T,t+T]}$. H acts on level sets, so all the $H-$orbits of these loops will be null homotopic in $X_{[t-T,t+T]}$, so $X_t \hookrightarrow X_{[t-T,t+T]}$ induces the trivial map on π_1. Corollary 19 says that the induced map on π_1 is an epimorphism, so $X_{[t-T,t+T]}$ is simply connected. But there are infinitely many vertices lying outside $X_{[t-T,t+T]}$, so if L is not simply connected then $X_{[t-T,t+T]}$ can't be either. □

Remark: In [BG] the authors use Brown's criterion to provide a shorter proof of the main theorem, but do not determine the exact homotopy type of the sublevel sets.

7 Applications

The number of citations of [BB] speaks to the utility of combinatorial Morse theory as a tool in geometric group theory, and it would be well beyond the scope of this essay to give a complete survey of the various directions the theory has grown in. For example, see[Bra99], where the author constructs hyperbolic groups with finitely presented non-hyperbolic subgroups, among many other interesting applications.
Both applications below will involve the Poincare homology sphere, so for completeness' sake we define it.

Definition. The alternating group A_5 embeds in $SO(3)$ via its action on the icosahedron. In particular, it is a discrete subgroup that acts freely and properly discontinuously on $SO(3)$. The *Poincare homology sphere* is the manifold $SO(3)/A_5$.

This is a homology sphere, i.e. a topological space with the same homology as a sphere, and the only homology $3-$sphere with finite fundamental group. For other descriptions of the Poincare homology sphere, see [KS].

7.1 Eilenberg-Ganea versus Whitehead

In [BB] the authors show that at most one of the Eilenberg-Ganea conjecture and the Whitehead conjecture is true using their main theorem, which remains one of the milestones in the history of these notoriously difficult conjectures that have resisted mathematicians' efforts to solve them for decades. See the appendix for a brief introduction to the conjectures.
Given a vertex $v \in X$ we identify the complex L with the descending link $Lk_{\downarrow}(v; X)$. Thus each point $x \in L$ determines a unique geodesic $g_x \subset K_v$ through v. Starting at any vertex v which is above a level set X_t, i.e. $f(v) > t$, the geodesic g_x eventually reaches X_t, which gives a way of projecting links to a level set.

Definition. Let $t \in \mathbb{R}$ and $v \in X$ be a vertex such that $f(v) > t$, and $M \subset L$ be a subcomplex of $L = Lk_{\downarrow}(v, X)$. Define the shadow of M on X_t to be the set

$$S_{v,M} = \bigcup \{g_x : x \in M\} \cap X_t$$

Proposition 32. *Note that $S_{v,L}$ is homeomorphic to L by the map which sends each simplex $\sigma \subset L$ to $S_{v,\sigma} \subset S_{v,L}$. Give L a metric by requiring that each 2-simplex is an equilateral triangle with side lengths equal to $|f(v) - t|$ and then taking the induced path metric. With this metric, the homeomorphism is a quasi-isometry with constants which are independent of $|f(v) - t|$.*

Proof. Metrize $S_{v,L}$ by restricting that of K_v. The quasi-isometry constants are then bounded by a multiple of the cardinality of the 1-skeleton of L. But K_v is a convex subset of X by Lemma 27. Hence we obtain the same quasi-isometry inequality for $S_{v,L}$ as a subset of X. □

In particular, by taking $|f(v) - t|$ to be arbitrarily large, one obtains arbitrarily large copies of L quasi-isometrically embedded into X_t.

Theorem 33. *Let L be a flag triangulation of a spine [7] of the Poincare homology sphere. Then either H_L is a counterexample to the Eilenberg-Ganea conjecture or there is a counterexample to the Whitehead conjecture.*

Proof. The associated X is 3-dimensional so the level set X_0 is 2-dimensional. Since L is acyclic X_0 is acyclic as well, and since H_L acts freely and cocompactly on X_0 it follows that H_L has cohomological dimension 2. If H_L does not have geometric dimension 2 then the Eilenberg-Ganea conjecture is false.
Suppose H_L does have geometric dimension 2. Then there exists a contractible

[7] It is often ambiguous what spine means in different contexts in geometry and topology. Here it just means the $2-$skeleton.

2-complex Y on which H_L acts freely, faithfully, properly, and cellularly. Since H_L is FP_∞ by the Main Theorem it is finitely generated and we may assume it acts co-compactly on the 1-skeleton of Y. Since Y is contractible there exists an H_L-equivariant quasi-isometry $\phi : X_t \to Y$.

Metrize X_t and $\phi(X_t)$ by H_L-equivariant path-metrics. H_L then acts on both properly and cocompactly (even though H_L may not act on Y cocompactly) by isometries, so is quasi-isometric to both. In particular, X_t and $\phi(X_t)$ are quasi-isometric to each other.

Thus point pre-images of ϕ will have diameters bounded by the quasi-isometry constants. Denote the restriction $\phi|S_{v,L}$ by ϕ_v.

By the previous proposition, we may choose a vertex $v \in X$ so that $|f(v)-t|$ is large in comparison with the quasi-isometry constants (since there are arbitrarily large copies of L). For each point $x \in S_{v,L}$ the geodesic $[x, v]$ determines a unique minimal simplex $\sigma \subset L = Lk_\downarrow(v, X)$ containing it. The pre-image $\phi_v^{-1}(\phi_v(x))$ will be contained in the contractible (by Proposition 12) subcomplex $S_{v,St'(\sigma,L)}$. For vertices v with $|f(v) - t|$ large enough one can define a left homotopy inverse to ϕ_v by taking each vertex of $\phi_v(S_{v,L})$ to a point of its ϕ_v-pre-image, and extending over skeleta. This is well defined up to homotopy since X is the universal cover so is simply-connected. By functoriality the map on π_2 induced by inclusion is split injective so $\pi_2(S_{v,L}) \subset \pi_2(\phi_v(S_{v,L}))$. However, L is the $2-$skeleton, in particular homotopic to the homology sphere with a point removed, so in the universal cover S^3 this corresponds to removing 120 points. Letting S denote the set of removed points,

$$\pi_2(S_{v,L}) \cong \pi_2(L) = \pi_2(\tilde{L}) \cong H_2(S^3 \backslash S) \neq 0$$

Hence $\phi_v(S_{v,L}) \subset Y$ is a non-aspherical subset of an aspherical $2-$complex. $\square$

7.2 An infinite relation gap

The material here follows [Ha].

The group ring $\mathbb{Z}G$ admits a natural ring homomorphism to $\mathbb{Z}$ by sending all the formal symbols $g \mapsto 1$. The kernel of this homomorphism is known as the *augmentation ideal* and is denoted IG. Given a short exact sequence of groups

$$1 \to N \to F(S) \to G \to 1$$

where $F(S)$ is a free group on the set S, one obtains a short exact sequence of $\mathbb{Z}G-$modules

$$0 \to N/[N, N] \to \oplus_{s \in S} \mathbb{Z}G\{s\} \to IG \to 0$$

In particular, if G has an $n-$generator $m-$relator presentation, then so does IG. Let $d(G)$ be the minimal number of generators needed to generate G, and $d_G(IG)$ be the minimal number of generators of IG as a left $\mathbb{Z}G-$ module. The difference $d(G) - d_G(IG)$ is called the generation gap. In [CGK] the authors show that arbitrarily large generation gaps can occur for finite groups.

F acts on N by conjugation since N is normal, and N acts trivially by conjugation on its abelianisation N_{ab}, so the action of Γ on N_{ab} by conjugation is well-defined.

As a consequence, $\min\{k | \exists s_1 \dots s_k \in F, N = \langle\langle s_1 \dots s_k \rangle\rangle\}$ is an upper bound on the rank of R_{ab} as a $\mathbb{Z}\Gamma$ module. Let $d_F(N)$ denote the minimal number of generators needed to normally generate N, and $d_G(N/[N, N])$ the minimal number of generators of $N/[N,N]$ as a left $\mathbb{Z}G-$ module. The next natural question one can ask is what values the *relation gap* $d_F(N) - d_G(N/[N,N])$ can take, but for finite groups this is an open problem. For finitely generated infinite groups Morse theory gives a way of creating a group with infinite relation gap. We will need a result of homological algebra:

Proposition 34. *Let* $\Gamma = F/R$*, where* F *is a free group of finite rank* m*. Then* Γ *is of type* FP_2 *if and only if* R_{ab} *is finitely generated as a* $\mathbb{Z}\Gamma$ *module*

Let L be a flag triangulation of a spine of the Poincare homology sphere as before. By the main theorem, H_L is of type FP (hence FP_2), but since L isn't simply connected H_L is not finitely presented. H_L is of type FP_2 so R_{ab} is finitely generated as a $\mathbb{Z}\Gamma$ module by the proposition, it isn't of type F_2, so any set of relations is infinite and hence the relation gap is infinite. Note that this construction works for any acyclic non-simply-connected finite flag complex L of dimension 2.

Appendix

A The Eilenberg-Ganea conjecture

Definition. The *cohomological dimension* $cd_{\mathbb{Z}}(G)$ of a group G with $\mathbb{Z}$ coefficients is the smallest number n such that there exists a projective resolution

$$0 \to P_n \to \cdots \to P_0 \to \mathbb{Z} \to 0$$

over the trivial $\mathbb{Z}G$-module $\mathbb{Z}$ where each P_i is a finitely generated projective $\mathbb{Z}G$-module, and infinity if no such resolution exists.

Definition. The *geometric dimension* $gd(G)$ of a group G is the smallest n such that there exists an $n-$dimensional $K(G, 1)$.

The universal cover of the $K(G, 1)$ gives rise to a projective resolution so $gd(G) \geq cd(G)$. The reverse inequality is almost always true, and is the Eilenberg-Ganea theorem [EG]:

Theorem 35. *Let G be a finitely presented group and $n \geq 3$. Suppose $cd_{\mathbb{Z}}(G) \leq n$. Then there exists an $n-$dimensional aspherical CW complex X such that $\pi_(X) = G$.*

When $n = 1$, it was proved in [Sta] that

Theorem 36. *Every finitely generated group of cohomological dimension one is free.*

(Finitely generated free groups act on trees, the universal cover of an appropriate wedge of circles.) Hence this leaves open only the case $n = 2$, the Eilenberg-Ganea conjecture.

B The Whitehead conjecture

The treatment in this section follows [Ro].
A $2-$dimensional complex K is aspherical if and only if $\pi_2(K)$ vanishes: the universal cover $\tilde{K}$ is again a $2-$dimensional complex whose π_1 and π_2 also vanish, so by the Hurewicz theorem all its homology groups vanish in dimensions 1 and 2. Since $\tilde{K}$ is a 2-complex all its higher homology groups vanish, so again by Hurewicz so do all its homotopy groups.

Conjecture 37. *If L is an aspherical $2-$complex, then so is any subcomplex $K \subset L$.*

The Whitehead conjecture was at least partly motivated by the study of knot complements $S^3 \backslash K$ (where K is a smooth embedding of S^1), which are now fairly well understood.

Theorem 38. *Conjecture 37 implies that knot complements are aspherical.*

Proof. Glue a (thickened) meridian disk into $S^3 \backslash K$ to get a 3-ball which collapses to an aspherical 2-complex L. If Conjecture 37 were true then L has to be aspherical. $\square$

The asphericity of knot complements has since been shown in [Pa] using techniques of 3-manifold theory. There have been partial results:

Theorem 39. *The Whitehead conjecture is true if either*

- *K has at most one 2-cell [Co], or*
- *$\pi_1(L)$ is finite and non-trivial, abelian, or free*

However, stronger results that would imply the Whitehead conjecture, such as Wise's conjecture, are known to be false [Fi], and some partial results have been proved showing where to look for counterexamples. First note that if $K \subset L$ with $\pi_2(L) = 0$ is a counterexample it can be assumed that:

- L is obtained from K by attaching 2-cells: we can always add the 1-skeleton of L to K without changing its asphericity).
- K is finite: If K were infinite and non-aspherical, then restrict K to the image of a nontrivial map $f : S^2 \to K$
- L is contractible: take instead the universal cover $\tilde{L}$

Theorem 40. *[Ho] If the Whitehead conjecture is false then there exists a counterexample $K \subset L$ such that either:*

1. *L is finite and contractible and $K = L \backslash e$ for one 2-cell e.*
2. *L is the union of an infinite chain of finite non-aspherical subcomplexes $K = K_0 \subset K_1 \subset K_2 \subset \dots$ where each $K_i \subset K_{i+1}$ is nullhomotopic*

Theorem 41. *[Lu] If the Whitehead conjecture is false then there is a counterexample of the second type.*

Note that the answer to the corresponding problem of whether $\pi_n(L) = 0$ implies $\pi_n(K) = 0$ when $K^n \subset L^n$ are $n-$dimensional complexes for $n \neq 2$ is known: $\pi_1(K^1) \to \pi_1(L^1)$ is always injective. On the other hand, for dimension $n \geq 3$, it is false. Consider $K = S^{n-1}$ and $L = D^n$. For $n = 3$, the Hopf map gives a non-trivial element of $\pi_3(S^2) \equiv \mathbb{Z}$ (by the long exact sequence for the fibration $S^1 \to S^3 \to S^2$), and similarly it was checked that for $n \geq 4, \pi_n(S^{n-1}) \equiv Z/(2)$, [Hu].
So only the 2-dimensional case, namely Whitehead's Conjecture, remains open.

C Further facts about Right-Angled Artin Groups

The material in this section follows [Wi].
If $C \cong [-1, 1]^n$, then a *midcube* $M \subseteq C$ is the intersection of C with $\{x_i = 0\}$ for some i.

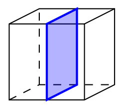

Now if X is a non-positively curved cube complex, and M_1, M_2 are midcubes of cubes in X, we say $M_1 \sim M_2$ if they have a common face, and extend this to an equivalence relation. The equivalence classes are *immersed hyperplanes*. We usually visualize these as the union of all the midcubes in the equivalence class.

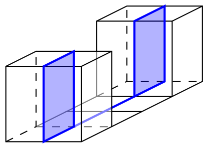

Note that in general, these immersed hyperplanes can have self-intersections, hence the word ”immersed”. Thus, an immersed hyperplane can be thought of as a locally isometric map $H \looparrowright X$, where H is a cube complex.
In general, these immersed hyperplanes can have several “pathologies”. The first is self-intersections, as we have already met. The next problem is that of *orientation*, or *sidedness*. For example, we can have a (closed) Mobius band. This is bad, for the reason that if we think of this as a $(-1, 1)$-bundle over H, then it is non-orientable, and in particular, non-trivial.
In general, there could be self intersections. So we let N_H be the pullback interval bundle over H. That is, N_H is obtained by gluing together $\{M \times (-1, 1) \mid M \text{ is a cube in } H\}$. Then we say H is *two-sided* if this bundle is trivial, and *one-sided* otherwise.

Sometimes, we might not have self-intersections, but something like this:

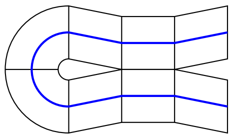

This is a *direct self-osculation* (when H is two sided and a pair of points in the same component of the boundary have the same image in X). We can also have *indirect osculations* (different components of the boundary) that look like this:

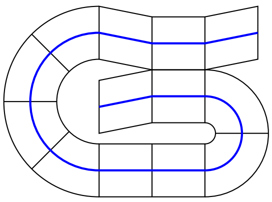

Finally, we have *inter-osculations*, which look roughly like this:

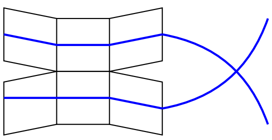

While in the figures above the cubes meet along an edge, meeting along a vertex would also be considered an osculation.

Definition. A cube complex is *special* if its hyperplanes do not exhibit any of the following four pathologies:

- One-sidedness
- Self-intersection
- Direct self-osculation

- Inter-osculation

Example. A cube is a special cube complex.

Example. If X is special then so is any covering space of X.

Example. If $X = \mathcal{S}_N$ is a Salvetti complex, then it is a special cube complex. Parallelism in the cubes of X preserves orientations of 1-cells, from which it follows that every hyperplane is two-sided. If a hyperplane H_a were to self-intersect, it would follows that some square has every edge glued to a, which does not occur in the construction of X. If H_a directly self-osculates then it follows that H_a is dual to two distinct edges incident at the same vertex; but each hyperplane is dual to a unique edge. If H_a and H_b inter-osculate then it follows that a and b both bound a square and do not bound a square, a contradiction.

This shows that all subgroups of right-angled Artin groups are fundamental groups of special cube complexes (by the Galois correspondence). The key theorem is the following:

Theorem 42 (Haglund, Wise)**.** *[HW] If X is a special cube complex, then there exists a graph N and a local isometry of cube complexes*

$$\varphi_X : X \looparrowright \mathcal{S}_N.$$

Corollary 43. *The characteristic map lifts to an isometric embedding $\tilde{\phi} : \tilde{X} \to \widetilde{\mathcal{S}_{\mathbb{H}(X)}}$. In particular, it induces an injective homomorphism $\phi_* : \pi_1 X \hookrightarrow A_{\mathbb{H}(X)}$.*

Some group theoretic facts about right-angled Artin groups are known, such as they are linear [DJ], residually finite (for any non-trivial element g, there is a homomorphism to a finite group such that g isn't in the kernel), Hopfian (any surjective homomorphism from the group to itself is an isomorphism) and so on. These then give control over the fundamental group. Equally important is the following:

Corollary 44. *A group G is a subgroup of a right-angled Artin group if and only if G is the fundamental group of a (not necessarily compact) special cube complex.*

References

[Ba90] W. Ballman Singular spaces of nonpositive curvature. In Sur les groupes hyperboliques d'apr'es Mikhael Gromov (Bern, 1988), volume 83 of Progr. Math., pages 189–201. Birkh¨auser Boston, Boston, MA, 1990.

[BB] Bestvina, M., Brady, N. Morse theory and finiteness properties of groups. Invent math 129, 445–470 (1997). https://doi.org/10.1007/s002220050168

[BG] K.U. Bux, C. Gonzalez. The Bestvina–Brady Construction Revisited: Geometric Computation of Σ-Invariants for Right-Angled Artin Groups. Journal of the London Mathematical Society, Volume 60, Issue 3, December 1999, Pages 793–801

[BH] M. Bridson, A. Haefliger, Metric spaces of non-positive curvature, Grundlehren der Mathematischen Wissenschaften 319, Berlin: Springer-Verlag

[Bi76] R. Bieri, Normal subgroups in duality groups and in groups of cohomological dimension 2, Jour. Pure App. Algebra 7, 35-52 (1976)

[Br] Brown, K.S. The cohomology of groups

[Bra99] N. Brady, Branched Coverings of Cubical Complexes and Subgroups of Hyperbolic Groups, Journal of the London Mathematical Society, Volume 60, Issue 2, October 1999, Pages 461–480,

[CGK] J.Cossey, K.W. Gruenberg, L.G. Kovács, The presentation rank of a direct product of finite groups, Journal of Algebra, Volume 28, Issue 3, 1974, Pages 597-603, ISSN 0021-8693, https://doi.org/10.1016/0021-8693(74)90062-3.

[Co] Cockcroft, W. H., On two-dimensional aspherical complexes, Proc. London Math. Soc., (3) 4 (1954), 375–384

[DJ] M. Davis and T. Januszkiewicz. Right-angled artin groups are commensurable with right-angled coxeter groups. Journal of Pure and Applied Algebra, 153(3):229–235, 2000

[EG] Eilenberg, S.; Ganea, T. (1957). "On the Lusternik–Schnirelmann category of abstract groups". Annals of Mathematics. 2nd Ser. 65 (3): 517–518. doi:10.2307/1970062. MR 0085510.

[Fi] W. Fisher. Contractible complexes and non-positive immersions. arXiv:2210.02304 [math.GT]

[Ha] Harlander, J. (2017). The Relation Gap Problem. In W. Metzler, S. Rosebrock (Eds.), Advances in Two-Dimensional Homotopy and Combinatorial Group Theory (London Mathematical Society Lecture Note Series, pp. 128-148). Cambridge: Cambridge University Press. doi:10.1017/9781316555798.007

[Ho] J. Howie, Some remarks on a problem of J. H. C. Whitehead, Topology 22, 1983.

[Ho99] Howie, J. (1999). Bestvina-Brady groups and the plus construction. Mathematical Proceedings of the Cambridge Philosophical Society, 127(3), 487-493.

[Hu] S.-T. Hu. Homotopy theory. Academic Press, 1959.

[HW] F. Haglund and D. T. Wise. Special cube complexes. Geometric and Functional Analysis, 17(5):1551–1620, 2008.

[KS] R. Kirby, M. Scharlemann, Eight faces of the Poincaré homology 3-sphere. Geometric topology (Proc. Georgia Topology Conf., Athens, Ga., 1977), pp. 113–146, Academic Press, New York-London, 1979.

[Lu] E. Luft, On 2-dimensional aspherical complexes and a problem of J.H.C. Whitehead, Math. Proc. Cambridge Phil. Soc., 119 (1996), 493–495.

[Pa] C. D. Papakyriakopoulos. On Dehn's lemma and asphericity of knots. Annals of Mathematics, 66:1–26, 1957.

[Ro] S. Rosebrock, The Whitehead conjecture – an overview, Sib. Elektron. Mat. Izv. ‘ , 2007, Volume 4, 440–449

[Sta] J.R. Stallings, "On torsion-free groups with infinitely many ends", Annals of Mathematics 88 (1968), 312–334. MR0228573

[Wi] H. Wilton, Non-positively curved cube complexes