Cube complexes

Ng Chung Lok Andrew

1 Introduction

This expository essay will give a proof of Gromov's link condition for determining when a cube complex is negatively curved and show a couple of the many important consequences. Negative curvature (to be defined) is analogous to hyperbolicity, which is ubiquitous in geometry: e.g. a 'random' group is either hyperbolic, trivial, or $\mathbb{Z}/(2)$ [Gro93], and almost all oriented surfaces, including those with boundary components and punctures, admit hyperbolic metrics. The appendix sketches a connection to decision problems in group theory. Cube complexes (to be defined, but the reader's first guess at a definition won't be far off) also play an important role: the Wise conjecture, proved by Ian Agol in 2012, was the final piece needed to settle the last of Thurston's conjectures on 3-manifolds, and completed one of the most impressive decades-long programs in mathematics. As a result, 3-manifolds have a classification akin to that for surfaces. Part of what was so surprising about this is that Wise's conjecture was about cube complexes, which can be very different to 3-manifolds, but it turns out that there are bridges to be built between the two areas.
This essay follows the exposition of [Wi]. For ease of understanding, sometimes remarks which aim to build intuition will appear in place of an actual proof. The figures were made with Geogebra geometry. The assumed knowledge will be anything that appears in the undergraduate schedules for the Cambridge Mathematical Tripos. Here is some knowledge which isn't assumed.

1.1 Some useful definitions

Definition. A metric space X is said to be *geodesic* if every pair of points of X is joined by a geodesic (an isometrically embedded closed interval up to a uniform rescaling)

Definition. A metric space X is said to be *proper* if closed balls are compact.

Compare with the theorem that says in a topological vector space over $\mathbb{R}$ or $\mathbb{C}$ the closed unit ball is compact iff the vector space is finite dimensional.

Definition. (Compact-open topology) Let X and Y be two topological spaces, and let $C(X, Y)$ denote the set of all continuous maps between X and Y. Given a compact subset K of X and an open subset U of Y, let $V(K, U)$ denote the set of all functions $f \in C(X, Y)$ such that $f(K) \subseteq U$. Then the collection of all such $V(K, U)$ is a subbase (i.e. it generates) for the compact-open topology on $C(X, Y)$.

This is a common topology in analysis, topology, and geometry whose properties will not really be used beyond the definition and certainly can't be done justice here, so we content ourselves with the following remark: when both the domain and codomain have a metric structure, a sequence of functions converges in the topology precisely when the functions converge uniformly on every compact subset of the domain.

Definition. An action of a group G on a space X is said to be

- *free* if whenever $g \in G$ satisfies $g \cdot x = x$ for some $x \in X$, then $g = 1$
- *properly discontinuous* each $x \in X$ lies in some open set U such that $\{g \in G | g(U) \cap U \neq \varnothing\}$ is finite.[1]
- *cocompact* if the quotient of the space by the action is compact. [2]

2 CAT(0) spaces and groups

2.1 The CAT(κ) condition

A wide variety of conditions on metric spaces mimicking hyperbolicity have been given by mathematicians, most of them being equivalent to each other in some sense. One of the most common 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. [3]
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,

[1] This slightly unfortunate choice of terminology results in phrases such as a continuous, properly discontinuous action.
[2] The comedic geniuses of our time have sometimes called this an mpact action.
[3] Technically the condition described is known as the hyperbolic plane being δ-slim. See the appendix for more details.

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 $\kappa' > \kappa$.
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 that the egregious abuse of notation is not anything more serious. Also,

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 $\{x\}$. □

Definition. A group Γ that acts freely, properly discontinuously, and cocompactly by isometries on a proper CAT(0) space is called a CAT(0) group.

Some authors don't require the action to be free, so that Γ may have torsion. Some examples of CAT(0) groups:

- $\mathbb{Z}^n$ for any n (acting on $\mathbb{R}^n$)
- Free groups (acting on their Cayley complex)
- Direct products of free groups

Riemannian manifolds of non-positive sectional curvature are also CAT(0) spaces, so their fundamental groups are CAT(0), as are uniform lattices in semisimple Lie groups. Amazingly, a random group at density $d < \frac{1}{6}$ is CAT(0) too [HW04][OW09]. Hence while CAT(0) groups and spaces enjoy a large number of nice properties, they still include enough (important!) classes of groups to warrant study.
The next gluing lemma allows us to create new CAT(0) spaces easily.

Lemma 5 (Alexandrov)**.** *Suppose the triangles $\Delta_1 := \Delta(x, y, z_1), \Delta_2 := \Delta(x, y, z_2)$ both satisfy the CAT(0) condition and $y \in [z_1, z_2]$. Then $\Delta := \Delta(x, z_1, z_2)$ satisfies the CAT(0) condition.*

Here is one instance where the proof really is draw a picture and play with it.

Proof. Claim: The quadrilateral $\overline{Q}$ obtained by gluing the comparison triangles $\overline{\Delta_1}$ and $\overline{\Delta_2}$ along $[\overline{x}, \overline{y}]$ has a non-acute interior angle at $\overline{y}$.

Proof. Suppose not. Then there are $\overline{p_i} \in [\overline{y}, \overline{z_i}]$ with $[\overline{p_1}, \overline{p_2}] \cap [\overline{x}, \overline{y}] = \{\overline{q}\}$ with $\overline{q} \neq \overline{y}$. Since $y \in [z_1, z_2]$, a contradiction arises from

$$\begin{aligned} d(p_1, p_2) &= d(p_1, y) + d(y, p_2) \\ &= d(\overline{p_1}, \overline{y}) + d(\overline{y}, \overline{p_2}) \\ &> d(\overline{p_1}, \overline{q}) + d(\overline{q}, \overline{p_2}) \\ &\geq d(p_1, q) + d(q, p_2) \\ &\geq d(p_1, p_2) \end{aligned}$$

□

The comparison triangle for Δ is obtained by straightening $[\overline{z_1}, \overline{y}] \cup [\overline{y}, \overline{z_2}]$. There are number of cases to check. One of them is done below.
Suppose $p_i \in [x, z_i]$ and let $\overline{p_i}$ be the comparison points in $\overline{Q}$. In the worst case, the Euclidean geodesic $[\overline{p_1}, \overline{p_2}]$ isn't contained in $\overline{Q}$. Consider the geodesic $[\overline{p_1}, \overline{y}]$. During the process of straightening $[\overline{z_1}, \overline{y}] \cup [\overline{y}, \overline{z_2}]$, $d_{\mathbb{R}^2}(\overline{p_1}, \overline{y})$ doesn't decrease. Hence $d_X(p_i, y) \le d_{\mathbb{R}^2}(\overline{p_1}, \overline{y})$. Eventually $[\overline{p_1}, \overline{p_2}]$ will lie in the interior of $\overline{Q}$ so it suffices to consider that case. Let $\{\overline{q}\} = [\overline{p_1}, \overline{p_2} \cap [\overline{x}, \overline{y}] =$. Then

$$d(p_1, p_2) \le d(p_1, q) + d(q, p_2) \le d_{\mathbb{R}^2}(\overline{p_1}, \overline{q}) + d_{\mathbb{R}^2}(\overline{p_2}, \overline{q}) = d_{\mathbb{R}^2}(\overline{p_1}, \overline{p_2})$$

The last quantity only increases during the straightening process, so the CAT(0) inequality holds and the lemma is proved. □

Gluing constructions in mathematics are most useful for local-to-global properties, e.g. the gluing lemma of topology allows one to create global continuous functions from local information, and the construction of sheaves in algebraic geometry is similarly motivated. Alexandrov's lemma allows us to pass from local non-positive curvature to global non-positive curvature:

Lemma 6. *If X is proper, uniquely geodesic, and non-positively curved, then it is CAT(0).*

Note X must be uniquely geodesic if it is CAT(0).

Proof. Consider a triangle $\Delta := \Delta(x, y, z)$ contained in a closed ball $B := B_x(R)$ [4]. B is compact so there is an $\epsilon > 0$ such that $B_p(\epsilon)$ is CAT(0) for every $p \in B$. Let $\alpha = [y, z]$ and, for each t, let γ_t be the geodesic from x to $\alpha(t)$. If geodesics are unique, they vary continuously with their endpoints, so there is a $\delta > 0$ such that $d(\alpha_{t_1}(s), \alpha_{t_2}(s)) < \epsilon$, for all s, whenever $|t_1 - t_2| < \delta$.
Now divide δ into very small geodesic triangles, each contained in a ball of radius ϵ. By Alexandrov's lemma, it follows by induction on the number of triangles that the entire Δ satisfies the CAT(0) condition. □

2.2 Length Metrics

Definition. Let (X, d) be a metric space and let $\gamma : [a, b] \to X$ be a path. The length of γ is the quantity

$$l(\gamma) :== sup_{a=t_0<t_1<\cdots<t_k=b} \sum_{i=1}^{k} d(\gamma(t_{i-1}), \gamma(t_i))$$

where the supremum ranges over all finite partitions of $[a, b]$. If the length is finite, the path is called *rectifiable*.

Definition. A metric space is called a *length space* if the distance between any pair of points is equal to the infimum of the lengths of paths between them.

[4] It is an important theorem in mathematics that mathematicians can't agree on notation. Apologies to any readers who believe B should be reserved for open balls.

Any geodesic metric space is automatically a length space. In the exact same way that a Riemannian metric on a manifold induces a Riemannian metric on its universal cover, a covering space of a length space inherits an induced metric.

Definition. Let $p : \tilde{X} \to X$ be a covering map and X a length space. Then there is a unique length metric defined on $\tilde{X}$ that makes p a local isometry, defined by setting the length of a path $\tilde{\gamma}$ to be the length of $p \circ \tilde{\gamma}$. In particular, if X is complete then so is $\tilde{X}$

Alexandrov's lemma then implies

Proposition 7. *Suppose X_1, X_2 are locally compact, complete $CAT(0)$ spaces and Y is isometric to closed, convex subspaces of both X_1, X_2. Then $X_1 \cup_Y X_2$ equipped with the induced length metric is $CAT\ (0)$.*

Technically $X_1 \cup_Y X_2$ isn't a covering space, but the induced metric means whatever the reader thinks it means.

Corollary 8. *If Γ_1, Γ_2 are both $CAT(0)$ groups then so is their free product $\Gamma_1 * \Gamma_2$.*

Proof. When Y is a point, wedging gives a space whose fundamental group is $\Gamma_1 * \Gamma_2$, by the Seifert van-Kampen theorem. □

We continue by proving analogues of two important theorems of Riemannian geometry.

Theorem 9. *(Hopf-Rinow) Let X be a length space. If X is complete and locally compact then it is proper and geodesic.*

Proof. X is proper: consider the set of all r for which all closed balls of radius r are compact. Note that it suffices to show $B_{x_0}(r)$ is compact for all r for a fixed x_0. By local compactness, this set is open in $[0, \infty)$.
Let R be a limit point of this set, $x_0 \in X$, and (x_n) be a sequence of points in the closed ball $B_{x_0}(R)$. Let γ_n be a path from x_0 to x_n and for each $p \in \mathbb{N}$ let y_n^p be a point on the path with $d(x_0, y_n^p) < r - \frac{1}{2p}$ and $d(x_n, y_n^p) < \frac{1}{p}$.
Each $y_n^p \in B_{x_0}(r - \frac{1}{2p})$ which is compact by hypothesis, so (y_n^p) has a convergent subsequence for fixed p. One obtains sequences $\{n_i^1\} \supseteq \{n_i^2\} \supseteq \dots$ with y_n^p converging along n_i^p. Picking the j^{th} of the j^{th} sequence gives a sequences (n_k) with $(y_{n_k}^p)$ convergent, hence Cauchy, for all p. x_{n_k} is therefore also Cauchy, hence converges in X.
X is geodesic: let $\gamma_n : [0, 1] \to X$ be a sequence of paths from p to q whose lengths converge to $d(p, q)$. They can all be taken to have length at most $R = d(p, q) + 1$ and parametrised by arc length. Then for any $s, t \in [0, 1]$ we have $d(\gamma_n(s), \gamma_n(t)) < \frac{|s-t|}{R+1}$. By Arzela-Ascoli some subsequence converges to a path γ from p to q. By lower semicontinuity of length, γ is geodesic. □

The next theorem is a generalisation of the classical Cartan-Hadamard theorem, which says that the universal cover of a connected complete Riemannian manifold of non-positive sectional curvature is diffeomorphic to $\mathbb{R}^n$. The proofs were published by Ballman [Ba90], which we follow.

Theorem 10. *(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 $CAT(0)$.*

This is true without the locally compact hypothesis, but we won't need this generality. The following useful criterion is immediate upon unravelling the definitions.

Corollary 11. *A group Γ is $CAT(0)$ iff it is the fundamental group of a compact non-positively curved space.*

To prove the theorem we will need a couple of intermediate results.
Let $\gamma : [0,1] \to X$ be a path from p to q. Let γ be contained in a compact ball B and $\epsilon > 0$ such that $B_y(2\epsilon)$ is CAT(0) for every $y \in B$. Let N be such that $\gamma([\frac{i-1}{2N}, \frac{i+1}{2N}) \subseteq B_{\gamma(\frac{i}{2N}}(\frac{\epsilon}{2})$ for each integer i.
Define a process called *Birkhoff curve shortening* (this is a fanciful name for applying the greedy algorithm locally) as follows. Let $\beta^0(\gamma)$ be the curve obtained by replacing γ be the unique geodesic on each interval $[\frac{2i}{2N}, \frac{2i+2}{2N}]$, and $\beta^1(\gamma)$ be the curve obtained by replacing γ be the unique geodesic on each interval $[\frac{2i+1}{2N}, \frac{2i+3}{2N}]$. The Birkhoff curve-shortening map is defined to be $\beta(\gamma) := \beta^1 \circ \beta^0(\gamma)$. Note that

1. this is a continuous function of γ
2. γ is homotopic to $\beta(\gamma)$ (respecting endpoints), and they are equal iff γ is a local geodesic.
3. Set $\lambda(\gamma) := \sum_{i=1}^{2N+1} d(\gamma(\frac{i-1}{2N}, \gamma(\frac{i+1}{2N})))$. $\lambda(\beta(\gamma)) \leq \lambda(\gamma)$, with equality iff γ is a local geodesic parametrised by arc length.

Denote by $d_{sup}(\cdot,\cdot)$ the supremum metric on the space of continuous paths from p to q.

Lemma 12. *Let γ_1, γ_2 be continuous paths from p to q. If $d_{sup}(\gamma_1, \gamma_2) < \epsilon$, $d_{sup}(\beta(\gamma_1), \beta(\gamma_2)) \leq d_{sup}(\gamma_1, \gamma_2)$.*

Proof. Each component of the straightening happens with a single ball of radius 2ϵ. Apply convexity of the metric. □

The next result, together with lemma 6, prove the Cartan-Hadamard theorem.

Theorem 13. *Let X be a proper length space of non-positive curvature. Each homotopy class of paths from p to q contains exactly one local geodesic.*

Proof. Existence: Let $\gamma : [0,1] \to X$ be a path from p to q. The distance between pairs of nearby points decreases when β is applied, so $\{\beta^n(\gamma)\}$ is equicontinuous. By Arzela-Ascoli, there is a curve γ_∞ from p to q and a convergent subsequence $\beta^{n_k}(\gamma) \to \gamma_\infty$. Since $\lambda(\beta^n(\gamma))$ is decreasing,

$$\lim_{n\to\infty} \lambda(\beta^n(\gamma)) = \lim_{k\to\infty} \lambda(\beta^{n^k}(\gamma)) = \lambda(\gamma_\infty)$$

$\lambda(\beta(\gamma_\infty)) = \lim_{n\to\infty} \lambda(\beta^{n+1}(\gamma)) = \lambda(\gamma_\infty)$, so γ_∞ is a local geodesic.
Uniqueness: let $\gamma_0, \gamma_1 : [0,1] \to X$ be local geodesics from p to q and let

$\gamma_s, s \in [0,1]$ be a homotopy between them. By compactness of the unit square there exist R, ϵ, N suitable for all the γ_s. Applying β^n, $\beta^n(\gamma_s)$ is a sequence of homotopies from γ_0 to γ_1. By the previous lemma $\{(s,t) \mapsto \beta^n(\gamma_s)(t)\}$ is equicontinuous, so there exists a subsequences n_k such that $\beta^{n_k}(\gamma_s)$ converges to a limiting homotopy $\tilde{\gamma_s}$, which is geodesic for each s by the argument given in the existence part of the proof.
For sufficiently close s_1, s_2, $d_{sup}(\tilde{\gamma_{s_1}}, \tilde{\gamma_{s_2}}) < \epsilon$ so $t \mapsto d(\tilde{\gamma_{s_1}}(t), \tilde{\gamma_{s_2}}(t))$ is a locally convex function. A locally convex function over the reals is globally convex, so in fact $\tilde{\gamma_{s_1}} = \tilde{\gamma_{s_2}}$, hence $\gamma_0 = \gamma_1$. □

3 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. This requires a couple of steps. In this section we give a proof of the link condition for any Euclidean complex.

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.
The torus is the quotient of $\mathbb{R}^2$ by the action of $\mathbb{Z}^2$, and $\mathbb{Z}^2$ is already known to be CAT(0). By the corollary to Cartan-Hadamard, to exhibit a CAT(0) group it suffices to exhibit a compact metric space with non-positively curved metric, a torus for example.

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 $\epsilon > 0$. [5]

The link 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 $Lk(x_0)$

[5] Another 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

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

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

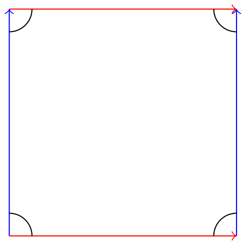

so is S^1. At the corner of a cube, the link is also (homeomorphic to) S^1 but looks a bit different:

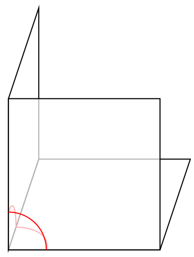

If one were to flatten the cube however, it wouldn't be possible to do so isometrically while keeping the S^1 intact. This is because the total length of the link is too short, or phrased differently, the angle is too small. This is a sign of positive curvature, and in fact the angle will be a useful 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 $\angle_{x_0}$.

Theorem 14. *(Gromov's link condition) Any Euclidean complex X is non-positively curved iff $Lk(x_0)$ is $CAT(1)$ for every $x_0 \in X$.*

The proof is put the next three lemmas together. First, the concept of a cone is defined (it likely means what the reader thinks of when imagining a cone), and the significance of the first lemma is just that the cone is geodesic.

Definition. Let L be a metric space. A (Euclidean) cone on L, denoted by $C_\epsilon L$, is the metric space associated to the pseudometric space $L \times [0, \epsilon]$, where we define the pseudometric by $d((x,s),(y,t))^2 = s^2+t^2-2st\cos\min\{\pi, d(x,y)\}$. Two distinct points are at distance zero iff $s = 0 = t$.

After crushing $L \times \{0\}$ to a point, this becomes the cone point and the space is a metric space. Checking that the triangle inequality holds is omitted.

Lemma 15. *Let $x, y \in L$ be two points at distance less than π. For any $s, t > 0$, there is a bijection between the set of geodesic segments joining x to y in L and the set of geodesics joining (x,s) to (y,t) in $C_\epsilon L$.*

Proof. Consider a geodesic $[x, y]$ in L. Together with the cone point, it spans a subcone $C_\epsilon[x, y]$ which is isometric to a Euclidean cone. The unique Euclidean geodesic from (x, s) to (y, t) is then the geodesic corresponding to $[x, y]$.
Conversely, consider a geodesic $[(x, s), (y, t)]$ in $C_\epsilon L$. If the cone point is in the geodesic then the distance between them is $s+t$, which implies that $d(x, y) \geq \pi$. Otherwise, there is a well defined projection of the geodesic to L. Let (z, r) be any point on $[(x, s), (y, t)]$. It is enough to prove that $d(x, y) \geq d(x, z) + d(z, y)$; the converse is the triangle equality, and it then follows that the image is a geodesic. To see this, simply note that the length of the projection of a geodesic is equal to the angle in the corresponding comparison triangle. But the comparison triangle for $[(x, s), (y, t)]$ can be obtained by straightening the comparison triangles for $[(x, s), (z, r)]$ and $[(y, t), (z, r)]$, and when we do so the angle at the cone point increases. $\square$

Note that if $d(x, y) \geq \pi$ then the path through the cone point is a geodesic.

Lemma 16. *If $x_0 \in X$ then there exists $\epsilon > 0$ such that the closed ball $B_{x_0}(\epsilon)$ is convex and isometric to $C_\epsilon Lk(x_0)$.*

Proof. Let $\{\Sigma_i\}$ be the set of open simplices whose closures contain x_0. Their union U is an open neighbourhood of x_0. For each $y \in U$ there is a well defined geodesic $[x_0, y]$ and so a continuous map $\pi : U - \{x_0\} \to Lk(x_0)$. Let $\gamma : [0, 1] \to U - \{x_0\}$ be any local geodesic, and suppose that $\pi \circ \gamma$ is non-constant. Consider the triangle $\Delta := \Delta(\gamma(0), x_0, \gamma(1))$. Let $0 = t_0 < t_1 < \cdots < t_n = 1$ be so that $\gamma|_{[t_k,t_{k+1}]}$ is equal to a component of the intersection of the image of γ with the interior of Δ_{i_k} . Then $\Delta(\gamma(t_{k-1}), x_0, \gamma(t_k))$ is isometric to a Euclidean triangle, which we shall denote by $\overline{\Delta_k}$. Arranging the $\overline{\Delta_k}$ side by side in $\mathbb{R}^2$, we obtain a comparison triangle $\overline{\Delta}$ for Δ. There is a distance-non-increasing continuous map $\overline{\Delta} \to \Delta$. Let 2ϵ be smaller than the minimal distance from x_0 to a simplex not contained in U. Then every point of every geodesic in $B_{x_0}(\epsilon)$ is contained in U, and so $B_{x_0}(\epsilon)$ is convex. In particular, the induced metric on $B_{x_0}(\epsilon)$ is a length metric. By construction it agrees with the cone metric on the interior of each cell. As both metrics are length metrics, it follows that they agree. $\square$

Lemma 17. *The cone $C_\epsilon Lk(x_0)$ is $CAT(0)$ iff $Lk(x_0)$ is $CAT(1)$.*

This is a special case of a slightly more general result known as Berestovskii's theorem. See [BH].

Proof. Retain the notation introduced in the proof of the previous lemma.
Suppose $Lk(x_0)$ is CAT(1). Let $\Delta := \Delta(x, y, z)$ be a triangle in $C_\epsilon L$. Note that it suffices to prove the inequality in the case in which one of the comparison points is a vertex. There are three cases to consider. In the first, x_0 is on the boundary of Δ. Cutting Δ into pieces, we may assume that $x_0 = x$. As in the proof of convexity above, the map $\overline{\Delta} \to \Delta$ does not increase distance and so we are done.
We may therefore assume that x_0 is not contained in an edge of Δ. Let $p = x$ and $q \in [y, z]$, where $[y, z]$ is the side of shortest length. In the second case, $\pi(\Delta)$ has perimeter at least 2π.

Consider the comparison triangles $\overline{\Delta}(x_0, x, y), \overline{\Delta}(x_0, y, z), \overline{\Delta}(x_0, z, x)$, which exist by case 1. Glue them along $[\overline{x_0}, \overline{y}]$ and $[\overline{x_0}, \overline{z}]$ and let $\overline{x_1}, \overline{x_2}$ be the unglued

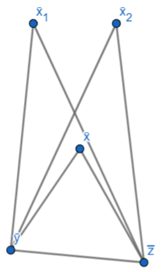

Figure 1: The case $\pi(\Delta)$ has perimeter at least 2π. Note that the perimeter of a triangle in the link is the sum of angles at x_0 and angles in Euclidean space are not less than the corresponding ones in the metric space, so $\overline{x_0}$ should be in the interior of the other comparison triangles.

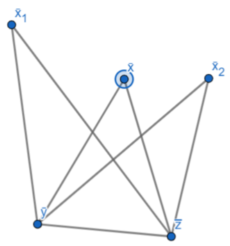

Figure 2: The case $\pi(\Delta)$ has perimeter less than 2π.

vertices of $\overline{\Delta}(x0, x, y)$ and $\overline{\Delta}(x_0, z, x)$. Let $\overline{x}$ be the point of intersection of the circle based at $\overline{y}$ of radius $d(x, y)$ and the circle based at $\overline{z}$ of radius $d(x, z)$ on the same side of $[\overline{y}, \overline{z}]$ as the triangles (they intersect by the triangle inequality). The comparison triangle for Δ is constructed by moving $\overline{x_1}$ to $\overline{x}$ along the first circle and $\overline{x_2}$ to $\overline{x}$ along the second circle. For any $\overline{r} \in [\overline{x_0}, \overline{y}]$ we have $d(\overline{x}, \overline{r}) \geq d(\overline{x_1}, \overline{r})$, and similarly for $\overline{r} \in [\overline{x_0}, \overline{z}]$. After straightening, $d((\overline{x}, \overline{q}) = d(\overline{x}, \overline{r}) + d(\overline{r}, \overline{q})$ for some $\overline{r}$, and the second term has remained constant during the straightening process, so $d(x, q) \leq d((\overline{x}, \overline{q})$.

The last case is when $\pi(\Delta)$ has perimeter less than 2π. An inspection of Figure 2 shows that the proof for the previous case will not work.

Consider a geodesic $\gamma = [x, q]$. Then $\pi \circ \gamma$ is a geodesic in $Lk(x_0)$, which we can develop as follows. Let $0 = t_0 < t_1 < \cdots < t_n = 1$ be so that $\pi \circ \gamma|_{[t_k, t_{k+1}]}$ is equal to a component of the intersection of the image of γ with the interior of $\Delta_{i_k} \cap Lk(x_0)$. Then $\pi \circ \gamma(t_{k-1}, t_k)$ defines a set of directions, which are equal to the intersection of Δ_{i_k} with a 2-dimensional Euclidean plane. Let $\tilde{\Delta_k}$ denote the Euclidean segment spanned by this set of directions. Gluing the $\tilde{\Delta_k}$ together, we obtain a Euclidean segment $\tilde{\Delta}$, and we may consider points $\tilde{x_0}$ at the cone point of the segment, $\tilde{x}$ on one edge with $d(\tilde{x}, \tilde{x_0}) = d(x, x_0)$ and $\tilde{q}$ on the other edge with $d(\tilde{q}, \tilde{x_0}) = d(q, x_0)$. To finish the proof, we compare $\tilde{\Delta}$ with the comparison triangle $\overline{\Delta}$. The CAT(1) hypothesis applied to $\pi(\Delta)$ tells us that $\angle_{\tilde{x_0}}(\tilde{q},\tilde{x}) \leq \angle_{\overline{x_0}}(\overline{x},\overline{q})$. Therefore

$$\begin{aligned} d(x,q)^2 &= d(\tilde{q},\tilde{x})^2 \\ &= d(\tilde{x},\tilde{x_0})^2 + d(\tilde{q},\tilde{x_0})^2 - 2d(\tilde{x},\tilde{x_0})d(\tilde{q},\tilde{x_0})\cos\angle_{\tilde{x_0}}(\tilde{q},\tilde{x}) \\ &\leq d(\overline{x},\overline{x_0})^2 + d(\overline{q},\overline{x_0})^2 - 2d(\overline{q},\overline{x_0})d(\overline{x},\overline{x_0})\cos\angle_{\overline{x_0}}(\overline{x},\overline{q}) \\ &= d(\overline{x},\overline{q})^2 \end{aligned}$$

as required.
For the converse assertion, note that if $Lk(x_0)$ fails to be CAT(1) then the inequality in the final calculation fails, and so the oringal triangle Δ did not satisfy the CAT(0) condition. □

3.1 Injectivity radius and systole

We will need a technical lemma establishing when local to global lifting holds for general CAT(κ) spaces, analogous to Lemma 6.

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 18. *Let X be a compact geodesic metric space of curvature at most κ. Then X is not CAT(κ) 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.*

Proof. If X contains an isometrically embedded circle of length less than $2D_{(\kappa)}$, picking any 3 points on the circle and looking at the comparison triangle shows X isn't CAT(κ), since in the Riemannian setting circles aren't triangles.
Now suppose X isn't CAT(κ). Define a *digon* to be the union of two distinct geodesic segments $[x,y]$ and $[x,y]'$, called its sides, joining two points $x, y \in X$. These are called the vertices of the digon. Then the injectivity radius r satisfies $0 < r < D_\kappa$. The proof of the right hand inequality is essentially Alexandrov's lemma, since any triangle of perimeter less than $2r$ satisfies the CAT(κ) condition. The proof consists of showing there is a digon of perimeter $2r$ and this digon is an isometrically embedded circle.
By definition of r, there are pairs of distinct geodesics $[x_n,y_n],[x_n,y_n]'$ whose lengths tend to r. By the usual compactness and Arzela-Ascoli argument, some subsequence converges to a pair of geodesics $[x,y],[x,y]'$ with $d(x,y) = r$.
Claim: these geodesics are distinct.

Proof. Suppose not. Then for large n, the midpoints m_n and m'_n of $[x,y]$ and $[x,y]'$ are close enough together, so that $\Delta(x_n,m_n,m'_n)$ and $\Delta(y_n,m_n,m'_n)$ are both of perimeter less than $2r$, hence satisfy the CAT(κ) condition. The comparison triangle $\overline{\Delta}(x_n, m_n, y_n)$ is degenerate, so by Alexandrov's lemma it follows that the comparison triangles $\Delta(x_n, m_n, m_n')$ and $\Delta(y_n, m_n, m_n')$ are also degenerate, so $m_n = m_n'$. □

To show that $[x, y] \cup [x, y]'$ is an isometrically embedded circle, just apply the above argument to any pair of points z, z' that are distance r in the digon but distance less than r in X. □

3.2 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. In the case of cube complexes, they are additionally *all-right spherical* simplicial complexes, meaning every edge has length $\frac{\pi}{2}$.

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 19. *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[6]. Also note the following: the proof of the Link Condition above goes through to show L is locally CAT(1) iff the link of every vertex is CAT(1). 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 18, 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

[6]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.

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 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. $\square$

This lets us produce potentially very many CAT(0) groups, by just writing down random cube complexes and computing their fundamental groups. However, how do we know whether they are different?

4 Special cube complexes

4.1 Right-angled Artin groups

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 Σ be the unique flag complex with $\Sigma^{(1)} = N$. A_N is also denoted by A_Σ.

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* $D(\Sigma)$ of a simplicial complex Σ 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 Σ.
- The simplices are those of the form $\langle v_{i_0}^\pm, \ldots, v_{i_k}^\pm \rangle$, where $\langle v_{i_0}, \ldots, v_{i_k} \rangle \in \Sigma$.

Note that

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

Lemma 20. *The link of the unique vertex x_0 of S_Σ is isomorphic to $D(\Sigma)$*

Proof. For $v \in \Sigma^{(0)}$, 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 Σ iff the only face of $[0,1]^{\Sigma^{(0)}}$ 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 $\{v_0 \ldots v_n\}$. □

The double of a flag complex is flag so by the link condition S_Σ is non-positively curved and A_Σ is CAT(0). Thus, right-angled Artin groups and their Salvetti complexes give examples of non-positively curved spaces with very general links. It turns out that their subgroups display interesting homological behaviour (see [BB]), and also

Theorem 21. *Right-angled Artin groups embed into $GL_n\mathbb{Z}$ (where n depends on N).*

Unfortunately exploring these in detail would be beyond the scope and space of this essay. Instead, we will conclude by showing that subgroups of right-angled Artin groups are precisely fundamental groups of 'nice' cube complexes.

4.2 Hyperplanes and their pathologies

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

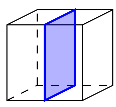

Now if X is a non-positively curved cube complex, and M_1, M_2 are midcubes of cubes in X, we say $M_1 \sim M_2$ if they have a common face, and extend this to an equivalence relation. The equivalence classes are *immersed hyperplanes*. We usually visualize these as the union of all the midcubes in the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

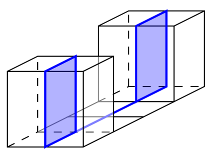

Note that in general, these immersed hyperplanes can have self-intersections, hence the word ”immersed”. Thus, an immersed hyperplane can be thought of as a locally isometric map $H \looparrowright X$, where H is a cube complex.
In general, these immersed hyperplanes can have several “pathologies”. The 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 $\{M \times (-1, 1) \mid M$ 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:

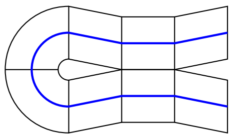

This is a *direct self-osculation* (when H is two sided and a pair of points in the same component of the boundary have the same image in X). We can also have *indirect osculations* (different components of the boundary) that look like this:

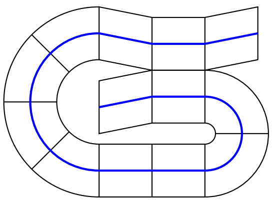

Finally, we have *inter-osculations*, which look roughly like this:

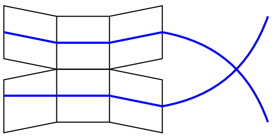

Definition. A cube complex is *special* if its hyperplanes do not exhibit any of the following four pathologies:

- 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. Traditionally, the way to exhibit a surface as a cube complex is to first tile it by right-angled polygons, so that every vertex has degree 4, and then the dual exhibits the surface as a cube complex. The advantage of this approach is that the hyperplanes are exactly the edges in the original tiling! From this, one checks that we in fact have a special curve complex.

This is one example, but it is quite nice to have infinitely many. Since covers of special things are special we already have infinitely many special cube complexes. But we want others.
For each 1-cell e of X there is a unique dual hyperplane, which we will denote by H_e. Two 1-cells that intersect the same hyperplane are called parallel. Hence, hyperplanes correspond exactly to parallelism classes of 1-cells.

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

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

To prove the theorem we need some setup. Define the *hyperplane graph* of a cube complex X, denoted by $\mathbb{H}(X)$, to be a graph whose vertex set is the set of hyperplanes of X; two hyperplanes are joined by an edge if and only if they intersect.
Let X be a cube complex with all walls two-sided. Fix a choice of transverse orientation on each hyperplane H of X, and for x_0 the unique vertex of $\mathcal{S}_{\mathbb{H}(X)}$, fix an identification $Lk(x_0) \cong D(\mathbb{H}(X))$. We can now define a *characteristic map*

$$\phi : X \to \mathcal{S}_{\mathbb{H}(X)}$$

as follows

1. The 0-skeleton $X^{(0)}$ is sent to the unique 0-cell of $\mathcal{S}_{\mathbb{H}(X)}$

2. An oriented 1-cell e of X determines a unique hyperplane H_e with a choice of transverse orientation. The hyperplane H_e corresponds to a vertex in $\mathbb{H}(X)$, which determines a 1-cell $\hat{e}$ of $\mathcal{S}_{\mathbb{H}(X)}$. If the orientation of e coincides with the fixed transverse orientation of H_e then orient $\hat{e}$ pointing from H_e^- to H_e^+; otherwise, give it the reverse orientation. The map ϕ sends e to $\hat{e}$, preserving orientations.

3. A higher dimensional cube C in X, spanned by edges $e_1, \ldots, e_n$, is sent by ϕ to the unique cube in $\mathcal{S}_{\mathbb{H}(X)}$ spanned by $\phi(e_1), \ldots, \phi(e_n)$, preserving orientations. Note that the necessary cube exists because $H_{e_1}, \ldots, H_{e_n}$ pairwise intersect and so span a simplex in $\mathbb{H}(X)$.

The part which requires thought is the second point. If one wants to define a local isometry $\phi : X \to \mathcal{S}_{\mathbb{H}(X)}$ it is natural to require $n-$cells to be sent to $n-$cells for every n. There is only one way to map the $0-cells$ and having decided how the $1-$cells are mapped the choices for the rest of the cells follow naturally.
Now it suffices to show that the characteristic map $\phi_X : X \to \mathcal{S}_{\mathbb{H}(X)}$ is a local isometry.

Proof. Every point has a neighbourhood isometric to the cone on the link so it is sufficient to check that ϕ_X induces injection on links of vertices $x \in X$. Indeed, because no hyperplanes self-intersect, no pair of distinct vertices at distance $\frac{\pi}{2}$ in $Lk_X(x)$ are identified in $Lk_{\mathcal{S}_{\mathbb{H}(X)}}$; likewise, because no hyperplanes self-osculate, no pair of distinct vertices at distance greater than $\frac{\pi}{2}$ in $Lk_X(x)$ are identified by ϕ. Since no pairs of hyperplanes inter-osculate, it follows that there is no pair of vertices $u, v \in Lk_X(x)$ with $d(u, v) > \frac{\pi}{2}$ but $d(\phi(u), \phi(v)) = \frac{\pi}{2}$. It now follows by induction on the length of the shortest non-injective path that ϕ is injective on $Lk_X(x)$, as required. □

Corollary 23. *The characteristic map lifts to an isometric embedding $\tilde{\phi} : \tilde{X} \to \tilde{\mathcal{S}_{\mathbb{H}(X)}}$. In particular, it induces an injective homomorphism $\phi_* : \pi_1 X \hookrightarrow A_{\mathbb{H}(X)}$.*

Proof. Consider the lift $\tilde{\phi} : \tilde{X} \to \tilde{\mathcal{S}_{\mathbb{H}(X)}}$. Let γ be a geodesic path in $\tilde{X}$ such that $\tilde{\phi} \circ \gamma$ is a loop in $\tilde{\mathcal{S}_{\mathbb{H}(X)}}$. Then $\tilde{\phi} \circ \gamma$ is a local geodesic, and so must be constant, because $\tilde{\mathcal{S}_{\mathbb{H}(X)}}$ is CAT(0) and so each based homotopy class contains a unique locally geodesic representative. In particular, the action of $\pi_1(X)$, when pushed forward by ϕ_*, is free on $\tilde{\mathcal{S}_{\mathbb{H}(X)}}$. □

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 24. *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.*

We have come full circle back to the introduction. Wise proved that a large class of hyperbolic 3-manifold groups are virtually the fundamental groups of compact special cube complexes (this means that they have a finite index subgroup which is the fundamental group of some compact special cube complex; more generally a group *virtually* has some property P if a finite index subgroup has that property). This linked two of Thurston's questions on 3-manifold topology and attracted the attention of Agol, who was among the first to realise the importance of and eventually proved one of Wise's conjectures that answered the questions.

Appendix

In these largely proof-free appendices various interesting related topics are introduced.

A Indecisiveness

For many students who are just beginning to learn abstract mathematics, group theory is one of the topics found most difficult. After learning lots more group theory, the student who still thinks group theory is hard is the one who understands group theory the best.
Groups often arise naturally, e.g. as a result of the Seifert-van Kampen theorem, as presentations of the form $< S|R >$, where S is a set of generators and R is a set of words in the alphabet S which act as relators. A natural question arises: given a word in the generators, is it possible to determine algorithmically whether the word is in fact the trivial element of the group? This is important in topology: given a word representing an element of the fundamental group of some space, whether it is trivial is equivalent to whether the loop is null-homotopic. Naively, one might hope that by applying the relations in some clever way this question would be decided. Unfortunately:

Theorem 25 (Novikov-Boone)**.** *There exists a group given by finitely many generators and relators (said to be finitely presented) for which there is no algorithm deciding whether two words u, v represent the same element in the group.*

It gets worse.

Definition. A *Markov property P* of groups is a property such that

1. P is preserved under group isomorphism.
2. There exists a finitely presentable group A_+ with property P.
3. There exists a finitely presentable group A_- which cannot be embedded as a subgroup in any finitely presentable group with property P.

Intuitively, most useful, interesting properties preserved under group isomorphism are exhibited by some finitely presentable group, so of course,

Theorem 26 (Adian-Rabin)**.** *Let P be a Markov property of finitely presentable groups. Then there does not exist an algorithm that, given a finite presentation $G = \langle X \mid R \rangle$, decides whether or not the group G defined by this presentation has property P.*

As a consequence, one cannot decide whether a given group is trivial, finite, abelian, finitely generated free or nilpotent, finitely presented solvable, torsion-free, or amenable. The despair-inducing list goes on, and, for better or for worse, so do the connections with geometry and topology.
Dimension four holds a special place in geometric topology for being particularly difficult to understand. It is the unique dimension for which the same topological manifold can admit infinitely non-diffeomorphic smooth structures. It is also the first dimension where attempts at complete classification of manifolds of that dimension are doomed to fail.

Theorem 27. *Let $n \geq 4$ and $G = \langle S \mid R \rangle$ be a finitely-presented group. Then we can construct a closed, smooth, orientable $n-$manifold M such that $\pi_1 M \cong G$.*

Proof. Let $S = \{a_1, \dots, a_m\}$ and $R = \{r_1, \dots, r_n\}$. We start with

$$M_0 = \#_{i=0}^{m}(S^1 \times S^{n-1}).$$

Note that when we perform this construction, as $n \geq 3$ so S^{n-1} is simply connected, we have

$$\pi_1 M_0 \cong F_m$$

by Seifert–van Kampen theorem. We now construct M_k from M_{k-1} such that

$$\pi_1 M_k \cong \langle a_1, \dots, a_m \mid r_1, \dots, r_k \rangle .$$

We realize r_k as a loop in M_{k-1}. Because $n \geq 3$, we may assume (by diff geo and alg top continuous maps are homotopic to smooth ones) that this is represented by a smooth embedded map $r_k : S^1 \to M_{k-1}$.

We take N_k to be a smooth tubular neighbourhood of r_k. Then $N_k \cong S^1 \times D^{n-1} \subseteq M_{k-1}$. Note that $\partial N_k \cong S^1 \times S^{n-2}$.

Let $U_k = D^2 \times S^{n-2}$. Notice that $\partial U_k \cong \partial N_k$. Since $n \geq 4$, we know U_k is simply connected. So we let

$$M'_{k-1} = M_k \setminus \mathring{N}_k,$$

a manifold with boundary $S^1 \times S^{n-2}$. Choose an orientation-reversing [7] diffeomorphism $\varphi_k : \partial U_k \to \partial M'_{k-1}$. Let

$$M_k = M'_{k-1} \cup_{\varphi_k} U_k.$$

Then by applying Seifert van Kampen repeatedly, we see that

$$\pi_1 M_k = \pi_1 M_{k-1} / \langle\langle r_k \rangle\rangle ,$$

as desired. □

Thus, classifying 4-manifolds would solve the word problem, which is a contradiction. Hence 4-manifold topologists have had to adjust their expectations. Even smooth manifolds, said to be nice topological spaces, are hard to understand.
Now for some positive results.

Greendlinger's lemma

The material here is based on [LS]

Definition. Let $G = \langle S|R \rangle$ be a group presentation where $R \subseteq F(X)$ is a set of freely reduced and cyclically reduced words in the free group $F(X)$ such that R is *symmetrized*, that is, closed under taking cyclic permutations and inverses.

A nontrivial freely reduced word u in $F(X)$ is called a *piece* with respect to the presentation if there exist two distinct elements r_1, r_2 in R that have u as maximal common initial segment.

[7] this is just to make sure the oriented manifolds glue together compatibly

Note that any presentation can be symmetrized by just adding in the cyclic permutations and inverses without altering the isomorphism type of the group.

Definition. Let $0 < \lambda < 1$. The presentation is said to satisfy the $C'(\lambda)$ small cancellation condition if whenever u is a piece w.r.t. the presentation and u is a subword of some $r \in R$, then $|u| < \lambda |r|$. Here $|v|$ is the length of a word v.

A piece is a basically a subword that appears twice (possibly in the same relator). Note that small cancellation is a property of a presentation, not a group, e.g. $\langle a, b | ab^{-1} \rangle$ is a small cancellation presentation for $\mathbb{Z}$ but $\langle a, b | ab^{-1}, ab^{-1}ab^{-1} \rangle$ is not.

Lemma 28 (Greendlinger). *Let G be a group presentation as above satisfying the $C'(\lambda)$ small cancellation condition where $0 \leq \lambda \leq 1/6$. Let $w \in F(X)$ be a nontrivial freely reduced word such that $w = 1$ in G. Then there is a subword v of w and a defining relator $r \in R$ such that v is also a subword of r and such that $|v| > (1 - 3\lambda)|r|$*

Note that the assumption $\lambda \leq \frac{1}{6}$ implies that $(1 - 3\lambda) \geq \frac{1}{2}$, so that w contains a subword more than half the length of some defining relator. In some sense this captures a lot of what can go wrong if one tries to cleverly apply the relators to decide whether a word is trivial: unless the correct choices are made, which may be hard if there are lots of choices, the length of the word might go up and down and never actually end up at the identity. Greendlinger's lemma shows that for a certain class of groups this doesn't happen, and these groups end up being quite important. In particular, the following algorithm, known as the Dehn algorithm, solves the word problem for $C'(\frac{1}{6})$ groups.
Dehn's algorithm:
Given a freely reduced word w on $X^{\pm 1}$, construct a sequence of freely reduced words $w = w_0, w_1, w_2, \dots,$ as follows.
Suppose w_j is already constructed. If it is the empty word, terminate the algorithm. Otherwise check if w_j contains a subword v such that v is also a subword of some defining relator $r = vu \in R$ such that $|v| > |r|/2$. If no, terminate the algorithm with output w_j. If yes, replace v by u^{-1} in w_j, then freely reduce, denote the resulting freely reduced word by w_{j+1} and go to the next step of the algorithm.
Note that we always have $|w_0| > |w_1| > |w_2| > \dots$ which implies that the process must terminate in at most $|w|$ steps. Moreover, all the words w_j represent the same element of G as does w and hence if the process terminates with the empty word, then w represents the identity element of G.

Hyperbolicity and the word problem

Definition. Let S be a generating set for a group G. The Cayley graph of G with respect to S is the 1-skeleton of the universal cover of any presentation complex for G with generators S. Equivalently, $Cay_S(G)$ is the graph with vertex set G and with an edge joining g_1 and g_2 for each $s \in S^{\pm 1}$ such that $g_1 = g_2$.

From now let P be presentation for some group G.

Definition. For a word $w \in F(S)$ that represents the trivial element in G, define $Area_P(w)$ to be the minimal N such that

$$w = \prod_{k=1}^{N} g_k r_{j_k} g_k^{-1}$$

that is, the area is the smallest number of conjugates of relators needed to prove that w represents the trivial element.

If one attaches a 2-cell for every relator to the corresponding loops in the Cayley graph, then the area is the area bounded by the loop

Definition. The *Dehn function* $\delta_P : \mathbb{N} \to \mathbb{N}$ is defined by

$$\delta_P(n) = max_{l_S(w) \leq n} Area_P(w)$$

where w ranges over words which represent the trivial element.

These are in fact intrinsic properties of a group which don't depend on the specific presentation, if one allows a slightly coarser (but still sensible!) notion of equivalence of functions and metric spaces than same order of growth or isometry.

Definition. A map of metric spaces $f : X \to Y$ is a (λ, ϵ)-*quasi-isometric embedding* if for all $x_1, x_2 \in X$,

$$\frac{1}{\lambda} d_X(x_1, x_2) - \epsilon \leq d_Y(f(x_1), f(x_2)) \leq \frac{1}{\lambda} d_X(x_1, x_2) + \epsilon$$

If it is also *quasi-surjective*, i.e. for all $y \in Y$ there is some $x \in X$ such that $d_Y(f(x), y) \leq \epsilon$, then f is said to be a *quasi-isometry* and X, Y are said to be quasi-isometric.

This is basically up to bi-Lipschitz functions with a translation.
For functions, the new equivalence relation is given as follows: $f \preceq g$ iff there is a constant C such that $f(n) \leq Cg(Cn + C) + Cn + C$, and $f \simeq g$ iff $f \preceq g$ and $g \preceq f$, and the equivalence class of the Dehn function is invariant up to quasi-isometry.

Lemma 29. *The word problem is solvable iff the Dehn function is computable.*

Definition. Let M be a Riemannian manifold. For a null-homotopic loop γ, $Area_M(\gamma)$ is the infimal area of a filling disc for γ. The filling function $Fill_M : [0, \infty) \to [0, \infty)$ is defined by

$$Fill_M(x) := sup_{l(\gamma) \leq x} Area_M(\gamma)$$

Theorem 30. *If M is a closed Riemannian manifold then $Fill_M \simeq \delta_{\pi_1(M)}$*

Manifolds of non-positive curvature in fact satisfy a quadratic isoperimetric inequality which means that their filling function is computable, so their fundamental groups have solvable word problem. Hyperbolicity, or at least non-negative curvature, is good.

B Gromov-Hyperbolicity

Throughout let X be a geodesic metric space.

Definition. A geodesic triangle is said to be *δ-slim* if each side is contained in the δ neighbourhood of the union of the other two sides.

Definition. A metric space X is said to be *δ-hyperbolic* or *Gromov hyperbolic* if every geodesic triangle is $\delta-$slim

As mentioned before the hyperbolic plane is $\delta-$slim. A tree is trivially $0-$slim, and is as hyperbolic as any metric space can get. Any definition of hyperbolicity should include trees, and groups acting on trees give rise to a rich and beautiful theory known as Bass-Serre theory. See for instance Serre's fantastic book *Trees* for details.
CAT(-1) spaces are hyperbolic, but note that δ-hyperbolicity only constrains large triangles since triangles of diameter less than δ is trivially $\delta-$slim, while CAT(κ) is a restriction on the 'small' triangles too.

Theorem 31. *Let X, Y be geodesic metric spaces. If X is hyperbolic and Y is quasi-isometric to X, then Y is hyperbolic too.*

This is proved using the next lemma. As a consequence, many graphs, which are quasi-isometric to trees, are hyperbolic.
A $(\lambda, \epsilon)-$quasi-geodesic in X is a $(\lambda, \epsilon)-$quasi-isometric embedding of an arc into X.

Lemma 32 (Morse)**.** *Suppose X is δ-hyperbolic. Any $(\lambda, \epsilon)-$quasi-geodesic from x to y is contained in the $R-$neighbourhood of a geodesic from x to y, where R is constant dependent on λ, ϵ, and δ.*

An important consequence of the theorem is the following:

Example. If Σ is a closed surface of constant Gaussian curvature -1 then the universal cover of Σ is $\mathbb{H}^2$, so for any finite generating set, the Cayley graph of $\pi_1\Sigma$ is quasi-isometric to the hyperbolic plane and hence is Gromov-hyperbolic

Definition. A finitely generated group is called (word-)hyperbolic if its Cayley graph is Gromov-hyperbolic.

Any CAT(-1) group is Gromov-hyperbolic by above, but whether the converse holds is an open question (it isn't true that hyperbolic spaces are necessarily CAT(-1)).
The theorem implies that a group which is quasi-isometric to a hyperbolic group is hyperbolic. Another result is that a group is hyperbolic iff its Dehn function is linear. Neither of these are true for CAT(0) groups. As far as group theory is concerned, Gromov-hyperbolicity appears to be the most natural condition for now, which explains why CAT(-1) groups receive relatively little attention.

C Artin and Coxeter groups

The material here is based on [Ch].
Coxeter, a giant of geometry, was interested in a type of generalised reflection

group defined as follows:

$$W = \langle s_1 \dots s_n | s_i^2 = 1, (s_i s_j)^{m_{ij}} = 1 \rangle$$

for some $m_{ij} = m_{ji}$ integers at least 2 or ∞, in which case we omit the relation between s_i and s_j. This can be summarised in a symmetric matrix M so W is also denoted C_M. The reason they are considered generalised reflection groups is

Theorem 33. *Define a bilinear form on $\mathbb{R}^n$ by $\langle e_i, e_j \rangle = -\cos \frac{\pi}{m_{ij}}$. Setting $r_i(x) = x - 2\langle e_i, x \rangle e_i$ defines a faithful representation of C_M, making them linear over $\mathbb{R}$.*

By making all $m_{ij}, i \neq j$ either 2 or ∞ one obtains what is known as a right-angled Coxeter group which is linear over $\mathbb{Z}$ by looking at the coefficients in the definition of the representation above. It turns out that every right-angled Coxeter group is virtually a subgroup of a right-angled Artin group.
A general Artin group is defined as

$$A = \langle s_1 \dots s_n | (s_i s_j)^{m_{ij}} = 1 \text{for all} i \neq j \rangle$$

which the reader will notice looks very similar to the definition of a Coxeter group. To each matrix M one can associate a Coxeter group and an Artin group with a natural surjective homomorphism from the latter to the former. There are many more connections between the two classes of groups.
An Artin group is said to be *spherical* or of *finite type* if the associated Coxeter group is finite and *non-spherical* or *infinite type* otherwise. Spherical Artin groups are better understood. They have the following properties:

- linear [Bi] [CW]
- torsion-free [BK]
- infinite cyclic centre [BK]
- have solvable word and conjugacy problem (the conjugacy problem is the question of determining whether two elements are conjugate) [BK] [Ch92]

None of the above are known to hold for general Artin groups, but the word problem *is* solvable for right-angled Artin groups. However it remains to be seen whether Artin groups, spherical or not, are CAT(0).

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] M. Bestvina and N. Brady. Morse theory and finiteness properties of groups. Inventiones Mathematicae, 129:445–470, 1997

[Bi] S. Bigelow, Braid groups are linear J. Amer. Math. Soc. 14 (2001), no. 2, 471–486

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

[BK] E. Brieskorn and K. Saito, Artin-gruppen und Coxeter-gruppen, Invent. Math. 17 (1972), 245-271.

[Ch92] R. Charney, Artin groups of finite type are biautomatic, Math. Annalen 292 (1992), 671-683.

[Ch] R. Charney, An introduction to right-angled Artin groups

[CW] A. Cohen and D. Wales, Linearity of Artin groups of finite type, Isreal J. of Math. 131 (2002), 101–123

[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

[Gro93] M. Gromov. Asymptotic invariants of infinite groups. In Geometric group theory, Vol. 2 (Sussex, 1991), pages 1–295. Cambridge Univ. Press, Cambridge, 1993.

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

[HW04] C. Hruska and D.T. Wise. Axioms for finiteness of cubulations. Preprint, 2004.

[LS] R. C. Lyndon and P. Schupp, Combinatorial group theory. Reprint of the 1977 edition. Classics in Mathematics. Springer-Verlag, Berlin, 2001. ISBN 3-540-41158-5.

[OW09] Y.Ollivier and D.T. Wise, Cubulating random groups at density less than $\frac{1}{6}$

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