

CONSTRUCTING NON-POSITIVELY CURVED SPACES AND GROUPS

JON MCCAMMOND¹

ABSTRACT. The theory of non-positively curved spaces and groups is tremendously powerful and has enormous consequences for the groups and spaces involved. Nevertheless, our ability to construct examples to which the theory can be applied has been severely limited by an inability to test – in real time – whether a random finite piecewise Euclidean complex is non-positively curved. In this article I focus on the question of how to construct examples of non-positively curved spaces and groups, highlighting in particular the boundary between what is currently do-able and what is not yet feasible. Since this is intended primarily as a survey, the key ideas are merely sketched with references pointing the interested reader to the original articles.

Over the past decade or so, the consequences of non-positive curvature for geometric group theorists have been thoroughly investigated, most prominently in the book by Bridson and Haefliger [26]. See also the recent review article by Kleiner in the Bulletin of the AMS [59] and the related books by Ballmann [4], Ballmann-Gromov-Schroeder [5] and the original long article by Gromov [48]. In this article I focus not on the consequences of the theory, but rather on the question of how to construct examples to which it applies. The structure of the article roughly follows the structure of the lectures I gave during the Durham symposium with the four parts corresponding to the four talks. Part 1 introduces the key problems and presents some basic decidability results, Part 2 focuses on practical algorithms and low-dimensional complexes, Part 3 presents case studies involving special classes of groups such as Artin groups, small cancellation groups and ample-twisted face pairing 3-manifolds. In Part 4 I explore the weaker notion of conformal non-positive curvature and introduce the notion of an angled n -complex. As will become clear, the topics covered have a definite bias towards research in which I have played some role. These are naturally the results with which I am most familiar and I hope the reader will pardon this lack of a more impartial perspective.

Part 1. Negative curvature

Although the definitions of δ -hyperbolic and CAT(0) spaces / groups are well-known, it is perhaps under-appreciated that there are at least seven potentially distinct classes of groups which all have some claim to the description “negatively-curved.” Sections 1 and 2 briefly review the definitions needed in order to describe these classes, the known relationships between them, and the current status of the (rather optimistic) conjecture that all seven of these classes are identical. Since the emphasis in this article is on the construction of examples, finite piecewise

Date: August 9, 2004.

¹Partially supported by NSF grant no. DMS-0101506

Euclidean complexes play a starring role. Section 3 describes how, in theory, one can decide whether such a complex is non-positively curved and Section 4 concludes with a few comments about the length spectrum of such a complex.

1. VARIETIES OF NEGATIVE CURVATURE

In this section I define the basic objects of study and highlight the distinctions between classes of groups which are often treated interchangeably by new-comers to the field. Because the main definitions and results in this section are well-known, the presentation is brief and impressionistic. A more detailed overview covering similar territory can be found in the expository article by Martin Bridson [23] or the one by Ruth Charney [35]. The most basic objects are δ -hyperbolic spaces and $\text{CAT}(\kappa)$ spaces. Recall that a *geodesic* is a length minimizing curve and that a *geodesic metric space* is a metric space in which every pair of points is connected by a geodesic.




FIGURE 1. A triangle which is not quite δ -thin.

Definition 1.1 (δ -hyperbolic). A geodesic metric space X is δ -hyperbolic if there is a fixed $\delta \geq 0$ such that for all points $x, y, z \in X$ and for all geodesics connecting x, y , and z and for all points p on the chosen geodesic connecting x to y , the distance from p to the union of the other two geodesics is at most δ .




FIGURE 2. A geodesic triangle and the corresponding triangle in \mathbb{R}^2 .

Definition 1.2 (CAT(κ)). A geodesic metric space X is called (globally) CAT(0) if for all points $x, y, z \in X$, for all geodesics connecting x , y , and z and for all points p on the chosen geodesic connecting x to y , $d(p, z) \leq d(p', z')$ where the second distance is calculated in the corresponding configuration in the Euclidean plane. A *corresponding configuration* is the unique configuration (up to isometry) of labeled points in \mathbb{R}^2 where the distance in X between labeled points matches the Euclidean distance between the corresponding labeled points in \mathbb{R}^2 . See Figure 2 for an illustration. The conditions CAT(1) and CAT(-1) are defined similarly using \mathbb{S}^2 and \mathbb{H}^2 , respectively - with restrictions on x , y , and z in the spherical case, since not all spherical comparison triangles are constructible. More generally, CAT(κ) is defined using the complete simply-connected 2-dimensional space with constant curvature κ . See [26] for further details.

As should be clear from the definition, increasing the value of κ weakens the conditions that need to be satisfied. In particular, the following result is immediate.

Theorem 1.3 (Increasing κ). *If X is a CAT(κ) space, then for all $k' \geq k$, X is also a CAT(κ') space. In particular, every CAT(-1) space is a CAT(0) space.*

Since hyperbolic n -space, \mathbb{H}^n , is a common inspiration for both of these theories of negative curvature, it is not too surprising that it is both δ -hyperbolic and a CAT(-1) space. Notice that the definition of δ -hyperbolic only implies that the large scale curvature is negative: we get no information about local structure. On the other hand, CAT(0) and CAT(-1) spaces have good local curvature properties which can be formalized as follows.

Definition 1.4 (Locally CAT(κ)). A complete geodesic metric space is *locally* CAT(κ) if every point has a neighborhood which satisfies the CAT(κ) condition in the induced local metric.

Note that the completeness of the metric is assumed primarily to rule out such pathologies as the lateral surface of a right circular cone minus its apex. This is a perfectly good locally CAT(0) space (except for completeness), but its universal cover is not even a geodesic metric space. Assuming completeness excludes such oddities. The connection between this local CAT(κ) definition and the previous global definition is now straightforward: it is known as the Cartan-Hadamard Theorem. See [26, p. 193].

Theorem 1.5 (Local-to-global). *If X is a complete geodesic metric space and $\kappa \leq 0$ is a fixed constant, then the space X is CAT(κ) if and only if X is locally CAT(κ) and simply-connected.*

Notice that one consequence of this lemma is that a local examination of a compact complex is sufficient to establish that its universal cover satisfies the global version. In keeping with recent practice, the phrases “locally CAT(0)” and “non-positively curved” are used interchangeably. Similarly, the phrase “locally CAT(-1)” is often equated with the phrase “negatively curved”, but since I am trying to highlight the various classes of spaces and groups with a legitimate claim to this title, the phrase “negatively curved” is never used in this sense in this article.

Shifting attention from spaces to groups, consider those groups which act properly and cocompactly by isometries on these types of spaces. Following Cannon [33] I call such an action *geometric*. Of the seven classes of groups to be delineated, one acts geometrically on δ -hyperbolic spaces and the other six on CAT(κ)

spaces with various restrictions. A group G is *hyperbolic* (or word-hyperbolic or Gromov-hyperbolic) if for some δ it acts geometrically on some δ -hyperbolic space. Similarly, a group G is a $\text{CAT}(\kappa)$ *group* if it acts geometrically on some $\text{CAT}(\kappa)$ space. Deciding whether a group is word-hyperbolic is made easier by the following result.

Lemma 1.6 (Testing hyperbolicity). *If A is a finite generating set for a group G , then G is word-hyperbolic if and only if there exists a δ such that the Cayley graph of G with respect to A is δ -hyperbolic.*

Thus hyperbolicity can be detected from a close examination of the metric space associated with any finite generating set. By way of contrast, showing that a group is $\text{CAT}(\kappa)$ first requires the construction of an appropriate $\text{CAT}(\kappa)$ space and there is no natural candidate with which to begin. This lack of a natural space on which the group should act is one of the main sources of difficulty for proving (and especially for disproving) that a group is a $\text{CAT}(\kappa)$ group.

The first type of restriction on $\text{CAT}(\kappa)$ groups I wish to consider involves the existence of flats in the spaces upon which they act. A *flat* is an isometric embedding of a Euclidean space \mathbb{R}^n with $n > 1$. Because an isometrically embedded flat plane would violate the thin triangle condition, every $\text{CAT}(-1)$ space is also a $\text{CAT}(0)$ space with no flats. Also, every $\text{CAT}(-1)$ space is δ -hyperbolic. The strongest result along these lines is the flat plane theorem. See [26, Theorem 3.1 on p. 459] or [22].

Theorem 1.7 (Flat Plane Theorem). *If X is a $\text{CAT}(0)$ -space and G is a group which acts geometrically on X , then X is δ -hyperbolic (and G is hyperbolic) if and only if X does not contain an isometrically embedded copy of the Euclidean plane.*

Notice that the quotient of X by its G -action is a non-positively-curved space with fundamental group G and universal cover X . Conversely, if we start with a compact non-positively curved space with fundamental group G , then G acts geometrically on its universal cover. The rough algebraic equivalent to the existence of an isometrically embedded flat plane is a rank 2 free abelian subgroup. The following is a small portion of the result known as the flat torus theorem ([26, Theorem 7.1 on p. 244]).

Theorem 1.8 (Flat Torus Theorem). *Let X be a $\text{CAT}(0)$ -space and let G be a group which acts geometrically on X . If G contains a $\mathbb{Z} \times \mathbb{Z}$ subgroup then X contains an isometrically embedded flat plane.*

An obvious question at this point is whether the converse also holds. Unfortunately this question is wide open.

Problem 1.9 (Flat vs. $\mathbb{Z} \times \mathbb{Z}$). *If X is a $\text{CAT}(0)$ -space and a group G acts geometrically on X , is it true that there is an isometrically embedded flat plane in X if and only if there is a $\mathbb{Z} \times \mathbb{Z}$ subgroup in G ?*

A good test case is a compact non-positively curved 2-complex constructed entirely out of unit Euclidean squares. Problem 1.9 remains open even in this highly restricted setting and it is not even clear that the answer should be yes. One phenomenon which frustrates many naive approaches to this problem is the following construction due to Dani Wise [86].

Theorem 1.10 (Isolated aperiodic flats). *There exists a compact piecewise Euclidean non-positively curved space constructed entirely out of unit Euclidean squares which contains an aperiodic flat that is not the limit of periodic flats.*

A *periodic flat* is one on which a $\mathbb{Z} \times \mathbb{Z}$ subgroup of G acts cocompactly. Alternatively, a flat is *aperiodic* if no such subgroup exists. The final assertion of the theorem essentially means that there does not exist a sequence of periodic flats whose intersections with the aperiodic flat increase so as to exhaust the aperiodic flat in the limit. The specific construction can be found in Section 7 of [85] or in [86]. Wise's example shows that attempts to prove that the existence of a flat in X implies the existence of a $\mathbb{Z} \times \mathbb{Z}$ subgroup in G cannot start by closely examining a randomly chosen flat in X with the hope of showing that it is periodic – or even that there is a periodic flat nearby. The route from the flat to the subgroup, if indeed it exists, is thus likely to start from a flat which has been carefully selected to satisfy additional conditions.

The final type of restriction I wish to consider severely constrains the type of spaces on which the group G can act. In particular, suppose we restrict our attention to the spaces constructed out of polyhedral pieces taken from one of the three standard constant curvature models: the n -sphere \mathbb{S}^n , Euclidean n -space \mathbb{R}^n , or hyperbolic n -space \mathbb{H}^n . A rough definition of these M_κ -complexes is given below. See [26] for a more precise definition.

Definition 1.11 (M_κ -complexes). A *piecewise spherical / euclidean / hyperbolic complex* X is a polyhedral complex in which each polytope is given a metric with constant curvature $1 / 0 / -1$ and the induced metrics agree on overlaps. In the spherical case, the cells must be convex polyhedral cells in \mathbb{S}^n in the sense that they must be embeddable in an open hemisphere of \mathbb{S}^n . The generic term is M_κ -complex, where κ is the curvature common to each of its cells.

It is an early foundational result of Bridson that compact M_κ -complexes are indeed geodesic metric spaces [21, 24]. In fact, Bridson showed that noncompact M_κ -complexes are also geodesic metric spaces so long as one assumes that there are only finitely many isometry types of cells. For ease of exposition, we restrict our attention to compact complexes. A second key result is that one can check whether an M_κ -complex ($\kappa \leq 0$) is non-positively curved by checking for the existence of short geodesic loops in the links of cells.

Definition 1.12 (Links). Let X be an M_κ -complex. The *link of a point* x in X is the set of unit tangent vectors at x . Notice that the link of x comes equipped with a natural piecewise spherical structure. The *link of a cell* is the set of unit tangent vectors at one of its points, but restricted to those which are orthogonal to all of the tangent vectors lying in the cell. For example, the link of a point in the interior of an edge in a tetrahedron is a spherical lune, but the link of the edge which contains it is a metric circular arc whose length is equal to the dihedral angle along this edge.

Definition 1.13 (Geodesics). A *piecewise geodesic* in X is a path $\gamma : [a, b] \rightarrow X$ where $[a, b]$ can be subdivided into a finite number of subintervals so that the restriction of γ to each closed subinterval is a path lying entirely in some closed cell σ of X . There is a further restriction that this portion of the path is the unique geodesic connecting its endpoints in the metric of σ . A *local geodesic* is a local

isometric embedding of an interval into X and a *closed geodesic loop* is a local isometric embedding of a metric circle into X . In an M_κ -complex, the structure of a local geodesic (or a closed geodesic loop) is always that of a piecewise geodesic, and to test whether a piecewise geodesic is a local geodesic it is sufficient to check whether at each of the transition points, the “angles are large”, meaning that the distance between the in-coming tangent vector and the out-going tangent vector is at least π in the link of this point.

Using this language the precise statement is the following.

Theorem 1.14 (Gromov’s link condition). *For $\kappa \leq 0$, an M_κ -complex is locally $\text{CAT}(\kappa) \Leftrightarrow$ the link of each vertex is globally $\text{CAT}(1) \Leftrightarrow$ the link of each cell is an piecewise spherical complex which contains no closed geodesic loop of length less than 2π .*

Thus, showing that piecewise Euclidean complexes are non-positively curved hinges on showing that piecewise spherical complexes have no short closed geodesic loops. I return to this theme in Section 3. Finally, notice that a compact M_κ -complex can be described using only a finite amount of data. This observation is important since it makes M_κ -complexes suitable for computer investigation. More specifically, it is sufficient to give the curvature κ , the simplicial structure of (a simplicial subdivision of) the complex together with the exact length of each of its edges. This is sufficient because the shape of each simplex (of curvature κ) is completely determined by its edge length data. This section concludes with an exercise.

Exercise 1.15 (Feasible edge lengths). What restrictions on a set of edge lengths ℓ_{ij} , $1 \leq i, j \leq n$, are necessary in order for a nondegenerate piecewise Euclidean n -simplex to be constructible with these lengths? The answer is classical, but the necessary and sufficient conditions are not necessarily easily reconstructed from first principles. Historically, the full answer consists of a set of equations (due to Cayley) and some inequalities (due to Menger) [69], although an alternative description in terms of quadratic forms was found by Schoenberg [79].

2. CURVATURE CONJECTURE

The primary goal of this section is to state and discuss the curvature conjecture. In order to state the conjecture seven potentially distinct classes of groups need to be defined.

Definition 2.1 (7 classes of groups). Consider the set of all groups G which act geometrically on

1. δ -hyperbolic spaces,
2. $\text{CAT}(-1)$ spaces,
3. $\text{CAT}(0)$ spaces with no isometrically embedded flat planes,
4. $\text{CAT}(0)$ spaces with no $\mathbb{Z} \times \mathbb{Z}$ subgroups in G ,
5. piecewise hyperbolic $\text{CAT}(-1)$ spaces,
6. piecewise Euclidean $\text{CAT}(0)$ spaces with no isometrically embedded flat planes, or
7. piecewise Euclidean $\text{CAT}(0)$ spaces with no $\mathbb{Z} \times \mathbb{Z}$ subgroups in G .

Notice that first two classes define word-hyperbolic groups and $\text{CAT}(-1)$ groups, respectively. The known relationships between these classes are summarized in Figure 3. In the figure, PH and PE stand for piecewise hyperbolic and piecewise Euclidean, respectively.




FIGURE 3. The known relationships between various classes of groups

Although the following is formulated as a conjecture, I should note that many experts in the area firmly believe the conjecture is false.

Conjecture 2.2 (Curvature conjecture). *The seven classes of groups described above are identical.*

Conjecture 2.2 is a natural analogue of Thurston's hyperbolization conjecture. In particular, if the hyperbolization conjecture is true then this conjecture is true when restricted to 3-manifolds and their fundamental groups. As we noted above, many experts expect the general conjecture to be false. A resolution either way would be a major advance for geometric group theory and its quest to understand the nature of negative curvature.

Perhaps the most likely source of a counterexample to Conjecture 2.2 is a quotient of a cocompact lattice in a high-dimensional quaternionic hyperbolic space by a long relation, the idea being that the non-existence of a $\text{CAT}(0)$ space on which this group acts geometrically can be established by a suitable extension of super-rigidity. Several groups of researchers have been working towards resolving this conjecture either positively or negatively (Mineyev-Yu [70], Burger-Monod [27, 28], etc.). Nevertheless, even should this approach yield a counterexample, the more general question of *which* word-hyperbolic groups can act geometrically on $\text{CAT}(0)$ complexes will remain wide open. It is at least conceivable that essentially the only counterexamples are ones constructible from quaternionic hyperbolic spaces. As (admittedly rather weak) evidence for this, there is the well-known propensity of the quaternions and octonions to be involved whenever there are only a few exceptions to an otherwise cleanly stated classification theorem [2, 41].

One aspect of Figure 3 which might be surprising at first is the lack of a relationship between piecewise hyperbolic $\text{CAT}(-1)$ groups and piecewise Euclidean $\text{CAT}(0)$ groups. A naive approach would be to simply replace each hyperbolic simplex with the Euclidean simplex which has the same edge lengths. This procedure does tend to lengthen piecewise geodesics in vertex links (where the break points are described using barycentric coordinates), but the problem is that the status of such a path as a local geodesic is not stable under this inflation. In particular, there might exist a short piecewise geodesic loop in the hyperbolic version which is not a local geodesic, but which becomes a slightly longer (although still less than 2π) local geodesic in the Euclidean version. Thus, the path that shows the Euclidean

version is not CAT(0) may not even be considered in the hyperbolic version since it is not a local geodesic. Another more serious problem is that the effect of this replacement on the length of piecewise geodesics in links of edges and of other higher dimensional faces is completely unclear: some may lengthen, others may shrink, depending on the intricate details of the shapes of the original hyperbolic simplices. A partial positive result in this direction has been shown by Charney, Davis and Moussong [34].

Theorem 2.3 (hyperbolic vs. PE CAT(0)). *If M is a compact n -dimensional manifold which supports a hyperbolic metric (i.e. a Riemannian metric with constant sectional curvature -1), then M also supports a metric which is piecewise Euclidean and locally CAT(0).*

Their proof looks at the orbit of a point in the hyperboloid model of hyperbolic space under the action of the fundamental group of M and then considers the piecewise Euclidean boundary of the convex hull of these points. They then quotient this piecewise Euclidean complex by the piecewise linear action of the fundamental group. Notice, however, that this proof scheme is highly dependent on the fact that M is a constant curvature manifold. The corresponding result for compact n -manifolds with *variable* negative Riemannian curvature or a locally CAT(-1) piecewise hyperbolic structure is still open. In the other direction, there is also a partial result, at least in low-dimensions.

Proposition 2.4 (No tight loops). *If M is a 2-dimensional piecewise Euclidean CAT(0) complex, and there does not exist a closed geodesic loop in the link of a vertex in M which has length exactly 2π , then M can also be given a piecewise hyperbolic metric which is CAT(-1).*

The argument is easy and goes roughly as follows. Give each simplex the metric of a simplex with identical edge lengths but with constant curvature ϵ where ϵ is negative but close to 0. By choosing ϵ sufficiently close to 0, the length of the closed geodesics in the vertex links (which are metric graphs in this case) do not shrink below 2π . Thus by Gromov's criterion (Theorem 1.14), the result is CAT(ϵ). Rescaling now produces a metric of curvature -1 . The restriction that M be 2-dimensional is crucial. As far as I can tell from the literature, the following stronger conjecture has neither been established nor disproved.

Conjecture 2.5 (No tight loops). *Suppose M is a piecewise Euclidean CAT(0) complex, and there does not exist a closed geodesic loop of length exactly 2π in the link of any simplex in M , is it true that this particular cell structure for M can also be given a piecewise hyperbolic metric which is CAT(-1)?*

As mentioned above, the potential problems lie in the links of faces other than vertices since a simple substitution of slightly hyperbolic simplices for the Euclidean ones has unpredictable effects on the metrics in these links. Note also that the no-tight-loops restriction in Proposition 2.4 is absolutely necessary for this result to hold in dimension 2. Specifically, Noel Brady and John Crisp [12] and Misha Kapovich [58] have produced several examples with the following properties.

Theorem 2.6 (PH CAT(-1) vs. PE CAT(0)). *There exists groups which acts geometrically on a 3-dimensional piecewise hyperbolic CAT(-1) spaces, and on a 2-dimensional piecewise Euclidean CAT(0) spaces, but not on any 2-dimensional piecewise hyperbolic CAT(-1) space.*




FIGURE 4. A one-relator group satisfying Theorem 2.6.

The simplest example of such a group is the one-relator group shown in Figure 4. The moral here is that although it is conceivable that the class of piecewise hyperbolic $CAT(-1)$ groups and the class of piecewise Euclidean $CAT(0)$ groups without flat planes are identical, a change in cell structure – and dimension! – is sometimes necessary in order to pass between them. In other words, higher dimensions are sometimes necessary in order to smooth things out. In particular, the space which proves a group is $CAT(0)$ might have a higher dimension than its geometric dimension [12, 25] and, if $CAT(-1)$ is desired, another space of even higher dimension might be required. Despite the exceptions under special circumstances noted above (Theorem 2.3 and Proposition 2.4) the behavior noted by Brady-Crisp and Kapovich is likely to be fairly common. For an explicit illustration where higher dimensions are useful see Section 9. On the other hand, the passage to higher-dimensions is not without perils of its own, as we shall see.

If our goal is to create complexes with good local curvature properties on which a particular word-hyperbolic group can act, there is at least one obvious candidate: the Rips complex, or some variant on it. Let $P_d(G, A)$ denote the flag complex on the graph whose vertices are labeled by G and which has an edge connecting g and h iff gh^{-1} is represented by a word of length at most d in the generators A . Recall that a *flag complex* is a simplicial complex which contains a simplex if and only if it contains the 1-skeleton of that simplex. This complex has extremely nice properties for large values of d [1].

Theorem 2.7 (Rips complex). *If G is word-hyperbolic and d is large relative to the hyperbolicity constant δ , then the Rips complex $P_d(G, A)$ is contractible, finite dimensional, and G acts discretely and compactly on $P_d(G, A)$.*

As defined above, the Rips complex is a simplicial complex with no natural metric. One approach to the curvature conjecture would be to try and add a metric to the Rips complex and then to show that the result is a $CAT(0)$ space. Unfortunately, this can be quite difficult even in the simplest of cases. Let G be a word-hyperbolic group and suppose we carefully pick a generating set A , pick a d very large and declare each simplex in $P_d(G, A)$ to be a regular Euclidean simplex with every edge length 1. Is the result a $CAT(0)$ space? Is this true when G is free and A is a basis? Although it is embarrassing to admit, I believe the answer is that no one knows, even in the case of a free group generated by a basis. The moral here is that our ability to test whether a compact constant curvature metric space is locally $CAT(0)$ or locally $CAT(-1)$ is *very* primitive.

Finally, there is another class of groups which deserves to be highlighted in connection with the seven classes listed above, even though they are not negatively

curved. These are the groups which act geometrically on CAT(0) spaces with isolated flats.

Remark 2.8 (Spaces with isolated flats). A CAT(0) space X has the *isolated flats property* if there is a collection \mathcal{F} of isometrically embedded Euclidean spaces (each of dimension at least 2) such that this collection is maximal and isolated. It should be maximal in the sense that every flat in X lies in a k -neighborhood of a flat in \mathcal{F} , and isolated in that k -neighborhoods of distinct flats in \mathcal{F} intersect in bounded sets whose diameter is a function of k alone. This definition was introduced and extensively studied by Chris Hruska in his dissertation [52], although it was implicit in earlier works by other authors [57, 85]. There are a number of results for hyperbolic groups which readily extend to groups acting geometrically on CAT(0) spaces with isolated flats, but which do not extend to CAT(0) groups in general. For example, if G is such a group (and it also has the closely related relative fellow traveler property), then there is a well-defined notion of a quasiconvex subgroup in G which is independent of the CAT(0) space with isolated flats on which G acts and there is a well-defined boundary up to homeomorphism. Both of these results were shown by Gromov for hyperbolic groups but fail for general CAT(0) groups.

3. DECIDABILITY

When constructing examples of non-positively curved spaces, attention has naturally focused on the class of constant curvature complexes. This attention is partly due, no doubt, to the reduction allowed by Gromov's link condition (Theorem 1.14) and to the fact that these complexes can be described using only a finite amount of data. But this leads to a fundamental question. Let X be a finite M_κ -complex, say piecewise Euclidean. Is there an algorithm to determine whether this metric space is non-positively curved? The answer, fortunately, is yes, but the proof is not as straightforward as one might hope. In [43] Murray Elder and I proved the following result.

Theorem 3.1 (CAT(κ) is decidable). *Given a compact M_κ -complex, there is an explicit algorithm which decides whether it is locally CAT(κ). In particular, it is possible to determine if a finite piecewise Euclidean complex is non-positively curved.*

The word ‘‘algorithm’’ in the theorem should really be placed in quotation marks since the algorithm we described in the course of the proof is not one which can be used in practice except in the most trivial of situations where the answer is already obvious. The difficulty is that the proof relies, ultimately, on a result of Tarski's about the decidability of the first order theory of the reals. In outline, the proof can be described as follows. First, reduce the problem to the existence of certain types of paths in special piecewise spherical complexes that we call circular galleries. Next, encode the existence of these paths into a real semi-algebraic set. Recall that a *real semi-algebraic set* is a subset of Euclidean space which can be described using Boolean combination of polynomial equations and inequalities. The name is derived from the fact that these subsets are also boolean combinations of (projections of) real algebraic varieties. Finally, Tarski's theorem about the decidability of first order sentences over the reals implies that there is an algorithm which decides whether a real semi-algebraic set is empty or not. The rest of this section describes the proof in more detail. The key concept is that of a gallery.

Definition 3.2 (Galleries). Let X be a piecewise constant curvature complex and let γ be a local geodesic in X . Consider the sequence of open cells through which γ passes. This sequence is essentially what we call a *linear gallery*. If γ is a closed geodesic, then this sequence is given a cyclic ordering and the sequence is called a *circular gallery*. Instead of a precise definition, consider the example shown in Figure 5. The figure in the upper left shows a very simple complex (which is essentially the boundaries of two Euclidean tetrahedra identified along one edge) along with a local geodesic path from x to y . On the upper right is the corresponding linear gallery. The lower left is result of concatenating the closed cells through which the local geodesic passes rather than the open cells, and the difference between these two is the *boundary* of the gallery. The circular version constructs a similar complex homotopy equivalent to a circle.




FIGURE 5. A 2-complex, a linear gallery, its closure and its boundary.

The reasoning now goes as follows. If there is a cell in X whose link is a piecewise spherical complex containing a closed geodesic loop of length less than 2π , then this geodesic determines a circular gallery. Moreover, since X is finite, this circular gallery involves only finitely many shapes. There is a result of Bridson [21] which asserts that in this context the length of the geodesic is quasi-isometric to the number of simplices in the circular gallery and that the quasi-isometry constants only depends on the shapes of the cells involved. Thus, if we are able to (1) check whether any particular piecewise spherical circular gallery contains a short closed geodesic loop with winding number 1 and (2) determine whether any of the closed geodesic loops found remain locally geodesic when the circular gallery is immersed back into the cell link from which it came, then we are done. Carrying out this procedure for each of the finite number of circular galleries containing a bounded number of simplices completes the test. To highlight the dependence of the argument on the quasi-isometry constant derived from the shapes, consider the situation where the metric is not given in advance.

Remark 3.3 (Testing for the existence of a CAT(0) metric). If X is a finite cell complex endowed with a piecewise Euclidean metric then the argument outlined




FIGURE 6. A 1-complex, a linear gallery and its model space.

above can determine whether or not this particular metric is non-positively curved. It remains an open question, however, (in dimensions greater than 2) whether there is an algorithm to determine whether a finite simplicial complex (with no pre-assigned metric) supports a metric of non-positive curvature. The problem is that without knowing in advance what the metric is, it is not clear what the quasi-isometry constants are, and thus it is not clear how large the circular galleries are which need to be considered. In particular, there is no *a priori* bound on the combinatorial length of the galleries we need to check. Dimension 2 is special in this regard because its vertex links are metric graphs. Thus, the only closed circular galleries which need to be considered are the finite number of simple closed loops in these graphs – regardless of the metric. Therefore, a (highly impractical) algorithm does exist in dimension 2.

The remainder of the argument is a complicated conversion process which encodes the existence of a closed geodesic loops in a circular gallery into a Boolean combination of polynomial equations and inequalities. The full details can be found in [43], but as an indication of the issues involved, here are two brief remarks. First, even determining whether a piecewise geodesic in a piecewise spherical complex is longer than π is a non-trivial task using only polynomials. Analytically, one would simply think of the transition points as unit vectors and add up the arccosines of the appropriate dot products. Since arccosines are prohibited, we create instead a 2-dimensional “model space” such as the one pictured in Figure 6 and then test whether the second coordinate of the final point plotted is positive.

The second remark is that to conclude that a particular piecewise geodesic is a local geodesic necessarily involves an induction on dimension: it is a local geodesic if and only if at each transition point the distance between in the in-coming tangent vector and the out-going tangent vector are distance at least π apart in the link

of this point. To show this one would need to show that for every possible linear gallery (up to a certain combinatorial length) connecting these two points in this lower dimensional piecewise spherical complex and for every possible local geodesic in each of these galleries, the distance (calculated using a 2-dimensional model space) is at least π . By an induction on dimension, one can assume that the appropriate logical combination of equations and inequalities exists to test each of these conditions, completing the proof. The result is that by inducting through dimensions it is possible to construct a real semi-algebraic set whose points are in one-to-one correspondence with the closed geodesic loops of length less than 2π in the specific circular gallery under consideration. Tarski's theorem can then be used to test whether this set is empty or not.




FIGURE 7. A regular triangular prism.

Remark 3.4 (Why is this so hard?). At this point it would be natural to wonder whether it is really necessary to use real semi-algebraic sets in order to decide this type of problem. Perhaps there is an elementary solution waiting to be discovered. Perhaps. On the other hand, the problem of determining whether a complex is $CAT(0)$ involves finding 1-dimensional paths in high dimensional piecewise spherical complexes, and problems involving high codimension (where “high” typically means codimension 2 or more) can often be surprisingly hard. As an example, consider the question of finding the unit volume Euclidean 3-polytope with the smallest 1-skeleton (where the size of the 1-skeleton is measured by adding up all of the edge lengths). The answer, in this case, is not known, although the solution is probably the regular triangular prism shown in Figure 7 (all edges have equal length, all faces are regular, and rescaled to have unit volume). This is known as the Melzak problem and it is a classical problem in geometric analysis which has stubbornly resisted all attempts to resolve it. For a summary of recent partial progress see [76].

4. LENGTH SPECTRUM

The ability (in theory) to test whether a finite constant curvature complex is non-positively curved is not the only result which can be shown using the methodology described in the previous section. A similar approach allows one to prove that length spectrum of such a complex must be discrete.

Definition 4.1 (Length spectrum). Let X be a compact M_κ -complex and let x and y be points in X . The lengths of all the local geodesics from x to y is called the length spectrum from x to y . The lengths of all the closed geodesic loops in X is simply its *length spectrum*.

In [26], Bridson and Haefliger use a relatively easy compactness argument to establish the discreteness of the length spectrum between specific points. Noel Brady and I were recently able to show that the same restriction holds for lengths of all closed geodesic loops as well [13]. The proof, on the other hand, requires a type of argument which is reminiscent of the proof of Theorem 3.1.

Theorem 4.2 (Length spectrum). *The length spectrum of a compact M_κ -complex is discrete.*

As one might expect, the compactness hypothesis can be weakened to finitely many isometry types of cells without altering the proof. A sketch of the proof goes as follows. If the length spectrum is not discrete then we can find a sequence of closed geodesic loops in X whose lengths are distinct but bounded. By passing to a subsequence we can assume that all of these loops determine the same circular gallery (this reduction uses the quasi-isometry constant alluded to above). By establishing that the function which calculates the length of certain piecewise geodesic paths is analytic, one can use the Morse-Sard theorem to show that the length spectrum for this particular circular gallery has measure zero in \mathbb{R} . Moreover, by encoding these piecewise geodesics as a real semi-algebraic set, one can show that the topology on the set of closed geodesic loops has only finitely many connected components. Combining these two ingredients shows that the length spectrum of this particular circular gallery is finite, providing the necessary contradiction.

One reason for proving a result of this type has to do with locally geodesic surfaces. Call an isometric immersion $f : D \rightarrow X$ of a metric polyhedral surface *locally geodesic* if for all points $d \in D$, the link of d is sent to a local geodesic in the link of $f(d)$. In a 2-dimensional piecewise Euclidean complex, every null-homotopic curve bounds a locally geodesic surface, but this fails in dimensions 3 and higher. This is also one of the key reasons why some theorems in dimension 2 fail to generalize easily to higher dimensions. Chris Hruska, for example, proves in [52] that in dimension 2 the isolated flats property is equivalent to the relative fellow travel property and to the relatively thin triangles property. He conjectures, moreover, that the same result should hold in all dimensions. In his proof, he heavily uses the fact that the curvature of points in locally geodesic surfaces mapped into finite 2-dimensional complexes is bounded away from 0. By the *curvature of a point* in the interior of D we mean 2π minus the length of its link. See also Section 12 where a more general definition of curvature is discussed. An immediate corollary of Theorem 4.2 is that this portion of his proof immediately generalizes to arbitrary dimensions.

Corollary 4.3 (Quantizing curvature). *If $D \rightarrow X$ is a locally geodesic surface in a compact piecewise Euclidean complex X , then those points in the interior of D with nonzero curvature have curvatures uniformly bounded away from 0. Moreover, this bound depends only on X and not on D . As a consequence, if the total amount of negative (or positive) curvature is known, then there is a bound on the number of points in D whose curvature is less than (or more than) zero.*

Part 2. Algorithms

In this part the focus shifts from theoretical results to practical algorithms which can actually be implemented on computers to obtain results. In Section 5 a practical algorithm in dimension 3 is discussed and in Section 6 this algorithm

is used to solve a conjecture about 3-manifolds. Finally, in Section 7 the inherent difficulties in higher dimensions are considered.

5. ALGORITHM IN DIMENSION 3

As pointed out above, the use of Tarski’s theorem in the course of the proof makes the implementation of the “algorithm” described in Theorem 3.1 impractical if one is really interested in actually carrying out a test on a particular concrete example. Thus, it makes sense to continue to search for alternative approaches in those special cases where real algebraic geometry can be avoided. In [45] Murray Elder and I were able to find a rather elementary geometric argument when the complexes under consideration are restricted to dimension at most 3.

Theorem 5.1 (Practical algorithm). *There exists an elementary practical algorithm which decides whether or not a 3-dimensional piecewise Euclidean complex is non-positively curved.*

As with the word “algorithm” in Theorem 3.1, the word “practical” in Theorem 5.1 should really be in quotation marks since the procedure is reasonably fast for each circular gallery, but the number of circular galleries which need to be tested in even a modest size example can quickly grow to unmanageable levels. The results described in Section 6 illustrate this phenomenon. The heart of the argument uses only elementary 3-dimensional geometry and has been implemented in both GAP (Groups, Algorithms, Programming) and PARI (a number theory package). The packages are called `cat.g` and `cat.gp`, and are available to download [44].

The key to the argument is the restricted nature of circular galleries in low dimensions. If X is a 3-dimensional piecewise Euclidean complex, the only links of cells which need to be examined are the links of vertices. The piecewise spherical complexes which result are 2-dimensional and the circular galleries to which they lead can be classified as either annular galleries, Möbius galleries or necklace galleries.

Definition 5.2 (Types of galleries). If a hypothetical geodesic does not pass through a vertex of the link, then the circular gallery it produces will be homeomorphic to either an annulus or a Möbius strip, whereas those which pass through a vertex look like a beaded necklace. We call these *annular galleries*, *Möbius galleries*, and *necklace galleries*, respectively. The portion of a necklace gallery from one vertex to the next is called a *bead*.

Annular galleries and Möbius galleries can be “cut open” and then developed onto a 2-sphere in an essentially unique way up to isometry. Once this is done, an easy construction using dot products and cross products enables one to test whether or not the original circular gallery contained a short closed geodesic loop. Similarly, each bead in a necklace gallery can be developed onto a 2-sphere in an essentially unique way and there is an easy procedure to check whether there is a local geodesic in this bead which connects its endpoints. These processes are illustrated in Figures 8, 9 and 10. See [45] for further details.

As mentioned above, the main reason why the algorithm described is not completely practical is because of the sheer number of circular galleries which need to be examined in any reasonable example. In practice, this search can be restricted to unshrinkable geodesics. A geodesic is *unshrinkable* if there does not exist a homotopy through rectifiable curves of non-increasing length to a curve of strictly




FIGURE 8. An annular gallery, cut open and developed.




FIGURE 9. A Möbius gallery, cut open and developed.




FIGURE 10. A necklace gallery with one of its beads developed.

shorter length. One corollary of a result established by Brian Bowditch in [11] is the following.

Theorem 5.3 (Unshrinkable geodesics). *If an M_κ -complex (with $\kappa \leq 0$) contains a cell whose link contains a closed geodesic loop of length less than 2π , then it also contains a cell whose link contains an unshrinkable closed geodesic loop of length less than 2π .*

Using a slightly more delicate argument, we can then show that when testing whether an M_κ -complex X is non-positively curved it is sufficient to test whether there are piecewise spherical circular galleries arising from the links of cells in X which contain closed geodesic loops that can neither be shrunk nor homotoped until it meets the boundary of its gallery without increasing length. All annular galleries, Möbius galleries of length at least π , and beads of length at least π are shrinkable in this expanded sense. Figure 11 illustrates the main ideas. The left




FIGURE 11. Shrinking annular galleries and long beads.

hand side shows how a geodesic in a cut-open and developed annular gallery can be homotoped through latitude lines to a strictly shorter curve, while the right hand side shows how a geodesic in a bead of length at least π can be homotoped in a non-length-increasing way to a piecewise geodesic which intersects the boundary of the bead. This reduction simplifies the search for short geodesics immensely.

This section concludes with two examples showing the type of results produced by the software implementing this algorithm.




FIGURE 12. 5 Möbius galleries.




FIGURE 13. 2 non-trivial beads.

Example 5.4 (Regular tetrahedra). Let X be a piecewise Euclidean complex built entirely out of regular tetrahedra with each edge length equal to 1. Each of the vertex links in X is piecewise spherical 2-complexes constructed out of equilateral spherical triangles with side length $\pi/3$. The software produces a list of circular galleries constructed out of triangles of this type which can carry short closed geodesic loops. The task of determining whether any of these circular galleries can actually be immersed into any of the vertex links of X and, more importantly, whether any of the closed geodesic loops they carry remain local geodesics under these immersions has not been implemented. Currently, the researcher must check these conditions by hand.

In the situation under consideration, the circular galleries carrying unshrinkable short closed geodesics consist of 22 necklace galleries and the 5 Möbius galleries shown schematically in Figure 12. In order to produce the Möbius strips the thick edges on the left and the right should be identified with a half-twist. The 22

necklaces are those which can be strung together using the one edge type (which is a spherical arc of length $\pi/3$) and the two non-trivial beads shown in Figure 13. If we call these beads types A , B , and C , respectively, then the 22 necklaces which are can strung together to contain a closed geodesic loop of length less than 2π are described by the sequences $A, A^2, A^3, A^4, A^5, B, BA, BA^2, BA^3, BA^4, B^2, B^2A, B^2A^2, BABA, B^3, C, CA, CA^2, CA^3, CB, CBA$, and C^2 . The final result is that X is non-positively curved if and only if it does not contain a vertex whose link contains an isometric immersion of one of these 5 Möbius galleries or one of these 22 necklaces galleries where the short closed geodesic loop it contains survives as a local geodesic under this immersion.



FIGURE 14. A Möbius gallery to avoid.



FIGURE 15. A non-trivial bead to avoid.

Example 5.5 (Coxeter shapes). Let X be a piecewise Euclidean 3-complex built entirely out of metric tetrahedra with 4 edges of length $\sqrt{3}$ and two non-adjacent edges of length 2. Tetrahedra of this type arise in the theory of Euclidean Coxeter groups. In particular, there is a regular tiling of \mathbb{R}^3 using these tetrahedra. In the standard notation it is an \tilde{A}_3 shape. The vertex links in X will be piecewise spherical 2-complexes built out of isometric isosceles spherical triangles. This is because all four vertices of our standard tetrahedron have isometric links. The software in this case produces a single Möbius gallery and a list of 19 necklace galleries. The Möbius gallery is shown in Figure 14.

The 19 necklaces are again strung together using 3 beads. In this case, there are two types of edges (of different lengths) in the vertex links and a single non-trivial bead shown in Figure 15. Labeling these A , B , and C as before, the 19 necklace galleries which can be strung together to contain a closed geodesic loop of length less than 2π are those described by the sequences $A^2, A^4, A^6, A^2B, A^2B^2, A^2B^3, ABAC, A^2C, A^2C^2, A^4B, CA^4 B, B^2, B^3, B^4, B^5, C, C^2$, and C^3 . The final result is that X is non-positively curved if and only if it does not contain a vertex whose link contains an isometric immersion of this Möbius gallery or one of these 19 necklaces galleries where the short closed geodesic loop it contains survives as a local geodesic under this immersion.

6. 3-MANIFOLD RESULTS

In this section we discuss how the software described in the previous section can be used in conjunction with a sort of spherical small cancellation theory to prove a result about 3-manifolds originally conjectured by Bill Thurston. The theorem is a result is about $5/6^*$ manifolds.

Definition 6.1 ($5/6^*$ -triangulations). Let M be a closed triangulated 3-manifold. The triangulation is a $5/6$ -triangulation if every edge has degree 5 or 6 (where the degree of an edge equals the combinatorial length of its link). The triangulation is a $5/6^*$ -triangulation if, in addition, each 2-cell contains at most one edge of degree 5.

The conjecture that Thurston makes is that every closed 3-manifold with a $5/6^*$ -triangulation has a word-hyperbolic fundamental group. In [46] Murray Elder, John Meier and I proved a slightly stronger result.

Theorem 6.2. *Every $5/6^*$ -triangulation of a closed 3-manifold M admits a piecewise Euclidean metric of non-positive curvature with no isometrically embedded flat planes in its universal cover. Thus $\pi_1 M$ is word-hyperbolic.*

The proof involves a mixture of traditional combinatorial group theory and computations carried out by the software `cat.g`. The combinatorial group theory portion is described first. The link of a vertex in M is a triangulated 2-sphere in which every vertex has valence 5 or 6 and no two vertices of valence 5 are connected by an edge. Thus, the dual is a tiling of the 2-sphere by pentagons and hexagons in which every vertex has valence 3 and no two pentagons have an edge in common. Since the smallest tiling with these properties is that of a soccer ball, we call these *soccer tilings* and a simply-connected subdiagram of a soccer tiling homeomorphic to a disc is called a *soccer diagram*. A typical example is shown in Figure 16.




FIGURE 16. A soccer diagram.

Curvature considerations imply that every soccer tiling has exactly 12 pentagons. If γ is a simple closed curve in the 1-skeleton of a soccer tiling, cutting along this curve splits the 2-sphere into two soccer diagrams, at least one of which has at most 6 pentagons. Since every vertex in a soccer tiling has valence 3, the vertices on the boundary of a soccer diagram have valence either two or three. The boundary vertices of valence 2 are called *right turns* and those of valence 3 are called *left turns*. The terminology arises from imagining that every corner of each polygon has been assigned an angle of $2\pi/3$ and we are traversing the boundary cycle clockwise. An analysis of soccer diagrams reminiscent of small cancellation theory yields the following key result.

Lemma 6.3 (Special diagrams). *The only soccer diagrams D with $\partial D \leq 14$, at most six pentagons, and no three consecutive right turns are the two diagrams shown in Figure 6.*




FIGURE 17. The only two soccer diagrams satisfying the conditions of Lemma 6.3

This combinatorial result reappears at the end of the argument. Returning to the triangulated 3-manifold M , we assign a piecewise Euclidean metric to M as follows. Each edge of degree 5 is assigned a length of 2 and each edge of degree 6 is assigned a length of $\sqrt{3}$. The corresponding metrics on the 2-cells and 3-cells are the unique Euclidean metrics with these specific edge lengths. Since no tetrahedron in M contains edges of degree 5 which are adjacent, there are exactly three types of metric tetrahedra which can be found in M . These are shown in Figure 18. The thicker edges are the ones of length 2. Notice that the tetrahedron on the left is regular and the one on the right is the Coxeter shape considered in Example 5.5.




FIGURE 18. The 3 metric tetrahedra which arise in a $5/6^*$ 3-manifold.

It is easy to show that the links of edges are metric circles with length at least 2π by calculating the sizes of the various dihedral angles. Thus, the only question which remains is whether the links of vertices contain closed geodesic loops of length less than 2π . Since annular galleries are shrinkable and Möbius galleries cannot immerse into 2-spheres, only necklace galleries need to be examined. Using the software `cat.g`, we calculated the list of beads which can be used to string together a necklace gallery containing a short closed geodesic loop. Using the simplified search algorithm that only looks for beads whose geodesic has length less than π , it takes less than an hour to produce the list of 75 beads. In the process, the program examines over 116,000 circular galleries. Before restricting the search to unshrinkable geodesics, the program needed to search through more than 3 million circular galleries and it found more than a thousand beads which could be used. Moreover, performing this search required over 2 months of computation time, and this in a problem which only involves three metric tetrahedra! The problem under consideration is thus only barely feasible using current technology. Slight modifications are likely to quickly lead to intractable complications. In this instance, the output of 75 beads is small enough to be successfully analyzed. In fact, all 75 beads in the output belong to one of the three combinatorial types shown in Figure 19. More specifically, the output consisted of 4 edges of varying length, 26 beads consisting of exactly 2 triangles as




FIGURE 19. The three combinatorial types of beads in the output.

shown in the middle figure, and 45 beads with 4 triangles as shown in the figure on the right. The different beads of each type were combinatorially identical but metrically distinct.

Labeling the beads of each type as type A , B , and C , it was possible to calculate the length of the shortest geodesic in each case. The results are given in Table 1. Using this explicit output, we can string together a rough list of all possible necklace types.

Type	Minimum length
A	$.302\pi$
B	$.5\pi$
C	$.833\pi$

TABLE 1. The minimum geodesic length for each type.

If M is not non-positively curved, then there is a vertex link in M which contains a short closed geodesic loop. This loop determines one of the combinatorial types of necklaces galleries in our list. Moreover, it is possible to perturb this loop slightly so that it misses all of the vertices in the link. The circular gallery determined by this new path will be an annular gallery, and it is straightforward to argue using the list of combinatorial types that the annular gallery which results contains at most 14 triangles. As a result, the dual of this annular gallery is a simple closed curve of length at most 14 in the 1-skeleton of a soccer tiling of \mathbb{S}^2 , and an analysis of dihedral angles shows that no three consecutive left or right turns can occur as this path is traversed.

Using Lemma 6.3 we can now conclude that if a short geodesic exists, this process leads to an approximation of this geodesic which bounds one of the two diagrams shown in Figure 6 in the dual soccer tiling. A careful inspection of the possible left/right sequences coming from our three explicit (combinatorial) beads shows this is impossible. This contradiction proves that vertex links cannot contain short closed geodesic loops and that M is locally $CAT(0)$. Finally, the argument that the universal cover of M contains no isometrically embedded flat planes is relatively easy, but omitted. See [46] for details.

Remark 6.4 (The role of computers). One aspect of this proof scheme that I find particularly appealing is the way in which computer computations and traditional proofs are intimately intertwined. After a metric is assigned, the computer performs a series of calculations, the output provides leads to an investigation of a particular type of disc diagram, and finally, following this analysis, the consequences must be

compared back with the computer output in order to reach the final contradiction. It does not seem too far-fetched to predict that in the near future, this type of human-computer interaction will be increasingly common.

The ubiquity of $5/6^*$ triangulations is suggested by the following result obtained in collaboration with Noel Brady and John Meier.

Theorem 6.5 (Bounding edge degrees). *Every closed orientable 3-manifold has a triangulation in which each edge has degree 4, 5, or 6.*

The proof essentially uses the universality of the figure eight knot complement and a careful triangulation of various 3-manifolds with boundary that we use as building blocks. See [15] for details.

Remark 6.6 (Foams). As a final remark, we note that the structures that originally prompted Thurston to consider these types of 3-manifolds are *positively curved* 3-manifolds called foams. A *foam* is a $5/6$ -triangulation in which no 2-cell contain more than one edge of degree 6. The name and the conditions are derived from the structure of the bubbles in chemical foams. See [81] for details and further references.

7. HIGHER DIMENSIONS

Although the focus of this part has been on low-dimensional complexes, there are, of course, some classes of higher dimensional complexes where easy algorithms are known which check non-positive curvature. Most of these follow from either Gromov's lemma or its generalization known as Moussong's lemma [71]. Recall that a simplicial complex is *flag* if every 1-skeleton of a simplex is actually filled with a simplex, and a metric complex is *metric flag* if every 1-skeleton of a metrically feasible simplex is filled with a simplex.

Lemma 7.1 (Gromov's lemma). *If every edge in a piecewise spherical complex K has length $\pi/2$, then K is globally CAT(1) if and only if it is a flag complex. In particular, a piecewise Euclidean complex built out of unit cubes (of various dimensions) is non-positively curved if and only if its vertex links are flag complexes.*

Lemma 7.2 (Moussong's lemma). *Suppose every edge in a piecewise spherical complex K has length at least $\pi/2$. Then K is globally CAT(1) if and only if it is a metric flag complex.*

As a consequence of his lemma, Moussong was able to establish that all Coxeter groups are CAT(0) groups. By way of contrast, no such easily checked combinatorial conditions have been found for any class of high dimensional piecewise Euclidean *simplicial* complexes. One might initially harbor the hope that every CAT(-1) group could be made to act geometrically on a non-positively piecewise Euclidean cube complex. Graham Niblo and Lawrence Reeves [74] have shown that this is not the case.

Theorem 7.3 (Cube complexes are insufficient). *There exist locally CAT(-1) Riemannian manifolds which are not homotopy equivalent to any finite dimensional, locally CAT(0) cube complex.*

In other words, the class of groups which act geometrically on CAT(0) cube complexes is strictly smaller than the class of CAT(0) groups. As an aside, we

note that the manifolds alluded to in the theorem are constructed using cocompact lattices in quaternionic hyperbolic space and that they only exist in (real) dimension 8 and higher. As a consequence of Theorem 7.3, working with general simplicial complexes is probably inevitable when attacking questions such as the curvature conjecture.

Finally, Section 3 concluded with a remark trying to explain why testing for short closed geodesics in high dimensional complexes is difficult. After seeing the 3-dimensional algorithm (Section 5), one might think that searching only for *unshrinkable* closed geodesics might be easier. Unfortunately, even this restricted task is likely to remain intractable.

Remark 7.4 (Why is testing shrinkability hard?). Consider a sequence of n all-right spherical tetrahedra (i.e. all edges have length $\frac{\pi}{2}$), and let \mathcal{G} be the circular gallery formed by identifying an edge in one tetrahedron to an edge in the next in such a way that the two edges used in each tetrahedron are non-adjacent. In this example, every piecewise geodesic determined by selecting one point on each of the shared edges and connecting them has the same length, namely $\frac{n\pi}{2}$. A closer examination would also show that all of these loops are closed geodesics. It is also instructive to consider a slight modification of this example where every edge which is not shared between two tetrahedra has length $\frac{\pi}{2} + \epsilon$. When ϵ is small and positive \mathcal{G} contains no closed geodesic loops, and when ϵ is small and negative \mathcal{G} contains exactly one. This extreme sensitivity to shape is one difficulty inherent in problems of this type.

More generally, consider the same combinatorial configuration but with a more random metric. Even if one were able to find a closed geodesic loop in this complex with winding number 1, it would not necessarily be easy to determine whether it was a shrinkable geodesic. Since each transition point has essentially one degree of freedom, the question of shrinkability can be reframed as a question about whether the Hessian of the multivariable function which calculates length from the positions of the transition points is positive definite or not. A moment's reflection indicates that the Hessian in this case is a tri-diagonal real symmetric matrix since the local distance calculation only depends on the neighboring transition points. Next, one can determine whether the matrix is positive definite by calculating the determinants of its principal minors. One more recoding transforms this sequence of determinants into a sequence of inequalities involving the initial portions of a finite continued fraction. Since continued fractions are notoriously sensitive to small perturbations and rather chaotic in their dependence on the precise numerators and denominators, we can conclude that whether or not a closed geodesic loop in this complex is shrinkable can be quite sensitive to the precise lengths of the edges of these simplices. Under these circumstances, it is highly unlikely that an elementary extension of the algorithm described in Section 5 exists, and even if it does, the restriction of the search to only unshrinkable geodesics is no longer likely to be feasible in practice.

Part 3. Special classes of groups

Since high-dimensional complexes are necessary in order to fully address questions such as the curvature conjecture, and since arbitrary metric simplicial complexes in high dimensions are hard to work with, one option is to restrict our

attention to those special classes of groups where special types of complexes can be used. In this part, I discuss recent progress along these lines for three such classes of groups. Section 8 considers the relationship between non-positive curvature and Artin groups, Section 9 does the same for small cancellation groups, and Section 10 considers non-positive curvature and the fundamental groups of ample twisted face-pairing 3-manifolds.

8. COXETER GROUPS AND ARTIN GROUPS

Ever since Moussong's dissertation it has been known that all Coxeter groups are CAT(0) groups. The situation for Artin groups, however, remains far from clear.

Definition 8.1 (Artin groups and Coxeter groups). Let Γ be a finite graph with edges labeled by integers greater than 1, and let $\langle a, b \rangle^n$ denote the length n prefix of $(ab)^n$. The *Artin group* A_Γ is the group generated by a set in one-to-one correspondence with the vertices of Γ with a relation of the form $\langle a, b \rangle^n = \langle b, a \rangle^n$ whenever a and b correspond to vertices joined by an edge labeled n . The *Coxeter group* W_Γ is the Artin group A_Γ modulo the additional relations $a^2 = 1$ for each generator a . There is also an alternative convention for associating diagrams with Coxeter groups and Artin groups which is derived from a consideration of the finite Coxeter groups. In that case, the graph Γ (using the above convention) is always a complete graph with most of the edges labeled 2 or 3. The alternate convention simplifies the diagram by removing all edges labeled 2 and leaving the label implicit for the edges labeled 3.




FIGURE 20. A labeled graph used to define a Coxeter group and an Artin group.

Example 8.2. Let Γ denote the labeled graph shown in Figure 20. The presentation of the Artin group A_Γ is $\langle a, b, c \mid aba = bab, ac = ca, bcbc = cbcb \rangle$ and the presentation $\langle a, b, c \mid aba = bab, ac = ca, bcbc = cbcb, a^2 = b^2 = c^2 = 1 \rangle$ defines the Coxeter group W_Γ .

Coxeter groups first arose in the classification of finite groups acting on Euclidean space which are generated by reflections (i.e. isometries fixing a codimension 1 hyperplane). In fact, the collection of finite Coxeter groups is the same as the collection of such finite reflection groups. The well-known classification of irreducible Coxeter groups divides them into type A_n ($n \geq 1$), B_n ($n \geq 2$), D_n ($n \geq 4$), E_6 , E_7 , E_8 , F_4 , H_3 , H_4 and $I_2(m)$ ($m \geq 2$). The corresponding diagrams (using the alternative convention) are shown in Figure 21. For more details see some of the standard references for Coxeter groups such as Bourbaki [10], Humphreys [54] or Kane [56].




FIGURE 21. Diagrams for the irreducible finite Coxeter groups.

An Artin group defined by the same labeled graph as a finite Coxeter group is called an *Artin group of finite-type*. There are a number of reasons for believing that all finite-type Artin groups might be CAT(0) groups. Most importantly, they act on several different spaces in ways which almost qualify. First, every finite-type Artin group is, essentially, the fundamental group of a locally Euclidean space derived from a complexified hyperplane arrangement. Being locally Euclidean, it is also locally CAT(0) but this space is neither compact nor complete. There is another space called the Deligne complex which is complete and homotopy equivalent to the one derived from the hyperplane complement, but the action of the Artin group on the universal cover of this new complex is not proper; there are infinite stabilizers. See [35] and the references therein for details. Moreover, Mladen Bestvina has used a weak version of non-positive curvature to show that finite-type Artin groups have essentially all of the expected group-theoretic consequences of non-positive curvature [8]. Finally, there is a finite Eilenberg-MacLane space for the braid groups

constructed independently by Tom Brady [17] and Daan Krammer [60, 61] and generalized to arbitrary Artin groups of finite-type by David Bessis [7] and by Tom Brady and Colum Watt [20]. If a piecewise Euclidean metric is assigned to one of these *Brady-Krammer complexes* then the corresponding Artin group will act geometrically on its universal cover. The only question is whether a metric can be assigned to these complexes which is non-positively curved. Thus, although Artin groups of finite-type do not yet qualify as $CAT(0)$ groups, given these different spaces and actions, they are extremely good candidates.




FIGURE 22. The (noncrossing) partition lattice for $n = 4$.

The general construction of the Brady-Krammer complexes is too complicated to review here, but it is possible to give a rough idea of their structure. In the case of the braid groups, the complex is closely connected with a well-known combinatorial object called the noncrossing partition lattice. For a survey of the connection between braid groups and noncrossing partitions see [64].

Definition 8.3 (Noncrossing partitions). A *noncrossing partition* is a partition of the vertices of a regular n -gon so that the convex hulls of the partitions are disjoint. One noncrossing partition σ is contained in another τ if each block of σ is contained in a block of τ . The noncrossing partition lattice for $n = 4$ is shown in Figure 22. The only “crossing partition” is the one inside the dashed line; the partition $\{\{1, 3\}, \{2, 4\}\}$.

More generally, there is a version of the noncrossing partition lattice for each (irreducible) finite-type Artin group. As an illustration, the partially ordered set for type F_4 is shown in Figure 23. The connection between the poset and the corresponding Brady-Krammer complex is that the geometric realization of this poset is a fundamental domain for the universal cover of the Brady-Krammer complex. In particular, every increasing path from the lowest vertex to the highest vertex in Figure 23 creates a distinct 4-simplex in the fundamental domain for the F_4 Brady-Krammer complex.




FIGURE 23. The poset used to construct the F_4 Brady-Krammer complex.

Although the Brady-Krammer complexes do not come with a metric preassigned, after examining their structure in low-dimensions, one is quickly led to consider a metric that Tom Brady and I have called the “natural metric.” The metric views the edges in a maximal chain as mutually orthogonal steps in a Euclidean space. This metric is natural in the sense that the metric it assigns to the geometric realization of a Boolean lattice is that of a Euclidean cube. See Figure 24. In other words, one of the nicest possible lattices is converted into one of the nicest possible shapes. Also notice that the link of the long diagonal in this n -cube is the Coxeter complex for a symmetric group.




FIGURE 24. The natural metric on the geometric realization of a Boolean lattice.

In several instances Tom Brady and I were able to show that the Brady-Krammer complexes with the natural metric were non-positively curved [16, 19, 18].

Theorem 8.4 (Low-dimensions). *The finite-type Artin groups with at most 3 generators are CAT(0)-groups and the Artin groups A_4 and B_4 are CAT(0) groups.*

When the Artin group has at most 3 generators, the complexes involved are very low dimensional and their links are easy to examine by hand. The A_4 and B_4 cases are slightly more complicated to carry out by hand, but the software `coxeter.g` (described below) is now able to quickly verify that there are no short closed geodesic loops in the various links in this case. Given this evidence, it made sense to conjecture that the Brady-Krammer complexes using the natural metric is non-positively curved for all Artin groups of finite type. Recently, however, a graduate student of mine, Woonjung Choi, disproved this conjecture. In fact, she was able to prove the following stronger result [40].

Theorem 8.5 (D_4 and F_4). *The Brady-Krammer complexes for D_4 and F_4 do not support any piecewise Euclidean metrics which respect the symmetries of their defining diagrams and at the same time are non-positively curved.*

By “respecting the symmetries” she means that all of the symmetries of the defining diagrams should be reflected in the metrics. Since the argument is the same in each case, both complexes are discussed simultaneously. The main idea behind the proof is to first use the cyclic center of these groups, together with the splitting theorem [26, Theorem 6.21 on p. 239], to force there to be a piecewise Euclidean structure on a well-defined 3-dimensional “cross-section” complex. Next, she determined which Euclidean metrics on the tetrahedra in the cross-section complex result in dihedral angles which ensure that the edge links (which are finite graphs) contain no closed geodesic loops of length less than 2π . Because of the sheer size of these computations, this portion of the proof was carried out by GAP using routines that she wrote specifically for this purpose. The symmetries of the defining diagrams were used to cut down on the complexity of the computation. This process drastically reduced the number of metrics which needed to be considered. Finally, for each feasible metric, the one 2-dimensional piecewise spherical vertex link was examined and found to always contains a short closed geodesic loop. Thus, every piecewise Euclidean metric on the cross section complex either produces a short closed geodesic loop in one of the edge links or in the vertex link. As a result, the cross section complex does not support any non-positively curved piecewise Euclidean metric, and this completes the proof.

Remark 8.6 (The software). The file `coxeter.g` is the name of a set of GAP routines used to examine Brady-Krammer complexes and they are available to download [39]. I initially developed some of these routines to test the curvature of the low-dimensional Brady-Krammer complexes exclusively using the natural metric. These early routines were extensively modified and extended by Woonjung Choi so that they are now able to first find the simplicial structure of the 3-dimensional cross-section complex, find representative vertex and edge links (up to automorphism), find the graphs representing the various edge links, find the simple cycles in these graphs, find the linear system of inequalities which need to be satisfied by the dihedral angles of the tetrahedra in order for these simple cycles to have length at least 2π , and finally, to simplify this large system of inequalities by removing redundancies.

When the software is run for the D_4 Artin group, the result is an initial 4-dimensional complex with 162 4-simplices, the cross section consists of 15 tetrahedra, and in the end the software produces a list of 13 simplified inequalities in 9 variables which needs to be analyzed. When the F_4 Artin group is considered, the

software finds 432 4-simplices in the original complex, 18 tetrahedra in the cross section, and produces a system of 27 simplified inequalities in 13 variables which needs to be analyzed.

A key result that Choi uses in her study of piecewise Euclidean metrics on Brady-Krammer complexes is the following fact about dihedral angle rigidity. It is almost assuredly classic, but I cannot find a reference to it in the literature.

Theorem 8.7 (Dihedral angle rigidity). *Let σ and τ be Euclidean n -simplices and let f be a bijection between their vertices. If the dihedral angle at each codimension 2 face of σ is at least as big as the dihedral angle at the corresponding codimension 2 face of τ , then σ and τ are isometric up to a scale factor.*

The proof is remarkably easy. Let \vec{u}_i and \vec{v}_i , $i = 1, \dots, n$ be the external unit normal vectors for the facets of σ and τ , numbered so that corresponding facets receive corresponding subscripts. It is a result due to Minkowski that there are positive numbers $a_i > 0$ such that $\sum_i a_i \vec{u}_i = \vec{0}$. In fact, the numbers a_i can be chosen to be the volume of the i^{th} facet. Noel Brady pointed out to me, that this can also be thought of as a consequence of the divergence theorem applied to this simplex. The proof consists of the following sequence of equations and inequalities.

$$0 = \left\| \sum_i a_i \vec{u}_i \right\|^2 = \sum_i \sum_j a_i a_j (\vec{u}_i \cdot \vec{u}_j) \geq \sum_i \sum_j a_i a_j (\vec{v}_i \cdot \vec{v}_j) = \left\| \sum_i a_i \vec{v}_i \right\|^2 \geq 0$$

The first inequality follows from the assumption on dihedral angles and the positivity of the a_i . Comparing the ends of the sequence we conclude that the inequality must be an equality and that each of the corresponding dihedral angles must be exactly equal in size. Finally, it is easy to show that the collection of dihedral angles determine the isometry type of the simplex up to rescaling.

Remark 8.8 (Non-positive curvature and Brady-Krammer complexes). A quick visual summary of the results to date is given in Figure 25. The circled cases are the ones which are CAT(0) groups (and all of these use the natural metric). The two which are boldly crossed out are the cases which do not support non-positively curved piecewise Euclidean in any reasonable sense. The higher rank examples which are lightly crossed out are also not going to support non-positively curved piecewise Euclidean metrics in any reasonable sense precisely because they contain the D_4 complex as part of their structure. And finally, there is a question mark next to the Artin group of type H_4 since this case is half-way analyzed, but the analysis is not yet complete.

Remark 8.9 (The H_4 complex). The case of H_4 is more difficult to analyze than the D_4 or F_4 cases because its defining diagram has no symmetries. This greatly increases the number of equations and variables involved in the computations. In particular, the software applied to the H_4 Artin group produces a list of 1350 4-simplices in the original complex, 23 tetrahedra in the cross section complex, and it produces a list of 638 simplified inequalities in 96 variables which needs to be analyzed. Although it is in principle feasible to plug these inequalities into linear programming software, neither Woonjung nor I have yet done so. The D_4 and F_4 cases produced systems which were small enough to analyze by hand; this system for H_4 is not.

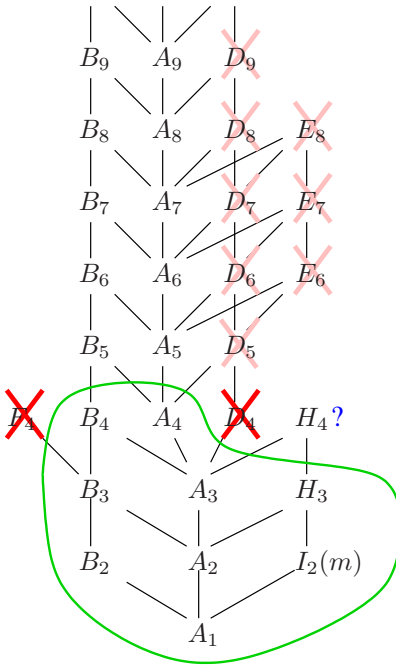


FIGURE 25. CAT(0) and Brady-Krammer complexes.

9. SMALL CANCELLATION GROUPS

The most successful reduction of curvature to combinatorial conditions is, of course, the case of complexes built entirely out of Euclidean cubes of various dimensions. Partly due to the ease with which curvature can be checked (Lemma 7.1), these *cube complexes* have become a favorite of geometric group theorists [9, 37, 72, 73, 77, 78]. In this section and the next I discuss how high dimensional cube complexes can be used to prove that various groups are CAT(0) groups. After a brief digression about cube complexes for Coxeter groups, the primary focus in this section is on cube complexes for small cancellation groups.

One source of additional interest in cube complexes is caused by the fact that groups acting on non-positively curved cube complexes have properties which do not hold for more general CAT(0) groups. For example, groups acting geometrically on non-positively curved cube complexes cannot have Kazhdan's property T (which is the property used by Niblo and Reeves [74] to establish Theorem 7.3). Thus, even for classes of groups such as Coxeter groups where they are already known to be CAT(0), it is still of interest to investigate which Coxeter groups are capable of acting geometrically on CAT(0) cube complexes. This question has been (essentially) completely answered by recent work of Graham Niblo and Lawrence Reeves [73] and by Ben Williams [82]. In particular, Niblo and Reeves use the Sageev cube construction ([77, 78]) to establish the following.

Theorem 9.1 (Coxeter groups and cube complexes). *If W is a finitely generated Coxeter group then W acts properly discontinuously by isometries on a locally finite, finite dimensional CAT(0) cube complex.*

As stated the action is not necessarily geometric since it need not be cocompact. On the hand, in his dissertation [82], Ben Williams was able to characterize exactly when this action is cocompact in terms of the subgroup structure.

Theorem 9.2 (Characterising cocompactness). *The action in the preceding theorem is cocompact if and only if for any triple p, q, r of positive integers, the Coxeter group W contains only finitely many conjugacy classes of subgroups isomorphic to the p, q, r triangle group $\langle a, b, c \mid a^2 = b^2 = c^2 = (ab)^p = (bc)^q = (ac)^r = 1 \rangle$.*

Despite this progress, it would still be of interest to determine more explicitly which Coxeter groups satisfy the conditions isolated by Williams.

The situation for small cancellation groups is much less satisfactory. Small cancellation groups are a relatively well-understood class of cohomological dimension 2 (mostly) word-hyperbolic groups whose definition and key properties have been known for nearly thirty years [63]. In fact, they formed one of the key models for the development of word-hyperbolic groups. Nonetheless, the relationship between small cancellation groups and the more local notions of curvature such as $\text{CAT}(0)$ has remained completely mysterious until recently. The basic definitions of small cancellation theory can be found in [63] or [67], but a rough description is included for the sake of completeness.

Definition 9.3 (Small cancellation groups). Let X be a finite combinatorial 2-complex. A *piece* in X is a path in the 1-skeleton which can be ϵ -pushed off the 1-skeleton in at least two distinct ways. In other words, this path must be liftable through the attaching maps of the 2-cells in more than one way. The complex X is called a $C(p)$ -complex if for each 2-cell in X its boundary cycle cannot be covered with fewer than p pieces. It is called a $T(q)$ -complex if there does not exist an immersed path in any vertex link with combinatorial length strictly between 2 and q . In other words, every 2-complex satisfies the condition $T(3)$ but the condition $T(4)$ represents an actual restriction.

There is also a metric version of the $C(p)$ condition known as $C'(\alpha)$. The complex is a $C'(\alpha)$ -complex if for each piece P in the boundary of a 2-cell R of X , the combinatorial length of P is less than α of the combinatorial length of the boundary of R . It is a basic result of small cancellation theory that $C'(1/4)$ - $T(4)$ -complexes and $C'(1/6)$ -complexes are word-hyperbolic. These two classes are called *metric small cancellation groups*. Sergei Ivanov and Paul Schupp have also characterized exactly when the non-metric small cancellation groups, that is the fundamental groups of $C(p)$ - $T(q)$ complexes where (p, q) is either $(3, 6)$, $(4, 4)$ or $(6, 3)$, are word-hyperbolic [55].

Noel Brady and I recently established a theorem concerning small cancellation groups and $\text{CAT}(0)$ structures [14], which has also been established by Dani Wise independently using different techniques [83]. The basic philosophy is described first since it is also used to prove the results in the next section.

Remark 9.4 (Philosophy). Let X be any finite combinatorial cell complex, let \mathcal{C} be the collection of maximal closed cells in its universal cover \tilde{X} , and let P be the partially ordered set of intersections of elements in \mathcal{C} . The poset P is called the *nerve* of \tilde{X} . See Figure 26 for a simple illustration.

The main idea is to replace each maximal cell in X with a high-dimensional cell so that they glue together nicely and the nerve of the resulting complex is identical.




FIGURE 26. A complex and its nerve.

In this new situation a *piece* is defined as a subcomplex of \tilde{X} which corresponds to an element of the nerve. In other words, a piece P is subcomplex which is an intersection of maximal cells in the universal cover. Notice that for small cancellation complexes this is a slightly modified notion of piece since arbitrary connected subcomplexes of pieces are no longer pieces. We need the following observation. If the pieces of X are contractible and the maximal cells embed in \tilde{X} , then X is homotopy equivalent to the geometric realization of the nerve quotiented by the group action. As a consequence, if we find two complexes with these properties and there is a correspondence between their fundamental groups and their nerves, then these complexes are homotopy equivalent to each other. Various small-cancellation-like conditions on X will guarantee both of these properties. For example, these properties will hold if the overlaps between maximal cells are “small” subcomplexes of its boundary (in some suitable sense) and the links are “large.”

In the case of small cancellation complexes, the idea is to replace each 2-cell and each piece with a high dimensional Euclidean cube, preserving the nerve and all of the other necessary properties described above. The following is a sample theorem along these lines. This result demonstrates how high-dimensional cube complexes really can be useful, even when studying objects whose cohomological dimension is quite small.

Theorem 9.5 (Cubes and $C'(1/4)$ - $T(4)$ groups). *Every $C'(1/4)$ - $T(4)$ group is the fundamental group of a compact high-dimensional non-positively curved cube complex.*

The key idea behind the proof is to subdivide the 1-skeleton of the standard 2-complex for G and then embed it into the 1-skeleton of a high-dimensional cube so that the pattern of intersections between the maximal cells in its universal cover, i.e. its nerve, remains the same. The proof proceeds roughly as follows. First subdivide every edge once so that every 2-cell has an even length boundary cycle. Next, identify each 2-cell R whose boundary cycle has length $2n$ with an n -dimensional cube. In addition, we select a path in the 1-skeleton of this n -cube which we think of as corresponding to the boundary of R . The choice is made so that any n consecutive edges of the path all travel in distinct directions and the antipodal edges are parallel. See Figure 27 for the path in a 3-cube which corresponds to an original hexagon.




FIGURE 27. The cube corresponding to a hexagon.

Finally, we glue these cubes together, exactly as the original 2-cells were glued, but with one crucial difference. If R_1 and R_2 shared a piece P of length m , then the cubes corresponding to R_1 and R_2 will be identified along the m -cubes spanned by P in R_1 and R_2 . For example, if P has length 2, then a square in one cube will be glued to a square in the other. It is now relatively easy to check that the result is a non-positively curved cube complex satisfying all the necessary conditions.

I should note that the hypothesis of the theorem can be weakened slightly without altering the proof: it is actually sufficient for the total length of any two consecutive pieces in R to be at most half the length of ∂R . Also, Dani Wise has extended this result to (a similar generalization of) $C'(1/6)$ groups. This case is more complicated and requires a more delicate construction. The next step would be, of course, to try and extend these results to all small cancellation groups, including the non-metric ones.

Conjecture 9.6 (Small cancellation groups and cubes). *All small cancellation groups are fundamental groups of compact high-dimensional non-positively curved cube complexes.*

The key difficulty here is to deal with the fact that 2-cells can now overlap a significant amount (even along more than half of their boundary cycles), so long as this large overlap is compensated by the existence of only very small overlaps along the remainder of the boundary path. A solution to this difficulty is likely to involve a more wholesale modification of the original 2-complex (perhaps along the lines described in Section 13 for one-relator groups) as a preparation for the transformation into a high-dimensional cube complex.

10. AMPLE TWISTED FACE-PAIRING GROUPS

The techniques and the philosophy described in the previous section are, of course, applicable more generally. As an illustration, Noel Brady and I have also shown that one of the ample twisted face-pairing 3-manifolds used as an illustration in [29] can be expanded in this way to 6-dimensional non-positively curved cube complex [14]. This brief tour of special cases concludes with a short account of this construction and its consequences.

Twisted face-pairing 3-manifolds were introduced by Jim Cannon, Bill Floyd and Walter Parry in [30], and their properties were further developed in [31] and [32]. Roughly speaking, one starts with a faceted 3-ball and a pairing of its faces.

Then, one subdivides the 1-skeleton of this 3-ball according to well-defined formulae and finally, one slightly alters the face-pairing maps by adding a small twist to each. Surprisingly, the result is always a 3-manifold. Moreover, in [29] they established that if the original faceted 3-ball satisfies certain elementary conditions that they call *ample*, then the 3-manifold constructed always has a word-hyperbolic fundamental group.

One of the primary motivations behind their construction was to construct explicit examples of word-hyperbolic 3-manifolds which were not already known to be hyperbolic (Kleinian) as a testing ground for various approaches to Thurston's geometrization conjecture. Having established that these 3-manifolds have word-hyperbolic fundamental groups, they have asked whether any of these constructions can be given a metric of non-positive curvature. The high-dimensional cube complex Noel Brady and I construct answers this question for at least one example. The natural conjecture to make is the following.

Conjecture 10.1 (Ample twisted 3-manifolds and $\text{CAT}(0)$). *If M is an ample twisted face-pairing 3-manifold, then $\pi_1 M$ is not only a word-hyperbolic group, but also the fundamental group of a compact non-positively curved cube complex of high dimension.*

Given the flexibility of our construction, it is already clear that Conjecture 10.1 does hold for an infinite number of ample twisted face-pairing 3-manifolds, and possibly for all of them. If one were even bolder, one could conjecture that for every compact, orientable hyperbolic 3-manifold M , there is a finite high-dimensional cube complex with fundamental group $\pi_1 M$. Rubinstein and Sageev's results on the k -plane property [77] are totally compatible with this conjecture, and, in fact, make it seem a bit more plausible. Moreover, this process of constructing high-dimensional non-positively curved cube complexes for word-hyperbolic 3-manifold groups might be interesting in another respect. Our cube complex construction forces the original 3-manifold to contain a π_1 -injective immersed surface. Thus, constructions of this type might be used as a way of approaching (at the very least) special cases of the virtual fibering conjecture.

Part 4. Combinatorial notions of curvature

In the last two parts I examined the possibility of restricting the complexes under consideration to low dimensions where good implementable algorithms exist, and of restricting the classes of groups under consideration so that special complexes such as cube complexes could be used. In this final part I consider a third possibility: weakening the notion of curvature itself. Rather than assigning piecewise Euclidean metrics to each of the polytopes in a complex, this more combinatorial version simply assigns values representing the "size" of each internal and external angle. Section 11 briefly reviews the various identities satisfied by these angles in a Euclidean polytope since this is the situation the more "conformal" version is modeled on. Next, Section 12 discusses the 2-dimensional combinatorial Gauss-Bonnet theorem and proposes a number of ways in which this theorem could be extended to higher dimensions. Finally, Section 13 discusses the notion of conformally $\text{CAT}(0)$ 2-complexes, sectional curvature, and the potential relationship between these concepts, special polyhedra and the coherence of one-relator groups.

11. ANGLES IN POLYTOPES

Every mathematician learns early on that the sum of the internal angles in a triangle equals π and that sum of the external angles of a convex polygon is 2π , but far fewer people learn the corresponding relations in higher dimensional Euclidean polytopes. For example, how many students learn that the sum of the dihedral angles in a tetrahedron can vary depending on its shape? This section is a short digression into the various identities among the internal and external angles of a Euclidean polytope particularly in high dimensions. All of this material is well-known and classical [49, 68, 80], but perhaps less well-known in our own subfield than it should be. The main goal of this section is to review the background material necessary to state and prove a fairly general high-dimensional combinatorial Gauss-Bonnet type theorem in the next section. The first concept to review is the precise definition of normalized internal and external angles in a Euclidean polytope.




FIGURE 28. Internal and external angles.

Definition 11.1 (Internal and external angles). Let P be a Euclidean polytope and let F be a face of P . The *internal angle*, $\alpha(F, P)$, is the proportion of unit vectors perpendicular to F which point into P (i.e. the measure of this set of vectors divided by the measure of the sphere of the appropriate dimension). The *external angle*, $\beta(F, P)$, is the proportion of unit vectors perpendicular to F so that there is a hyperplane with this unit normal which contains F and the remainder of P is on the other side. Also, by convention $\alpha(P, P) = \beta(P, P) = 1$. As an example, consider the vertex in the triangle shown in Figure 28. The internal angle (marked by α and the external angle (marked by β) are the radian measures of these angles divided by 2π (which is the normalization). The angles considered here are always normalized in this sense, since in higher dimensions it is the normalized versions which satisfy the most elegant relations.




FIGURE 29. A regular tetrahedron.

Example 11.2 (Angles in a regular tetrahedron). Consider the regular tetrahedron shown in Figure 29. Let t denote the tetrahedron itself containing a face f which contains an edge e which contains a vertex v . Both $\alpha(v, e) = \frac{1}{2}$ and $\beta(v, e) = \frac{1}{2}$ since each contains a single unit vector and the 0-sphere consists of two unit vectors. Similarly, $\alpha(e, f) = \beta(e, f) = \alpha(f, t) = \beta(f, t) = \frac{1}{2}$. The value of $\alpha(v, f) = \frac{1}{6}$ since $\frac{\pi}{3}$ is one-sixth of a circle. Also, since the dihedral angles along the edges are $\cos^{-1}(\frac{1}{3})$ and the area of the vertex link is $3 \cos^{-1}(\frac{1}{3}) - \pi$, the values of $\alpha(e, t)$ and $\alpha(v, t)$ are $\frac{1}{2\pi} \cos^{-1}(\frac{1}{3})$ and $\frac{3}{4\pi} \cos^{-1}(\frac{1}{3}) - \frac{1}{4}$ respectively. The normalized external angles are $\beta(e, t) = \frac{1}{2} - \frac{1}{2\pi} \cos^{-1}(\frac{1}{3})$, $\beta(v, f) = \frac{1}{3}$ and $\beta(v, t) = \frac{1}{4}$.

From the definition of external angles and the fact that a hyperplane approaching a polytope from a generic direction first encounters the polytope at a vertex, the following result is immediate.

Theorem 11.3 (Summing external vertex angles). *If P is a Euclidean convex polytope then the result of summing the normalized external angles $\beta(v, P)$ over all the vertices $v \in P$ is 1.*

This result is one explanation for the relative simplicity of the external angles of the regular tetrahedron as compared with its internal angles. In order to easily describe the other identities which hold between the internal and external angles in a Euclidean polytope a short digression into combinatorics is necessary.

Definition 11.4 (Face lattice of a polytope). Associated to every Euclidean polytope P is a partially ordered set called its face lattice. The *face lattice* of a polytope is its poset of faces under inclusion. Conventionally, the empty set and the whole polytope are included as faces. When these elements are included the poset is a lattice in the combinatorialist's sense (i.e. every pair of elements has a unique least upper bound and a unique greatest lower bound). Notice that the set of all normalized internal angles of a polytope can be viewed as a single real-valued function on those pairs of elements in its face lattice which are comparable in the poset ordering. I call this function α . The normalized external angles can be viewed as a similar function which I call β . There is a small issue about extending the values of α and β to pairs of faces where the smaller face is the empty set since these values were not defined above. The conventional assignments, derived from thinking of a Euclidean polytope as a rescaled version of a very small spherical polytope, are as follows.

$$\alpha(\emptyset, F) = \begin{cases} 1 & \text{if } F = \emptyset \\ \frac{1}{2} & \text{if } \dim F = 0 \\ 0 & \text{if } \dim F > 0 \end{cases} \quad \beta(\emptyset, F) = \begin{cases} 1 & \text{if } F = \emptyset \\ \frac{1}{2} & \text{if } \dim F \geq 0 \end{cases}$$

The justifications are that if F is a very small spherical polytope of dimension n , then as seen from the origin $\alpha(\emptyset, F)$ is one-half of a 0-sphere when F is a vertex and nearly zero when F is higher dimensional. Similarly, the external angle is one-half of a 0-sphere when F is a vertex and nearly one-half when F is higher dimensional.

The collection of all functions of this type is known as the incidence algebra of the poset and the algebraic structure of this object is quite useful for expressing identities involving the internal and external angle functions α and β .

Definition 11.5 (Incidence algebra of a poset). Let P be a finite partially ordered set whose elements are labeled by the set $\{1, 2, \dots, n\}$. The set of $n \times n$ matrices




FIGURE 30. A sample poset.

over the reals with $a_{ij} \neq 0$ only when $i \leq j$ in P is called the *incidence algebra of P* , or $I(P)$. Notice that for any finite poset P the ordering of the vertices can be chosen so it is consistent with its partial order. For example, the numbering of the vertices in Figure 30 is consistent. Using such an ordering, the incidence algebra is a set of upper triangular matrices. Alternatively, and more abstractly, the elements of $I(P)$ can also be thought of as functions $f : P \times P \rightarrow \mathbb{R}$ where $f(x, y) = 0$ whenever $x \not\leq y$ in P . Because of this possibility, elements of $I(P)$ are often referred to as functions rather than matrices. Finally, matrix multiplication defines a corresponding multiplication of functions, called their *convolution*.

Inside the incidence algebra, there are three elements that are particularly crucial: the delta function, the zeta function and the Möbius function.

$$\zeta = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \quad \mu = \begin{bmatrix} 1 & -1 & -1 & -1 & 2 \\ 0 & 1 & 0 & 0 & -1 \\ 0 & 0 & 1 & 0 & -1 \\ 0 & 0 & 0 & 1 & -1 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

FIGURE 31. The functions ζ and μ for the sample poset.

Definition 11.6 (δ , ζ , and μ). The identity matrix is the *delta function* where $\delta(x, y) = 1$ iff $x = y$. The *zeta function* is the function $\zeta(x, y) = 1$ if $x \leq y$ in P and 0 otherwise (i.e. 1's wherever possible). The *Möbius function* is the matrix inverse of ζ . Note that $\mu \cdot \zeta = \zeta \cdot \mu = \delta$. The matrices that represent ζ and μ for the sample poset are given in Figure 31.

Remark 11.7 (Euler characteristics). There is a close connection between the values of the Möbius function and various Euler characteristics. In particular, the value of $\mu(x, y)$ is the reduced Euler characteristic of the geometric realization of the portion of the poset that lies strictly between x and y . For example, the value of the Möbius function on the interval of the sample poset from vertex 1 to vertex 5 is 2 (i.e. the value in the upper left corner of the matrix in Figure 31). The portion of the poset between the first and fifth elements is the three incomparable elements labeled 2, 3, and 4, the geometric realization is three points with the discrete topology, and the reduced Euler characteristic of this topological space is also 2.

The connection between Möbius functions and Euler characteristics enables a quick calculation of the Möbius function for the face lattice of a polytope.

Lemma 11.8 (Möbius functions for polytopes). *If F and G are faces of a polytope P with $F \leq G$, then the value of the Möbius function of the face lattice evaluated on the interval (F, G) is $\mu(F, G) = (-1)^{\dim G - \dim F}$.*

The idea is that the geometric realization of the portion of the face lattice between F and G is topologically an n -sphere, where n is determined by the difference in dimension between F and G . Thus, $\mu(F, G)$ is the reduced Euler characteristic on this n -sphere. This connection between Möbius functions and Euler characteristics was also recently used by myself and John Meier to calculate the ℓ^2 Betti numbers of the pure symmetric automorphism groups [65].




FIGURE 32. The poset of hypertrees on a set of size 4.

Remark 11.9 (ℓ^2 Betti numbers). Elsewhere in these proceedings, Wolfgang Lück describes various theorems about ℓ^2 Betti numbers which make them easier to calculate [62]. For the groups that John Meier and I studied, we first used a spectral sequence argument to show that all of the ℓ^2 Betti numbers were trivial - except for the one in top dimension. As a consequence, this final Betti number must be the Euler characteristic of the group. Next, we show this value must equal the reduced Euler characteristic of the fundamental domain of the space which is in turn a geometric realization of the poset of hypertrees. The hypertree poset for $n = 4$ is shown in Figure 32. The size of the hypertree posets grows quite rapidly. In fact, the poset for $n = 5$ already has 311 elements. Using the connection between reduced Euler characteristics and Möbius functions, the final calculation could be completed using standard generating function techniques from enumerative combinatorics. The moral here is that combinatorial ideas such as Möbius functions of posets can be relevant, even to seemingly distant areas such as the calculation of ℓ^2 Betti numbers.

The final identities I wish to highlight (Sommerville's identities and McMullen's identity) involve the convolution product of elements already defined, although for Sommerville's identities, it is convenient to introduce a signed version of the internal and external angles. Let $\bar{\alpha}(F, G)$ denote a *signed normalized internal angle* where the sign is given by $\mu(F, G)$. In other words $\bar{\alpha}(F, G) = \mu(F, G)\alpha(F, G)$. There is also a *signed normalized external angle*, $\bar{\beta}$ defined as $\bar{\beta}(F, G) = \mu(F, G)\beta(F, G)$. Signed normalized external angles resurface in the discussion of vertex curvature in Section 12.

Theorem 11.10 (Sommerville's identity). *If P is a Euclidean polytope and α , $\bar{\alpha}$, β , $\bar{\beta}$ and μ are defined as above, then $\mu \cdot \alpha = \bar{\alpha}$ and $\beta \cdot \mu = \bar{\beta}$. In other words, for all faces $F \leq H$ in P , the sum of $\mu(F, G)\alpha(G, H)$ over all G with $F \leq G \leq H$ is equal to $\mu(F, H)\alpha(F, H)$. Similarly, for all faces $F \leq H$ in P , the sum of $\beta(F, G)\mu(G, H)$ over all G with $F \leq G \leq H$ is equal to $\beta(F, H)\mu(F, H)$.*

Since the values of μ are simply signs, Sommerville's identity asserts that summing up the internal angles $\alpha(G, H)$ over all G between F and H with the appropriate signs equals the value of $\alpha(F, H)$ with the appropriate sign. The proof of this result is an easy application of inclusion-exclusion. The second identity listed ($\beta \cdot \mu = \bar{\beta}$) is derived from the first by applying Sommerville's identity to the dual polytopal cone.

Example 11.11. To illustrate the first identity, consider a regular tetrahedron, let F be one of vertices and H be the tetrahedron itself. Using v , e , f and t to denote arbitrary vertices, edges, faces and the tetrahedron itself, the sum under consideration is

$$\mu(v, v) \cdot \alpha(v, t) + 3\mu(v, e) \cdot \alpha(e, t) + 3\mu(v, f) \cdot \alpha(f, t) + \mu(v, t) \cdot \alpha(t, t)$$

where the non-trivial coefficients account for the fact that there are three edges and three faces which contain the vertex and are contained in the tetrahedron. Numerically this works out to be

$$1 \cdot \left(\frac{3}{4\pi} \cos^{-1} \left(\frac{1}{3} \right) - \frac{1}{4} \right) - 3 \cdot \left(\frac{1}{2\pi} \right) + 3 \cdot \frac{1}{2} - 1 \cdot 1,$$

which simplifies to $-\frac{3}{4\pi} \cos^{-1} \left(\frac{1}{3} \right) + \frac{1}{4} = -\alpha(v, t)$ as claimed.

Finally, the most interesting angle identity is the nonlinear identity relating internal and external angles discovered by Peter McMullen [68].

Theorem 11.12 (McMullen's identity). *If P is a Euclidean polytope and α , β and ζ are defined as above, then $\alpha \cdot \beta = \zeta$. In other words, for all faces $F \leq H$ in P , the sum of $\alpha(F, G)\beta(G, H)$ over all G with $F \leq G \leq H$ is equal to 1.*

The proof of this remarkable identity is a surprisingly short half page argument. First reduce to the case where F is a vertex and view the neighborhood of F in H as a polytopal cone. The idea is to look at a product of this particular polytopal cone with its dual in \mathbb{R}^{2n} (where n is the dimension of H) and then to integrate the function $f(\vec{x}) = \exp(-\|\vec{x}\|^2)$ over this space in two different ways. As an immediate corollary the identities $\mu \cdot \alpha \cdot \beta = \bar{\alpha} \cdot \beta = \delta$ and $\alpha \cdot \beta \cdot \mu = \alpha \cdot \bar{\beta} = \delta$ must also hold. In other words $\bar{\alpha}$ and β are inverses of each other and α and $\bar{\beta}$ are inverses. Once again, the regular tetrahedron provides a good illustration.

Example 11.13. Let F be one of the vertices of a regular tetrahedron H . Using the notational abbreviations defined above, the sum under consideration is

$$\alpha(v, v) \cdot \beta(v, t) + 3\alpha(v, e) \cdot \beta(e, t) + 3\alpha(v, f) \cdot \beta(f, t) + \alpha(v, t) \cdot \beta(t, t)$$

where the coefficients account for the three edges and three faces between F and H as before. Numerically this works out to be

$$1 \cdot \frac{1}{4} + 3 \cdot \frac{1}{2} \cdot \left(\frac{1}{2} - \frac{1}{2\pi} \cos^{-1} \left(\frac{1}{3} \right) \right) + 3 \cdot \frac{1}{6} \cdot \frac{1}{2} + \left(\frac{3}{4\pi} \cos^{-1} \left(\frac{1}{3} \right) - \frac{1}{4} \right) \cdot 1$$

which simplifies to 1 as claimed.

12. COMBINATORIAL GAUSS-BONNET

In this section I discuss two seemingly different well-known combinatorial versions of the Gauss-Bonnet formula and a very flexible common generalization which is being introduced here for the first time. The first version is a formula due to Cheeger, Müller and Schrader which starts with a piecewise Euclidean complex and assigns a curvature to each vertex (based on certain external angles of the Euclidean polytopes involved) so that the sum of these vertex curvatures is the Euler characteristic of the complex. The second version (rediscovered by a number of different researchers over the years) starts with a polygonal 2-complex, randomly assigns real numbers as external “angles,” and then calculates vertex and face curvatures so that the sum of all the vertex and face curvatures is the Euler characteristic. Finally, using the identities from the previous section, I show how both versions are special cases of a very general procedure for assigning random “angles” to high-dimensional polyhedral cell complexes and then calculating combinatorial curvatures which add up to the Euler characteristic of the complex. The first step is to introduce the two versions which are already well-known.

The version of the combinatorial Gauss-Bonnet theorem due to Cheeger, Müller and Schrader is essentially a quick calculation [38]. See also the paper by Charney and Davis [36] where they review this calculation and then use the resulting formula to find a combinatorial analogue of the Hopf conjecture in the context of non-positively curved piecewise Euclidean n -manifolds. Let X be a piecewise Euclidean complex. In the following calculation P denotes an arbitrary (nonempty) cell of X and v an arbitrary vertex. By convention, when both variables are mentioned below a summation sign, it is the first listed which is varying over the course of the summation. Using Theorem 11.3 and an inversion of the order of summation, the Euler characteristic of X can be rewritten as follows.

$$\begin{aligned}
 \chi(X) &= \sum_P (-1)^{\dim P} \\
 &= \sum_P \sum_{v \in P} (-1)^{\dim P} \beta(v, P) \\
 &= \sum_v \sum_{P \ni v} (-1)^{\dim P} \beta(v, P) \\
 &= \sum_v \sum_{P \ni v} \bar{\beta}(v, P) = \sum_v \kappa(v)
 \end{aligned}$$

where the *vertex curvature* $\kappa(v)$, as might be guessed from the final equality, is defined by summing the signed normalized external angles $\bar{\beta}(v, P)$ over all polytopes P containing v . Notice that if the portion of X containing the vertex v has a unique maximum element P , then by Sommerville’s identity (Theorem 11.10) the absolute value of the vertex curvature $\kappa(v)$ will be the normalized external angle $\beta(v, P)$.

The second combinatorial version of the Gauss-Bonnet formula is specifically 2-dimensional, but where the angles are now randomly assigned real numbers rather than numbers derived from a piecewise Euclidean metric.

Definition 12.1 (Angled 2-complex). An *angled 2-complex* X is one where arbitrary real numbers are assigned as values for the normalized external angles $\beta(v, f)$ for each vertex v in a face f , but the smaller intervals retain their conventional values. Thus $\beta(v, v) = \beta(e, e) = \beta(f, f) = 1$ and $\beta(v, e) = \beta(e, f) = \frac{1}{2}$ for each

vertex v contained in an edge e contained in a face f . We can then define values for the internal angle function α based on the goal of preserving the relationship $\alpha \cdot \beta = \zeta$. This forces $\alpha(v, v) = \alpha(e, e) = \alpha(f, f) = 1$, $\alpha(v, e) = \alpha(e, f) = \frac{1}{2}$ and $\alpha(v, f) = \frac{1}{2} - \beta(v, f)$ as one would expect from the Euclidean situation. In any angled 2-complex, one can define a vertex curvature $\kappa(v)$ as Cheeger, Müller and Schrader did above. In addition define a *face curvature* $\kappa(f)$, designed as a correction term, which measures how far the sum of the external vertex angles is from 1. Concretely, this means that $\kappa(v)$ and $\kappa(f)$ are defined by the following formulas.

$$\begin{aligned}\kappa(v) &= \beta(v, v) - \sum_{e \ni v} \beta(v, e) + \sum_{f \ni v} \beta(v, f) = 1 - \frac{1}{2} \text{degree}(v) + \sum_{f \ni v} \beta(v, f) \\ \kappa(f) &= 1 - \sum_{v \in f} \beta(v, f)\end{aligned}$$

Summing $\kappa(v)$ over all vertices of X and considering the result term by term, it is easy to see that this sum counts the number of vertices minus the number of edges plus the sum of all the assigned external angles. Similarly, summing $\kappa(f)$ over all faces of X counts the number of faces minus the sum of all the assigned external angles. Combining these two counts yields the following theorem, versions of which have been proven by Lyndon [63], Gersten [47], Pride [75], Ballmann and Buyalo [3], and myself and Wise [67], among others. See Section 4 of [67] for a more detailed discussion of its history.

Theorem 12.2 (2-dimensional combinatorial Gauss-Bonnet theorem). *If X is an angled 2-complex and vertex and face curvatures are defined as above, then*

$$\sum_v \kappa(v) + \sum_f \kappa(f) = \chi(X)$$

Typically, the sum in these references is actually equal to $2\pi\chi(X)$ instead of $\chi(X)$ since in 2-dimensions angles are not normally normalized. As seen in the previous section normalization is crucial when extending these types of results to higher dimensions.

The final goal of this section is to place both of these results in a broader context. The main idea is the following. Let X be a finite polyhedral cell complex and consider the zeta function of its face lattice (including the empty simplex). If α and β are any two elements of its incidence algebra such that $\alpha \cdot \beta = \zeta$ then the reduced Euler characteristic of X can be rewritten as follows. The variables P and Q here represent (possibly empty) cells of X and since they are being identified with

particular elements of the face poset, inequalities are used to denote inclusions.

$$\begin{aligned}
\tilde{\chi}(X) &= \sum_{P \geq \emptyset} (-1)^{\dim P} = \sum_{P \geq \emptyset} (-1)^{\dim P} \zeta(\emptyset, P) \\
&= \sum_{P \geq \emptyset} (-1)^{\dim P} \left(\sum_{Q \in [\emptyset, P]} \alpha(\emptyset, Q) \beta(Q, P) \right) \\
&= \sum_{Q \geq \emptyset} (-1)^{\dim Q} \alpha(\emptyset, Q) \left(\sum_{P \geq Q} (-1)^{\dim P - \dim Q} \beta(Q, P) \right) \\
&= \sum_{Q \geq \emptyset} (-1)^{\dim Q} \alpha(\emptyset, Q) \kappa^\uparrow(Q)
\end{aligned}$$

where $\kappa^\uparrow(Q)$ is defined as the obvious signed sum implicit in the final equality. Call this reformulation the *general combinatorial Gauss-Bonnet formula*. For later reference, here is a formal statement of what this calculation shows.

Theorem 12.3 (General combinatorial Gauss-Bonnet). *If X is a finite polyhedral cell complex and α and β are any two elements of the incidence algebra of the face poset of X such that their product $\alpha \cdot \beta$ is the zeta function ζ of this poset, then*

$$\tilde{\chi}(X) = \sum_{P \geq \emptyset} (-1)^{\dim P} \alpha(\emptyset, P) \kappa^\uparrow(P)$$

where $\kappa^\uparrow(P)$ is found by summing $(-1)^{\dim Q - \dim P} \beta(P, Q)$ over all cells Q containing P .

Because of the similarities with angled 2-complexes, a polyhedral cell complex together with a particular factorization of its zeta function into $\alpha \cdot \beta$ is called an *angled n -complex*. The resemblance between $\alpha \cdot \beta = \zeta$ and McMullen's identity is, of course, not accidental. I claim that both of the results given above can be reinterpreted in this manner. For the Cheeger-Müller-Schrader result, the values of α and β are the standard ones, most of the values of $\alpha(\emptyset, F)$ are zero, and most of the values of $\beta(\emptyset, F)$ are $\frac{1}{2}$. In particular, using the values of $\alpha(\emptyset, F)$ and $\beta(\emptyset, F)$ established above, the general formula reduces to the following.

$$\begin{aligned}
\tilde{\chi}(X) = -1 + \chi(X) &= \sum_{Q \geq \emptyset} (-1)^{\dim Q} \alpha(\emptyset, Q) \kappa^\uparrow(Q) \\
&= (-1) \kappa^\uparrow(\emptyset) + \frac{1}{2} \sum_v \kappa^\uparrow(v) \\
&= \sum_{Q \geq \emptyset} (-1)^{\dim Q} \beta(\emptyset, Q) + \frac{1}{2} \sum_v \kappa^\uparrow(v) \\
&= -1 + \frac{1}{2} \sum_{Q > \emptyset} (-1)^{\dim Q} + \frac{1}{2} \sum_v \kappa^\uparrow(v) \\
&= -1 + \frac{1}{2} \chi(X) + \frac{1}{2} \sum_v \kappa^\uparrow(v)
\end{aligned}$$

Cancelling the negative ones, collecting both Euler characteristic terms on one side and factoring out the one-half yields the previous formula. Alternatively, a slight

change in the conventions for the values of $\alpha(\emptyset, F)$ and $\beta(\emptyset, F)$ makes the derivation even quicker and more immediate. Let $\alpha(P, Q)$ and $\beta(P, Q)$ denote the normalized internal and external angles as originally defined when P and Q are not empty, but when $P = \emptyset$ set

$$\alpha(\emptyset, F) = \begin{cases} 1 & \text{if } F = \emptyset \\ 1 & \text{if } \dim F = 0 \\ 0 & \text{if } \dim F > 0 \end{cases} \quad \beta(\emptyset, F) = \begin{cases} 1 & \text{if } F = \emptyset \\ 0 & \text{if } \dim F \geq 0 \end{cases}$$

It is an easy exercise to show that this new extension of α and β to the empty simplex preserves the identity $\alpha \cdot \beta = \zeta$. Thus, the general combinatorial Gauss-Bonnet formula is applicable and the result is the calculation of Cheeger, Müller and Schrader with no rearrangement necessary.

$$\begin{aligned} \tilde{\chi}(X) = -1 + \chi(X) &= \sum_{P \geq \emptyset} (-1)^{\dim P} \alpha(\emptyset, P) \kappa^\uparrow(P) \\ &= (-1)\kappa^\uparrow(\emptyset) + \sum_v \kappa^\uparrow(v) \\ &= \sum_{P \geq \emptyset} (-1)^{\dim P} \beta(\emptyset, P) + \sum_v \kappa^\uparrow(v) \\ &= -1 + \sum_v \kappa^\uparrow(v) \end{aligned}$$

The derivation of the 2-dimensional version is even easier. If X is an angled 2-complex, then the values of $\beta(v, f)$ have been assigned at random, and remaining values of $\beta(P, Q)$ and $\alpha(P, Q)$ with Q and P nonempty were either chosen to be the standard ones or derived from the conditions embedded in the equation $\alpha \cdot \beta = \zeta$. For the extension to the empty simplex there is again a degree of choice, but the extension which forces the more recently defined choices for $\alpha(\emptyset, P)$ (i.e. equal to 1 when P is the empty set or a vertex and 0 otherwise) is one feasible possibility. Once made, this choice forces the values of $\beta(\emptyset, P)$ to be 1 when P is empty, 0 when P is a vertex or an edge, and equal to 1 minus the sum of its external angles $\beta(v, f)$ when P is a face f . Notice in particular, that $\beta(\emptyset, f)$ equals the face curvature $\kappa(f)$ and $\kappa^\uparrow(v)$ equals the vertex curvature $\kappa(v)$ as previously defined. Applying Theorem 12.3 to this factorization of the zeta function yields the result recorded in Theorem 12.2:

$$\begin{aligned} \tilde{\chi}(X) = -1 + \chi(X) &= \sum_{P \geq \emptyset} (-1)^{\dim P} \alpha(\emptyset, P) \kappa^\uparrow(P) \\ &= (-1)\kappa^\uparrow(\emptyset) + \sum_v \kappa^\uparrow(v) \\ &= \sum_{P \geq \emptyset} (-1)^{\dim P} \beta(\emptyset, P) + \sum_v \kappa^\uparrow(v) \\ &= -1 + \sum_f \beta(\emptyset, f) + \sum_v \kappa^\uparrow(v) \\ &= -1 + \sum_f \kappa(f) + \sum_v \kappa(v) \end{aligned}$$

As these two examples illustrate, the technique of rewriting the reduced Euler characteristic described in Theorem 12.3 is extremely flexible. Various factorizations of the zeta function are likely to be useful, but particularly those which produce lots of zeros and for which many terms are known to have the same sign. Further explorations of the possibilities are certainly needed.

13. CONFORMAL CAT(0) AND SECTIONAL CURVATURE

In this final section I discuss two uses of the 2-dimensional version of the combinatorial Gauss-Bonnet theorem (Theorem 12.2). The first involves the notion of conformally CAT(0) 2-complexes as developed by Gersten [47] and investigated further by Pride [75], Corson [42], Huck and Rosebrock [53] and others. The second application uses a more recent notion due to Dani Wise [84] that he calls non-positive section curvature. After showing that these two notions are equivalent for the class of special polyhedra, the section concludes with some speculations about a possible relationship between these notions and a long-standing open question about one-relator groups.

Definition 13.1 (Conformally CAT(0) 2-complexes). Let X be an angled 2-complex as defined in the previous section and notice that the links of the vertices in X are metric graphs. If the face curvatures are non-positive and the vertex links are CAT(1), then X is called *conformally* CAT(0). Concretely, the second condition means that every simple closed curve in a vertex links should have internal angles which add up to at least 1 (or 2π in the un-normalized version). If the face curvatures are negative and the vertex links are CAT(1) then X is called *conformally* CAT(-1).

Conformally CAT(0) 2-complexes have a number of nice properties. Steve Gersten has shown, for example, that conformally CAT(0) complexes are aspherical [47] and Jon Corson has shown that conformally CAT(-1) complexes (with all internal angles positive) are word-hyperbolic [42]. It is an interesting question whether either of these results can be extended to higher dimensions.

Problem 13.2. *Is there a higher dimensional analogue of a conformally CAT(0) or CAT(-1) 2-complex? More concretely, are there any reasonably general conditions on the angles assigned in an angled n -complex X which would allow one to conclude that X is aspherical or that $\pi_1 X$ is word hyperbolic?*

The distinction between 2-dimensional CAT(0) spaces and 2-dimensional conformal CAT(0) spaces is witnessed by the standard 2-complex for the Baumslag-Solitar groups $BS(n, m) = \langle a, b | a^m b = b a^n \rangle$. Assigning internal and external angles so that the defining relator looks like a rectangle with vertical b 's and horizontal a 's easily shows that this complex is conformally CAT(0). See Figure 33 for an illustration. On the other hand, an easy argument using translation length shows that $BS(n, m)$ is a CAT(0) group if and only if $n = m$.

The second concept is of more recent vintage: non-positive sectional curvature.

Definition 13.3 (Non-positive sectional curvature). Let X be an angled 2-complex. A *section* of a vertex v in X is a based immersion $(S, s) \rightarrow (X, v)$ so that the image of the inclusion of the link of s into the link of v is a non-trivial finite spurless connected subgraph. Non-trivial means the link contains at least one edge and spurless means that the link contains no vertices of degree 1. Given a section




FIGURE 33. A portion of universal cover of the Baumslag-Solitar group $BS(1, 2)$.

$(S, s) \rightarrow (X, v)$ we can pull back the angle assignments and speak of the vertex curvature at s . If every face of X has non-positive face curvature and every section $(S, s) \rightarrow (X, v)$ has non-positive vertex curvature (calculated at s), then X has is said to have *non-positive sectional curvature*.

Notice that given a simple closed curve in the link of v , it is easy to construct a section with this loop as the image of its vertex link. Thus, it is immediate that non-positive sectional curvature implies conformally $CAT(0)$. It is also easy to construct examples which distinguish between the two since conformally $CAT(0)$ only checks simple closed loops and non-positive sectional curvature adds restrictions on many additional subgraphs of the vertex link. One situation where the distinction between the two concepts disappears is when X is a special polyhedron.

Definition 13.4 (Special polyhedra). Let X be a 2-complex in which the link of every point is either a circle, a theta graph, or the complete graph on 4 vertices. See Figure 34. The collection of points with each of these link types define the intrinsic 2-skeleton, 1-skeleton and 0-skeleton of X , respectively. Complexes with this property are called *closed fake surfaces*. If, in addition, the components of the intrinsic 2-skeleton are discs, then X is called a *special polyhedron*.




FIGURE 34. The three possible local structures in a special polyhedron.

Closed fake surfaces and special polyhedra are ubiquitous in at least one sense: every compact 2-complex is simple homotopy equivalent to a special polyhedron. See [50] for an in-depth discussion of these ideas. The simplicity of the vertex links in special polyhedra makes it easy to show the following.

Lemma 13.5. *If X is an angled 2-dimensional special polyhedron, then X is conformally $CAT(0)$ if and only if X has non-positive sectional curvature.*

One direction is immediate. In the other direction, note that the only section of a vertex v not tested by the conformal CAT(0) condition is the one derived from the entire link. Adding up the known inequalities for each of the simple closed curves in the link of length 3 one finds that twice the sum of all six angles at v is at least twice the value required to make the vertex curvature non-positive. This completes the proof.

Dani Wise has been able to show that non-positive sectional curvature has strong group-theoretic consequences [84]. Recall that a group is *coherent* if every finitely generated subgroup is also finitely presented. Examples of coherent groups include free groups, surface groups and fundamental groups of compact 3-manifolds. The standard example of a group which is not coherent is the direct product of two non-abelian free groups. Wise proves the following.

Theorem 13.6 (Non-positive sectional curvature and coherence). *If X is an angled 2-complex with non-positive sectional curvature, then $\pi_1 X$ is coherent.*

The proof uses Howie towers [51] and the 2-dimensional combinatorial Gauss-Bonnet theorem. The types of sublinks used to define sectional curvature are precisely those that need to be considered as one lifts various complexes and subcomplexes through the tower. Using the 2-dimensional combinatorial Gauss-Bonnet theorem, one can then show that certain intermediate complexes have non-positive Euler characteristics which, moreover, are strictly increasing as the process continues. It is precisely this tension between increasing Euler characteristics and non-positivity which forces the process to terminate at a complex which then witnesses the finite presentability of the subgroup. The general strategy is similar to the earlier work by myself and Wise on coherence using the idea of perimeter reduction [66], but the construction here are much more flexible in spirit. As an immediate corollary of Lemma 13.5 and Theorem 13.6 one has the following.

Corollary 13.7. *If X is a 2-dimensional special polyhedron with a conformal CAT(0) structure, then $\pi_1 X$ is coherent.*

Finally, this section concludes with some speculations about a possible relationship between the general results list above and the coherence of one-relator groups. Consider the following two conjectures.

Conjecture 13.8 (Coherent). *Every one-relator group is coherent.*

Conjecture 13.9 (Conformally CAT(0)). *Every one-relator group is either the fundamental group of a 2-dimensional special polyhedron with a conformal CAT(0) structure, or it can be built from such a one-relator group using very simple constructions.*

The first conjecture was asked as a question by Gilbert Baumslag as early as his 1974 survey of one-relator groups [6]. The second is being proposed here as a natural geometric conjecture. Although it is being stated in a slightly vague manner, the idea is that in light of Corollary 13.7, Conjecture 13.9 can be used as a possible line of attack on Conjecture 13.8. There is a much more detailed and precise version of Conjecture 13.9 which I have not stated formally since its development is not suitable for a survey article.

To give at least some idea of the general strategy, consider the following observations. First, it is well-known that every one-relator group is a subgroup of a

2-generator one-relator group so that in order to prove the coherence of all one-relator groups, it is sufficient to consider only those with 2-generators. Next, these groups are particularly nice from the perspective of conformally CAT(0) structures on special polyhedra, since their Euler characteristics are zero. As a result, the only feasible conformal CAT(0) structures are those where the face curvatures are 0 and the vertex curvatures are 0. In other words, all the inequalities become equalities. Another less obvious consequence of the tightness of the constraints is that the pairs of non-adjacent angles in the link of a vertex must have the exact same angles assigned. Thus, if we let n denote the number of vertices in the complex, there are exactly $3n$ variables which determine the angle values. There are n linear equations associated with the face curvatures and n linear equations associated with the curvature of the vertex links. From a linear algebra perspective, so long as the system of $2n$ linear equations in $3n$ variables is consistent, there should be at least an n -dimensional family of solutions. Finally, the process of converting the standard 2-complex for a one-relator group into a special polyhedron is quite flexible. It is possible to insist that the result has no polygons with one or two sides and no “untwisted” triangles. Even at this point there are a number of moves which can be done which only slightly change the complex while significantly changing the system of linear equations. The net result is that the n vertex equations are always consistent with each other and the n face equations are almost always consistent each other.

Concretely, Noel Brady and I have also tried our hand at finding conformal CAT(0) structures for a half-a-dozen examples of 2-generator one-relator groups chosen at random. Starting from a randomly chosen relator of length 12 or so, we haphazardly converted it to a special polyhedron using the pineapple and banana tricks among others [50]. Once we had a special polyhedron, we then worked to remove all monogons and bigons. At this point we stopped and found the linear system of equations which needed to be satisfied and we tested whether this system had a solution. In all cases, it did, and in all but one case, the resulting angle assignment could be extended to a piecewise Euclidean metric on the 2-complex. In other words, the results were typically CAT(0) in addition to being conformally CAT(0). The positive evidence is, admittedly, rather meager at the moment. A computer program to check all of the 2-generator one-relator groups out to a modest size would be more convincing and is currently under development. A general program for investigating conformal CAT(0) structures on special polyhedra for 2-generator one-relator groups is also being actively pursued. In addition to being a reasonable approach to the coherence question for one-relator groups, even partial progress on this geometric conjecture would certainly help to explain why one-relator groups have such a tendency to act like non-positively curved groups, even though it is well-known that they are not non-positively curved in general.

REFERENCES

- [1] J. M. Alonso and et al. Notes on word hyperbolic groups. In *Group theory from a geometrical viewpoint (Trieste, 1990)*, pages 3–63. World Sci. Publishing, River Edge, NJ, 1991. Edited by H. Short.
- [2] John C. Baez. The octonions. *Bull. Amer. Math. Soc. (N.S.)*, 39(2):145–205 (electronic), 2002.
- [3] W. Ballmann and S. Buyalo. Nonpositively curved metrics on 2-polyhedra. *Math. Z.*, 222(1):97–134, 1996.

- [4] Werner Ballmann. *Lectures on spaces of nonpositive curvature*, volume 25 of *DMV Seminar*. Birkhäuser Verlag, Basel, 1995. With an appendix by Misha Brin.
- [5] Werner Ballmann, Mikhael Gromov, and Viktor Schroeder. *Manifolds of nonpositive curvature*, volume 61 of *Progress in Mathematics*. Birkhäuser Boston Inc., Boston, MA, 1985.
- [6] Gilbert Baumslag. Some problems on one-relator groups. In *Proceedings of the Second International Conference on the Theory of Groups (Australian Nat. Univ., Canberra, 1973)*, pages 75–81. Lecture Notes in Math., Vol. 372. Springer, Berlin, 1974.
- [7] David Bessis. The dual braid monoid. Preprint available at arXiv:math.GR/0101158.
- [8] Mladen Bestvina. Non-positively curved aspects of Artin groups of finite type. *Geom. Topol.*, 3:269–302 (electronic), 1999.
- [9] Louis J. Billera, Susan P. Holmes, and Karen Vogtmann. Geometry of the space of phylogenetic trees. *Adv. in Appl. Math.*, 27(4):733–767, 2001.
- [10] Nicolas Bourbaki. *Éléments de mathématique*. Masson, Paris, 1981. Groupes et algèbres de Lie. Chapitres 4, 5 et 6. [Lie groups and Lie algebras. Chapters 4, 5 and 6].
- [11] B. H. Bowditch. Notes on locally $\text{cat}(1)$ spaces. In *Geometric group theory (Columbus, OH, 1992)*, pages 1–48. de Gruyter, Berlin, 1995.
- [12] Noel Brady and John Crisp. On the dimensions of $\text{CAT}(0)$ and $\text{CAT}(-1)$ complexes with the same hyperbolic group action. Preprint available at <http://aftermath.math.ou.edu/~nbrady/papers/>.
- [13] Noel Brady and Jon McCammond. The length spectrum of compact, constant curvature complex is discrete. Preprint 2004. Available at <http://www.math.ucsb.edu/~mccammon/papers/>.
- [14] Noel Brady and Jon McCammond. Nonpositive curvature and small cancellation groups. In preparation.
- [15] Noel Brady, Jon McCammond, and John Meier. Bounding edge degrees in triangulated 3-manifolds. Available at <http://www.math.ucsb.edu/~mccammon/papers/>. To appear in *Proceedings of the American Mathematics Society*.
- [16] Thomas Brady. Artin groups of finite type with three generators. *Michigan Math. J.*, 47(2):313–324, 2000.
- [17] Thomas Brady. A partial order on the symmetric group and new $K(\pi, 1)$'s for the braid groups. *Adv. Math.*, 161(1):20–40, 2001.
- [18] Thomas Brady and Jonathan P. McCammond. Four-generator artin groups of finite-type. In preparation.
- [19] Thomas Brady and Jonathan P. McCammond. Three-generator Artin groups of large type are biautomatic. *J. Pure Appl. Algebra*, 151(1):1–9, 2000.
- [20] Thomas Brady and Colum Watt. $K(\pi, 1)$'s for Artin groups of finite type. Preprint available at arXiv:math.GR/0102078.
- [21] Martin R. Bridson. Geodesics and curvature in metric simplicial complexes. In *Group theory from a geometrical viewpoint (Trieste, 1990)*, pages 373–463. World Sci. Publishing, River Edge, NJ, 1991.
- [22] Martin R. Bridson. On the existence of flat planes in spaces of nonpositive curvature. *Proc. Amer. Math. Soc.*, 123(1):223–235, 1995.
- [23] Martin R. Bridson. Non-positive curvature in group theory. In *Groups St. Andrews 1997 in Bath, I*, volume 260 of *London Math. Soc. Lecture Note Ser.*, pages 124–175. Cambridge Univ. Press, Cambridge, 1999.
- [24] Martin R. Bridson. On the semisimplicity of polyhedral isometries. *Proc. Amer. Math. Soc.*, 127(7):2143–2146, 1999.
- [25] Martin R. Bridson. Length functions, curvature and the dimension of discrete groups. *Math. Res. Lett.*, 8(4):557–567, 2001.
- [26] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.
- [27] M. Burger and A. Iozzi. Boundary maps in bounded cohomology. Appendix to: “Continuous bounded cohomology and applications to rigidity theory” [Geom. Funct. Anal. **12** (2002), no. 2, 219–280; MR 2003d:53065a] by Burger and N. Monod. *Geom. Funct. Anal.*, 12(2):281–292, 2002.
- [28] M. Burger and N. Monod. Continuous bounded cohomology and applications to rigidity theory. *Geom. Funct. Anal.*, 12(2):219–280, 2002.

- [29] J. W. Cannon, W. J. Floyd, and W. R. Parry. Ample twisted face-pairing 3-manifolds. Available at <http://www.math.vt.edu/people/floyd/research/>.
- [30] J. W. Cannon, W. J. Floyd, and W. R. Parry. Introduction to twisted face-pairings. *Math. Res. Lett.*, 7(4):477–491, 2000.
- [31] J. W. Cannon, W. J. Floyd, and W. R. Parry. Twisted face-pairing 3-manifolds. *Trans. Amer. Math. Soc.*, 354(6):2369–2397 (electronic), 2002.
- [32] J. W. Cannon, W. J. Floyd, and W. R. Parry. Heegaard diagrams and surgery descriptions for twisted face-pairing 3-manifolds. *Algebr. Geom. Topol.*, 3:235–285 (electronic), 2003.
- [33] James W. Cannon. The theory of negatively curved spaces and groups. In Tim Bedford, Michael Keane, and Caroline Series, editors, *Ergodic theory, symbolic dynamics, and hyperbolic spaces*, Oxford Science Publications, pages 315–369, New York, 1991. Oxford University Press.
- [34] R. Charney, M. Davis, and G. Moussong. Nonpositively curved, piecewise Euclidean structures on hyperbolic manifolds. *Michigan Math. J.*, 44(1):201–208, 1997.
- [35] Ruth Charney. Metric geometry: connections with combinatorics. In *Formal power series and algebraic combinatorics (New Brunswick, NJ, 1994)*, volume 24 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 55–69. Amer. Math. Soc., Providence, RI, 1996.
- [36] Ruth Charney and Michael Davis. The Euler characteristic of a nonpositively curved, piecewise Euclidean manifold. *Pacific J. Math.*, 171(1):117–137, 1995.
- [37] Indira Chatterji and Graham Niblo. From wall spaces to CAT(0) cube complexes. Available at <http://arxiv.org/abs/math.GT/0309036>.
- [38] Jeff Cheeger, Werner Müller, and Robert Schrader. On the curvature of piecewise flat spaces. *Comm. Math. Phys.*, 92(3):405–454, 1984.
- [39] Woonjung Choi. Software package `coxeter.g`. Available to download from <http://www.math.tamu.edu/~wjchoi/>.
- [40] Woonjung Choi. *The existence of metrics of nonpositive curvature on the Brady-Krammer complexes for finite-type Artin groups*. PhD thesis, Texas A&M University, in preparation. Expected date of graduation May 2004.
- [41] John H. Conway and Derek A. Smith. *On quaternions and octonions: their geometry, arithmetic, and symmetry*. A K Peters Ltd., Natick, MA, 2003.
- [42] Jon Michael Corson. Conformally nonspherical 2-complexes. *Math. Z.*, 214(3):511–519, 1993.
- [43] Murray Elder and Jon McCammond. CAT(0) is an algorithmic property. Available at <http://www.math.ucsb.edu/~mccammon/papers/>. To appear in *Geometriae Dedicata*.
- [44] Murray Elder and Jon McCammond. Software packages `cat.g` and `cat.gp`. Available at <http://www.math.ucsb.edu/~mccammon/software/>.
- [45] Murray Elder and Jon McCammond. Curvature testing in 3-dimensional metric polyhedral complexes. *Experiment. Math.*, 11(1):143–158, 2002.
- [46] Murray Elder, Jon McCammond, and John Meier. Combinatorial conditions that imply word-hyperbolicity for 3-manifolds. *Topology*, 42(6):1241–1259, 2003.
- [47] S. M. Gersten. Reducible diagrams and equations over groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 15–73. Springer, New York, 1987.
- [48] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.
- [49] Branko Grünbaum. *Convex polytopes*, volume 221 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 2003. Prepared and with a preface by Volker Kaibel, Victor Klee and Günter M. Ziegler.
- [50] Cynthia Hog-Angeloni and Wolfgang Metzler, editors. *Two-dimensional homotopy and combinatorial group theory*, volume 197 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 1993.
- [51] James Howie. On pairs of 2-complexes and systems of equations over groups. *J. Reine Angew. Math.*, 324:165–174, 1981.
- [52] G. Christopher Hruska. *Nonpositively curved spaces with isolated flats*. PhD thesis, Cornell University, 2002.
- [53] Günther Huck and Stephan Rosebrock. Weight tests and hyperbolic groups. In *Combinatorial and geometric group theory (Edinburgh, 1993)*, volume 204 of *London Math. Soc. Lecture Note Ser.*, pages 174–183. Cambridge Univ. Press, Cambridge, 1995.
- [54] James E. Humphreys. *Reflection groups and Coxeter groups*. Cambridge University Press, Cambridge, 1990.

- [55] S. V. Ivanov and P. E. Schupp. On the hyperbolicity of small cancellation groups and one-relator groups. *Trans. Amer. Math. Soc.*, 350(5):1851–1894, 1998.
- [56] Richard Kane. *Reflection groups and invariant theory*. CMS Books in Mathematics/Ouvrages de Mathématiques de la SMC, 5. Springer-Verlag, New York, 2001.
- [57] M. Kapovich and B. Leeb. On asymptotic cones and quasi-isometry classes of fundamental groups of 3-manifolds. *Geom. Funct. Anal.*, 5(3):582–603, 1995.
- [58] Misha Kapovich. An example of a 2-dimensional hyperbolic group which can't act on 2-dimensional negatively curved complexes. Preprint, Utah, 1994.
- [59] Bruce Kleiner. Review of ballmann, bridson-haefliger and eberlein. *Bulletin of the American Mathematical Society*, 39(2):273–279, 2002.
- [60] Daan Krammer. The braid group B_4 is linear. *Invent. Math.*, 142(3):451–486, 2000.
- [61] Daan Krammer. Braid groups are linear. *Ann. of Math. (2)*, 155(1):131–156, 2002.
- [62] Wolfgang Lueck. L^2 -Invariants from the Algebraic Point of View.
- [63] Roger C. Lyndon and Paul E. Schupp. *Combinatorial group theory*. Springer-Verlag, Berlin, 1977. Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 89.
- [64] Jon McCammond. Noncrossing partitions in surprising locations. Preprint 2003. Available at <http://www.math.ucsb.edu/~mccammon/papers/>.
- [65] Jon McCammond and John Meier. The hypertree poset and the ℓ^2 -betti numbers of the motion group of the trivial link. Available at <http://www.math.ucsb.edu/~mccammon/papers/>. To appear in *Mathematisches Annalen*.
- [66] Jonathan P. McCammond and Daniel T. Wise. Coherence, local quasiconvexity, and the perimeter of 2-complexes. Preprint 2003.
- [67] Jonathan P. McCammond and Daniel T. Wise. Fans and ladders in small cancellation theory. *Proc. London Math. Soc. (3)*, 84(3):599–644, 2002.
- [68] P. McMullen. Non-linear angle-sum relations for polyhedral cones and polytopes. *Math. Proc. Cambridge Philos. Soc.*, 78(2):247–261, 1975.
- [69] Karl Menger. New foundations of euclidean geometry. *American J. of Math.*, 53:721–745, 1931.
- [70] Igor Mineyev and Guoliang Yu. The Baum-Connes conjecture for hyperbolic groups. *Invent. Math.*, 149(1):97–122, 2002.
- [71] G. Moussong. *Hyperbolic Coxeter Groups*. PhD thesis, The Ohio State University, 1988.
- [72] G. A. Niblo and L. D. Reeves. The geometry of cube complexes and the complexity of their fundamental groups. *Topology*, 37(3):621–633, 1998.
- [73] G. A. Niblo and L. D. Reeves. Coxeter groups act on CAT(0) cube complexes. *J. Group Theory*, 6(3):399–413, 2003.
- [74] Graham Niblo and Lawrence Reeves. Groups acting on CAT(0) cube complexes. *Geom. Topol.*, 1:approx. 7 pp. (electronic), 1997.
- [75] S. J. Pride. Star-complexes, and the dependence problems for hyperbolic complexes. *Glasgow Math. J.*, 30(2):155–170, 1988.
- [76] Simon Rentzmann. *The Melzak problem for triangulated convex polytopes*. PhD thesis, Texas A&M University, 2000.
- [77] Hyam Rubinstein and Michah Sageev. Intersection patterns of essential surfaces in 3-manifolds. *Topology*, 38(6):1281–1291, 1999.
- [78] Michah Sageev. Codimension-1 subgroups and splittings of groups. *J. Algebra*, 189(2):377–389, 1997.
- [79] I. J. Schoenberg. Remarks to Maurice Fréchet's article "Sur la définition axiomatique d'une classe d'espace distanciés vectoriellement applicable sur l'espace de Hilbert" [MR1503246]. *Ann. of Math. (2)*, 36(3):724–732, 1935.
- [80] D. M. Y. Sommerville. *An introduction to the geometry of n dimensions*. Dover Publications Inc., New York, 1958.
- [81] John Sullivan. New tetrahedrally close-packed structures. Preprint 2000.
- [82] B. T. Williams. *Two topics in geometric group theory*. PhD thesis, University of Southampton, 1999.
- [83] Daniel T. Wise. Cubulating small cancellation groups. Preprint 2003.
- [84] Daniel T. Wise. Sectional curvature and local quasiconvexity. In preparation.
- [85] Daniel T. Wise. *Non-positively curved squared complexes, aperiodic tilings, and non-residually finite groups*. PhD thesis, Princeton University, 1996.

- [86] Daniel T. Wise. A flat plane that is not the limit of periodic flat planes. *Algebr. Geom. Topol.*, 3:147–154 (electronic), 2003.

DEPT. OF MATH., U. C. SANTA BARBARA, SANTA BARBARA, CA 93106
E-mail address: `jon.mccammond@math.ucsb.edu`