

Curvature Testing in 3-Dimensional Metric Polyhedral Complexes

Murray Elder and Jon McCammond

CONTENTS

- 1. Introduction
- 2. Preliminaries
- 3. Galleries
- 4. Annular Galleries
- 5. Möbius Galleries
- 6. Disc Galleries
- 7. Necklace Galleries
- 8. Enumerating Galleries
- 9. Testing a Complex
- 10. The Software
- 11. Examples
- Conclusion
- Acknowledgements
- References

In a previous article, the authors described an algorithm to determine whether a finite metric polyhedral complex satisfied various local curvature conditions such as being locally $CAT(0)$. The proof made use of Tarski's theorem about the decidability of first order sentences over the reals in an essential way, and thus it was not immediately applicable to a specific finite complex. In this article, we describe an algorithm restricted to 3-dimensional complexes which uses only elementary 3-dimensional geometry. After describing the procedure, we include several examples involving Euclidean tetrahedra which were run using an implementation of the algorithm in GAP.

1. INTRODUCTION

In this article we describe an algorithm to determine whether or not a finite 3-dimensional M_κ -complex is locally $CAT(\kappa)$. The procedure is based purely on elementary 3-dimensional geometry, and has considerable computational advantages over the algorithm for M_κ -complexes of arbitrary dimension which was described by the authors in [Elder and McCammond 01]. After describing the procedure, we include several examples involving Euclidean tetrahedra which were run using an implementation of the algorithm in GAP. The implementation is available from the authors upon request.

The article is structured as follows: Section 2 contains a brief review of polyhedral and comparison geometry which is included for completeness, but can easily be skipped by readers familiar with the area. In Section 3, we review the notion of a gallery as introduced in [Elder and McCammond 01]. A gallery in our sense is a generalization of the notion of gallery used in the study of Coxeter complexes.

Sections 4–7 contain the core of the argument. In these sections we describe how to detect a closed geodesic of length less than 2π in each of the four types of

2000 AMS Subject Classification: Primary 20F65; Secondary 20F67

Keywords: non-positive curvature, $CAT(0)$

circular galleries which can occur in dimension 3: annular, Möbius, disc and necklace. In Section 8 we describe a technique for enumerating the precise finite list of circular galleries which can be built from a specified list of building blocks and contain a closed geodesic of length less than 2π . We then show in Section 9 how to apply this list of forbidden galleries to test whether a particular finite complex is locally $\text{CAT}(\kappa)$.

In Section 10 we describe the GAP routines mentioned above; in Section 11 we give three examples of the results this software has produced, and in Section 12 we describe some directions for future research.

2. PRELIMINARIES

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 in which we are primarily interested, will be cell complexes constructed out of convex polyhedral cells in \mathbb{H}^n , \mathbb{E}^n or \mathbb{S}^n . We will start by describing these spaces, then quickly review their relationship with comparison geometry. Details can be found in [Ballman 95], [Bridson 91] and [Bridson, Haefliger 99].

Definition 2.1. (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 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. For 0-cells, $\sigma^\circ = \sigma$, by definition and $\partial\sigma$ is empty.

Definition 2.2. (M_κ -complexes.) An \mathbb{H} -complex [E-complex, 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 2.3. (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 2.4. (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 2.5. (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 .

Definition 2.6. (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 the 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 [Bridson, Haefliger 99, Theorem 7.19].

Definition 2.7. (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 2.8. (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 2.9. (Local geodesics.) A piecewise geodesic γ in an M_κ -complex 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 $\text{LINK}(x, K)$. Similarly, a *closed geodesic* is a loop which is a local geodesic at each point.

The size and length of local geodesics are closely related. Bridson proves that for every $\ell > 0$ there exists 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 [Bridson 91, Theorem 1.11] [Bridson, Haefliger 99, Theorem I.7.28]. We can in fact compute such an N directly. See Section 8.

Metric polyhedral complexes are particularly useful in the creation of metric spaces of non-positive curvature.

Definition 2.10. (Globally CAT(κ)). 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 CAT(κ)*, 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 CAT(κ)* if every point in K has a neighborhood which is globally CAT(κ). Locally CAT(0) spaces are often referred to as *non-positively curved*.

Theorem 2.11. *In a globally 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 CAT(κ) can be reduced to local properties and how local properties can be reduced to the existence of geodesics in \mathbb{S} -complexes.

Theorem 2.12. *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 2.13. *If K is an M_κ -complex, then K is locally CAT(κ) if and only if the link of each vertex in K is globally CAT(1) if and only if the link of each cell of K is an \mathbb{S} -complex which contains no closed geodesic 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.

3. GALLERIES

As a consequence of Theorem 2.13, the main goal of the algorithm will be to detect short closed geodesics in the link of a cell. In a 3-dimensional complex the link of a 3-cell is empty, the link of a 2-cell is a discrete set of points (and thus never contains a short closed geodesic), and the link of a 1-cell is a finite metric graph which is easy to test for short closed loops. Thus, the entire focus will be on the link of a 0-cell which is a 2-dimensional spherical complex.

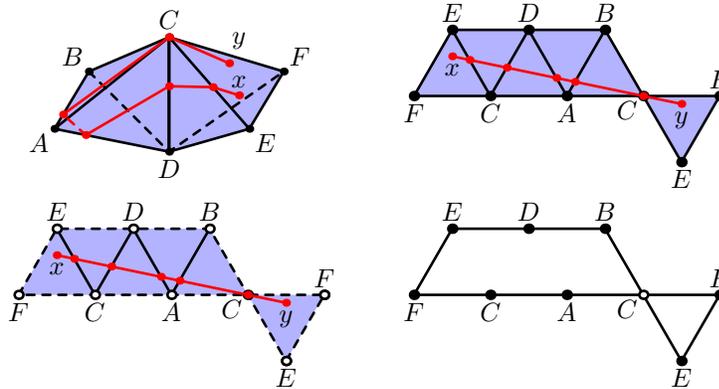


FIGURE 1. The 2-complex and linear gallery described in Example 3.1.

Let K be a 2-dimensional \mathbb{S} -complex, let γ be a local geodesic in K , and consider the sequence of open cells through which γ passes (we will assume throughout that K has been subdivided so that all of its cells are proper spherical simplices). This sequence is essentially what we call a *linear gallery*. If γ is instead a closed geodesic, then this sequence is given a cyclic ordering and the sequence is called a *circular gallery*. Before giving a precise definition, we begin with an example.

Example 3.1. Let K be the 2-dimensional \mathbb{S} -complex formed by attaching the boundaries of two regular spherical tetrahedra along a 1-cell. The complex K is shown in the upper left corner of Figure 1 (where the spherical nature of the 2-cells has been left to the reader’s imagination). 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 linear gallery determined by γ is shown in the upper right corner, its interior in the lower left corner, and its boundary in the lower right of Figure 1.

We will now give a semi-formal definition of a gallery. For precise definitions and proofs of the properties listed see [Elder and McCammond 01].

Definition 3.2. (Linear gallery determined by a geodesic.) Let \mathcal{S} be a finite set of proper spherical simplices, and assume that if σ is a shape in \mathcal{S} then every face of σ is also in \mathcal{S} . As noted in the previous section, σ° will denote the open cell corresponding to a closed cell $\sigma \in \mathcal{S}$.

Let K be an \mathbb{S} -complex such that each simplex in K is isometric to a cell in \mathcal{S} . If γ is a local geodesic in K , then the sequence $\{\sigma_i^\circ\}_{i=1}^k$ of open cells through which γ passes, corresponds to a sequence $\{\sigma_i\}_{i=1}^k$ of closed

cells in \mathcal{S} . The cells σ_1 and σ_k are called the *start cell* and *end cell* \mathcal{G} , respectively. Moreover, because γ is a local geodesic, it is easy to see that the dimensions of consecutive cells in this sequence are distinct and that these dimensions alternate, going up and down as the sequence progresses. If the dimension of a cell σ_i is larger than that of σ_{i-1} and σ_{i+1} (if these exist), then σ_i is called a *top cell*. If its dimension is smaller than that of its neighbors, then it is a *bottom cell*. Notice also that each bottom cell can be identified with a particular cell in the boundary of each neighboring top cell in a unique way which is determined by the path γ .

The *linear gallery determined by γ* is technically the complex \mathcal{G} which results from taking the cells σ_i , $i = 1, \dots, k$ and gluing the bottom cells into the boundaries of the neighboring top cells according to these identifications. Figure 1 should help to make this rough description precise.

Notice that even though cells in K may be traversed more than once by different portions of γ (such as triangle CEF in the example), the corresponding cells in the gallery are not identified. The *interior of \mathcal{G}* is the union of the open cells σ_i° , while the *boundary of \mathcal{G}* is the complement of the interior in \mathcal{G} . We note two properties of linear galleries which are essentially immediate from the definition.

Lemma 3.3. *If K is an \mathbb{S} -complex and \mathcal{G} is a linear gallery determined by a local geodesic γ in K , then the interior of \mathcal{G} immerses into K and retracts onto the lift of γ to \mathcal{G} .*

Note that \mathcal{G} itself may not immerse into K . In the example, the map $\mathcal{G} \rightarrow K$ is not a local embedding at the vertex labeled A .

Definition 3.4. (Circular gallery determined by a geodesic.) Let K be an \mathbb{S} -complex built out of the shapes in \mathcal{S} , and let γ be a closed geodesic in K . There is a *circular gallery determined by γ* whose definition is essentially the same as that of a linear gallery, only the cells are given a cyclic ordering. In dimension 2, typical circular galleries will be annuli and Möbius strips. The full list of possibilities are given in Definition 3.7.

The definitions of top cell, bottom cell, interior, and boundary are as above. As with linear galleries, we record the following properties.

Lemma 3.5. *If K is an \mathbb{S} -complex and \mathcal{G} is a circular gallery determined by a closed geodesic in K , then the interior of \mathcal{G} immerses into K and retracts onto the loop which is the lift of γ to \mathcal{G} .*

Definition 3.6. (Linear and Circular Galleries.) More generally, a linear gallery, or a circular gallery, is any complex built out of the shapes in \mathcal{S} which has this type of linear/circular ordering to its cells. See [Elder and McCammond 01] for precise details.

We say that a linear gallery \mathcal{G} *contains a geodesic* γ if γ is a local geodesic which starts in the start cell, ends in the end cell, and remains in the interior of \mathcal{G} throughout. In particular, the sequence of open cells through which γ passes must be exactly the sequence which determined \mathcal{G} in the first place. Similarly, a circular gallery \mathcal{G} *contains a closed geodesic* γ if the cyclic sequence of open cells through which γ passes is the same as the sequence which determines \mathcal{G} . As a result, γ will remain in the interior of \mathcal{G} throughout.

A gallery *occurs in a complex* K if there is a cellular map $\mathcal{G} \rightarrow K$ which restricts to an immersion on the interior of \mathcal{G} . Notice that a local geodesic in \mathcal{G} need not be sent to a local geodesic in K under this map. Finally, note that a circular gallery can be viewed as a linear gallery where the start and end cells have been identified by an isometry.

Definition 3.7. (Types of circular galleries.) In the 3-dimensional case there are four distinct types of circular galleries. If all of the bottom cells are edges, then the circular gallery is either an annulus, a Möbius strip, or a disc, and we refer to these as *annular galleries*, *Möbius galleries*, and *disc galleries*, respectively. A disc gallery is created precisely when all of the cells in the gallery share a common 0-cell.

On the other hand, if at least one of the bottom cells is a vertex, then the gallery can be split uniquely into a

sequence of linear galleries whose start and end cells are vertices and which have no bottom cells which are in-between vertices. In this case, the gallery is a “necklace” of these *vertex-to-vertex* pieces, and we will refer to it as a *necklace gallery*.

The next four sections will examine circular galleries of each type in turn. The goal will be to determine which circular galleries could have been determined by a closed geodesic (of length less than 2π) in some complex K . Equivalently, we will seek to determine which circular galleries \mathcal{G} themselves contain such a closed geodesic (in the sense of Definition 3.6).

4. ANNULAR GALLERIES

If \mathcal{G} is an annular gallery, then by definition all of its top cells are 2-cells and all of its bottom cells are 1-cells. Let $\{f_1, e_1, f_2, e_2, \dots, f_n, e_n\}$ denote the sequence of cells in \mathcal{G} where the e_i are the bottom cells and the f_i are the top cells. In order to determine whether there exists a closed geodesic in \mathcal{G} of length less than 2π , we will “cut open” \mathcal{G} and then develop the resulting linear gallery onto the 2-sphere. In particular, let \mathcal{G}' be the linear gallery whose sequence of cells is $\{f_1, e_1, f_2, e_2, \dots, f_n, e_n, f'_1\}$, where f'_1 is simply another copy of f_1 . The idea is that \mathcal{G}' is a linear gallery such that identifying f_1 and f'_1 by an isometry should convert \mathcal{G}' into the circular gallery \mathcal{G} . We say that \mathcal{G}' is obtained by *cutting open* \mathcal{G} .

Next, we define a map ϕ from \mathcal{G}' to \mathbb{S}^2 cell by cell. To start, pick an arbitrary isometric embedding $f_1 \rightarrow \mathbb{S}^2$. This restricts to a map from $e_1 \rightarrow \mathbb{S}^2$. Then, we define an isometric embedding $f_2 \rightarrow \mathbb{S}^2$ which extends the map $e_1 \rightarrow \mathbb{S}^2$ and which is an immersion on the interior of the linear gallery $\{f_1, e_1, f_2\}$. Note that there is only one such map. Continuing in this way, we define a map $\phi : \mathcal{G}' \rightarrow \mathbb{S}^2$ which is an immersion on the interior of \mathcal{G}' and which is uniquely determined once the initial map $f_1 \rightarrow \mathbb{S}^2$ has been chosen. In other words, it is unique up to an isometry of \mathbb{S}^2 . This unique map ϕ is called the *developing map for \mathcal{G}'* , and the following is one of its key properties.

Lemma 4.1. *Let \mathcal{G} is an annular gallery, let \mathcal{G}' be a linear gallery obtained by cutting open \mathcal{G} , and let $\phi : \mathcal{G}' \rightarrow \mathbb{S}^2$ be the developing map. If \mathcal{G} contains a closed geodesic γ whose lift to \mathcal{G}' is denoted γ' , then $\phi(\gamma')$ is an arc lying in a great circle of \mathbb{S}^2 . Moreover, the unique isometry of \mathbb{S}^2 which takes the image of the start cell of \mathcal{G}' to the image of the end cell (according to the way they are identified*

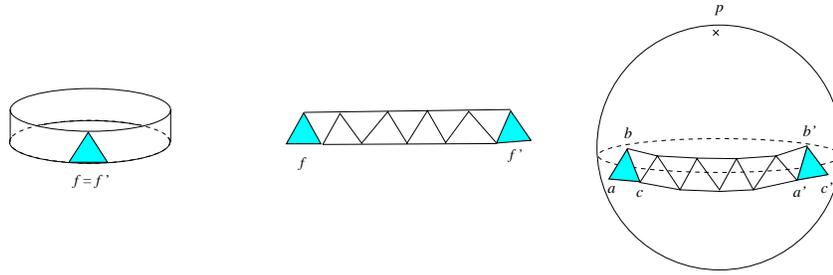


FIGURE 2. An annular gallery, cut open and developed.

to produce \mathcal{G}) will be a rotation around the line through the origin perpendicular to the plane containing this great circle.

Proof: Using the definitions, it is easy to see that the image of γ' must be a local geodesic in \mathbb{S}^2 . For the second assertion, note that the geodesic passes through the start and end cells of \mathcal{G}' in the exact same way, since under the isometric gluing to produce the circular gallery \mathcal{G} , it forms a closed geodesic. Thus, in \mathbb{S}^2 , the images of the start and end cells lie across the great circle containing this geodesic in the same way. A rotation of \mathbb{S}^2 around the pole corresponding to this great circle is clearly an isometry of \mathbb{S}^2 which takes $\phi(f_1)$ to $\phi(f'_1)$ in the correct way, and because these cells are non-degenerate proper spherical triangles, this is the unique isometry of \mathbb{S}^2 with this property. \square

If γ' and \mathcal{G}' are as described above, then there are exactly two position vectors in \mathbb{S}^2 which are perpendicular to the great circle containing $\phi(\gamma')$. The pole of γ' is selected from these two options by orienting γ' (and its image under ϕ) from start cell to end cell and then choosing the pole by the right-hand rule.

We are now ready to describe a procedure to determine whether \mathcal{G} contains a closed geodesic γ of length less than 2π . Continuing with the notation established above, let $a, b,$ and c be the position vectors in \mathbb{S}^2 of the 0-cells of f_1 under ϕ and let a', b' and c' be the corresponding position vectors of the images of the 0-cells of f'_1 . See Figure 2 for an illustration.

Step 1: By Lemma 4.1, if \mathcal{G} contains a closed geodesic in its interior, then there is a position vector p on \mathbb{S}^2 such that the angle between p and a is the same as the angle between p and a' . The set of such points is $\{x \in \mathbb{S}^2 : x \cdot a = x \cdot a'\} = \{x \in \mathbb{S}^2 : x \cdot (a - a') = 0\}$. Similar restrictions hold for b, b' and c, c' . Thus the pole, if it exists, will be orthogonal to the span of the vectors $(a - a'), (b - b')$ and

$(c - c')$. The possible dimensions of $\text{SPAN}\{(a - a'), (b - b'), (c - c')\}$ are 0, 1, 2, or 3.

Step 2: If the dimension is 3, then the pole p cannot exist and \mathcal{G} contains no closed geodesics. If the dimension is 2, then there is a unique line through the origin which is perpendicular to this span and thus a unique great circle which could contain the image of this hypothetical geodesic. We will examine this case in Step 3. If the dimension is 1, then there is a unique plane through the origin which is perpendicular to the vectors $a - a', b - b'$ and $c - c'$. Also, since a and a' lie on \mathbb{S}^2 , this plane must bisect the line segment connecting a to a' and similarly for the other two pairs. This shows that a reflection through this plane is an orientation-reversing isometry of \mathbb{S}^2 which takes $\phi(f_1)$ to $\phi(f'_1)$, contradicting Lemma 4.1. Thus, there is no closed geodesic in \mathcal{G} in this case. Finally, if the dimension is 0, then $a = a', b = b',$ and $c = c'$. If \mathcal{G} contains a closed geodesic with these properties then $\phi(\gamma')$ will travel around the 2-sphere at least once, and the length of the closed geodesic in \mathcal{G} cannot be less than 2π . This shows that a 2-dimensional span is the only case of interest.

Step 3: When the span is 2-dimensional, it is easy to calculate the position vector $p \in \mathbb{S}^2$ which is the pole of the great circle containing $\phi(\gamma')$, if it exists. In particular, p is a scalar multiple of $(a - a') \times (b - b')$ or $(a - a') \times (c - c')$, whichever is nonzero. If this great circle passes through the image of the interior of \mathcal{G}' , then, working backwards, we can create a geodesic in \mathcal{G}' which glues together to form a closed geodesic in \mathcal{G} . Conversely, if this great circle passes through the image of the boundary of \mathcal{G}' then \mathcal{G} does not contain a closed geodesic, since the only possible pole for its image has been eliminated. Notice that the length condition is also easy to check by looking at the image of \mathcal{G}' under the developing map.

Step 4: To check whether the image of the boundary of \mathcal{G}' intersects the great circle, it is sufficient to check that for every bottom cell e_i of \mathcal{G}' , the vertices of this

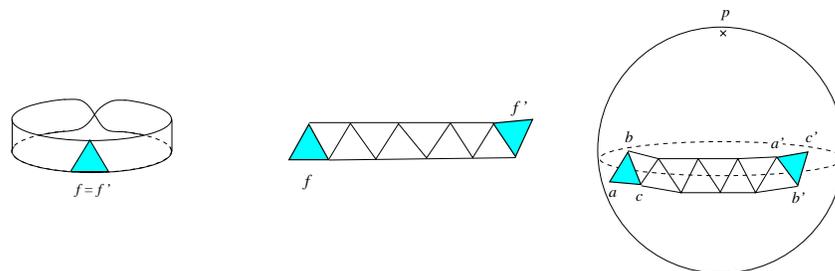


FIGURE 3. A Möbius gallery, cut open and developed.

edge are sent to different sides of the great circle under consideration. In other words, if the vertices of e_i are v and v' , we simply check whether $p \cdot \phi(v)$ and $p \cdot \phi(v')$ have opposite signs. Dot products equal to 0 are not allowed since this would indicate that the supposed geodesic contains a boundary point of \mathcal{G}' .

To summarize, testing an annular gallery \mathcal{G} for a closed geodesic of length less than 2π involves developing the cut-open linear gallery \mathcal{G}' onto \mathbb{S}^2 , and then checking the dimension of the span of $a - a'$, $b - b'$ and $c - c'$. If this dimension is 2, we calculate the only possible pole of the image of the hypothetical geodesic, and check whether this possibility is feasible or not. Otherwise we conclude that \mathcal{G} cannot contain such a closed geodesic. Notice that each of these steps involves only elementary computations.

5. MÖBIUS GALLERIES

If \mathcal{G} is a Möbius gallery, then by definition all of its top cells are 2-cells and all of its bottom cells are 1-cells. Let $\{f_1, e_1, f_2, e_2, \dots, f_n, e_n\}$ denote the sequence of cells in \mathcal{G} where the e_i are the bottom cells and the f_i are the top cells. As in the annular case, we define a linear gallery \mathcal{G}' whose cell sequence is $\{f_1, e_1, f_2, e_2, \dots, f_n, e_n, f'_1\}$, where f'_1 is another copy of f_1 , and we define a developing map $\phi : \mathcal{G}' \rightarrow \mathbb{S}^2$. The only difference is that, in this case, because identifying f_1 with f'_1 produces a Möbius strip, the triangles $\phi(f_1)$ and $\phi(f'_1)$ will occur in \mathbb{S}^2 with opposite orientations. As before, the developing map is unique up to an isometry of \mathbb{S}^2 . Its key properties are summarized in the following lemma.

Lemma 5.1. *Let \mathcal{G} be an Möbius gallery, let \mathcal{G}' be a linear gallery obtained by cutting open \mathcal{G} , and let $\phi : \mathcal{G}' \rightarrow \mathbb{S}^2$ be the developing map. If \mathcal{G} contains a closed geodesic γ whose lift to \mathcal{G}' is denoted γ' , then $\phi(\gamma')$ is an arc lying in a great circle of \mathbb{S}^2 . Moreover, the unique isometry of \mathbb{S}^2 which takes the image of the start cell of \mathcal{G}' to the image*

of the end cell (according to the way they are identified to produce \mathcal{G}) will be a rotation around the line through the origin perpendicular to the plane containing this great circle followed by a reflection through this plane.

Proof: Using the definitions, it is easy to see that the image of γ' must be a local geodesic in \mathbb{S}^2 . For the second assertion, note that the geodesic passes through the start and end cells of \mathcal{G}' in the exact same way, since under the isometric gluing to produce the circular gallery \mathcal{G} , it forms a closed geodesic. Thus, in \mathbb{S}^2 , the images of the start and end cells lie across the great circle containing this geodesic in the same way, but flipped. A rotation of \mathbb{S}^2 around the pole corresponding to this great circle followed by a reflection through the plane containing the great circle is clearly an isometry of \mathbb{S}^2 which takes $\phi(f_1)$ to $\phi(f'_1)$ in the correct way, and because these cells are non-degenerate proper spherical triangles, this is the unique isometry of \mathbb{S}^2 with this property. \square

We are now ready to describe a procedure to determine whether \mathcal{G} contains a closed geodesic γ of length less than 2π . The pole of γ and the position vectors a, b, c, a', b' , and c' are defined as in the annular case. See Figure 3 for an illustration.

Step 1: By Lemma 5.1, if \mathcal{G} contains a closed geodesic in its interior, then there is a position vector p on \mathbb{S}^2 such that the angle between p and a equals the angle between $-p$ and a' , and similarly for the other two pairs b, b' and c, c' . One way to see this is to think of the great circle as the equator and to realize that a point a at some latitude in the northern hemisphere will be taken by the isometry of \mathbb{S}^2 described in the lemma to a point in the southern hemisphere with the same latitude. Thus p , if it exists, must be in the set $\{x \in \mathbb{S}^2 : x \cdot a = -x \cdot a'\} = \{x \in \mathbb{S}^2 : x \cdot (a + a') = 0\}$. In other words, p will be orthogonal to the span of the vectors $a + a'$, $b + b'$ and $c + c'$. The possible dimensions of $\text{SPAN}\{(a + a'), (b + b'), (c + c')\}$ are 0, 1, 2, or 3.

Step 2: If the dimension is 3, then the pole p cannot exist and \mathcal{G} contains no closed geodesics. If the dimension is 2, then there is a unique line through the origin which is perpendicular to this span and thus a unique great circle which could contain the image of this hypothetical geodesic. We will examine this case in step 3A. If the dimension is 1, then there is a unique plane through the origin which is perpendicular to the vectors $a + a'$, $b + b'$ and $c + c'$. If \mathbb{S}^2 is oriented so that this plane contains the equator, then the fact that a and a' are unit vectors whose sum is vertical, implies that they lie at the same latitude, and their longitudes differ by exact π radians. Similar arguments apply to the pairs b, b' and c, c' . Thus a rotation through π radians around the line perpendicular to this plane is an orientation-preserving isometry of \mathbb{S}^2 which takes $\phi(f_1)$ to $\phi(f'_1)$, contradicting Lemma 5.1. Thus, there is no closed geodesic in \mathcal{G} in this case. Finally, if the dimension is 0, then $a = -a'$, $b = -b'$, and $c = -c'$ and the triangle $\phi(f_1)$ is antipodal to the triangle $\phi(f'_1)$. This case will be examined in Step 3B.

Step 3A: When the span is 2-dimensional, it is easy to calculate the position vector $p \in \mathbb{S}^2$ which is the pole of the great circle containing $\phi(\gamma')$, if it exists. In particular, p is a scalar multiple of $(a+a') \times (b+b')$ or $(a+a') \times (c+c')$, whichever is nonzero. Once the only possible pole, p , has been found, we check whether this determines a feasible geodesic as in the annular case. See Step 4 in Section 4.

Step 3B: When the span is 0-dimensional, the pairs a, a', b, b' and c, c' are antipodal and the unique isometry of \mathbb{S}^2 taking $\phi(f_1)$ to $\phi(f'_1)$ is the antipodal map. The main difficulty in this case is that when a solution exists, it is underdetermined. Concretely, the isometry sending the start to the end cell does not help narrow the search for poles of geodesics in \mathcal{G} , so we must take a different approach. The approach we adopt is to determine directly the set of all poles of all possible geodesics contained in the interior of \mathcal{G}' . The vertices of \mathcal{G}' can be divided into two classes which we will call *upper and lower vertices*. Let e_1 be the first bottom cell of \mathcal{G}' and arbitrarily call one of its vertices an upper vertex and the other one a lower vertex. The next bottom cell e_2 shares a vertex in common with e_1 since f_2 is a triangle. Label the unlabeled vertex so that e_2 has one upper and one lower vertex as well. Continuing in this way through all of the bottom cells of \mathcal{G}' labels all but two of its vertices as upper or lower. The two remaining are in the triangles f_1 and f'_1 . Label these according to the opposite of the label of the vertex in the other triangle with whichn it gets identified when forming the circular gallery \mathcal{G} . This

switch in label is due to the twist in the Möbius strip. In order for p to be the pole of a great circle which contains the image of a geodesic contained in \mathcal{G}' , p must satisfy the equations $p \cdot \phi(v) > 0$ for each upper vertex and $p \cdot \phi(v) < 0$ for each lower vertex in \mathcal{G}' . Conversely, each point p which satisfies these conditions will determine a closed geodesic in \mathcal{G} since the antipodal nature of $\phi(f_1)$ and $\phi(f'_1)$ ensures that the great circle corresponding to p passes through the two triangles in the same way. As a result, the preimage of this great circle in \mathcal{G}' will form a closed geodesic in \mathcal{G} under the identification map.

Another way to see this is to consider the double cover \mathcal{G}^2 of the original circular gallery \mathcal{G} . In general, developing maps from circular galleries to \mathbb{S}^2 do not exist because the final 2-cell is placed in the wrong position to complete the cycle, but in this case, there is a well-defined developing map from \mathcal{G}^2 to \mathbb{S}^2 which is an immersion on the interior of this annulus. We illustrate this in Figure 4. Notice that a hypothetical geodesic in \mathcal{G} will lift to a geodesic twice as long in \mathcal{G}^2 and immerse to a path which wraps an integral number of times around a great circle in \mathbb{S}^2 . We will only be interested in the cases where the geodesic wraps around exactly once, since this is the only case where the original geodesic has length less than 2π . In this situation, we can identify the geodesic with the great circle to which it is sent.

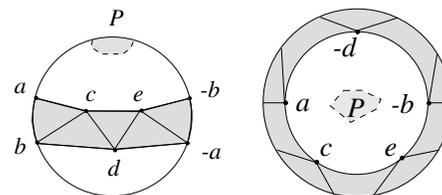


FIGURE 4. Side view and top view of an antipodal situation.

For each bottom cell in \mathcal{G}^2 , the hypothetical geodesic must separate its endpoints in \mathbb{S}^2 . More specifically, one boundary cycle of the annulus must be sent to one side of the great circle and the other boundary cycle must be sent to the other. A position vector p is a pole of a great circle which satisfies these conditions if $p \cdot \phi(v) > 0$ for all v in one boundary cycle of \mathcal{G}^2 . This system is equivalent to the previous system since roughly half of the vertices in this boundary cycle are sent to points antipodal to the images of the lower vertices of \mathcal{G}' . See Figure 4.

In either formulation this linear system is easy to solve, and the solution set is either empty, or an open spherical cell. If we orient \mathbb{S}^2 so that the image of some closed geodesic in \mathcal{G}^2 is sent to the equator, then the solution

set of P will be found hovering around one of the poles. Thus, we call this the *polar region*. See Figure 4 for an illustration. For each $p \in P$, the 2-sphere can be tilted so that p is the north pole and the equator corresponds to a closed geodesic of length 2π in the interior of \mathcal{G}^2 , and thus to a closed geodesic of length π in the interior of \mathcal{G} .

To summarize, testing a Möbius gallery \mathcal{G} for a closed geodesic of length less than 2π involves developing the cut-open linear gallery \mathcal{G}' onto \mathbb{S}^2 , and then checking the dimension of the span of $a + a'$, $b + b'$ and $c + c'$. If this dimension is 2, we calculate the only possible pole of the image of the hypothetical geodesic, and check whether this possibility is feasible or not. If this dimension is 0, we calculate a feasibility region for the pole and check whether this is empty or not. In all other cases, we conclude that \mathcal{G} cannot contain such a closed geodesic. As in the annular case, each of these steps involves only elementary computations.

6. DISC GALLERIES

If \mathcal{G} is a disc gallery, then by definition all of its bottom cells are edges which share a common vertex, v . In this short section we show that if a disc gallery contains a short closed geodesic then the link of v is not CAT(1).

Let \mathcal{G} be a disc gallery with common vertex v , let \mathcal{G}' be a linear gallery obtained by cutting open \mathcal{G} , and let $\phi : \mathcal{G}' \rightarrow \mathcal{S}^2$ be the developing map. It is easy to see that the unique isometry of \mathcal{S}^2 which takes the image of the start cell of \mathcal{G}' to the image of the end cell of \mathcal{G}' (according to the way they are identified to form \mathcal{G}) must be a rotation of \mathcal{S}^2 about the origin and through v . As a result, the only possible geodesic in \mathcal{G} will lift to a path in \mathcal{G}' , which is sent to a path in the 2-sphere and the image of this lifted path will be entirely contained in the great circle perpendicular to v . When this path in the great circle stays inside the image of \mathcal{G}' , the arc length of the path inside a particular spherical triangle is the same as the angle of this triangle at v . Thus a closed geodesic of length ℓ in \mathcal{G} immediately leads to a closed geodesic of length ℓ in the link of v . In particular, a 2-dimensional \mathcal{S} -complex which contains a disc gallery with a short closed geodesic will not even be locally CAT(1). And since the complete process of checking curvature conditions on a complex will be done inductively by dimension, disc galleries will not need to be examined at all.

7. NECKLACE GALLERIES

If \mathcal{G} is a necklace gallery, then by definition at least one of its bottom cells is a vertex. Let $\mathcal{G}_1, \dots, \mathcal{G}_k$ be the

unique linear vertex-to-vertex galleries into which \mathcal{G} can be decomposed. Instead of analyzing whether \mathcal{G} contains a closed geodesic, it is sufficient to analyze whether each \mathcal{G}_i contains a geodesic from start vertex to end vertex. If all of the \mathcal{G}_i contain a geodesic, then these geodesic pieces can be concatenated to produce a closed geodesic in \mathcal{G} . Whether this closed geodesic in \mathcal{G} remains geodesic when it is included in a 2-dimensional \mathcal{S} -complex is a different issue entirely, and one which will be addressed in Section 9.

Vertex-to-vertex galleries come in two types: either \mathcal{G}_i has a top cell which is an edge, or all of its top cells are 2-cells and all of its bottom cells (excluding the start and end cells) are edges. In the former situation, \mathcal{G}_i is a linear gallery whose sequence of cells is $\{v, e, v'\}$ and it is clear that this gallery contains a geodesic from v to v' , namely, the edge e itself. In the latter case, \mathcal{G}_i is a linear gallery whose sequence of cells is $\{v, f_1, e_1, \dots, e_{n-1}, f_n, v'\}$ for some $n \geq 1$. As in the previous sections, there is a developing map $\phi : \mathcal{G}_i \rightarrow \mathbb{S}^2$ which starts with an isometric embedding $f_1 \rightarrow \mathbb{S}^2$ and proceeds through the gallery, all the while requiring that the map be an immersion on the interior of the gallery. Such a map exists and it is unique up to an isometry of \mathbb{S}^2 . The key property of this developing map is the following.

Lemma 7.1. *Let \mathcal{G}_i be a vertex-to-vertex gallery which contains a 2-cell and let $\phi : \mathcal{G}_i \rightarrow \mathbb{S}^2$ be the developing map. If \mathcal{G}_i contains a geodesic γ in its interior from start vertex to end vertex, then $\phi(\gamma)$ is an arc lying in a great circle of \mathbb{S}^2 .*

Proof: Once again, it is easy to see that the image of γ must be a local geodesic in \mathbb{S}^2 . \square

We are now ready to describe a procedure to determine whether \mathcal{G}_i contains a geodesic. The pole of γ will be defined as in the annular case. Let a and a' denote the position vectors of the images of the start and end vertices, respectively. See Figure 5 for an illustration.

Step 1: By Lemma 7.1, if \mathcal{G}_i contains a geodesic from start cell to end cell, then there is a position vector p on \mathbb{S} such that p is the pole of a great circle containing a and a' . In other words, p is orthogonal to the span of a and a' . The possible dimensions of $\text{SPAN}\{a, a'\}$ are 1 and 2.

Step 2: If this span is 2-dimensional, then there is a unique line through the origin which is perpendicular to this span and thus a unique great circle which could contain the image of this hypothetical geodesic. We will examine this case in Step 3A. If this span is 1-dimensional,

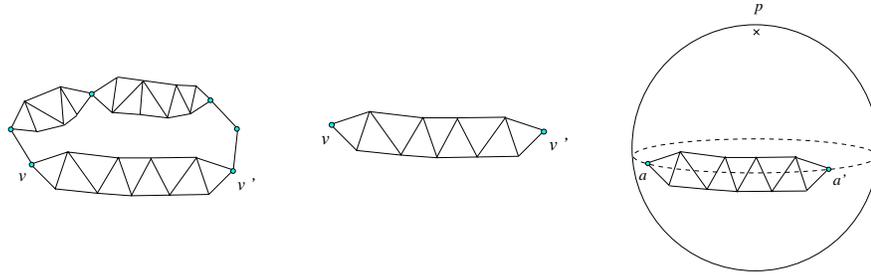


FIGURE 5. A developed vertex-to-vertex gallery.

then $a = \pm a'$. If $a = a'$ then any geodesic in \mathcal{G}_i will travel around the 2-sphere at least once and the length of the geodesic cannot be less than 2π . The case where $a = -a'$ will be examined in Step 3B.

Step 3A: When the span is 2-dimensional, the only possible pole is a scalar multiple of $a \times a'$ and we check whether this possibility is feasible as in the previous cases. For later use, we note that it is easy to calculate the unit tangent vector at a which represents the direction with which the geodesic between a and a' leaves a , and the unit tangent vector at a' which represents the direction with which the geodesic between a and a' approaches a' .

Step 3B: When a and a' are antipodal, we proceed as in the antipodal Möbius case. More specifically, we label the vertices of \mathcal{G}_i as upper or lower as in that case, starting with e_1 and continuing through all of the bottom cells of \mathcal{G}_i . The start and end vertices are neither upper nor lower since their images are supposed to lie on the great circle containing the hypothetical geodesic. The polar region is defined as in the Möbius case except that we can add the restriction that $p \cdot a = p \cdot a' = 0$. This restricts our polar region from the start to the great circle perpendicular to a and a' . If this 1-dimensional polar region is non-empty it will be an open interval in this great circle. If we identify this great circle with the set of unit tangent vectors at a , then the elements in this region can be thought of as the directions it is possible to leave a , remain in the interior of (the image of) the gallery \mathcal{G}_i and arrive safely at a' . Notice also that the length of any such geodesic will always be π since a and a' are antipodal.

Step 4: The output for each vertex-to-vertex gallery will include the length of the geodesic, and the incoming and outgoing angle of the geodesic (or range of angles for the antipodal case). The first will be used immediately; the second will be used in Section 9.

The necklace gallery \mathcal{G} will contain a closed geodesic of length less than 2π if and only if each of its vertex-to-vertex pieces contain a geodesic and the sum of the lengths of these geodesics is less than 2π . As described above, each of these steps involves only elementary computations.

8. ENUMERATING GALLERIES

Let \mathcal{S} be a finite set of proper spherical simplices (of dimension at most 2), and assume that if σ is a shape in \mathcal{S} then every face of σ is also in \mathcal{S} . In this section we show how to construct a finite list which consists precisely of the circular galleries built out of the shapes in \mathcal{S} which contain closed geodesics of length less than 2π and which can occur in locally CAT(1) complexes built out of the shapes in \mathcal{S} . The latter restriction merely rules out the need to examine disc galleries. The steps in the procedure will be interspersed with comments.

Step 1: Since \mathcal{S} is a finite set, there exists an n such that each angle in each spherical triangle in \mathcal{S} has radian measure at least π/n . Moreover, it is easy to calculate such an n explicitly given the triangles.

Let \mathcal{G} be a circular gallery which is not a disc gallery. If the sequence of cells for \mathcal{G} contains a subsequence of the form $\{e_1, f_2, e_2, \dots, f_k, e_k\}$ where the f_i are triangles which are top cells, the e_i are edges which are bottom cells, and all of the e_i share a common vertex v in \mathcal{G} , then this portion of \mathcal{G} is said to be *spinning around* v . Notice that if $k > n$ then when this linear portion of \mathcal{G} is developed on \mathbb{S}^2 (with developing map ϕ) there is no local geodesic in \mathbb{S}^2 which starts in the interior of $\phi(e_1)$, passes through the interiors of $\phi(e_i)$ and $\phi(f_i)$ in order and then ends in the interior of $\phi(e_k)$. This is because the cells involved are proper and the angle at $\phi(v)$ between $\phi(e_1)$ and $\phi(e_k)$ is greater than π . Thus, we only need to consider circular galleries which do not spend too much time spinning around a single vertex.

Step 2: Let L_1 be a list of linear galleries \mathcal{G} built out of cells in \mathcal{S} whose sequence of cells is of the form $\{e_1, f_2, e_2, f_3, e_3\}$, where the f_i are triangles which are top cells, the e_i are edges which are bottom cells, and \mathcal{G} does *not* spin around a vertex. The list L_1 is clearly finite since \mathcal{S} is finite. For each \mathcal{G} in L_1 the images of e_1 and e_3 under the developing map ϕ are disjoint, and we can calculate the minimum distance between $\phi(e_1)$ and $\phi(e_3)$. In particular, it is easy to show that the arc exhibiting the minimum distance either connects an endpoint of one edge to an endpoint of the other edge, or it connects an endpoint of one edge to a point in the interior of the other edge where the arc itself is perpendicular to the other edge. Each of these eight distances can be calculated explicitly by elementary means. Let δ be the minimum of these minimum distances for the linear galleries in L_1 .

Let \mathcal{G} be a linear gallery which contains a geodesic γ or a circular gallery which contains a closed geodesic γ . If \mathcal{G} contains a linear portion \mathcal{G}' which is in the list L_1 , then the portion of γ which lies in \mathcal{G}' has length at least δ . Thus, if \mathcal{G} contains at least ℓ/δ such linear portions which are disjoint, then the length of γ will be longer than ℓ .

Step 3: Let L_2 be the list of annular and Möbius galleries built out of shapes in \mathcal{S} which do not contain $2\pi/\delta$ disjoint linear portions which occur in the list L_1 . Because of the bound on the number of edges which can spin around a vertex, this list is finite, and it must also contain every annular or Möbius gallery which contains a geodesic of length less than 2π . Using the procedures in the previous sections, we can narrow this list so it contains exactly those annular and Möbius galleries built out of shapes in \mathcal{S} which do in fact contain geodesics of length less than 2π .

Step 4: Let L_3 be the list of vertex-to-vertex galleries built out of shapes in \mathcal{S} which do not contain $2\pi/\delta$ disjoint linear portions which occur in the list L_1 . This list is also finite, and it must contain every vertex-to-vertex gallery which contains a geodesic of length less than 2π . Using the procedure in Section 7 we can narrow this list so it contains exactly the vertex-to-vertex galleries built out of shapes in \mathcal{S} which do in fact contain geodesics of length less than 2π . Moreover, the procedure also outputs the length of the geodesic(s) contained in each such gallery.

Step 5: Using the finite list of vertex-to-vertex galleries produced by Step 4 (and the lengths of their geodesics), we can easily produce a finite list of all necklace

galleries which contain a closed geodesic of length less than 2π .

Since every circular gallery is either annular, Möbius, disc, or necklace, the union of the lists produced by Steps 3 and 5 form a nearly complete list of all circular galleries built out of shapes in \mathcal{S} which contain a closed geodesic of length less than 2π . The only exceptions are those contained in disc galleries and, as mentioned above, this will be ruled out in other ways. This proves that such a list can be produced mechanically, but the procedure outlined above initially produces finite lists which are prohibitively large in actual computations. Thus, in the GAP implementation mentioned in the introduction, a completely different method of enumerating potential galleries has been utilized. The linear galleries are enumerated in a depth-first way by adding triangles onto the end of the linear gallery under consideration. We stop adding triangles as soon as we can show that the gallery is too long to contain a geodesic of length less than 2π . At this point we backtrack by removing a triangle from the end and try to add a different triangle. Since the size and shape of the spherical triangles being added will vary greatly, the number of triangles in a gallery when it becomes too long will vary quite a bit from one linear gallery to another, and as a result, many fewer galleries will be examined during the construction of the final list.

9. TESTING A COMPLEX

Given a finite set of 2-dimensional spherical shapes \mathcal{S} , we have shown how to create the finite list L of annular, Möbius and necklace galleries built out of these shapes which contain closed geodesics of length less than 2π . In this section we show how to use this list to test whether a particular finite, 2-dimensional, locally CAT(1) \mathbb{S} -complex K built out of the shapes in \mathcal{S} , contains a closed geodesic of length less than 2π .

The first thing to note is that if K contains such a closed geodesic, then it determines a circular gallery \mathcal{G} in L , and as a result there is a cellular map from \mathcal{G} to K which is an isometry on each cell and an immersion on the interior of \mathcal{G} . Since the list L is finite and the complex K is finite, there are only a finite number of such maps from $\mathcal{G} \in L$ to K and we can easily determine which ones have these properties. For annular and Möbius gallery in L the mere existence of such a map is sufficient.

Lemma 9.1. *Let K be a 2-dimensional \mathbb{S} -complex and let \mathcal{G} be either an annular or Möbius gallery which contains*

a closed geodesic γ or a vertex-to-vertex gallery which contains a geodesic γ . If there is a cellular map $\phi : \mathcal{G} \rightarrow K$ which is an isometry on each cell and an immersion on the interior of \mathcal{G} , then the image of γ under ϕ is a local geodesic in K .

Proof: If \mathcal{G} is a 1-dimensional vertex-to-vertex gallery, then this result is immediate so assume \mathcal{G} is 2-dimensional. Since K itself is 2-dimensional, $\phi(\gamma)$ will clearly be locally geodesic on the interior of its 2-cells. Let x be a point in \mathcal{G} where γ crosses a bottom cell e . The link of x in \mathcal{G} is a circle of length 2π and the incoming and outgoing unit tangent vectors to γ at x will be antipodal points in this circle since γ is a locally geodesic in \mathcal{G} at x . On the other hand, the link of $\phi(x)$ in K will be a finite graph with two vertices and a number of arcs connecting one vertex to the other, each of length π . The fact that ϕ is an isometry on cells and an immersion on the interior of \mathcal{G} ensures that $\text{LINK}(x, \mathcal{G})$ isometrically embeds in $\text{LINK}(\phi(x), K)$ and that the images of the incoming and outgoing unit tangent vectors remain a distance π apart under this embedding. \square

When \mathcal{G} is a necklace gallery with geodesic γ , the situation is more complicated. By Lemma 9.1, the portion of the geodesic in the interior of each vertex-to-vertex gallery is sent to a local geodesic in K , but the lemma is silent about the places where γ passes through a vertex v of \mathcal{G} .

If \mathcal{G} contains a single closed geodesic of length less than 2π then we merely need to construct the metric graph $\text{LINK}(\phi(v), K)$, find the points in this link corresponding to the incoming and outgoing unit tangent vectors of $\phi(\gamma)$ at $\phi(v)$, and then check whether they are a distance of at least π apart. If this is true for every vertex v in γ , then $\phi(\gamma)$ is a local geodesic in K and conversely.

On the other hand, if \mathcal{G} contains an antipodal vertex-to-vertex gallery then it will contain a continuum of closed geodesics and a different procedure is needed. Since geodesics in antipodal vertex-to-vertex galleries all have length π , \mathcal{G} contains at most one such gallery. Call this gallery \mathcal{G}_1 . If v is a vertex in γ which is not an end cell of \mathcal{G}_1 , then the check can proceed as above. If v is one of the end cells of \mathcal{G}_1 but not both, then we can construct the metric graph $\text{LINK}(\phi(v), K)$ and identify the point corresponding to either the incoming or outgoing unit tangent vector of $\phi(\gamma)$ at $\phi(v)$. For the other point, all we know is that it lies somewhere in an open subinterval of an edge of this graph. Calculating the distance in the graph from the known point to each of the end-

points of the edge containing the unknown point, we can determine precisely the portion of this open subinterval which will make $\phi(\gamma)$ a local geodesic at $\phi(v)$. Once this information has been calculated for each end cell of \mathcal{G}_1 , we simply check whether the two feasible regions have any points in common. Finally, if \mathcal{G} is formed by identifying the end cells of \mathcal{G}_1 to a single vertex v , we construct the metric graph $\text{LINK}(\phi(v), K)$ as above. This time we only know the two open subintervals contained in distinct edges of $\text{LINK}(\phi(v), K)$ which correspond to the incoming and outgoing unit tangent vectors. After calculating the pairwise distances between the endpoints of these edges, we can determine whether there are points in these intervals which are a distance of at least π apart and which arise from the two ends of the same geodesic in \mathcal{G}_1 .

To summarize, given a map from $\mathcal{G} \in L$ to K which is an isometry on each cell and an immersion on the interior of \mathcal{G} , it is possible to check whether any of the short closed geodesics in \mathcal{G} have images which remain local geodesics in K . For annular and Möbius galleries no additional check is needed, while for necklace galleries, the situation at the vertices must be determined. Using this, it is possible to determine whether or not K contains a geodesic of length less than 2π .

10. THE SOFTWARE

As mentioned in the introduction, the algorithms described have been implemented on GAP for Euclidean tetrahedra. In this section we give a brief description of these software routines and the assumptions and limitations involved. In short, the user inputs a list of Euclidean tetrahedra by entering the six edge lengths of each shape, and the program outputs the list of annular, Möbius, and 2-dimensional vertex-to-vertex galleries which contain geodesics of length less than 2π that might arise in a 2-dimensional spherical link of a vertex in a complex built out of Euclidean tetrahedra with these metrics. Such a list is slightly smaller than the list produced by the arguments described in the previous sections. See Remark 10.1 below.

The first and most important restriction is that GAP only uses exact arithmetic: arbitrary precise real numbers are not implemented. We chose GAP for its speed and ease of use and determined that these benefits far outweighed the restriction imposed.

After loading the program, the user inputs a list of 6-tuples of rationals where each 6-tuple represents a Euclidean tetrahedron and the 6 rationals represent the squares of the 6 edge lengths. The set of tetrahedra which can

be described in this way is rather large since it includes every tetrahedron which ever arises in a simplicial decomposition of a rational polytope. The program then calculates all possible ways of orienting these tetrahedra in \mathbb{E}^3 so that one vertex is at the origin, a second is on the positive x -axis, a third is in the first quadrant of the xy -plane and the last has positive z -coordinate. For each orientation it calculates the coordinates of the vertices. Each coordinate lies in a quadratic extension of the rationals, and the field generated by all of these coordinates is an algebraic extension of the rationals of the form

$$\mathbb{F} = \mathbb{Q}(\sqrt{d_1}, \sqrt{d_2}, \dots, \sqrt{d_n})$$

where the integers d_i are distinct primes. All of the calculations made by the software will take place in this field.

When asked for a list of annular, Möbius and vertex-to-vertex galleries which contain short (closed) geodesics, the program starts with a tetrahedron placed in \mathbb{E}^3 with these coordinates and then starts gluing tetrahedra to its faces so that the link of the origin looks like a linear gallery developed on \mathbb{S}^2 . When we say that all of the calculations take place in \mathbb{F} , we primarily mean that all of the coordinates of vertices in tetrahedra which have been developed around the origin in this way will lie in \mathbb{F} . Finally, we should note that the speed of the program is greatly effected by the number of primes involved in the definition of \mathbb{F} since the degree of this extension over the rationals is 2^n where n is the number of primes. In practice, n greater than 3 or 4 is currently impractical.

The reason for developing tetrahedra around the origin instead of triangles on \mathbb{S}^2 is addressed by the following remark.

Remark 10.1. The spherical complexes of greatest interest are, of course, those which arise as a link of a vertex in a complex built out of a specified list of tetrahedra. Let \mathcal{S}_1 be a list of Euclidean tetrahedra and \mathcal{S}_2 be a list of the proper spherical cells which occur as corners of the shapes in \mathcal{S}_1 . The list of circular galleries containing geodesics of length less than 2π which can appear in the link of a vertex in a complex built out of the shapes in \mathcal{S}_1 is often strictly shorter than the list of circular galleries containing geodesics of length less than 2π which can appear in \mathbb{S} -complexes built out of the shapes in \mathcal{S}_2 . A 3-dimensional Euclidean example of this phenomenon is given in Example 11.2 below. The reason is that a pair of spherical triangles can meet along an edge in such a complex if and only if they are corners of tetrahedra where the corresponding faces are isometric.

The current implementation searches for and finds only the smaller list which is, of course, of greater interest to researchers. Part of the reason for this is that the number of things to be checked grows exponentially with the length of the galleries. To illustrate this point, consider the computations described in Example 11.3 below. Even though only 3 tetrahedral shapes are involved, they can be arranged in \mathbb{E}^3 (as described above) in 10 distinct ways. As we are developing a linear gallery there is also a choice at each stage of adding the new tetrahedra onto the left or right face of the previous one. Thus, when trying to construct all linear galleries of size n , one is lead to consider in the neighborhood of $(20)^n$ different galleries where n is often in the mid-teens. This type of exponential behavior is often the rule in problems of this type and the technical name is a *combinatorial explosion*. Strategies for keeping the search tightly focused become crucial in such contexts.

11. EXAMPLES

In this section we will give three examples which illustrate both the procedures described in this article and the type of results which can be expected from the software described in Section 10.

Example 11.1. (Regular tetrahedra.) Our first example will involve complexes built out of regular Euclidean tetrahedra. Notice that any 3-dimensional simplicial complex can be given a metric of this type simply by assigning each simplex a regular Euclidean metric with all edge lengths equal to the same constant. Since the dihedral angle in a regular tetrahedra is slightly less than $\frac{2\pi}{5}$ (actually $\arccos(1/3) \approx .39\pi$), the link of each edge must be a graph which does not contain any cycles of length less than 6. The link of a vertex is a spherical 2-complex made up of equilateral triangles with edge length $\frac{\pi}{3}$.

The results are shown in Figures 6, 7 and 8 where the reader is left to imagine that the triangles pictured are spherical triangles with side length $\pi/3$. In Figure 6, the heavy edges on the left and right of the gallery depicted need to be identified to form an annular gallery. Similarly, in Figure 7 the heavy edges on the left and right are to be identified (with a half-twist) to form a Möbius gallery. In Figure 8, the geodesic connects the vertex on the far left to the vertex on the far right.

As can be seen from the illustrations, there are 10 annular galleries to avoid, 8 Möbius galleries to avoid, and 6 vertex-to-vertex galleries which are 2-dimensional. With the addition of the one vertex-to-vertex gallery

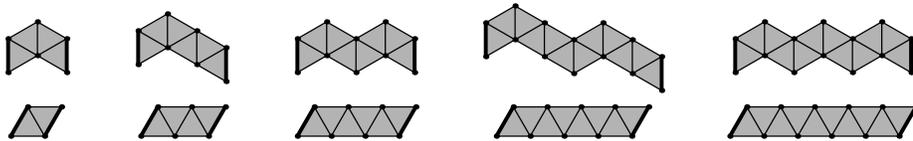


FIGURE 6. Annular galleries in the regular case.

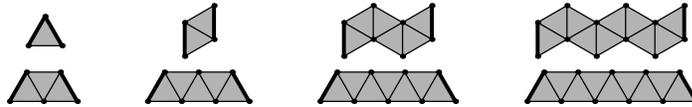


FIGURE 7. Möbius galleries in the regular case.

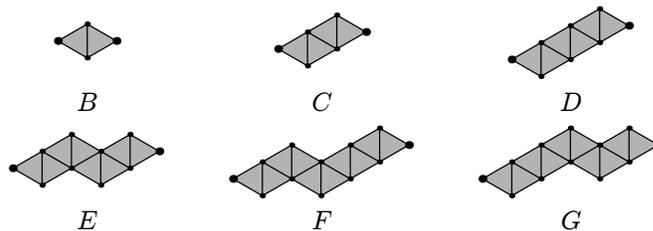


FIGURE 8. Vertex-to-vertex galleries in the regular case.

of the form $\{v, e, v'\}$, which we will denote A , these 7 vertex-to-vertex galleries can be combined in 29 distinct ways to form necklace galleries which must be avoided. Specifically, the 29 necklace galleries which string together vertex-to-vertex galleries in the sequences: $A, A^2, A^3, A^4, A^5, B, BA, BA^2, BA^3, BA^4, B^2, B^2A, B^2A^2, BABA, B^3, C, CA, CA^2, CA^3, CB, CBA, C^2, D, DA, DA^2, DB, E, EA$, and F , form necklace galleries which contain closed geodesics of length less than 2π and these are the only ones.

Example 11.2. (Coxeter shapes.) Our second example involves complexes in which every tetrahedra has a Euclidean metric such that four of its edges have length $\sqrt{3}$ and two non-adjacent edges have length 2. This shape is the isometry type of the Coxeter cell in a Coxeter complex for the \tilde{A}_3 Coxeter group. Complexes of this type have also arisen in the work of Tom Brady and the second author on CAT(0) structures for various Artin groups of finite type. In such a complex, the link of a length 2 edge is a graph where each arc has length $\pi/2$, and the link of a length $\sqrt{3}$ edge is a graph where each arc has length $\pi/3$. Thus these graphs must have no cycles of length less than 4 or 6, respectively. The link of a vertex is a 2-dimensional \mathbb{S} -complex built out of triangles whose angles are $\pi/2, \pi/3$, and $\pi/3$. If a 2-dimensional \mathbb{S} -complex K built out of this triangle arises as a link of a vertex

in a 3-complex C built out of these Euclidean tetrahedra, then the vertices of K can be colored depending on whether they arose from edges of length 2 or $\sqrt{3}$ in C . In the triangle with angles $\pi/2, \pi/3$, and $\pi/3$, the corners with angle $\pi/3$ will be one color and the corner with angle $\pi/2$ will be another. Thus if K is the link of a vertex in C , all of the triangles meeting at a vertex in K will have the same angle incident at that vertex. The list of circular galleries described below will only include those which satisfy this extra restriction.

The list of annular, Möbius, and 2-dimensional vertex-to-vertex galleries which can be built out of this spherical triangle and which contain a geodesic of length less than 2π are shown in Figures 9, 10, and 11. The conventions are the same as in the figures for Example 11.1. Once again, the figures should be viewed as representing spherical triangles. The angles in the triangles which look like $\pi/2$ angles are in fact $\pi/2$ angles, while the angles which look like $\pi/4$ angles represent $\pi/3$ angles. Thus, in the third figure of Figure 11 both sides connecting the specified end cells are actually geodesics. The application of this list to the study of finite-type Artin groups will be developed in [Brady and McCammond 02].

As can be seen from the illustrations, there are 2 annular galleries to avoid, 4 Möbius galleries to avoid, and 3 vertex-to-vertex galleries which are 2-dimensional. This time there are two vertex-to-vertex galleries of



FIGURE 9. Annular galleries in the Coxeter case.



FIGURE 10. Möbius galleries in the Coxeter case.

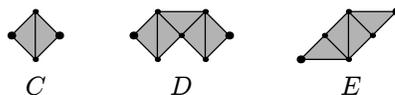


FIGURE 11. Vertex-to-vertex galleries in the Coxeter case.

the form $\{v, e, v'\}$ of length $\arccos(1/\sqrt{3}) \approx .304\pi$ and $\arccos(1/3) \approx .392\pi$. We will denote these A and B , respectively. These 5 vertex-to-vertex galleries can be combined in 26 distinct ways to form necklace galleries which contain a geodesic of length less than 2π and which can occur in a vertex link of a complex built out of tetrahedra with this metric. Specifically, the 26 necklace galleries which string together vertex-to-vertex galleries in the sequences: $A^2, A^4, A^6, A^2B, A^2B^2, A^2B^3, ABAC, A^2C, A^2C^2, A^2D, A^2E, A^4B, CA^4, B, B^2, B^3, B^4, B^5, BE, B^2E, C, C^2, C^3, CD, D$, and E form necklace galleries which contain closed geodesics of length less than 2π , and these are the only ones.

Example 11.3. Our final example is slightly more complicated and is included to show how the software may be of use in the investigation of standard topics in geometric group theory and low dimensional topology. As a special case of the (weak) hyperbolization conjecture that might be more tractable, Thurston has proposed the following:

Conjecture 11.4. (Thurston.) *Let M be a compact triangulated 3-manifold without boundary such that the link of every edge is either a 5-cycle or a 6-cycle and such that any two edges whose links are 5-cycles do not lie in a common 2-cell. Under these hypotheses, the fundamental group of M should be word-hyperbolic in the sense of Gromov.*

The authors, in collaboration with John Meier, have been investigating the following stronger conjecture.

Conjecture 11.5. *Let M be as above. If the edges whose links are 5-cycles are assigned a length of $\sqrt{3}$, the other edges are assigned a length of 2, and the 2-cells and 3-cells are given Euclidean metrics corresponding to these edge lengths, then the resulting metric on M should be*

locally CAT(0). Moreover, since M under this metric will contain no immersed flat planes, this would imply that $\pi_1 M$ is word-hyperbolic.

It is easy to check that each tetrahedra in M is assigned one of three metrics, and that the links of edges in M are circles of length at least 2π . Thus it only remains to investigate the links of vertices in complexes which can be built out of these three types of tetrahedra. Using the software developed, we are currently determining the list of annular and vertex-to-vertex galleries which need to be avoided. Möbius galleries are not being searched since the link of each vertex is a triangulated 2-sphere and Möbius strips cannot immerse into 2-spheres. By current estimates the software will take approximately two months to determine this list running on a 500Mhz Dell PC running Linux. At the time of this writing, the search is 44% complete. In contrast, the lists given in our first two examples took 5 seconds and 3 seconds respectively. The time difference is the result of the combinatorial explosion described in Section 10.

Given this list, which we expect to contain on the order of 100,000 galleries, there are various ways to show that most of these possibilities cannot occur in triangulated 2-spheres such as the ones which will occur as links of vertices in M . If all the items can be eliminated, Conjecture 11.5 will become a theorem. And if not, there will at least be a partial result which states that if M has a triangulation of this type which avoids the galleries remaining in the list, then it will be CAT(0) and $\pi_1 M$ will be word-hyperbolic.

12. CONCLUSION

In this final section we discuss two limitations of the elementary approach adopted in this article and some open problems which seek to address these limitations.

Limitation 1: Dimension. The procedures described for determining whether particular circular galleries contain short closed geodesics rely implicitly on the fact that the set of vectors orthogonal to a plane in \mathbb{E}^3 is a line. If we try to mimic these procedures for \mathbb{S} -complexes of dimension $n > 2$, circular galleries \mathcal{G} can still be cut open and developed on \mathbb{S}^n . A hypothetical geodesic in \mathcal{G} can still, in many instances at least, be sent to an arc lying in a great circle of \mathbb{S}^2 , but the set of all great circles in \mathbb{S}^n is more complicated than the set of great circles in \mathbb{S}^2 . The set of great circles in \mathbb{S}^n is equivalent to the set of 2-planes in \mathbb{E}^{n+1} which is the Grassmannian manifold $G_{n+1,2}$. The special case $n = 2$ is easier to work with because of the well-known homeomorphisms $G_{n,k} \cong G_{n,n-k}$ (via orthogonal complements) and $G_{n,1} \cong \mathbb{R}P^{n-1}$ (i.e., real projective space). Despite the fact that $G_{n,2}$ is slightly more complicated than $G_{3,2} \cong \mathbb{R}P^2$, there should exist an elementary geometric procedure in higher dimensions. By results in [Elder and McCammond 01], it is at least known that an algorithm exists. The only question is whether there is an elementary, geometric algorithm.

Problem 12.1. Find an elementary geometric algorithm to determine whether a finite \mathbb{S} -complex K contains a closed geodesic of length less than 2π when the dimension of K is greater than 2.

Limitation 2: Metrics. Some of the complexes which arise naturally in geometric group theory and low dimensional topology come equipped with a natural metric, but many others do not. The procedures described here require that the metric be given from the start. A natural question is whether a particular simplicial complex will support some piecewise Euclidean metric of

non-positive curvature. This algorithm is not able to address this question since galleries without metrics cannot be uniquely developed onto the 2-sphere. This leads to the following research direction.

Problem 12.2. Given a 3-dimensional complex without a specified metric, find an algorithm to determine whether it supports a piecewise Euclidean [or piecewise hyperbolic] metric of non-positive [resp. negative] curvature.

ACKNOWLEDGEMENTS

This work was supported under NSF Grant No. DMS-99781628.

REFERENCES

- [Ballman 95] W. Ballmann. *Lectures on spaces of nonpositive curvature*, With an appendix by Misha Brin. Birkhäuser Verlag, Basel, 1995.
- [Bowditch 95] B. Bowditch. “Notes on locally CAT(1) spaces.” In *Geometric group theory (Columbus, OH, 1992)*, pp. 1–48, Walter de Gruyter, Berlin, 1995.
- [Bridson 91] M. R. Bridson. “Geodesics and curvature in metric simplicial complexes.” In *Group theory from a geometrical viewpoint (Trieste, 1990)*, pp 373–463, World Sci. Publishing, River Edge, NJ, 1991.
- [Bridson, Haefliger 99] M. R. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*. Springer-Verlag, Berlin, 1999.
- [Brady and McCammond 02] T. Brady and J. McCammond. “Four generator Artin groups of finite type.” In preparation.
- [Elder and McCammond 01] M. Elder and J. McCammond. “CAT(0) is an algorithmic property.” Preprint 2001.

Murray Elder, Department of Mathematics, Texas A&M University, College Station, TX 77843
(murray.elder@math.tamu.edu)

Jon McCammond, Department of Mathematics, Texas A&M University, College Station, TX 77843
(jon.mccammond@math.tamu.edu)

Received January 27, 2001; accepted October 3, 2001.