GEODESIC COMPLEXITY

DONALD M. DAVIS

ABSTRACT. Geodesic complexity has to do with motion planning when you require the paths to be minimal geodesics. We survey many of the interesting results obtained by David Recio-Mitter. Then we focus on two cases, the cube and the 2-point ordered configuration space of a star graph, with a fairly thorough sketch of the proof of the geodesic complexity of each.

1. INTRODUCTION

Michael Farber's concept ([11]) of the topological complexity, $\mathsf{TC}(\mathsf{X})$, of a topological space X involves finding the smallest partition of $X \times X$ into sets S on which you can make a continuous choice of a path from x_0 to x_1 for all $(x_0, x_1) \in S$. In many cases, the paths which appear in an optimal solution are somewhat indirect or roundabout.

In [15], David Recio-Mitter had the idea of requiring that, if X is a metric space, the paths be minimal geodesics; i.e. shortest paths between their endpoints. He defined the *geodesic complexity*, $\mathrm{GC}(X)$, for a metric space X as follows.

Definition 1.1. *If X is a metric space, $\mathrm{GC}(X)$ is the smallest integer k such that there is a partition of $X \times X$ into locally compact sets E_i, $0 \le i \le k$, such that for each i there is a continuous function $\sigma_i : E_i \to P(X)$, where $P(X)$ is the space of all paths in X, and for all $(x_0, x_1) \in E_i$, $\sigma_i(x_0, x_1)$ is a (minimal) geodesic from x_0 to x_1. Each σ_i is called a geodesic motion planning rule (GMPR).*

As we are only interested in minimal geodesics, we will omit the word "minimal."

In the examples that we will be discussing, the lengths of paths will be rather intuitive, but the formal definition of length is as follows.

Date: October 6, 2023.

Key words and phrases. Geodesic complexity, topological robotics, geodesics, configuration space.

2000 *Mathematics Subject Classification*: 53C22, 55M30, 55R80, 52B10.

Definition 1.2. *Let (X,d) be a metric space. The length of a path $\sigma:[0,1]\to X$ is defined as*

$$\ell(\sigma)=\sup_{0\le t_0\le\cdots\le t_N=1}\sum_{i=1}^{N}d(\sigma(t_{i-1}),\sigma(t_i)),$$

where the supremum is taken over all finite partitions of $[0,1]$. A path σ is a geodesic if $\ell(\sigma)\le\ell(\tau)$ for all paths τ with the same endpoints as σ.

Clearly $\mathsf{TC}(\mathsf{X})\le \mathrm{GC}(\mathsf{X})$. In this survey, we will discuss many examples in which equality occurs, and several in which $\mathrm{GC}(X)$ is strictly greater than $\mathsf{TC}(\mathsf{X})$. The simplest example of equality is the sphere S^n with the usual Euclidean metric. The standard motion planning rules (e.g., [11, Theorem 8]) use paths which are geodesics, and so we have the elementary result that with the Euclidean metric,

$$\mathrm{GC}(S^n)=\mathsf{TC}(\mathsf{S^n})=\begin{cases}1 & n \text{ odd}\\ 2 & n \text{ even.}\end{cases}$$

Similarly, the n-torus $T^n=S^1\times\cdots\times S^1$, with n factors, in the product metric (called the flat metric in [15] since it is the metric when T^n is thought of as a quotient of $\mathbb{R}^n$) uses geodesics in its standard motion planning rules ([11, Theorem 12]) and hence satisfies

$$\mathrm{GC}(T^n_{\mathrm{flat}})=\mathsf{TC}(\mathsf{T^n})=\mathsf{n}.$$

This leads to two surprising examples of $\mathrm{GC}>\mathsf{TC}$ proved in [15].

In [15, Theorem 5.1] it is proved that, with $T=T^2$, the usual torus,

$$\mathrm{GC}(T_{\mathrm{emb}})=3>\mathrm{GC}(T_{\mathrm{flat}})=2. \tag{1.3}$$

Here T_{emb} is the usual embedded torus with the Euclidean metric,

$$T_{\mathrm{emb}}=\left\{(x,y,z)\in\mathbb{R}^3\,\middle|\,\left(\sqrt{x^2+y^2}-2\right)^2+z^2=1\right\}.$$

The delicate proof of (1.3) relied on the analysis in [14] of cut loci of points in the embedded torus. The concept of the cut locus of a point of a metric space is intimately related to the study of geodesic complexity.

Definition 1.4. *The cut locus of a point P in a metric space is the closure of the set of points Q for which there is more than one geodesic from P to Q.*

The other interesting related example in [15] is the construction of a metric g on S^{n+1} for which $\mathrm{GC}(S^{n+1}, g) \geq n$, hence yielding examples in which the difference $\mathrm{GC} - \mathsf{TC}$ is arbitrarily large. We sketch the construction. Embed the torus T^n in S^{n+1} with a neighborhood N_2 homeomorphic to $T^n \times (-1, 1)$. Give the subset N_1 which corresponds to $T^n \times (-\frac{1}{2}, \frac{1}{2})$ under the homeomorphism the product metric $T^n_{\mathrm{flat}} \times (-D, D)$, where D is the length of the longest geodesic in T^n_{flat}. Then geodesics between points of $T^n \times \{0\}$ must stay in N_1 and hence be geodesics in S^{n+1}. Use the usual metric on S^{n+1} outside N_2, and construct a metric on $N_2 - N_1$ to yield a metric g on all of S^{n+1}. Then $\mathrm{GC}(S^{n+1}, g) \geq \mathrm{GC}(T^n_{\mathrm{flat}}) = n$.

Another interesting example in [15] is the Klein bottle K with its metric as a quotient of a square. In [15, Theorem 4.5] direct proofs of both inequalities of

$$\mathrm{GC}(K) = 4$$

are given; i.e., it did not rely on the known result that $\mathsf{TC}(\mathsf{K}) = \mathsf{4}$ to obtain the lower bound. The proof in [3] that $\mathsf{TC}(\mathsf{K}) \geq \mathsf{4}$ involved quite complicated algebraic calculations, and, of course, implies that $\mathrm{GC}(K) \geq 4$, but the proof in [15] is of a very different, geometric, nature. The proof that we give of Theorem 3.6 is modeled on it. In [10], we adapted this proof to show that for the n-dimensional Klein bottle K^n, $\mathrm{GC}(K^n) = 2n$. It is interesting that $\mathsf{TC}(\mathsf{K^n})$ is not known. It is $\geq n + 2$; a recent `arXiv` posting ([4]) argues that it is $\leq \frac{3}{2}n + 2$.

In the seminal paper [15], Recio-Mitter showed that if C denotes the (boundary of the) cube, then $\mathrm{GC}(C) \geq 3$. He called it the flat sphere, since it is a metric on a space homeomorphic to S^2. Since $\mathsf{TC}(\mathsf{S^2}) = \mathsf{2}$, this is another instance of $\mathrm{GC} > \mathsf{TC}$. In Section 3, we sketch the proof in [6] that, in fact,

$$\mathrm{GC}(C) = 4. \tag{1.5}$$

The geodesic complexity of 2-point configuration spaces has also been studied, with $\mathrm{GC} = \mathsf{TC}$ in all cases known to the author. If X is a topological space, the ordered n-point configuration space $F(X, n)$ is defined by

$$F(X, n) = \{(x_1, \ldots, x_n) \in X^n : x_i \neq x_j \text{ if } i \neq j\},$$

while the unordered configuration space $C(X,n)$ is the quotient of $F(X,n)$ by the action of the symmetric group. Then $\mathrm{GC}(F(X,n))$ (resp. $\mathrm{GC}(C(X,n))$) can be interpreted to be the number of geodesic motion planning rules (minus 1) required to move n tiny robots without collision between any points of X if you care (resp. don't care) about which robot will be at each point.

In [5], geodesics in $F(\mathbb{R}^n,2)$ and $C(\mathbb{R}^n,2)$ were studied. There is a subtlety that $F(\mathbb{R}^n,2)$ is not geodesically complete.

Definition 1.6. *A metric space X is geodesically complete if there is at least one minimal geodesic between any two points of X.*

If the linear path in $\mathbb{R}^n$ from a_1 to b_1 and the one from a_2 to b_2 intersect for the same parameter value t, then there is no geodesic in $F(\mathbb{R}^n,2)$ from (a_1,a_2) to (b_1,b_2) since the linear path between them is not a path in $F(\mathbb{R}^n,2)$, but there are paths between them in $F(\mathbb{R}^n,2)$ arbitrarily close to the linear one. So we replace $F(\mathbb{R}^n,2)$ by the homotopy equivalent space $F_\varepsilon(\mathbb{R}^n,2)$ consisting of pairs (a_1,a_2) satisfying $d(a_1,a_2)\geq\varepsilon$, where d is the Euclidean metric in $\mathbb{R}^n$.

The metric on $F_\varepsilon(\mathbb{R}^n,2)$ is induced from the Euclidean metric on $\mathbb{R}^n\times\mathbb{R}^n$, while the metric on $C(\mathbb{R}^n,2)$ is

$$d(\{a_1,a_2\},\{b_1,b_2\})=\min\big(d\big((a_1,a_2),(b_1,b_2)\big),d\big((a_1,a_2),(b_2,b_1)\big)\big).$$

Both spaces are geodesically complete. The spaces $F_\varepsilon(\mathbb{R}^n,2)$ and $C(\mathbb{R}^n,2)$ have the homotopy type of S^{n-1} and real projective space RP^{n-1}, respectively, and it was shown in [5] that

$$\mathrm{GC}(F_\varepsilon(\mathbb{R}^n,2))=\mathsf{TC}(\mathsf{F}(\mathbb{R}^\mathsf{n},2))=\mathsf{TC}(\mathsf{S}^{\mathsf{n}-1})$$

and

$$\mathrm{GC}(C(\mathbb{R}^n,2))=\mathsf{TC}(\mathsf{C}(\mathbb{R}^\mathsf{n},2))=\mathsf{TC}(\mathsf{RP}^{\mathsf{n}-1}),$$

by constructing explicit GMPRs.

It was shown in [13] that, if $n\notin\{1,3,7\}$, then $\mathsf{TC}(\mathsf{RP}^\mathsf{n})$ equals the immersion dimension of RP^n, the dimension of the smallest Euclidean space in which RP^n can be immersed, a number which is known for many values of n and unknown for many others. See [7] for tables. In [15] it was noted that, although the motion planning rules in RP^n described in [13] are not geodesic, they can easily be replaced by GMPRs.

It is interesting that the GMPRs in $C(\mathbb{R}^n, 2)$ do not rely on GMPRs in RP^{n-1}, but rather use just ordinary motion planning rules in RP^{n-1} to make continuous choices of geodesics in $C(\mathbb{R}^n, 2)$.

An important example of configuration spaces occurs when X is a graph G; i.e., a set of vertices and edges between some of them. Then motion planning rules in $F(G, n)$ can be thought of as moving n robots without collision along a system of wires, as in a warehouse. In [9], we proved that if G is a star graph (i.e., a graph with a single vertex and no cycles) with at least four edges, then

$$\mathrm{GC}(F_\varepsilon(G, 2)) = \mathsf{TC}(\mathsf{F}(\mathsf{G}, 2)) = 2, \tag{1.7}$$

and if G is a tree (i.e., a graph with no cycles), then

$$\mathrm{GC}(C(G, 2)) = \mathsf{TC}(\mathsf{C}(\mathsf{G}, 2)) = 2.$$

Here again, we need to require that the two points remain a distance ε apart in $F(G, 2)$ in order that geodesics always exist. We will discuss this in more detail in Section 2.

In the final two sections, we focus on the cube and the 2-point ordered configuration space of a star graph, because the two cases differ in many respects: configuration space versus familiar geometric space, $\mathrm{GC} = \mathsf{TC}$ in one but not the other, ε needed in one, direct argument for lower bound needed in the other. We give explicit GMPRs, of very different types, in the two, and give fairly complete sketches of proofs of (1.5) and (1.7).

2. Two robots moving on a star graph

In this section, we present many of the main ideas in [9] regarding $\mathrm{GC}(F_\varepsilon(G, 2))$ when G is a star graph with at least four edges. The three-edge case, called a Y graph, is also considered there, but we shall omit that. One issue is the metric, which is inherited from a metric on $G \times G$. Any tree graph has an obvious metric, given by shortest distance along the graph, but there are several well-known metrics on $G \times G$. Two are considered in [9]. The one that we will use here is the ℓ_2-metric, defined by

$$d_2((a_1, a_2), (b_1, b_2)) = \sqrt{d(a_1, b_1)^2 + d(a_2, b_2)^2}.$$

In $(G \times G, d_2)$, there is a unique geodesic from (a_1, a_2) to (b_1, b_2) given by uniform motion from a_1 to a_2 simultaneous with uniform motion from b_1 to b_2. That will often involve a collision, which must be avoided in $F(G, 2)$. Avoiding collision is usually accomplished by moving one particle onto another arm so that the other can pass. In such a case, there is no geodesic in $F(G, 2)$, since the motion onto the other arm can be arbitrarily small. Hence, the need for ε.

The other metric on $F(G, 2)$ used in [9] is the ℓ_1-metric, defined by

$$d_1((a_1, a_2), (b_1, b_2)) = d(a_1, b_1) + d(a_2, b_2).$$

Here it is the sum of the distances moved by the two particles that determines the length of a path in the configuration space, so simultaneous motion is not required, and there are many parametrizations yielding the same distance. In this survey, we will consider only the ℓ_2 metric.

In this paper, we diverge slightly from the treatment in [9] and depict the motion as of a disk of diameter ε (particle 1) and a square of side length ε (particle 2). This differs from the depiction in [9] (e.g., their Figure 1), where ε is unrelated to the depicted size of the point particles. Our model would be an accurate depiction if the arms did not make angles with one another. But because of the thickness of the disk and square outside of the graph, if they are on two edges near the vertex, they can intersect even though the distance along the graph between their centers is greater than ε. We should just think of the intersection with the graph of the disk and square. We can think of them as worms of length ε moving on the graph, and we will sometimes use the name "worm" to refer to the particles. See Figure 2.1 where the left diagram shows the ● and ■ overlapping even though the associated worms do not. You should think of the ● and ■ as visual approximations of the worms, which are hard to see. When a worm is close to a vertex, it grows temporary arms reaching onto all edges to a distance $\varepsilon/2$ from its center. This is illustrated in the right side of Figure 2.1.

Figure 2.1. Worms and overlapping approximation, and arms.

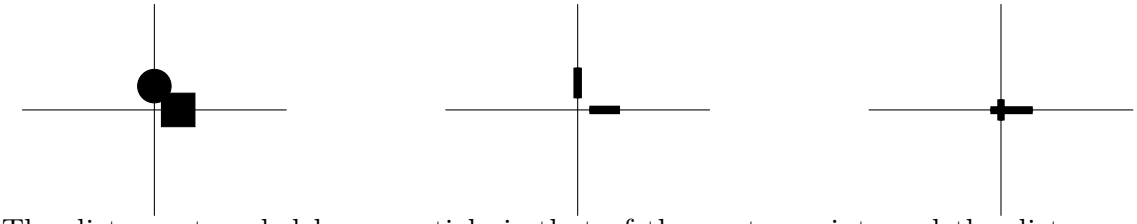

The distance traveled by a particle is that of the center point, and the distance traveled by the two particles in uniform motion between two pairs of points on the graph is $\sqrt{D_1^2 + D_2^2}$, where D_i is the distance traveled by particle i. The disk and square can be tangent to one another, but cannot overlap more than that along the graph. We will depict the first particle by a round black disk, and its destination by a round white disk, while the second particle and its destination will be depicted by black and white squares.

We introduce our ideas with a specific numerical example. Suppose $\varepsilon = 1$, so that the radius of the particles is 0.5. We consider the configuration on the left side of Figures 2.2 and 2.3, with distance of the center of each particle and of each destination from the vertex as indicated. We include only three arms since other arms are not needed in the geodesic motion. Clearly direct motion to the destination without collision is impossible, so one particle will have to move to the third arm to let the other pass by.

Figure 2.2. ■ moves onto third arm.

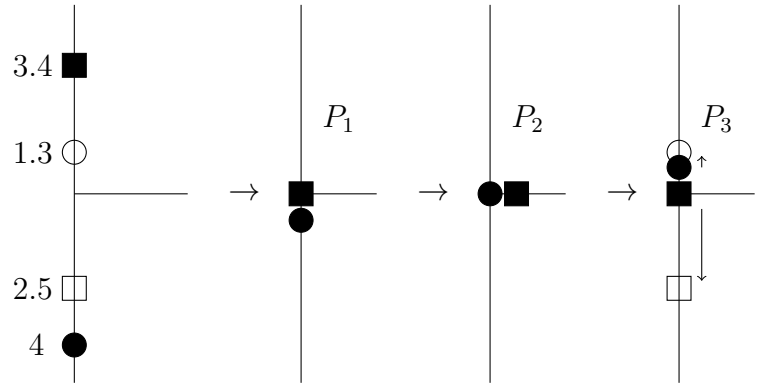

In Figure 2.2, we wish to move ■ onto the third arm so that ● can pass by. To move directly from the initial position to position P_2 with uniform motion of both particles, ■ would move 4.4 units while ● moves 4. So ■ is moving faster than ●,

and this uniform motion to P_2 cannot be done (the worms would overlap prior to arriving at P_2). So instead, they move uniformly to position P_1, and then at equal speeds from P_1 to P_2. Uniform motion from P_2 to the destination would have ■ moving 3.5 units, while ● moves only 1.3. So ■ is moving faster and would overtake ● immediately. Hence they move from P_2 to P_3 at equal speeds, and from there to their final destination. The total distance traveled is

$$\sqrt{3^2+3.4^2}+\sqrt{2}+\sqrt{2}+\sqrt{0.3^2+2.5^2}=9.88.$$

If instead, we wish to move ● onto the extra arm with direct uniform motion to P_5 in Figure 2.3, it would be moving faster than ■, and hence they would overlap prior to reaching P_5. So instead they first move to P_4 and then to P_5. From P_5 to the destination, ■ will be moving faster, so they can go directly from P_5 to their destination. The total distance is

$$\sqrt{4^2+2.4^2}+\sqrt{2}+\sqrt{2.3^2+2.5^2}=9.476,$$

so the latter path is the geodesic.

Figure 2.3. ● moves onto third arm.

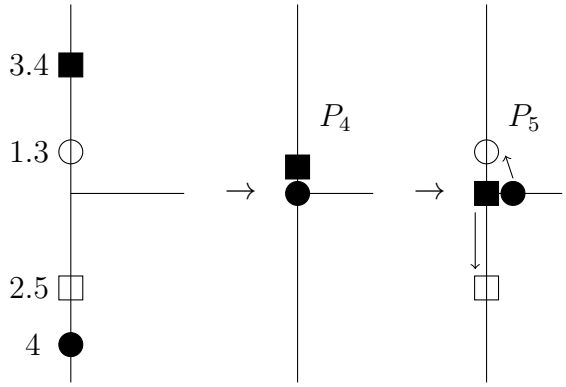

The configurations $((a_1,a_2),(b_1,b_2))$ in $F(G,2)$ are first partitioned into sets X_1, X_2, X_3, and X_4, defined by how many of the (open) arms contain at least one of a_1, a_2, b_1, and b_2. We will subdivide the sets X_i and group these into the three sets E_0, E_1, and E_2 on which there are GMPRs, as required for proving

$$\mathrm{GC}(F_\varepsilon(G,2))\le 2.$$

The lower bound,

$$2 = \mathsf{TC}(\mathsf{F}(\mathsf{G}, 2)) = \mathsf{TC}(\mathsf{F}_\varepsilon(\mathsf{G}, 2)) \leq \mathrm{GC}(\mathsf{F}_\varepsilon(\mathsf{G}, 2)),$$

was proved in [12].

The first set, E_0, is a subset of X_2^-, which is defined to be the subset of X_2 on which the relative orientation of a_1 and a_2 is opposite to that of b_1 and b_2. This means that direct motion from a_1 to b_1, and a_2 to b_2, will cause the particles to collide, and so motion of one onto a different arm to allow the other to pass is required. There are at least two arms onto which this extra motion could take place. In order to have a continuous selection of geodesics, we make a convention about this. We number the arms, from 1 to N (the total number of arms), and then say that, for configurations in X_2^-, the extra motion will be onto the unused arm with the smallest number. We compare the length of the path if a_1 moves onto the extra arm with that if a_2 moves onto it, taking into account the possible additional motion such as that illustrated in Figures 2.2 and 2.3. For each, this length will be a sum of two or more square roots of sums of squares. It is possible that the two sums in this comparison might be equal. The set E_0 consists of those in which they are not equal, which we call X_2^n. For these configurations, we choose the path with the shorter length. Although this wasn't really a "unique" geodesic, due to the choice of arm, it can be considered to be, by removal of the other arms, and it follows from [2, Lemma 3.12] that unique geodesics vary continuously with their endpoints. This defines the GMPR on E_0.

The second set with a GMPR, E_1, is the union $X_1^- \cup X_4^{\mathrm{eq}}$. Here X_1^- consists of configurations in X_1 for which the relative orientation of a_1 and a_2 is opposite to that of b_1 and b_2. For X_4 configurations, we consider direct motion from a_1 to b_1 and from a_2 to b_2. If this avoids collision, it is the geodesic, and will later be placed in set X_4^n. If not, we compare the path length when a_1 moves through the vertex first (pictured in Figure 2.4; the "if necessary" intermediate motion is done if ■ is moving faster than ● in the middle part of Figure 2.4, in which case the arrows in the middle diagram will not be followed.) with that when a_2 passes through first. The configuration is in X_4^{eq} if these two lengths are equal.

Figure 2.4. a_1 **moves through first**

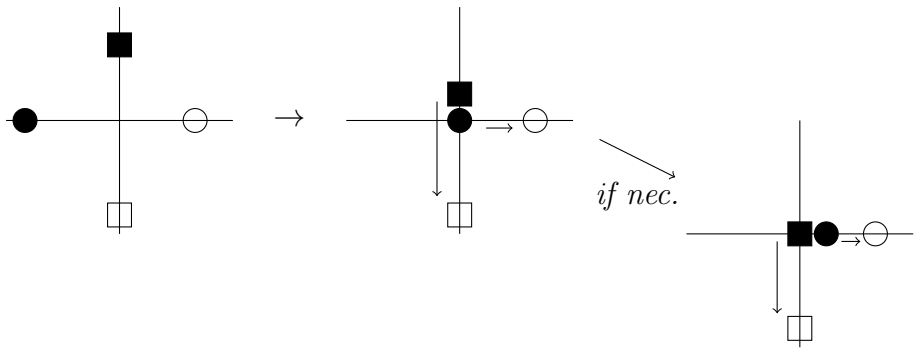

For configurations in X_4^{eq}, we choose the geodesic which moves a_1 through the vertex first. This is a GMPR on X_4^{eq}. For configurations in X_1^-, if they are on arm i, the particle closer to the vertex moves through onto arm $i+1$, while the other particle moves to the vertex. Then, at equal speed, they move into the position on the right side of Figure 2.5, using arm $i+2$, and then to their destination. This is a GMPR on X_1^-.

Figure 2.5. Motion in X_1^-

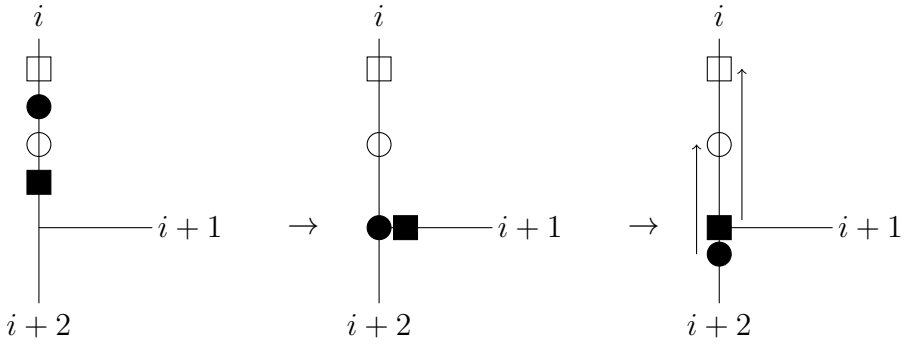

We have defined GMPRs on X_4^{eq} and X_1^-. These can be combined to yield a GMPR on their union, E_1, because the two sets are topologically disjoint.

Definition 2.6. *Two subsets of a topological space are topologically disjoint if no point of one is the limit of a sequence of points of the other.*

Continuous functions on topologically disjoint sets of a metric space can be combined to give a continuous function on the union. The limit of a sequence of configurations

in X_1 will have all four points on a single closed arm, while the limit of a sequence of configurations in X_4 cannot have more than one point on any open arm.

The third set with a GMPR consists of everything else, which is

$$E_2 = X_4^n \cup X_3 \cup X_2^+ \cup X_1^+ \cup X_2^{\text{eq}}.$$

We will show that configurations in the first four of these sets have unique geodesics. By [2, Lemma 3.12], unique geodesics vary continuously, so we have a GMPR on $X_4^n \cup X_3 \cup X_2^+ \cup X_1^+$. We will show that X_2^{eq} has a GMPR and is topologically disjoint from the four other sets. Thus we have a GMPR on all of E_2.

The sets X_1^+ and X_2^+ are the subsets of X_1 and X_2 in which the relative orientation of a_1 and a_2 agrees with that of b_1 and b_2. The direct path from (a_1, a_2) to (b_1, b_2) is clearly the unique geodesic on these. The set X_4^n is the complement in X_4 of the set X_4^{eq} considered earlier (e.g., Figure 2.4). The set X_4^n includes configurations in which direct motion avoids collision, as well as those in which the two paths, Figure 2.4 and its analogue, have differing total lengths. Choosing the shorter of the two unequal paths clearly gives a unique geodesic, yielding a GMPR on X_4^n by [2, 3.12].

The set X_2^{eq} equals $X_2 - X_2^n$, where X_2^n was described in E_0. There we explained the choice of the arm which will be used for passing. Now the two choices of which passes first have equal length. We choose the one in which a_1 passes first. That yields a GMPR on X_2^{eq}. That X_2^n is topologically disjoint from $X_4^n \cup X_3 \cup X_2^+ \cup X_1^+$ is quite clear, with the only slightly delicate point being that a sequence of X_3 configurations can approach an X_2 configuration with one of a_1, a_2, b_1, or b_2 at the vertex. However, such a configuration cannot have two equal geodesics because if a_i or b_i is at the vertex, one can verify that the path moving a_i onto the third arm is the shorter of the two.

Finally, we show that X_3 configurations have unique geodesics. This is the most complicated set, as there are three cases, and one of them has four subcases. The three cases are shown in the left side of Figure 2.7; they depend on whether the two on the same arm have the same color and whether they have the same shape. Distances from the vertex can vary greatly from what is pictured.

Figure 2.7. Three cases for X_3

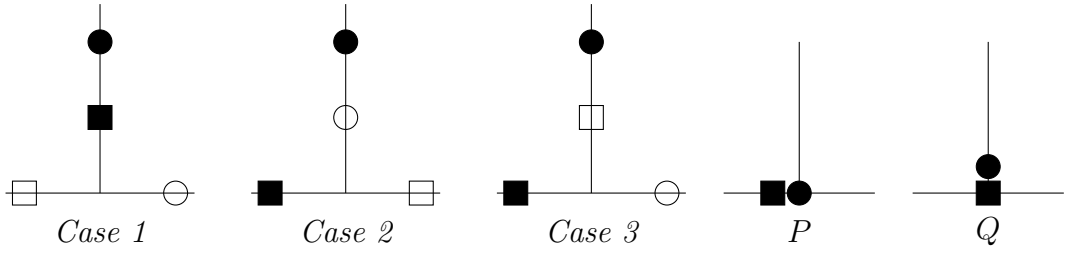

In Case 1, particles move uniformly toward their destination. However if, due to moving faster, the outer one overtakes the inner one, then, instead, they move uniformly to P, and then to their destination. This is a unique geodesic. In Case 2, they move uniformly toward their destination unless both ● and ○ are very close to the vertex, in which case the particles move uniformly to Q, and then on to their destination. This defines a unique geodesic.

Now we consider Case 3. Sometimes the particles can move uniformly to their destination without collision. If so, that's what they do. Subcase 3a, in the left side of Figure 2.8, is where they would collide, so instead move uniformly to A, from which they can move directly to their destination. Subcase 3b, in the right side of Figure 2.8, is where again they must move to A, but then, because in direct motion to the destination the trailing particle would have to move faster, they instead move at equal speed from A to B, before moving to their destination.

Figure 2.8. Two subcases for X_3

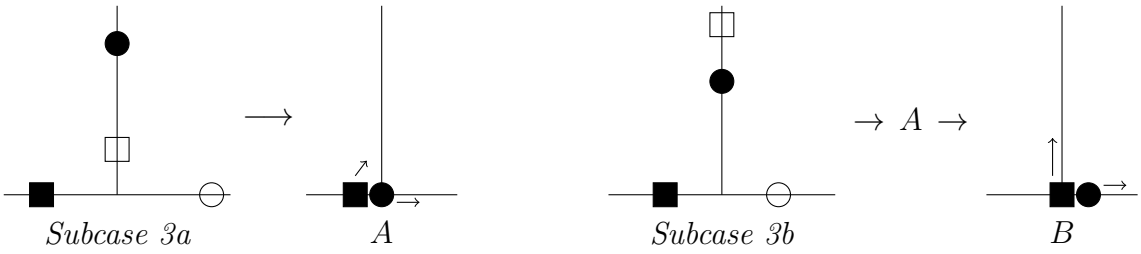

In the final subcase, 3c, illustrated in Figure 2.9, uniform motion to the destination would again result in a collision, as would uniform motion from configuration A of Figure 2.8 to the destination, so we move directly to B above and then to the destination. We could go to A as an intermediate, but that would be a longer path.

The difference between this and 3b is that in 3b you could not go directly to B because ● is moving faster. A different visualization of Case 3 can be seen in [9, Figure 5].

Figure 2.9. Subcase 3c

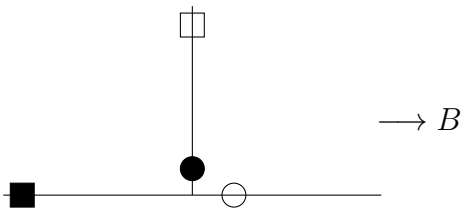

This completes our description of GMPRs on E_0, E_1, and E_2, yielding $\mathrm{GC}(F_\varepsilon(G,2)) \leq 2$.

3. Geodesic complexity of a cube

In this section, we sketch the proof from [6] that, with C a cube,

$$\mathrm{GC}(C) = 4.$$

The proof relies heavily on the analysis in [8] of cut loci (Definition 1.4) in the cube, and we begin with a discussion of this.

We begin with a visual example. Figure 3.1 shows a cube with the numbering of corner points that we will use, with the cut locus of the midpoint P of edge 5-8 shown in bold red. Our cut loci will usually be displayed as on the right side of the diagram.

Figure 3.1. Cut locus of P

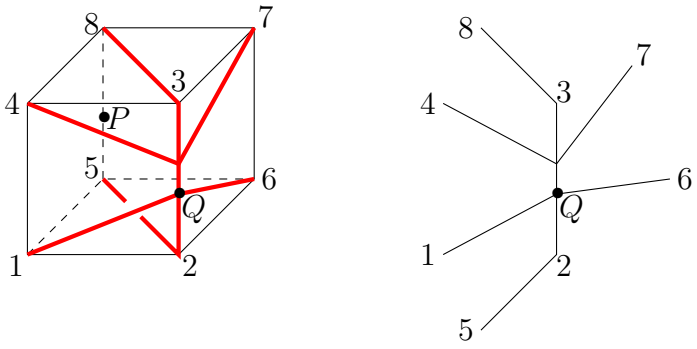

The path chosen from a point P to a point Q in its cut locus will be important. For example, if Q is the highlighted point in Figure 3.1, the reader can visualize the four types of paths from P to Q. For example, the type passing through the 1-4 region,

between the edges ending at 4 and at 1 in the diagram on the right side, will pass through the face on the left side of the cube and then through the front face, while the type which we see as passing through the 1-5 region will go down to the 1-5 edge of the cube, then along the bottom of the cube to the 1-2 edge, and then up along the front face to point Q.

The method that we use to find the cut loci is that of star unfolding and Voronoi diagrams developed in papers such as [1]. We illustrate with the same example we used in [8] and [6]. We unfold the cube, except for the 1234 face, and draw the shortest segment from P to each of the corner points. The diagram decomposes as eight polygons with P at a vertex of each. In a new drawing, we attach four of these polygons to the 1234 face, and then the other four polygons to now-exposed edges. This is the star unfolding of P. Eight of its sixteen vertices are different placements P_α of the point P in unfoldings of the cube, while the others correspond to corner points of the cube.

Figure 3.2. Star unfolding

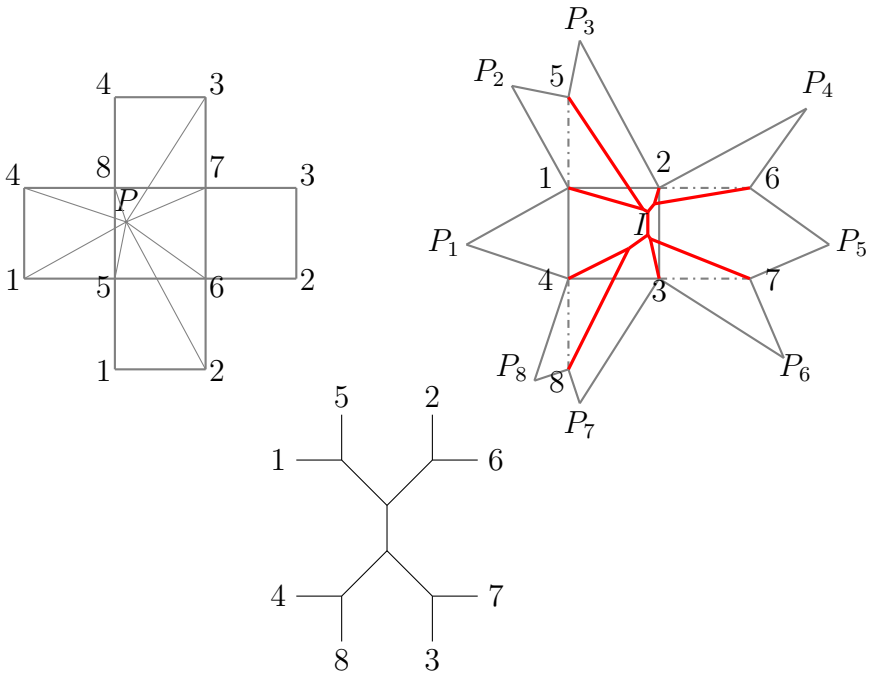

We partition the star unfolding into subsets which are closer to one P_α than to the others. This is called the Voronoi cell of P_α. The cut locus of P consists of

points bounding two Voronoi cells. This is called the Voronoi diagram. These will always be portions of the perpendicular bisector of the segment connecting some P_α and P_β. The delicate part is seeing how they intersect. There is a short segment in the diagram bounding the Voronoi cells of P_3 and P_5, and an almost-indiscernible segment bounding the cells of P_5 and P_7. We are concerned with the combinatorial structure of the cut locus, and so would draw this as in the bottom portion of Figure 3.2.

To visualize paths from P to points in its cut locus requires various unfoldings of the cube. We illustrate in Figure 3.3 the easiest of these, the paths from P_5 in the right diagram of Figure 3.2, which we will think of as from the 6-7 region in the schematic diagram at the bottom of Figure 3.2.

Figure 3.3. Paths from P to cut locus

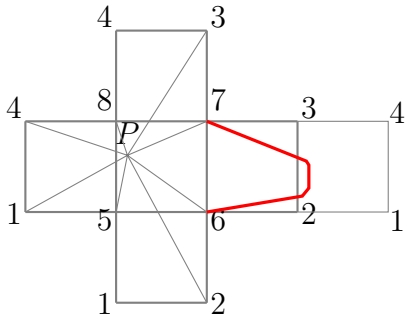

We explain in [8] how we were able to partition the face of the cube into 193 connected sets (of dimension 0, 1, and 2) on which the combinatorial structure of the cut locus is constant. These are displayed in Figure 3.4, while Figure 3.5 is a stretched blowup of the left side of the left quadrant. In [8] we present the combinatorial structure of the cut locus in each region and curve of Figure 3.5, as well as equations of the curves.

Figure 3.4. Regions in a face

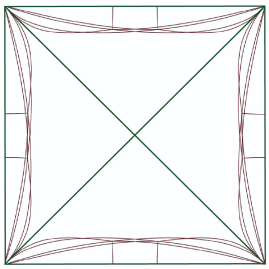

Figure 3.5. Regions in left quadrant

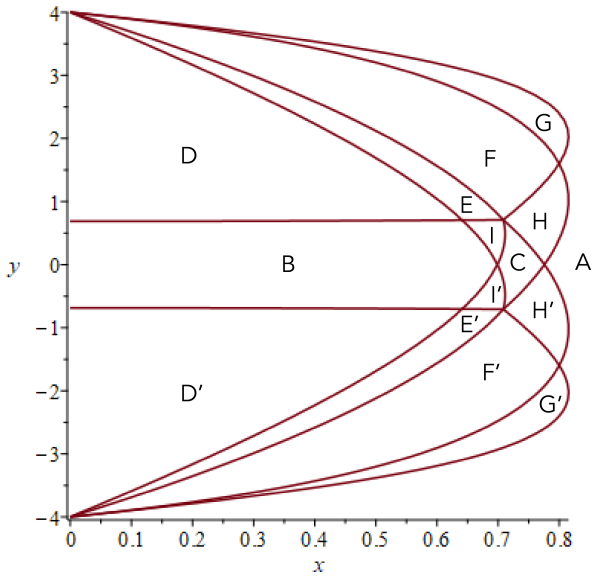

Now we sketch the proof that $\mathrm{GC}(C) \geq 4$ in the following result.

Theorem 3.6. *It is impossible to partition $C \times C$ as $E_0 \sqcup E_1 \sqcup E_2 \sqcup E_3$ with a GMPR σ_i on each E_i.*

Proof. Assume such a decomposition exists. In Figure 3.7, we display the combinatorial structure of the cut locus of the corner point V_8 (in the upper left corner in Figure 3.5) and of points in region D and on the curve DE bounding regions D and E. As points in region D or curve DE approach V_8, the vertical segment in the depiction of their cut locus in Figure 3.7 shrinks to a point, while as points in D approach DE, the vertical segment going up from edge 6 shrinks to a point.

Figure 3.7. Cut loci of V_8, DE, and D

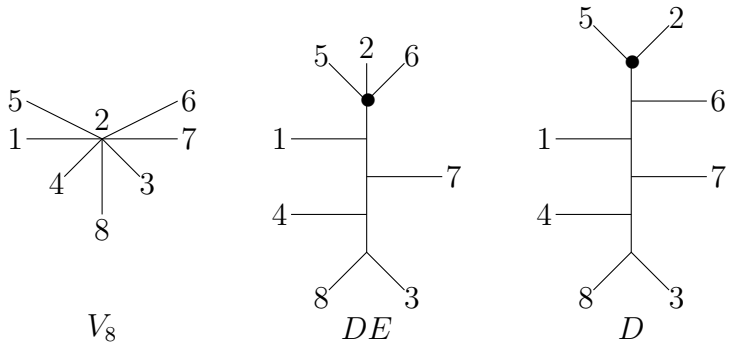

Let E_0 be the set containing (V_8, V_2). We assume that $\sigma_0(V_8, V_2)$ passes through the region bounded by edges 4 and 8 as it approaches V_2. If it passed through one of the upper regions, we would use a region different than D in our proof. For points P on the DE curve, let Q_P be the indicated point in the DE diagram. As P approaches V_8, if (P, Q_P) were in E_0, then $\sigma_0(P, Q_P)$ would have to approach $\sigma_0(V_8, V_2)$. But $\sigma(P, Q_P)$ cannot pass through the region between edges 4 and 8, so we conclude that there must be a sequence P_n approaching V_8 on DE such that (P_n, Q_{P_n}) are in the same set not equal to E_0. Let E_1 be the set containing these points. By restricting, we may arrange that $\sigma_1(P_n, Q_{P_n})$ all pass through the same region as they approach their Q_{P_n}, and we will assume this is the region between edges 6 and 7. If it was a different region, we would vary the argument, possibly using region E instead of D.

By a similar argument, there must be sequences $P_{n,m}$ in D approaching P_n such that, with $Q_{P_{n,m}}$ the indicated point in the D diagram, $(P_{n,m}, Q_{P_{n,m}})$ are all in the same set not equal to E_1, and, since a diagonal sequence $P_{n,n}$ approaches V_8, this set cannot be E_0, so must be a different set, E_2. We assume that $\sigma_2(P_{n,m}, Q_{P_{n,m}})$ pass through the region between edges 5 and 2. The same argument works if it was one of the other regions.

For a sequence of points $Q_{n,m,\ell}$ approaching $Q_{P_{n,m}}$ along the edge in the D diagram going down from the indicated point, $\sigma(P_{n,m}, Q_{n,m,\ell})$ must approach $Q_{n,m,\ell}$ from the left or right, so the limit of these cannot equal $\sigma_2(P_{n,m}, Q_{P_{n,m}})$. Hence there must be a sequence of these in the same set not equal to E_2. By utilizing diagonal subsequences, we deduce this set is not E_1 or E_0. So it must be E_3. By restricting, we may assume that all $\sigma_3(P_{n,m}, Q_{n,m,\ell})$ come from either the left side or right side, which we will say is the left side.

Choose points $Q_{n,m,\ell,k}$ in the complement of the cut locus of $P_{n,m}$ approaching $Q_{n,m,\ell}$ from the right side of the vertical cut-locus segment containing $Q_{n,m,\ell}$. By restricting, we may assume that all $(P_{n,m}, Q_{n,m,\ell,k})$ are in the same E_i. Then $\sigma_i(P_{n,m}, Q_{n,m,\ell,k})$ is the unique geodesic between these points, and must be coming from the region bounded by edges 2 and 6. Since $\sigma_3(P_{n,m}, Q_{n,m,\ell})$ comes from the opposite side, i cannot equal 3. By diagonal arguments, it cannot equal 0, 1, or 2. Therefore there must be another region E_4. ■

We complete the proof of (1.5) with the following. We will sketch the proof, referring the reader to [6] for quite a few details.

Theorem 3.8. *There is a decomposition*

$$C \times C = E_0 \sqcup E_1 \sqcup E_2 \sqcup E_3 \sqcup E_4$$

into locally compact sets E_i with a GMPR σ_i on each.

Proof. We start by letting E_0 be the set of pairs (P, Q) for which there is a unique geodesic from P to Q, and let $\sigma_0(P, Q)$ be that path.

Some adjacent regions in Figure 3.5 have the same combinatorial structure for the top halves of their cut loci, and so are combined, leading to the forms in Figure 3.9. The circles will be explained later. The A is for the big region in the left quadrant of Figure 3.4, and $\mathcal{E}$ is the left edge. The $*$ is vertex $I'CH'F'E'$ in Figure 3.5. It is approached by all the diagrams except $\mathcal{E}$, but its vertices will be in a separate E_i, so that compatibility is not a concern. The bottom halves of these cut loci have a similar form.

Figure 3.9. Top halves of cut loci

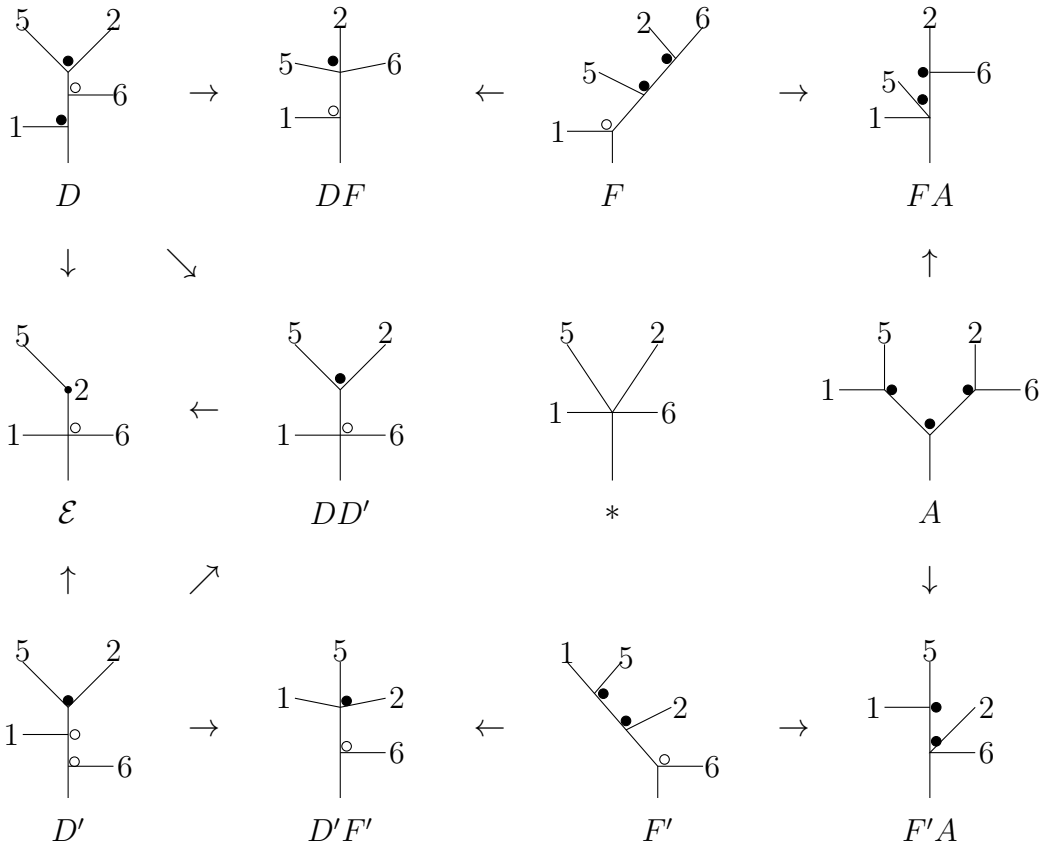

Our set E_1 consists of all (P, Q) such that Q is in the interior of an edge of the cut locus of P or is a degree-2 vertex, such as vertex 2 in $\mathcal{E}$ in Figure 3.9. The cut loci are all divided into a top half and bottom half, up to rotation. We choose an orientation of the cube and use that to consider a rotation of each cut locus around a central point between the two halves. We choose $\sigma_1(P, Q)$ to be the path in the direction of the rotation. See Figure 3.10. Regarding the segments connecting the two halves of cut loci: each edge of the cube bounds two quadrants, and all cut loci in those two quadrants have parallel connecting segments. Arbitrarily make a choice of direction for the approach to the connecting segments in each pair of quadrants. These quadrant pairs are separated by diagonals of faces. The connecting segments of the cut locus of points on a diagonal are a single degree-4 vertex, which is in E_4.

So compatibility of σ_1 on the connecting segments of the various quadrant pairs is not an issue.

Figure 3.10. Rotation direction for σ_1

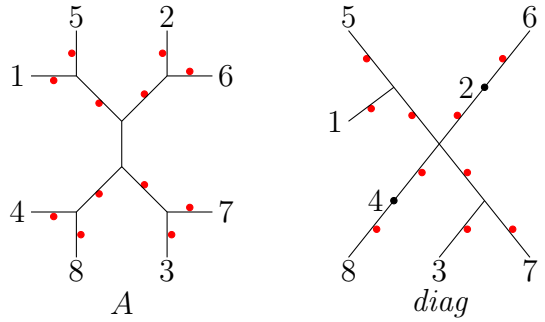

The set E_2 consists of the 56 points (P, Q) such that Q is the corner point of the cut locus of P of degree 5 or 6. Since this is a discrete set, the function σ_2 can be defined arbitrarily. Eight of these have P a corner point and Q the opposite corner point. The others are where P is the analogue of the point $*$ of Figure 3.9 in any of the 48 half-quadrants and Q the degree-5 vertex in its cut locus.

Degree-3 and -4 vertices of cut loci are placed in the final two sets, E_3 and E_4, but there is not a direct correspondence between the 3's and 4's. Refer to Figure 3.9, for top halves of cut loci in the left quadrant of the 5678 face. There, vertices which can be approached from the region between edges 2 and 5 are approached that way, and are in set E_3, indicated by $\bullet$, while others which can be approached from the region between edges 2 and 6 are approached that way, and are in E_4, indicated by $\circ$. Remaining cases, one each in the D, DF, and F diagrams, are placed as indicated there. One readily checks the continuity of these approaches within each set. It is not difficult to show that a similar grouping works on bottom halves of cut loci and in other quadrants. Detailed explanations are provided in [6]. ■

References

[1] P. Agarwal, B. Arnov, J. O'Rourke, and C. Schevon, *Star Unfolding of a Polytope with Applications*, SIAM J. Comput., **26** (1997) 1689–1713.

[2] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der mathematischen Wissenschaften **319** Springer-Verlag (1999).

[3] D. Cohen and L. Vandembroucq, *Topological complexity of the Klein bottle*, J. Appl and Comput Topology **1** (2017) 199–213.
[4] N. Daundkar and S. Sarkar, *LS-category and topological complexity of several families of fibre bundles*, arXiv 2302.00468.
[5] D. M. Davis, *Geodesics in the configuration spaces of two points in* $\mathbb{R}^n$, Tblisi Math Jour **14** (2021) 149–162.
[6] ———, *Geodesic complexity of the cube*, arXiv 2308.04316.
[7] ———, `www.lehigh.edu/~dmd1/immtable`
[8] D. M. Davis and M. Guo, *Isomorphism classes of cut loci on a cube*, Discrete Math **347** (2024), https://doi/10.1016/j.disc.2023.113709.
[9] D. M. Davis, M. Harrison, and D. Recio-Mitter, *Two robots moving geodesically on a tree*, Alg and Geom Topology **22** (2022) 785–814.
[10] D. M. Davis and D. Recio-Mitter, *Geodesic complexity of the n-dimensional Klein bottle*, New York Jour Math **27** (2021) 296–318.
[11] M. Farber, Topological complexity of motion planning, Discr Comp Geom **29** (2003) 211–221.
[12] ———, *Collision-free motion planning on graphs*, Algorithmic Foundations of Robotics VI (2005) Springer-Verlag 123–138.
[13] M. Farber, S. Tabachnikov, and S. Yuzvinsky, *Topological robotics: motion planning in projective spaces*, Int Math Res Note **34** (2003) 1853–1870.
[14] J. Gravesen, S. Markvorsen, R. Sinclair, and M. Tanaka, *The cut locus of a torus of revolution,* Asian J. Math **9** (2005) 103–120.
[15] D. Recio-Mitter, *Geodesic complexity of motion planning*, J. Appl. Comput. Topol, **5** (2021) 141–178.

Department of Mathematics, Lehigh University, Bethlehem, PA 18015, USA
Email address: `dmd1@lehigh.edu`