

CAT(0) IS AN ALGORITHMIC PROPERTY

MURRAY ELDER AND JON MCCAMMOND¹

ABSTRACT. In this article we give an explicit algorithm which will determine, in a discrete and computable way, whether a finite piecewise Euclidean complex is non-positively curved. In particular, given such a complex we show how to define a boolean combination of polynomial equations and inequalities in real variables, i.e. a real semi-algebraic set, which is empty if and only if the complex is non-positively curved. Once this equivalence has been shown, the main result follows from a standard theorem in real algebraic geometry.

CONTENTS

1. Polyhedral geometry	2
2. Comparison geometry	5
3. Galleries	7
4. Detecting piecewise geodesics	13
5. Detecting local geodesics	17
6. Coefficients	20
7. The Main Theorem	21
References	23

In this article we prove that for every finite collection of Euclidean polytopes, there exists a finite list of forbidden configurations which determine the non-positively curved complexes that can be built out of these shapes. More specifically, a complex built out of these shapes will be non-positively curved if and only if it avoids these forbidden configurations. The general result for arbitrary curvatures is as follows:

Theorem A. *If \mathcal{S} is a finite collection of shapes with curvature κ , then there exists a finite list of configurations \mathcal{C} such that an M_κ -complex K with $\text{SHAPES}(K) \subset \mathcal{S}$ is locally CAT(κ) if and only if K avoids all of the configurations in \mathcal{C} .*

¹Supported under NSF grant no. DMS-99781628

Date: February 8, 2008.

1991 *Mathematics Subject Classification.* 20F65.

Key words and phrases. non-positive curvature, CAT(0), decidability.

Notice, however, that Theorem A merely asserts the existence of a finite list of forbidden configurations that need to be avoided. It does not, in and of itself, show that there is a procedure for deciding whether a particular finite complex is locally $\text{CAT}(\kappa)$. The second half of the article is devoted to proving this stronger assertion. In particular we prove the following:

Theorem B. *There exists an algorithm which determines whether or not a finite M_κ -complex is locally $\text{CAT}(\kappa)$.*

Two special cases are of particular note. Theorem B shows that it is possible to test whether a finite, piecewise Euclidean complex is non-positively curved and to test whether a finite, piecewise hyperbolic complex is negatively curved.

The proof of Theorem B proceeds by converting curvature considerations into a real semi-algebraic set which is empty if and only if the complex is $\text{CAT}(\kappa)$. A *real semi-algebraic subset of \mathbb{R}^n* is a subset which can be described by a finite boolean combination of polynomial equations and inequalities in n real variables. The main property of real semi-algebraic sets we will need is that given a boolean combination of polynomial equations and inequalities in n variables there is an algorithm which decides whether this system has a solution. This is a special case of Tarski's theorem that the elementary theory of the reals is decidable. See [2] or [5] for details. Thus, once the conversion to real algebraic geometry has been completed, the existence of an algorithm will be immediate.

Finally, we should note that in dimension 3 there exists a particularly elementary procedure which avoids real algebraic geometry altogether. Since this algorithm differs substantially in flavor from the one given below, it will be described elsewhere. See [7] for details.

Overview of the Sections

In the first two sections we briefly review some basic definitions and results from polyhedral geometry, and comparison geometry. They are included for completeness and can readily be skipped by the reader familiar with these areas.

In Section 3 we introduce our main technical tool, that of a gallery, and establish its main properties. In Section 4 we analyze how to describe piecewise geodesics in galleries using polynomial equations and inequalities and in Section 5 we perform a similar analysis for local geodesics. These two sections represent the heart of the proof.

In Section 6 we show that the coefficients that arise in the encoding process are no worse than the original lengths in the original complex, and finally, in Section 7 we present the proof of the main theorem, Theorem B.

1. POLYHEDRAL GEOMETRY

We begin by reviewing the necessary background regarding polyhedral geometry. A metric space (K, d) is called a *geodesic metric space* if every

pair of points in K can be connected by a geodesic. The geodesic metric spaces we will be primarily interested in will be cell complexes constructed out of convex polyhedral cells in \mathbb{H}^n , \mathbb{E}^n or \mathbb{S}^n . After describing these spaces, we will quickly review their relationship with comparison geometry. Details can be found in [1], [3], [4], and [9].

Definition 1.1 (Polyhedral cells). A *convex polyhedral cell* in \mathbb{H}^n or \mathbb{E}^n is the convex hull of a finite set of points. The convex hull of $n + 1$ points in general position is an *n-simplex*. A *polyhedral cone* in \mathbb{E}^n is the positive cone spanned by a finite set of vectors. If the original vectors are linearly independent, it is a *simplicial cone*. A *cell (simplex) in \mathbb{S}^n* is the intersection of a polyhedral cone (simplicial cone) in \mathbb{E}^n with \mathbb{S}^n .

A spherical cell which does not contain a pair of antipodal points is *proper*. All Euclidean and hyperbolic cells are considered proper. Notice that every spherical cell can be subdivided into proper spherical cells by it cutting along the coordinate axes. If σ denotes a proper convex polyhedral cell, then σ° will denote its interior and $\partial\sigma$ will denote its boundary.

Definition 1.2 (M_κ -complexes). An \mathbb{H} -*complex* [\mathbb{E} -complex, \mathbb{S} -complex] is a connected cell complex K made up of proper hyperbolic [Euclidean, spherical] cells glued together by isometries along faces. A cell complex which has an \mathbb{H} -complex, an \mathbb{E} -complex, or an \mathbb{S} -complex structure, will be called a *metric polyhedral complex*, or M_κ -*complex* for short where κ denotes the curvature constant common to all of its cells. More generally, an M_κ -complex is formed from polyhedral cells with constant curvature κ .

Convention 1.3 (Subdivisions). As noted above, every spherical cell can be subdivided into proper spherical cells. Since proper spherical cells are required in an \mathbb{S} -complex, a subdivision will occasionally be necessary in order to convert a complex built out of pieces of spheres to be considered an \mathbb{S} -complex.

Definition 1.4 (Shapes). If K is an M_κ -complex, then the isometry types of the cells of K will be called the *shapes of K* and the collection of these isometry types will be denoted $\text{SHAPES}(K)$. When $\text{SHAPES}(K)$ is finite, K is said to have only *finitely many shapes*. Notice that since cells of different dimensions necessarily have different isometry types, finitely many shapes implies that K is finite dimensional. It does not, however, imply that K is locally finite. Notice also that if K is an M_κ -complex with only finitely many shapes, then there is a subdivision K' of K where the cells of K' are proper and simplicial, and $\text{SHAPES}(K')$ remains finite.

Definition 1.5 (Paths and Loops). A *path* γ in a metric space K is a continuous map $\gamma : [0, \ell] \rightarrow K$. A path is *closed* if $\gamma(0) = \gamma(\ell)$. A *loop* is a closed path where the basepoint $\gamma(0)$ has been forgotten. Technically, a loop is viewed as a continuous map from a circle to K . The circle of perimeter ℓ will be denoted C_ℓ .

Definition 1.6 (Piecewise geodesics). Let K be an M_κ -complex. A *piecewise geodesic* γ in K is a path $\gamma : [a, b] \rightarrow K$ 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 K and that this path is the unique geodesic connecting its endpoints in the metric of σ . The *length* of γ , denoted $\text{length}(\gamma)$, is the sum of the lengths of the geodesics into which it can be partitioned. A *closed piecewise geodesic* and a *piecewise geodesic loop* are defined similarly. The *intrinsic metric* on K is defined as follows:

$$d(x, y) = \inf\{\text{length}(\gamma) \mid \gamma \text{ is a piecewise geodesic from } x \text{ to } y\}$$

In general d is only a pseudometric, but when K has only finitely many shapes, d is a well-defined metric and (K, d) is a geodesic metric space [4, Theorem 7.19].

Convention 1.7 (Parameterizations). If $\gamma : [0, \ell]$ is a piecewise geodesic, we will always assume that the map γ has been reparameterized by arc-length. In particular, ℓ should be the length of γ and for all subintervals $[a, b]$ in $[0, \ell]$, the length of $\gamma([a, b])$ should be exactly $b - a$. Following the same convention, if a loop is a closed piecewise geodesic of length ℓ , then we will assume that its domain is a circle whose circumference is ℓ . We will use C_ℓ to denote such a circle and we will identify the points on C_ℓ with reals mod ℓ and in particular with the points $[0, \ell)$.

Definition 1.8 (Size). The *size* of a piecewise geodesic $\gamma : [0, \ell] \rightarrow K$ is the number of open cells of K through which $\gamma([0, \ell])$ passes, with multiplicities. Technically, the size of γ is the minimal number of subintervals (open, half-open, or closed) into which $[0, \ell]$ must be partitioned so that the image of each subinterval lies in a single open cell of K . Note that some of these subintervals may be single points. The fact that γ is a piecewise geodesic ensures that the size of γ is finite.

Definition 1.9 (Links). Let K be an M_κ -complex with only finitely many shapes and let x be a point in K . The set of unit tangent vectors to K at x is naturally an \mathbb{S} -complex called the *link of x in K* , or $\text{LINK}(x, K)$. If K has only finitely many shapes, then $\text{LINK}(x, K)$ has only finitely many shapes.

When x lies in the interior of a cell B of K , $\text{LINK}(x, B)$ is a sphere of dimension $k = \dim B - 1$ sitting inside $\text{LINK}(x, K)$. Moreover, the complex $\text{LINK}(x, K)$ can be viewed as a spherical join of \mathbb{S}^k and another \mathbb{S} -complex, denoted $\text{LINK}(B, K)$, which can be thought of as the unit tangent vectors to x in K which are orthogonal to B . The complex $\text{LINK}(B, K)$ is called the *link of the cell B in K* . Once again, if K has only finitely many shapes, then $\text{LINK}(B, K)$ has only finitely many shapes as well.

Definition 1.10 (Local geodesics). Let K be an M_κ -complex. A piecewise geodesic γ in K is called a *local geodesic* if for each point x on γ , the incoming and outgoing unit tangent vectors to γ at x are at a distance of at least π from each other in $\text{LINK}(x, K)$.

The size and length of local geodesics are closely related.

Theorem 1.11 (Bridson). *If K is an M_κ -complex with only finitely many shapes, then for every $\ell > 0$ there exist an integer $N > 0$, depending only on $\text{SHAPES}(K)$, such that every local geodesic of size at least N has length at least ℓ .*

For a proof of this result see [3, Theorem 1.11] or [4, Theorem I.7.28]. Both references state this theorem in a more general form using the concept of a taut- m -string. We merely need to note that local geodesics are taut- m -strings and the size of a local geodesic γ , as defined above, is within a constant factor of the integer m when γ is viewed as a taut- m -string.

Remark 1.12 (Constructing N). The *existence* of a constant such as N is important, but for our purposes, we need to know more. There needs to be an algorithm which constructs an integer N which will work solely from $\text{SHAPES}(K)$ and the real number ℓ . There is indeed such an algorithm, and we will return to this issue later in the article. See Lemma 5.3.

Notice that it is in fact sufficient to construct an N that will work for one particular value of ℓ . That is, if every local geodesic of size N has length at least ℓ , then every local geodesic of size $k \cdot N$ will have length at least $k \cdot \ell$. Thus for an arbitrary β , the least integer greater than $\frac{\beta}{\ell} \cdot N$ will work for β .

2. COMPARISON GEOMETRY

Metric polyhedral complexes are particularly useful in the creation of metric spaces of non-positive curvature. As in the previous section, details can be found in [1], [3], [4], and [9].

Definition 2.1 (Globally $\text{CAT}(\kappa)$). Let K be a geodesic metric space, let T be a geodesic triangle in K , and let $\kappa = -1$, [or 0, or 1]. A *comparison triangle* for T is a triangle T' in \mathbb{H}^2 , [or \mathbb{E}^2 or \mathbb{S}^2] with the same side lengths as T . Notice that for every point x on T , there is a corresponding point x' on T' . The space K is called *globally* $\text{CAT}(\kappa)$, if for any geodesic triangle T in K [of perimeter less than 2π when $\kappa = 1$] and for any points x and y on T , the distance from x to y in K is less than or equal to the distance from x' to y' in \mathbb{H}^2 [or \mathbb{E}^2 or \mathbb{S}^2]. Finally, a space K is called *locally* $\text{CAT}(\kappa)$ if every point in K has a neighborhood which is globally $\text{CAT}(\kappa)$. Locally $\text{CAT}(0)$ spaces are often referred to as *non-positively curved* and locally $\text{CAT}(-1)$ spaces are called *negatively curved*.

Theorem 2.2. *In a globally $\text{CAT}(0)$ space, every pair of points is connected by a unique geodesic, and a path is a geodesic if and only if it is a local geodesic.*

The next two results about M_κ -complexes show how global properties such as $\text{CAT}(\kappa)$ can be reduced to local properties and how local properties can be reduced to the existence of geodesics in \mathbb{S} -complexes.

Theorem 2.3. *Let K be an M_κ -complex which contains only finitely many shapes.*

- (1) *If K is an \mathbb{H} -complex, then K is globally CAT(-1) if and only if it is locally CAT(-1) and simply-connected.*
- (2) *If K is an \mathbb{E} -complex, then K is globally CAT(0) if and only if it is locally CAT(0) and simply-connected.*
- (3) *If K is an \mathbb{S} -complex, then K is globally CAT(1) if and only if it is locally CAT(1) and there are no geodesic cycles of length $< 2\pi$.*

Theorem 2.4. *If K is an M_κ -complex, then the following are equivalent:*

- (1) *K is locally CAT(κ).*
- (2) *The link of each vertex in K is globally CAT(1).*
- (3) *The link of each cell of K is an \mathbb{S} -complex which contains no closed geodesic cycle of length less than 2π .*

Thus showing that \mathbb{H} -complexes are CAT(-1) or that \mathbb{E} -complexes are CAT(0) ultimately depends on being able to show that various \mathbb{S} -complexes contain no short geodesic cycles. In order to study short geodesic cycles we introduce the concept of a configuration. Configurations are related to the “finite models” used by Bridson and Haefliger [4]. We will use configurations and the following theorem, which is a restatement of Theorem 1.11, to prove Theorem A.

Theorem 2.5. *If K is an \mathbb{S} -complex with only finitely many shapes, then there is a constant N depending only on $\text{SHAPES}(K)$, such that every local geodesic of length less than 2π has size less than N . In other words, if K contains a short closed geodesic γ , then γ is contained in a finite subcomplex of K which is the union of at most N cells from $\text{SHAPES}(K)$.*

Definition 2.6 (Configurations). A *configuration* C is a finite \mathbb{S} -complex C which contains at least one closed geodesic γ of length less than 2π . An M -complex K *contains a configuration* C if C isometrically embeds as a subcomplex of $\text{LINK}(B, K)$ for some cell B in K and under this embedding at least one of the short closed geodesics contained in C is sent to a short closed geodesic in $\text{LINK}(B, K)$. If K does not contain a configuration C , then K *avoids* this configuration.

If we fix the set of shapes under consideration and let the M_κ -complex vary, then Theorem 2.5 can be restated as an assertion about the existence of finite lists of forbidden configurations which characterize which of the complexes built out of these shapes are CAT(κ).

Theorem A. *If \mathcal{S} is a finite collection of shapes with curvature κ , then there exists a finite list of configurations \mathcal{C} such that an M_κ -complex K with $\text{SHAPES}(K) \subset \mathcal{S}$ is locally CAT(κ) if and only if K avoids all of the configurations in \mathcal{C} .*

Proof. By Theorem 2.4 it is sufficient to determine whether the link of a cell in K contains a closed geodesic of length less than 2π , and by Theorem 2.5

this closed geodesic must live in a finite subcomplex whose size is uniformly bounded by a constant which depends only on the set of shapes in \mathcal{S} . Since there are only a finite number of finite complexes of bounded size built out of the links of faces of shapes in \mathcal{S} (up to isometry), the list of configurations which will destroy the property of being locally $\text{CAT}(\kappa)$ is contained in a finite list of possibilities. Even though this does not determine which ones they are, this does show that the list of forbidden configurations is finite. \square

As we noted in the introduction, Theorem A merely asserts the existence of a finite list of forbidden configurations that need to be avoided but it does not provide an algorithm to construct the list. At present there are very few instances where the finite list for some collection of shapes is known explicitly. One of the main results along these lines is Moussong's Lemma which was first stated by Gromov ([8]) in the special case where all 1-cells in K have length $\pi/2$ and proved in general by Moussong in his dissertation [9].

Example 2.7 (Moussong's Lemma). Moussong's Lemma states that if K is a simplicial \mathbb{S} -complex with only finitely many shapes and all of the 1-cells of K have length at least $\frac{\pi}{2}$, then K is globally $\text{CAT}(1)$ if and only if K is a metric flag complex. A simplicial complex is called a *flag complex* if every 1-skeleton of an n -simplex is filled in with an n -simplex. A simplicial \mathbb{S} -complex is called a *metric flag complex* if every 1-skeleton of an n -simplex whose edge lengths are those of a possible proper spherical n -simplex is filled in with a copy of that n -simplex. This means, for example, that three 1-cells in an \mathbb{S} -complex each with length π need not and cannot bound a spherical triangle.

When all of the 1-cells have length $\pi/2$, the metric flag condition reduces to the (non-metric) flag condition. Moussong's lemma thus enables a researcher to easily determine whether an n -dimensional cubical complex with the usual metric has a locally $\text{CAT}(0)$ structure. The finite list of forbidden configurations in this case is the collection of empty simplices in dimensions 2 up to $n - 1$.

One way to paraphrase Theorem A is that for any finite collection of piecewise Euclidean or piecewise hyperbolic shapes, there is a finite description (similar to that in Moussong's lemma) of the combinatorial conditions needed in order for a complex built out of these shapes to be $\text{CAT}(0)$ or $\text{CAT}(-1)$. Such finite descriptions must exist by Theorem A, but most of them remain undiscovered.

3. GALLERIES

In this section we define one of our central concepts: that of a gallery. A gallery is a more precise tool than a configuration, as will become clear below. The term "gallery" has been borrowed from the study of Coxeter groups since a geodesic gallery in a Coxeter complex will be an example of

a linear gallery as defined below. The galleries defined here will be either linear or circular and we begin by defining the linear ones.

Definition 3.1 (Linear Galleries). Let K be a M_κ -complex, let $\gamma : [0, \ell] \rightarrow K$ be a piecewise geodesic in K , let k is the size of γ , and let $\{\sigma_i^\circ\}_{i=1}^k$ be the sequence of open cells that γ passes through. Finally, assume that for all $1 < i < k$, σ_i is either a common face of both σ_{i+1} and σ_{i-1} which are distinct or else both of these are incompatible faces of σ_i . In the former case, we say that σ_i is a *bottom cell*, and in the latter σ_i is an *top cell* in the sequence.

The *linear gallery* \mathcal{G} determined by γ is constructed inductively from this sequence of open cells. To start the induction, let \mathcal{G}_1 be a copy of the closed cell σ_1 , let α_1 be all of \mathcal{G}_1 , and let $\phi_1 : \alpha_1 \rightarrow \sigma_1$ be a specific isometry between α_1 and σ_1 . The cell α_i in \mathcal{G}_i , will always be isometric to σ_i by a specific isometry ϕ_i and the complex \mathcal{G}_{i+1} will always contain the complex \mathcal{G}_i as a subcomplex. We call α_i the *active cell in \mathcal{G}_i* . Next, assume that for some $i \geq 1$ we have defined the complex \mathcal{G}_i , its active cell α_i , and an isometric $\phi_i : \alpha_i \rightarrow \sigma_i$. The inductive step depends on whether σ_i is a top cell or a bottom cell.

- If σ_{i+1} is a bottom cell, then define $\mathcal{G}_{i+1} = \mathcal{G}_i$, define $\alpha_{i+1} = \phi_i^{-1}(\sigma_{i+1})$ and define ϕ_{i+1} as the restriction of ϕ_i to α_{i+1} .
- If σ_i is a top cell, then define let α_{i+1} be a copy of σ_{i+1} , let $\phi_{i+1} : \alpha_{i+1} \rightarrow \sigma_{i+1}$ be an isometry between the two, and let \mathcal{G}_{i+1} be the complex formed by gluing \mathcal{G}_i to α_{i+1} along their distinguished faces isometric to σ_i . Specifically identify $\phi_{i+1}^{-1}(\sigma_i) \subset \alpha_{i+1}$ with $\alpha_i \subset \mathcal{G}_i$ in the obvious fashion. Note that when \mathcal{G}_i is viewed as a subcomplex of \mathcal{G}_{i+1} , ϕ_{i+1} is an extension of ϕ_i .

The linear gallery \mathcal{G} is simply the final complex \mathcal{G}_k . The cells α_1 and α_k in \mathcal{G} are the *endcells* of \mathcal{G} . The *interior of \mathcal{G}* , denoted \mathcal{G}° , is the union of the *open cells* α_i , $i = 1, \dots, k$ in \mathcal{G} . The *boundary of \mathcal{G}* , denoted $\partial\mathcal{G}$, is the open cells in \mathcal{G} which are not in the interior. Notice that since the various maps ϕ_i are compatible on their overlaps, the gallery \mathcal{G} comes equipped with a map $\phi : \mathcal{G} \rightarrow K$ which is an isometry when restricted to any of the closed cells of \mathcal{G} .

Example 3.2. Let K be the 2-dimensional \mathbb{E} -complex formed by attaching the boundaries of two regular Euclidean tetrahedra along a 1-cell. The complex K is shown in the upper left corner of Figure 1. Let γ be the geodesic shown which starts at x travels across the front of K , around the back, over the top, and ends at y . The gallery determined by γ in shown in the upper right corner, its interior in the lower left corner and its boundary in the lower right of Figure 1.

Definition 3.3 (Circular Galleries). Let K be a M_κ -complex, let $\gamma : C_\ell \rightarrow K$ be a closed piecewise geodesic in K , let k is the size of γ , and let $\{\sigma_i^\circ\}_{i=1}^k$ be the sequence of open cells that γ passes through. Finally, assume that

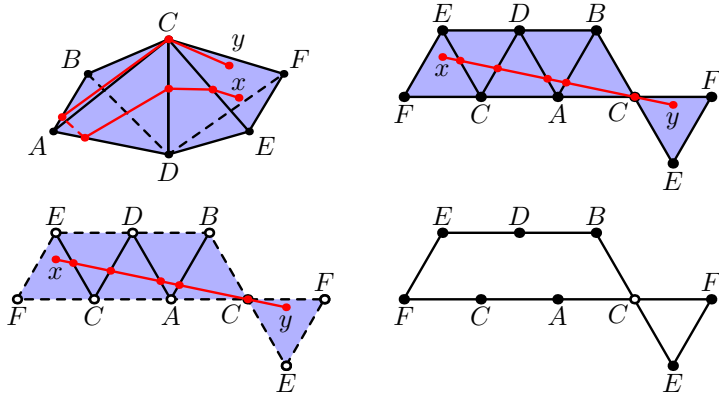


FIGURE 1. The 2-complex and linear gallery described in Example 3.2.

for all $1 \leq i \leq k$, σ_i is either a common face of both σ_{i+1} and σ_{i-1} which are distinct or else both of these are incompatible faces of σ_i where the subscripts are interpreted $\text{mod } k$. In the former case, we say that σ_i is a *bottom cell*, and in the latter σ_i is an *top cell* in the sequence.

As in Definition 3.1, the *circular gallery* \mathcal{G} determined by γ is constructed inductively from the sequence of open cells that γ passes through. Let a be an arbitrary point in C_ℓ , say one whose image lies in σ_1° . If we cut the circle at a we obtain a path γ' of length ℓ in K which is a piecewise geodesic and which happens to start and end at the same point, $\gamma'(a)$. Next, we construct a linear gallery \mathcal{G}' for the path γ' as above. Finally, we identify the endcells of \mathcal{G}' to form the circular gallery \mathcal{G} .

It should be clear that the end result is independent of the choice of a , and that the circular gallery \mathcal{G} comes equipped with a map $\phi : \mathcal{G} \rightarrow K$ which is an isometry when restricted to any of the cells of \mathcal{G} .

Definition 3.4 (Galleries in complexes). Let K be a M_κ -complex with only finitely many shapes. A complex \mathcal{G} with a map $\phi : \mathcal{G} \rightarrow K$ will be called a *linear gallery in K* if there is a piecewise geodesic γ satisfying the necessary restrictions which determines \mathcal{G} and ϕ . A *circular gallery in K* is defined analogously. We should note that when we speak of a linear gallery \mathcal{G} in K , we implicitly assume that \mathcal{G} has a specified pair of endcells, σ and τ , and a specified map $\phi : \mathcal{G} \rightarrow K$. From this information we can recover the interior \mathcal{G}° and the linear ordering on the open cells in \mathcal{G}° . Similarly, when we speak of a circular gallery \mathcal{G} in K , we implicitly assume that \mathcal{G} has a specified map $\phi : \mathcal{G} \rightarrow K$. From \mathcal{G} we can reconstruct its interior \mathcal{G}° and the circular ordering of open cells in \mathcal{G}° .

Remark 3.5 (Recognizing Galleries). Let $\phi : \mathcal{G} \rightarrow K$ be map from a finite complex to K which is an isometry when restricted to a closed cell of \mathcal{G} , and let σ and τ be two distinguished closed cells in \mathcal{G} . It is easy to determine

whether \mathcal{G} is a linear gallery with endcells σ and τ as determined by some piecewise geodesic in K . If \mathcal{G} is a gallery, then the maximal cells in the face lattice of \mathcal{G} are the top cells of the gallery. There needs to exist an ordering of these cells so that adjacent top cells have non-trivial intersections. These intersections are the bottom cells. If σ or τ is not in the list of top cells then the cells containing them are also bottom cells. At this point, we simply check whether the two bottom cells adjacent to a top cell are incompatible. If all of these conditions are true, then \mathcal{G} is a gallery determined by some piecewise geodesic. Since there are only finitely many maximal cells in \mathcal{G} and only finitely many linear ordering on this set, we can check all of the possibilities. A similar procedure will determine whether a map $\phi : \mathcal{G} \rightarrow K$ from a finite complex to K which is an isometry when restricted to a closed cell of \mathcal{G} is a circular gallery determined by some closed piecewise geodesic.

Lemma 3.6. *Let K be an M_κ -complex and let $\phi : \mathcal{G} \rightarrow K$ be the linear gallery determined by a piecewise geodesic $\gamma : [0, \ell] \rightarrow K$. The map γ lifts through ϕ in the sense that there is path $\gamma' : [0, \ell] \rightarrow \mathcal{G}$ with $\gamma = \phi \circ \gamma'$. Moreover, γ' is a piecewise geodesic in \mathcal{G} which starts in the one endcell of \mathcal{G} , ends in the other, and passes monotonically through the open cells in \mathcal{G}° .*

Proof. The notation of Definition 3.1 will be used without comment. Partition the path γ into open, closed and half-open subintervals so that the i -th subinterval is sent to the i -th open cell, σ_i° , in K that γ passes through. If σ_i is a top cell, then it contains the adjacent bottom cells in its boundary and thus σ_i contains the images of the $(i-1)$ -st and the $(i+1)$ -st subintervals if they exist. Thus all three subintervals can be simultaneously lifted to $\alpha_i \cap \mathcal{G}^\circ$ via the isometry ϕ_i^{-1} . Since these lifts agree on the bottom cells where they overlap, they combine to form the required path γ' . The second assertion is immediate from the definition of \mathcal{G} . \square

Remark 3.7 (Geodesics in galleries). Note that several piecewise geodesics may determine the same linear gallery \mathcal{G} in K . If γ is one of these paths which determine \mathcal{G} , then we say that \mathcal{G} *contains* γ . Notice in particular that if \mathcal{G} contains γ , then the lift γ' of γ to \mathcal{G} will be a piecewise geodesic in \mathcal{G}° which starts in one endcell of \mathcal{G} , ends in the other endcell of \mathcal{G} and which progresses monotonically from one endcell to the other. Conversely given a piecewise geodesic γ' in \mathcal{G} satisfying these restrictions, it is easy to see that its ϕ -image in K will be a piecewise geodesic γ which determines \mathcal{G} .

Lemma 3.8. *Let K be an M_κ -complex, and let $\phi : \mathcal{G} \rightarrow K$ be either the linear gallery determined by a geodesic $\gamma : [0, \ell] \rightarrow K$, or the circular gallery determined by a closed geodesic $\gamma : C_\ell \rightarrow K$. If γ' is the lift of γ to \mathcal{G} and α_i° is an open cell in \mathcal{G}° , then the inverse image of α_i° under γ' is a closed subinterval of $[0, \ell]$ or C_ℓ when α_i is a bottom cell and an open subinterval when α_i is a top cell.*

Proof. Notice that α_i° is a closed subspace of \mathcal{G}° when α_i is a bottom cell in \mathcal{G} and an open subspace when α_i is a top cell in \mathcal{G} . Thus the preimage

of α_i° is a closed [respectively open] subspace of $[0, \ell]$ or C_ℓ when α_i is a bottom [top] cell. Since γ' proceeds monotonically through the interiors of the α_i (Lemma 3.6), these open/closed subspaces can only be open/closed subintervals. \square

Lemma 3.9. *Let K be an M_κ -complex, and let $\phi : \mathcal{G} \rightarrow K$ be either the linear gallery determined by a geodesic $\gamma : [0, \ell] \rightarrow K$, or the circular gallery determined by a closed geodesic $\gamma : C_\ell \rightarrow K$. If γ' is the lift of γ to \mathcal{G} and α is a bottom cell in \mathcal{G} , then the inverse image of α° under γ' is a single point. The same is true for circular gallery.*

Proof. If α is a bottom cell whose preimage under γ' is not a single point, then by considering α and an adjacent top cell, it is easy to see that γ' (and hence γ) is not a local geodesic. This shows that preimages of the interiors of bottom cells must be single points. \square

Lemma 3.10. *Let K be an M_κ -complex and let $\phi : \mathcal{G} \rightarrow K$ be the linear gallery determined by a geodesic $\gamma : [0, \ell] \rightarrow K$. If γ' is the lift of γ to \mathcal{G} , and $L = \{0 = x_0 < x_1 < \dots < x_k = \ell\} \subset [0, \ell]$ is the list of point preimages of bottom cells under γ' together with the endpoints of the interval, then the image of γ' in \mathcal{G} is completely determined by the images of the points in L . Moreover, γ' is an embedding.*

Proof. If α denotes a top cell in \mathcal{G} , then the preimage of α under γ' is a closed subinterval $[x_i, x_{i+1}]$ for some i (Lemma 3.8 and Lemma 3.9). Since γ' is a local geodesic, and since the image of this subinterval lies completely in the closed cell α , this image must be the unique geodesic in α connecting $\gamma'(x_i)$ to $\gamma'(x_{i+1})$. This shows that γ' is an embedding and that its image is completely determined by the image of L in \mathcal{G} as required. \square

Lemma 3.11. *Let K be an M_κ -complex and let $\phi : \mathcal{G} \rightarrow K$ be the circular gallery determined by a closed geodesic $\gamma : C_\ell \rightarrow K$. If γ' is the lift of γ to \mathcal{G} , and $L = \{0 \leq x_0 < x_1 < \dots < x_k < \ell\} \subset [0, \ell)$ is the list of point preimages of bottom cells under γ' , then the image of γ' in \mathcal{G} is completely determined by the images of the points in L . Moreover, γ' is an embedding.*

Proof. If α denotes a top cell in \mathcal{G} , then the preimage of α under γ' is an closed subinterval $[x_i, x_{i+1}]$ for some i , or $[x_k, x_0 + \ell]$ (Lemma 3.8 and Lemma 3.9). Since γ' is a local geodesic, and since the image of this subinterval lies completely in the closed cell α , this image must be the unique geodesic in α connecting $\gamma'(x_i)$ to $\gamma'(x_{i+1})$. This shows that γ' is an embedding and that its image is completely determined by the image of L in \mathcal{G} as required. \square

Lemma 3.12. *Let K be an M_κ -complex, let $\phi : \mathcal{G} \rightarrow K$ be the linear gallery determined by a piecewise geodesic $\gamma : [0, \ell] \rightarrow K$, and let $\gamma' : [0, \ell] \rightarrow \mathcal{G}$ be the lift of γ to \mathcal{G} . The map ϕ restricted to \mathcal{G}° is an immersion.*

Proof. Let \mathcal{G} be the linear gallery determined by a geodesic γ and let $\phi : \mathcal{G} \rightarrow K$ be canonical map. The map ϕ is clearly an embedding when restricted

to each closed cell α_i in \mathcal{G} . For a point x in the interior of a top cell of \mathcal{G} , there is a neighborhood of x which lies completely in α_i° . Thus ϕ is an embedding in a neighborhood of x . If x is a point in the interior of a bottom cell α_i (which is not an endcell of \mathcal{G}), then there is a neighborhood of x which lies in $(\alpha_{i-1} \cup_{\alpha_i} \alpha_{i+1}) \cap \mathcal{G}^\circ$. The restrictions on γ now ensure that ϕ is an embedding when restricted to this subcomplex. Finally, when α_i is a bottom cell which is an endcell, there is neighborhood of x which lies in either $\alpha_{i+1} \cap \mathcal{G}^\circ$ or $\alpha_{i-1} \cap \mathcal{G}^\circ$. \square

Lemma 3.13. *Let K be an M_κ -complex, let $\phi : \mathcal{G} \rightarrow K$ be the linear [or circular] gallery determined by a local [closed] geodesic $\gamma : [0, \ell] \rightarrow K$, and let $\gamma' : [0, \ell] \rightarrow \mathcal{G}$ be the lift of γ to \mathcal{G} . There exists a deformation retraction from \mathcal{G}° to the image of γ' .*

Proof. To prove there is a deformation retraction from \mathcal{G} to γ' we first note that for each bottom cell of \mathcal{G} , we can use a radial retraction from α_i° to the unique point where γ' intersects α_i° (Lemma 3.9). Next for each top cell α_i there are two points which determine the intersection of γ' with α_i . These two points are either the two endpoints of γ' , an endpoint and a point in the interior of a boundary cell of α_i , or two points in the interiors of two incompatible boundary cells of α_i . In all three cases it is easy to see that there is a retraction of $\alpha_i \cap \mathcal{G}^\circ$ which is compatible with the deformations already defined on the interiors of the boundary cell(s) of α_i . \square

Definition 3.14 (Equivalent galleries). Let K be an M_κ -complex and let \mathcal{G} and \mathcal{G}' be two of its galleries, both linear or both circular. The galleries \mathcal{G} and \mathcal{G}' will be considered *equivalent* if there is an isometry between them which preserves their cell structures. They will be considered *identical* if they are equivalent and they are immersed in K in the same way. More specifically, if \mathcal{G} and \mathcal{G}' are immersed in K by maps ϕ and ϕ' , then there is a map $\rho : \mathcal{G} \rightarrow \mathcal{G}'$ showing they are equivalent with $\phi = \phi' \circ \rho$.

Lemma 3.15. *Let K be a M_κ -complex with $\text{SHAPES}(K)$ finite, let ℓ be a fixed real number. If a number N satisfying the conclusion of Theorem 1.11 can be constructed from $\text{SHAPES}(K)$ and ℓ , then there exists a finite, constructible list of linear galleries such that every geodesic in K of length at most ℓ determines a gallery equivalent to a linear gallery in this list. There is a similar list of circular galleries such that every closed geodesic in K of length at most ℓ determines a circular gallery equivalent to a gallery in this list.*

Proof. We prove the linear case; the circular case is analogous. Let Γ be the (infinite) collection of all geodesics in K of length at most ℓ . By Theorem 1.11 there is an N such that each γ in Γ passes through at most N open cells of K . As a result each γ in Γ determines a gallery formed by gluing together at most N closed cells, each of which is isometric to a closed cell in $\text{SHAPES}(K)$. Because $\text{SHAPES}(K)$ is finite, there are only a finite number of complexes which can be formed in this way, and by Remark 3.5 we can

$$\begin{array}{ccccc}
& & [0, \ell] & \rightarrow & \mathbb{S}^1 & \subset & \mathbb{R}^2 \\
& \swarrow & \downarrow & & & & \\
K & \leftarrow & \mathcal{G} & \rightarrow & \mathbb{S}^n & \subset & \mathbb{R}^{n+1}
\end{array}$$

FIGURE 2. Spaces and maps used in Theorem 4.1

determine which of these are linear galleries. The linear gallery determined by γ must be equivalent to one of them. \square

4. DETECTING PIECEWISE GEODESICS

In this section we discuss in detail how to search for piecewise geodesics of bounded length in a linear or circular gallery using boolean combinations of polynomial equations and inequalities.

Theorem 4.1. *If K is finite n -dimensional \mathbb{S} -complex and \mathcal{G} is a linear gallery in K , then we can construct a system of polynomial equations and inequalities which has a solution if and only if the gallery \mathcal{G} contains a piecewise geodesic of length less than π . As a consequence, there exists an algorithm to test whether \mathcal{G} contains such a piecewise geodesic.*

Proof. Let $\gamma : [0, \ell] \rightarrow K$ be a piecewise geodesic which is contained in \mathcal{G} , and let $\gamma' : [0, \ell] \rightarrow \mathcal{G}$ be its lift to \mathcal{G} . For each bottom cell α in \mathcal{G} , let x_α be the smallest value in the preimage of α° under γ' . Such a value exists by Lemma 3.8. Clearly the ordering of the x_α is consistent with the linear ordering of the bottom cells in \mathcal{G} . Next, define a list $L = \{0 = x_0 < x_1 < \dots < x_k = \ell\} \subset [0, \ell]$ which consists endpoints of the domain together with the collection of x_α defined above. Let α_1° and α_k° denote the open cells of \mathcal{G} containing $\gamma'(0)$ and $\gamma'(\ell)$ and let α_i° , $1 < i < k$ denote the open bottom cell of \mathcal{G} containing $\gamma'(x_i)$.

If there is an i such that γ' does not send this subinterval $[x_i, x_{i+1}]$ to the unique geodesic connecting $\gamma'(x_i)$ and $\gamma'(x_{i+1})$ in the unique top cell containing both points, then we can define a new path γ'' with this property which determines the same list L , but with a strictly shorter length. Thus, if \mathcal{G} contains a path of length less than π , it contains a path of this type. We may therefore assume that γ' has this shortest geodesic property for each i , and as a result that γ' intersects each bottom cell of \mathcal{G} in a unique point.

We now establish a system of equations and inequalities which will test whether the sum of the distances from x_i to x_{i+1} , $i = 0, \dots, k-1$ is less than π or not. In one sense all we need to do is to add up the lengths of the unique individual arcs connecting $\gamma'(x_i)$ to $\gamma'(x_{i+1})$, but we wish to do this with polynomial equations. Let V denote the set of 0-cells in \mathcal{G} and let E denote the set of 1-cells in \mathcal{G} . We will introduce the necessary variables and equations in three stages. The spaces and maps used are summarized in Figure 2.

Step 1: For every $v \in V$ we introduce a vector $\vec{u}_v = (u_{v,1}, \dots, u_{v,n+1})$ in \mathbb{R}^{n+1} and we add the equation $\vec{u}_v \cdot \vec{u}_v = 1$ to our system. In other words, we

introduce n new variables $u_{v,i}$, $i = 1, \dots, n+1$ and the polynomial equation $u_{v,1}^2 + u_{v,2}^2 + \dots + u_{v,n+1}^2 = 1$. The vector notation is simply a convenient shorthand. Next, for each 1-cell $e \in E$ connecting v to v' we add the equation $\vec{u}_v \cdot \vec{u}_{v'} = \cos(\theta)$ to our system, where θ is the arc length of e . Notice that the solutions to this system of equations defined so far are in one-to-one correspondence with the possible maps from \mathcal{G} to \mathbb{S}^n which restrict to an isometry on each closed cell of \mathcal{G} .

Step 2: For each x_i in L , we introduce a vector $\vec{y}_i = (y_{i,1}, \dots, y_{i,n+1})$, and add the equation $\vec{y}_i \cdot \vec{y}_i = 1$ which stipulates that \vec{y}_i has length 1. In addition, we add equations and inequalities which stipulate that \vec{y}_i is a positive linear combination of the vectors \vec{u}_v corresponding to the 0-cells v of α_i . When α_i is 0-dimensional the positive linear combination will reduce to a set of equations of the form $\vec{y}_i = \vec{u}_v$.

Notice that we are treating the position of $\gamma(x_i)$ in α_i° as unspecified. For example, if the start point $\gamma(x)$ lies in a 1-cell e in \mathcal{G} with 0-cells v and v' , then the equations added so far only state that the image of the endpoint x , \vec{y}_1 , is a positive linear combination of the vectors \vec{u}_v and $\vec{u}_{v'}$. They do not state what coefficients of that positive linear combination are. The variables and equations in steps 1 and 2 ensure that encoded in each solution to these system is the description of a distinct piecewise geodesic path between the endcells of \mathcal{G} .

Step 3: To check whether the total length of γ' is at least π , we use a 2-dimensional model space. For each point x_i in L we introduce a new vector $\vec{z}_i = (z_{i,1}, z_{i,2})$ in \mathbb{R}^2 starting with $\vec{z}_0 = (0, 1)$. Then we add the following equations:

$$\begin{aligned} \vec{z}_i \cdot \vec{z}_i &= 1, & \text{for } i = 0, 1, \dots, k \\ \vec{z}_{i-1} \cdot \vec{z}_i &= \vec{y}_{i-1} \cdot \vec{y}_i, & \text{for } i = 1, \dots, k \\ \det(\vec{z}_{i-1}, \vec{z}_i) &> 0, & \text{for } i = 1, \dots, k \end{aligned}$$

These equations stipulate that the length of \vec{z}_i is 1, that the angle between \vec{z}_i and \vec{z}_{i+1} equals the angle between \vec{y}_i and \vec{y}_{i+1} , and that $(\vec{z}_i, \vec{z}_{i+1})$ is a positively oriented frame for \mathbb{R}^2 . The third set of equations has this interpretation because the cells of \mathcal{G} are proper and thus all distances in these cells are less than π . Taken together, the second and third sets of equations ensure that \vec{z}_i is obtained from \vec{z}_{i-1} by a counterclockwise rotation through an angle equal to the distance from $\gamma(x_{i-1})$ to $\gamma(x_i)$ in \mathcal{G} . In other words the points \vec{z}_0, \vec{z}_1 , etc. are marching around the unit circle in \mathbb{R}^2 in a counterclockwise direction starting on the positive x -axis.

Step 4: Finally, since each angle is less than π , the length of the geodesic will be less than π if and only if each \vec{z}_i , $i > 0$, has a positive second coordinate. Thus we add the equations $z_{i,2} > 0$, for $i = 1, \dots, k$.

Combining all of the variables and equations introduced in each of these steps, we get a system of equations and inequalities which has a solution if and only if \mathcal{G} contains a piecewise geodesic of length less than π . The second assertion follows immediately from Tarski's theorem. \square

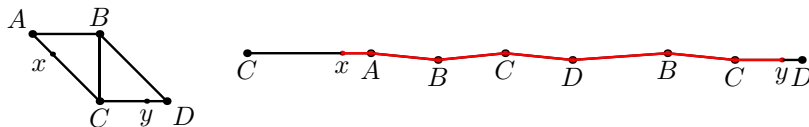


FIGURE 3. The 1-complex and gallery described in Example 4.2

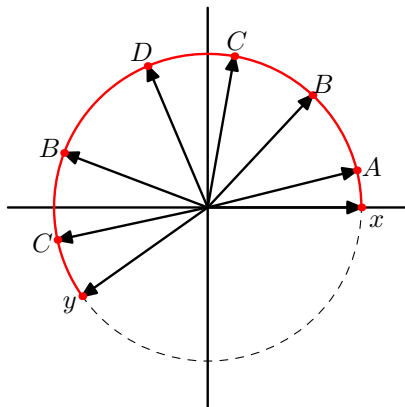


FIGURE 4. Model space determined by $\{x, A, B, C, D, B, C, y\}$.

Example 4.2. To illustrate the procedure described in Theorem 4.1 consider the 1-dimensional \mathbb{S} -complex on the left in Figure 3. The gallery determined by the sequence of points $\{x, A, B, C, D, B, C, y\}$ is shown on the right and the model space for this sequence is shown in Figure 4. The vectors in the model space are the vectors \vec{z}_i . In this example \vec{z}_6 and \vec{z}_7 have negative second coordinates, the total length is at least π , and thus this is not a solution to the system of equations and inequalities.

There is a similar theorem for circular galleries.

Theorem 4.3. *If K is a finite n -dimensional \mathbb{S} -complex and \mathcal{G} is a circular gallery in K , then we can construct a boolean combination of polynomial equations and inequalities which has a solution if and only if the gallery \mathcal{G} contains a closed piecewise geodesic of length less than 2π . As a consequence, there exists an algorithm to test whether \mathcal{G} contains such a closed piecewise geodesic.*

Proof. The proof is nearly identical to the one above, so only the differences will be noted. We begin noting that the circular gallery \mathcal{G} was constructed by first creating a linear gallery \mathcal{G}' and then identifying the two endcells of \mathcal{G}' by an isometry. We will need the unidentified linear gallery \mathcal{G}' to create some of our equations. By picking a point in top cell as a basepoint of the loop, we may assume that the endcells of \mathcal{G}' are top cells.

The notation is as before, but slightly modified since \mathcal{G} is circular. There is a loop $\gamma' : C_\ell \rightarrow \mathcal{G}$ and a path $\gamma'' : [0, \ell] \rightarrow \mathcal{G}'$. There is a list $L = \{0 < x_1 < x_2 < \dots < x_k < \ell\} \subset [0, \ell]$ where the point $\gamma'(x_i)$ is contained in the interior of a bottom cell of \mathcal{G}' denoted α_i . Moreover, γ' is a piecewise geodesic whose image is determined by the points $\gamma'(x_i)$. Let V denote the set of 0-cells in \mathcal{G}' , let E denote the set of 1-cells in \mathcal{G}' .

Step 1: Same as above with \mathcal{G} replaced by \mathcal{G}' . The solutions to this system of equations are in one-to-one correspondence with the possible maps from \mathcal{G}' to \mathbb{S}^n which restrict to an isometry on each closed cell of \mathcal{G}' .

Step 2: Same as above with the addition that since x_1 and x_k are two preimages of the same point in \mathcal{G} , we add equations that state that each coefficient variable used to describe \vec{y}_0 as a positive linear combination of the vectors \vec{u}_v corresponding to the 0-cells in α_1 , is equal to the corresponding coefficient variable used to describe \vec{y}_k as a positive linear combination of the vectors \vec{u}_v corresponding to the 0-cells in α_k . The correspondence of the 0-cells is determined by the isometry used to obtain \mathcal{G} from \mathcal{G}' . The variables and equations in steps 1 and 2 ensure that encoded in any solution to the system so far is the description of a piecewise geodesic loop which travels monotonically around \mathcal{G} .

Step 3: Same as above.

Step 4: When we were checking length less than π this only involved checking whether all the second coordinates were positive. This time the length calculation should determine whether the length is less than 2π which is more involved. For each $i = 2, \dots, k - 1$ add new variables p_i and q_i and equations of the form

$$\begin{aligned} p_i \vec{z}_i + q_i \vec{z}_{i+1} &= \vec{z}_1 \\ p_i < 0 \text{ or } q_i < 0 \end{aligned}$$

This is just a polynomial way of saying that \vec{z}_1 is not a nonnegative linear combination of \vec{z}_i and \vec{z}_{i+1} (which ensures that the total length is less than 2π). To see this note that if z_{i+1} ever comes back around the full circle in this model space, then z_1 will be a nonnegative linear combination of z_i and z_{i+1} , $i > 1$.

Combining all of the variables and equations introduced in each of these steps, we get a boolean combination of polynomial equations and inequalities which has a solution if and only if \mathcal{G} contains a piecewise geodesic loop of length less than 2π . The second assertion follows immediately from Tarski's theorem. \square

Remark 4.4 (Coefficients). The only equations in these systems with potentially non-rational coefficients are the equations of the form $\vec{u}_v \cdot \vec{u}_{v'} = \cos(\theta)$. Such an equation can be rewritten as the equation $(\vec{u}_v \cdot \vec{u}_{v'})^2 = \cos^2(\theta)$, and an inequality stipulating that $\vec{u}_v \cdot \vec{u}_{v'}$ has the appropriate sign. Thus, if $\cos(\theta)$ is at worst a square root of a rational number, using this

alternative ensures the systems use only rational coefficients. Such considerations can play a significant role when implementing algorithms, and we will explore the issue of coefficients in greater detail in Section 6.

5. DETECTING LOCAL GEODESICS

In this section we discuss how to search for local geodesics of bounded length in a linear or circular gallery using boolean combinations of polynomial equations and inequalities. The requirement that the paths be local geodesics instead of merely piecewise geodesics is the main complication. In the process we show how to bound the length of a linear gallery which can contain a short geodesic. The results in this section will be proved by induction on n . We begin with dimension $n = 1$.

Theorem 5.1. *If K is a finite n -dimensional \mathbb{S} -complex and \mathcal{G} is a linear gallery in K , then there exists a boolean combination of polynomial equations and inequalities which has a solution if and only if \mathcal{G} contains a local geodesic of length less than π whose image in K is also a geodesic. As a consequence there exists an algorithm which determines whether or not \mathcal{G} contains such a geodesic.*

Proof. When $n = 1$, the system created is the one created in the proof of Theorem 4.1. Because K and hence \mathcal{G} have dimension 1, the piecewise geodesics encoded in a solution to this system will be precisely the local geodesics contained in \mathcal{G} , and they will also remain local geodesics when immersed into K .

For $n > 1$, we may assume that results in this section have already been proved in lower dimensions. Let P denote the system created in the proof of Theorem 4.1. In essence, we will start with P and then add restrictions to require that any piecewise geodesics encoded in a solution to P will be locally geodesic at the transitions between the geodesic pieces. We will use the notations established in the proof of Theorem 4.1 without further comment.

Let α_i be a bottom cell of \mathcal{G} . Given any two points a and b in α_i° , the links $\text{LINK}(\phi(a), K)$ and $\text{LINK}(\phi(b), K)$ are canonically isometric. Let L_i be this well-defined \mathbb{S} -complex and note that its dimension is strictly less than n . In general the cells in L_i will not be proper so we assume that it has been suitably subdivided. Notice that the three points $\gamma(x_{i-1})$, $\gamma(x_i)$ and $\gamma(x_{i+1})$ determine two points in L_i . The complex L_i will have a finite number of galleries which could possibly contain a geodesic of length less than π by Lemma 5.6. Let G_i be the list of such galleries in L_i . For each gallery in G_i we write down the variables and equations necessary to the encode a generic path determined by crossing points into the gallery (this would be steps 1 and 2 in Theorem 4.1). Do this for each gallery in G_i and repeat this procedure for each of the links L_i . Each gallery in each link will use distinct variables.

For each i and for each gallery \mathcal{H}_j in G_i , we create a system $V_{i,j}$ which is a boolean combination of polynomial equations and inequalities. First we create the equations and inequalities needed to say that the three points $\gamma(x_{i-1})$, $\gamma(x_i)$ and $\gamma(x_{i+1})$ determine two points in L_i which lie in the endcells of \mathcal{H}_j . Then we add the boolean combination of polynomial equations and inequalities (which exist by the lower-dimensional versions of this theorem) which have a solution if and only if there is a local geodesic in \mathcal{H}_j connecting these exact points in its endcells.

Intuitively, the system $V_{i,j}$ has a solution if and only if the the piecewise geodesic in \mathcal{G} which is encoded in a solution of the system P , there is a path less than π contained in the gallery \mathcal{H}_j in the link of the point $\gamma(x_i)$ pushed into K . More informally, we might say that the solutions to $V_{i,j}$ encode the piecewise geodesics in \mathcal{G} which have a “kink” at $\phi(\gamma(x_i))$ as witnessed by a short path through gallery \mathcal{H}_j .

The final system is one which starts with P and subtracts off the solutions to each of the $V_{i,j}$. Clearly, this is a boolean combination of these solution sets. This final system will have a solution if and only if \mathcal{G} contains a piecewise geodesic of length less than π which does not fail to be locally geodesic in any of the finite number of possible ways which are available to it. In other words, if and only if \mathcal{G} contains a local geodesic of length less than π . Finally, the algorithmic assertion follows immediately from Tarski’s Theorem. \square

There is a similar theorem for circular galleries.

Theorem 5.2. *If K is a finite n -dimensional \mathbb{S} -complex and \mathcal{G} is a circular gallery in K , then there exists a boolean combination of polynomial equations and inequalities which has a solution if and only if \mathcal{G} contains a local geodesic of length less than 2π whose image in K is also a local geodesic. As a consequence there exists an algorithm which determines whether or not \mathcal{G} contains such a geodesic.*

Proof. The proof is identical to the one above, but starting with the system created in the proof of Theorem 4.3. Everything else is unchanged. \square

Next, we show how to construct the number N in dimension n .

Lemma 5.3. *If K be a n -dimensional \mathbb{S} -complex with finitely many shapes and $\ell > 0$ is a real number, then there is an algorithm which computes an integer N such that every local geodesic contained in a linear geodesic \mathcal{G} of size N in K has length at least ℓ .*

Proof. Note that Theorem 5.1 shows that it is possible to determine whether any specific linear gallery \mathcal{G} contains a local geodesic of length less than π . To find the smallest value of N which will work for $\ell = \pi$ we simply test a representative of each equivalence class of galleries of size 1, then a representative of each equivalence class of galleries of size 2, etc. That there is a finite constructible list of such galleries is guaranteed by Lemma 3.15. As

soon as we find an N such that all galleries of size N do not contain a local geodesic of less than length π , we are essentially done since by Remark 1.12 it is sufficient to find N for a single number ℓ . That such an N must exist follows from Theorem 1.11 and the fact that local geodesics contained in linear galleries of size N themselves have size N . Note that we are using the fact that geodesics “contained in” \mathcal{G} remain in the interior of the gallery (Remark 3.7). A geodesic between the endcells of \mathcal{G} which is allowed to contain points from the boundary of \mathcal{G} may very well have a smaller size. \square

Now that we have shown that the N in Theorem 1.11 is constructible, we can state the following corollaries.

Corollary 5.4. *Let K be a finite n -dimensional \mathbb{S} -complex and let ℓ be a fixed real number. Given open cells σ° and τ° in K there is a finite, constructible list of linear galleries from σ to τ , based only on $\text{SHAPES}(K)$ and ℓ , such that every geodesic from σ to τ of length at most ℓ determines a gallery in this list.*

Proof. Let N be a size which guarantees a length greater than ℓ and then construct all possible linear galleries with size at most N up to equivalence (Lemma 3.15), and then construct all possible maps from \mathcal{G} to K starting at σ . This is possible since N exists (Theorem 1.11), it is constructible (Lemma 5.3), and K is locally-finite. \square

Corollary 5.5. *Let K be a finite n -dimensional \mathbb{S} -complex and let ℓ be a fixed real number. There is a finite, constructible list of circular galleries such that every closed geodesic of length at most ℓ determines a circular gallery in this list.*

Proof. Let N be a size which guarantees a length greater than ℓ and then construct all possible circular galleries \mathcal{G} with size at most N up to equivalence (Lemma 3.15), and then construct all possible maps from \mathcal{G} to K . This is possible since N exists (Theorem 1.11), it is constructible (Lemma 5.3), and K is locally finite. \square

Finally, we prove a lemma to complete the induction.

Lemma 5.6. *Let K be a finite n -dimensional \mathbb{S} -complex. Given open cells σ° and τ° in K , there exists a boolean combination of polynomial equations and inequalities which has a solution if and only if there is a geodesic in K from some point in σ° to some point in τ° of length less than π . A similar statement holds when the points in σ° and τ° are specified.*

Proof. First note that any such geodesic determines a linear gallery. By Corollary 5.4 we can construct a finite list L which contains all linear galleries $\mathcal{G} \rightarrow K$ from σ to τ which could possibly be determined by a geodesic of length less than π . For each gallery in L , use Theorem 5.1 to create a system of polynomial equations which has a solution if and only if this particular gallery does not contain a piecewise geodesic of length less than π . If each

of these systems are created using distinct variables, then their union will be a finite system which has a solution if and only if there does not exist a path in K from σ° to τ° of length less than π . \square

6. COEFFICIENTS

The results in this section are a slight digression. Let K be a finite n -dimensional \mathbb{E} -complex. In this section we show that the type of numbers which show up as coefficients of the polynomials used to test whether K is locally CAT(0) are no worse than the types of numbers which show up as lengths of 1-cells in K . A precise version of this claim is contained in Theorem 6.6 where it is presented in terms of F -simplices.

Definition 6.1 (F -simplices). Let F be a subfield of \mathbb{R} and let σ be n -dimensional Euclidean simplex. If the square of the length of each 1-cell in σ lies in F , then σ is called an F -simplex.

Lemma 6.2. Let F be a subfield of \mathbb{R} , let $\sigma \subset \mathbb{R}^n$ be an k -dimensional Euclidean simplex with vertices v_i , $i = 0, \dots, k$, and let \vec{v}_i be the vector from v_0 to v_i . The following conditions are equivalent:

- (1) σ is an F -simplex,
- (2) $\vec{v}_i \cdot \vec{v}_j \in F$ for all i, j ,

Proof. (2 \Rightarrow 1) Since $\|\vec{v}_i\|^2 = \vec{v}_i \cdot \vec{v}_i$ we conclude that the square of the length of the edge connecting v_0 to v_i lies in F , and since

$$\|\vec{v}_j - \vec{v}_i\|^2 = (\vec{v}_j - \vec{v}_i) \cdot (\vec{v}_j - \vec{v}_i) = \|\vec{v}_i\|^2 + \|\vec{v}_j\|^2 - 2\vec{v}_i \cdot \vec{v}_j$$

we conclude that square of the length of the edge connecting v_i to v_j , $i, j > 0$, lies in F as well. (1 \Rightarrow 2) Since the square of the length of the edge connecting v_0 to v_i is $\|\vec{v}_i\|^2 = \vec{v}_i \cdot \vec{v}_i$, we conclude that $\vec{v}_i \cdot \vec{v}_j$ lies in F when $i = j$. Similarly, the square of the length of the edge connecting v_i to v_j is $\|\vec{v}_j - \vec{v}_i\|^2 = \|\vec{v}_i\|^2 + \|\vec{v}_j\|^2 - 2\vec{v}_i \cdot \vec{v}_j$. Since this result lies in F , -2 is in F , and the first two of its terms lies in F by assumption, we conclude that $\vec{v}_i \cdot \vec{v}_j \in F$ for $i \neq j$ as well. \square

Lemma 6.3. Let F be a subfield of \mathbb{R} and let $\sigma \subseteq \mathbb{R}^n$ be an k -dimensional F -simplex with one vertex, v_0 , at the origin. There is an orthogonal basis U of \mathbb{R}^n such that the squared length of each basis vector in U lies in F and all of the coordinates of all of the vertices of σ , relative to U , lies in F .

Proof. By replacing \mathbb{R}^n with the subspace spanned by the vectors \vec{v}_i if necessary, we may assume that $k = n$ and that collection of vectors $V = \{\vec{v}_i\}_{i=1}^n$ is a basis for \mathbb{R}^n . Next apply the Gram-Schmidt process to V without normalizing the final result. Specifically, let $\vec{u}_1 = \vec{v}_1$, and for $i > 1$ inductively define

$$\vec{u}_i = \vec{v}_i - \sum_{j=1}^{i-1} \frac{\vec{v}_i \cdot \vec{u}_j}{\vec{u}_j \cdot \vec{u}_j} \vec{u}_j.$$

An easy induction shows that \vec{u}_i is an F -linear combination of the vectors \vec{v}_j , $j \leq i$ so that the fraction in front of the vector \vec{u}_j in the above equation is an element of F . In particular, the change of basis matrix connecting U and V is an upper triangular matrix over F with 1's down the diagonal. As a result, the same is true about the change of basis matrix going in the other direction, thus showing that each \vec{v}_i is an F -linear combination of the vectors \vec{u}_j with $j \leq i$. \square

Lemma 6.4. *If F be a subfield of \mathbb{R} and σ is a F -simplex, then the square of each trigonometric function of each dihedral angle in σ lies in F .*

Proof. First note that it is sufficient to prove this for the square of the cosine of such an angle. Let $V = \{v_0, \dots, v_k\}$ be the set of vertices of σ . A dihedral angle between two faces of σ is determined by a subset $S \subset V$ and distinct $s, t \in V \setminus S$. In particular, there is a dihedral angle θ between the face determined by $S \cup \{s\}$ and the face determined by $S \cup \{t\}$. By renumbering if necessary we may assume that $S = \{v_0, \dots, v_{i-1}\}$, $s = v_i$, and $t = v_{i+1}$. Using the orthogonal basis U derived in Lemma 6.3, it is clear that $\cos^2(\theta)$ lies in F . Specifically, θ is angle between the projections of $vecv_i$ and \vec{v}_{i+1} into the subspace perpendicular to the subspace determined by S . In the basis U , this projection is accomplished by simply “zeroing out” the first $i-1$ coordinates. As a result these projections are still F -linear combinations of the basis vectors. That $\cos^2(\theta)$ lies in F now follows immediately. \square

One special case which deserves to be highlighted is where the field $F = \mathbb{Q}$. An angle θ for which $\cos^2(\theta)$ is rational is called *geodetic*. See [6] for a computational description of the space of geodetic angles.

Corollary 6.5. *If σ be a Euclidean \mathbb{Q} -simplex, then every dihedral angle in σ is geodetic.*

Notice that \mathbb{Q} -simplices are quite common. For example, every simplex in a simplicial decomposition of a rational polytope will be a \mathbb{Q} -simplex.

Theorem 6.6 (Coefficients). *Let K be a finite simplicial \mathbb{E} -complex and let F be a subfield of \mathbb{R} . If every simplex in K is an F -simplex, then the systems of polynomial equations and inequalities used to test whether K is locally CAT(0) have coefficients which belong to F .*

Proof. As we mentioned in Remark 4.4, the only coefficients used which are not integers are equal to $\cos^2(\theta)$ for some dihedral angle θ in some simplex of K . But by Lemma 6.4, this number lies in F . \square

Similar theorems could clearly be derived for other curvatures, but they will not be pursued here.

7. THE MAIN THEOREM

At this point the proof of Theorem B is nearly immediate.

Theorem 7.1. *There exists an algorithm which determines whether or not a finite \mathbb{S} -complex contains a closed geodesic of length less than 2π , and as a consequence there exists an algorithm which determines whether or not a finite \mathbb{S} -complex is globally CAT(1).*

Proof. A closed geodesic of length less than 2π in the complex would determine a circular gallery which contains it. By Corollary 5.5 there is a finite, constructible list of circular galleries in K which could possibly contain a closed geodesic of length less than 2π , and by Theorem 5.2 we can test each one for the presence of a short closed geodesic which survives in K . This proves the first assertion.

To prove the second let K be the finite \mathbb{S} -complex under consideration. By Theorem 2.3 and Theorem 2.4 it is sufficient to show that K contains no closed geodesic of length less than 2π and that for each cell B in K , $\text{LINK}(B, K)$ contains no closed geodesic of length less than 2π . Since K and $\text{LINK}(B, K)$ are finite \mathbb{S} -complexes for all cells B , the first part of the proof shows that these are checkable conditions. \square

Theorem B. *There exists an algorithm which determines whether or not a finite M_κ -complex is locally CAT(κ).*

Proof. For each vertex v in K , the link $\text{LINK}(v, K)$ is a finite \mathbb{S} -complex and by Theorem 7.1 we can check whether it is globally CAT(1). By Theorem 2.4 this is sufficient. \square

As can be seen from the proof, it is actually sufficient for the complex K to be locally-finite, so long as (1) only a finite list of finite \mathbb{S} -complexes occur as links of 0-cells up to isometry, and (2) there is a constructive procedure for enumerating this finite list.

Computational real algebraic geometry is able to show that testing curvature conditions is decidable and algorithmic, but the real algebraic sets described here quickly reach the limit of what is realistically computable using current techniques. For the working geometric group theorist who would like software which determines the local curvature of a finite M_κ -polyhedral complex, there are essentially two main options:

- use sophisticated techniques from computational real algebraic geometry to speed up the computations, or
- develop alternative, simpler algorithms which work directly with the geometry.

Both strategies are currently under investigation. The latter strategy, in particular, has already produced results in the form of a direct geometric algorithm to determine the local curvature of a 3-dimensional M_κ -complex using only elementary 3-dimensional geometry. This result can be found in [7]. At this point, the following question remains open:

Open Question 7.2. Is there a direct *geometric* algorithm which determines the local curvature of a finite M_κ -complex in dimensions greater than 3?

REFERENCES

- [1] W. Ballmann. *Lectures on spaces of nonpositive curvature*. Birkhäuser Verlag, Basel, 1995. With an appendix by Misha Brin.
- [2] J. Bochnak, M. Coste, and M. Roy. *Real Algebraic Geometry*. Springer-Verlag, Berlin, 1998.
- [3] M. 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.
- [4] M. R. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*. Springer-Verlag, Berlin, 1999.
- [5] B. F. Caviness and J. R. Johnson, eds. *Quantifier eliminations an cylindrical algebraic decomposition*. Springer-Verlag, Berlin, 1998.
- [6] J. H. Conway, C. Radin, and L. Sadun. On angles whose squared trigonometric functions are rational. *Discrete Comput. Geom.*, 22(3):321–332, 1999.
- [7] M. Elder and J. McCammond. Curvature testing in 3-dimensional metric polyhedral complexes. Preprint 2001.
- [8] M. Gromov. Hyperbolic groups. In *Essays in group theory*, pages 75–263. Springer, New York, 1987.
- [9] G. Moussong. Hyperbolic Coxeter groups. Ph.D. dissertation, The Ohio State University, 1988.

DEPT. OF MATHEMATICS, TUFTS UNIVERSITY, MEDFORD, MA 02155

E-mail address: Murray.Elder@tufts.edu

DEPT. OF MATH., UNIVERSITY OF CALIFORNIA, SANTA BARBARA, SANTA BARBARA, CA 93106

E-mail address: jon.mccammond@math.ucsb.edu