

Two robots moving geodesically on a tree

DONALD M. DAVIS
 MICHAEL HARRISON
 DAVID RECIO-MITTER

We study the geodesic complexity of the ordered and unordered configuration spaces of graphs in both the ℓ_1 and ℓ_2 metrics. We determine the geodesic complexity of the ordered two-point ε -configuration space of any star graph in both the ℓ_1 and ℓ_2 metrics and of the unordered two-point configuration space of any tree in the ℓ_1 metric, by finding explicit geodesics from any pair to any other pair, and arranging them into a minimal number of continuously-varying families. In each case the geodesic complexity matches the known value of the topological complexity.

[53C22](#), [55R80](#); [55M30](#), [68T40](#)

1 Introduction and statement of results

The topological complexity of a space X , introduced almost two decades ago by Farber [5], is a homotopy invariant of X which measures the complexity of the *motion planning problem* on X . For example, if X represents the space of all possible states of a robot arm, then $\text{TC}(X)$ measures the number of rules required to determine a complete algorithm which dictates how the robot arm will move from any given initial state to any given final state. The formal definition requires the notion of the *free path fibration* $PX \rightarrow X \times X$, which sends a path $\gamma : [0, 1] \rightarrow X$ to the pair $(\gamma(0), \gamma(1))$. The (*reduced*) *topological complexity* of X is then defined as the smallest number k for which there exists a decomposition

$$X \times X = \bigsqcup_{i=0}^k E_i$$

into Euclidean Neighborhood Retracts (ENRs) E_i , such that there exist local sections $s_i : E_i \rightarrow PX$ of the free path fibration. The sections s_i are the “rules” which specify, for any two points $(a, b) \in E_i \subset X \times X$, a path from a to b . The fact that s_i is a section ensures that the rules vary continuously with $(a, b) \in E_i$; here PX is considered with the compact-open topology.

Geodesic complexity, recently introduced by the third author [8], is a geometric counterpart to topological complexity defined for metric spaces (X, g) . The geodesic complexity also measures the complexity of the motion planning problem, except that all motions are required to follow minimizing geodesics in X .

Definition 1 Let (X, g) be a metric space. Let $\gamma : [0, 1] \rightarrow X$ be a path in X and let $\ell(\gamma)$ denote its length. We say that γ is a (*minimal*) *geodesic* if $\ell(\gamma) = g(\gamma(0), \gamma(1))$.

Remark We frequently drop the word “minimal,” but we follow the convention that “geodesic” always refers to a minimizing geodesic as defined above. See [8] for alternate equivalent definitions and some discussion on terminology in metric vs. riemannian geometry.

The formal definition of geodesic complexity is analogous to that of topological complexity. Let GX be the subspace of PX consisting of geodesics. Restricting the free path fibration $PX \rightarrow X \times X$ to GX yields a map $\pi : GX \rightarrow X \times X$. We note that the map π is no longer a fibration (except in the very special case when it is a homeomorphism), although it is sometimes a branched covering (see [8], Section 3).

Definition 2 The *geodesic complexity* $\text{GC}(X, g)$ of a metric space (X, g) is defined as the smallest number k for which there exists a decomposition

$$X \times X = \bigsqcup_{i=0}^k E_i$$

into ENRs E_i , such that there exist local sections $s_i : E_i \rightarrow GX$ of π . We refer to the collection $\{s_i\}$ as a *geodesic motion planner* and each s_i as a *geodesic motion planning rule* (GMPR).

By definition, $\text{TC}(X) \leq \text{GC}(X, g)$ for any metric g . In particular, the topological complexity of a space is a homotopy invariant, hence independent of the metric on X , but the geodesic complexity of a space genuinely depends on the metric. The third author showed in [8] that on each sphere S^n , $n \geq 3$, there exist two metrics with different geodesic complexity. Specifically, for every $k \in \mathbb{N}$, there exists a metric g on S^{k+2} with $\text{GC}(S^{k+2}, g) - \text{TC}(S^{k+2}) \geq k - 1$, so the gap between GC and TC may be arbitrarily large.

Besides trivial examples, the geodesic complexity has been computed for only a handful of spaces. In [8], the third author computes the geodesic complexity of the flat n -torus and the flat Klein bottle and gives lower bounds for the GC of the standard 2-torus and

for the boundary of the 3-cube in \mathbb{R}^3 which are larger than the TC of the respective spaces. In [4], the first and third authors compute the GC of the n -dimensional Klein bottles (see also [2]), for which the topological complexity is still unknown except in the case of $n = 2$ (the ordinary Klein bottle). In [3], the first author computes the GC of certain configuration spaces of \mathbb{R}^n .

Our present goal is to compute the geodesic complexity of configuration spaces of certain graphs. Configuration spaces of graphs are of central importance in topological robotics, since they model the situation of several robots moving along tracks, as in a warehouse ([7]). We consider the case of two points, either distinguished or indistinguishable, moving on a graph without colliding. We always assume that the graph G is a tree, and unless otherwise stated, we assume that G is not homeomorphic to an interval. We obtain explicit descriptions of the geodesics and optimal GMPRs on the configuration spaces of these graphs.

As a motivational example, let G be the figure-Y graph (three edges emanating from a single vertex) with its usual path metric d , and consider the space F_ε consisting of pairs of points of Y which are at least distance ε apart. A path in X from $a = (a_1, a_2)$ to $b = (b_1, b_2)$ may be thought of as the motion of two particles \bullet (the *first* particle) and \blacksquare (the *second* particle) in Y , beginning at a_1 and a_2 , ending at b_1 and b_2 , and staying at least ε apart throughout their trajectories (see Figure 1). Here ε is assumed to be small relative to the lengths of the arms.

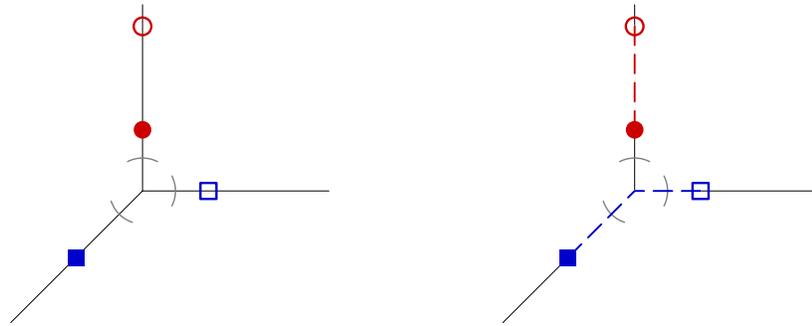


Figure 1: Left: The Y-graph with an ε -neighborhood of the vertex. The solid circle and square represent the initial points a_1 and a_2 ; the empty circle and square represent the destination points b_1 and b_2 . Right: A path in F_ε from a to b .

More formally, we define the *ordered two-point ε -configuration space* of a graph G with path metric d :

$$F_\varepsilon := F_\varepsilon(G, 2) = \text{Conf}_\varepsilon(G, 2) = \{(a_1, a_2) \in G \times G \mid d(a_1, a_2) \geq \varepsilon\},$$

and the *unordered two-point configuration space* of G :

$$C := C(G, 2) = F(G, 2)/\mathbb{Z}_2 = \{(a_1, a_2) \in G \times G \mid a_1 \neq a_2\} / [(a_1, a_2) \sim (a_2, a_1)].$$

The lack of ε in the unordered case will be explained shortly.

The product space $G \times G$ may be endowed with various natural metrics, any of which is inherited by F_ε and C . We focus our attention on the ℓ_1 and ℓ_2 metrics. Writing $a = (a_1, a_2)$ and $b = (b_1, b_2)$ for points $a, b \in G \times G$, the ℓ_1 and ℓ_2 metrics on $G \times G$ are defined:

$$\begin{aligned} \ell_1(a, b) &= |d(a_1, b_1) + d(a_2, b_2)|, \\ \ell_2(a, b) &= \sqrt{d(a_1, b_1)^2 + d(a_2, b_2)^2}. \end{aligned}$$

In Figure 1, it is easy to see that there is a unique geodesic from a to b in the ℓ_2 metric, obtained when the particles traverse their indicated paths at the appropriate relative speed.

Remark If G is a tree not homeomorphic to an interval, the ordered two-point configuration space $F(G, 2)$ is not geodesically complete in the ℓ_1 nor ℓ_2 metric. To see this, let v be a vertex of degree ≥ 3 with $a_2 = b_2 = v$ and such that a_1 and b_1 lie on different edges adjacent to v , both at equal distance d from v . By moving the second particle slightly onto a third edge, we can obtain a path in $F(G, 2)$ from (a_1, a_2) to (b_1, b_2) of length arbitrarily close to $2d$, but there is no path of length $2d$ in either metric. Therefore we replace $F(G, 2)$ by the geodesically complete space F_ε , which is a (\mathbb{Z}_2 -equivariant) deformation retract of $F(G, 2)$ and hence topologically equivalent to it (see Lemma 3).

Remark If G is a tree with a vertex v of degree ≥ 4 , the unordered two-point configuration space C is not geodesically complete in the ℓ_2 metric. To see this, suppose that a_1, a_2, b_1 , and b_2 lie on different edges adjacent to v , all equidistant from v at distance $\frac{\sqrt{2}}{2}d$. As above, there is a sequence of paths in C with length approaching $2d$, but there is no path of length $2d$. On the other hand, C is geodesically complete in the ℓ_1 metric for any tree G . The key difference is that the speed of travel is irrelevant in the ℓ_1 metric; only the total distance traveled by the particles is important. Thus there is no penalty for choosing a motion which first keeps one particle fixed while the other moves to its destination, and then moves the remaining particle to its destination. This strategy eliminates any issues caused by collisions which could potentially occur when both particles move simultaneously.

Our main results give the values of $\text{GC}(F_\varepsilon, \ell_i)$, $i \in \{1, 2\}$, and of $\text{GC}(C, \ell_1)$ for certain graphs G .

Theorem 1 *Let G be a star graph, with $k \geq 3$ edges emanating from a single vertex, and let F_ε be the ordered two-point ε -configuration space of G . Then*

- (a) *If $k = 3$, $\text{GC}(F_\varepsilon, \ell_i) = \text{TC}(F_\varepsilon) = 1$, for $i \in \{1, 2\}$.*
- (b) *If $k \geq 4$, $\text{GC}(F_\varepsilon, \ell_i) = \text{TC}(F_\varepsilon) = 2$, for $i \in \{1, 2\}$.*

Theorem 2 *Let G be a tree and let C be the unordered two-point configuration space of G . Then*

- (a) *If G is homeomorphic to an interval, then $\text{GC}(C, \ell_1) = \text{TC}(C) = 0$.*
- (b) *If G is the Y-graph, $\text{GC}(C, \ell_1) = \text{TC}(C) = 1$.*
- (c) *Otherwise, $\text{GC}(C, \ell_1) = \text{TC}(C) = 2$.*

For a tree G , the topological complexities of the configuration spaces $F(G, 2)$ and $C(G, 2)$ are well-known (see [6]); the lower bounds for GC follow from this and the fact that $F(G, 2)$ deformation retracts to F_ε (see Lemma 3). We show the upper bounds by constructing explicit geodesic motion planners. The general strategy for constructing geodesic motion planners on a metric space (X, g) is to analyze the structure of the *total cut locus*, which we define as follows:

$$\mathcal{C} = \{(a, b) \in X \times X \mid \text{there exist multiple geodesic paths from } a \text{ to } b\}.$$

On the complement of \mathcal{C} , the geodesic is uniquely determined, and there is a well-defined map $\mathcal{C}^c \rightarrow GX$ which sends a pair of points to the unique geodesic connecting them. If X is a proper metric space, this map is continuous (see [1], Corollary I.3.13), yielding a GMPR on \mathcal{C}^c . Thus to construct a geodesic motion planner on a metric space (X, g) , it suffices to find GMPRs over the total cut locus \mathcal{C} .

In the ℓ_1 metric on C , the only pairs of points in \mathcal{C}^c are those for which there is a path from the starting configuration to the ending configuration in which one particle does not move. Thus \mathcal{C} contains almost all of $C \times C$. In particular, the rule on \mathcal{C}^c does not contribute much to the geodesic motion planner, and so the proof of Theorem 2 still requires partitioning $C \times C$ into the appropriate number of ENRs, on each of which there is a continuous choice of a geodesic. This is the content of Section 3.

On the other hand, an essential part of the proof of the ℓ_2 case of Theorem 1 is a careful analysis of \mathcal{C} . To illustrate the difficulties which may arise when determining the total cut locus, we offer a second example, depicted in Figure 2. Due to the condition that the particles must stay distance ε apart throughout their trajectories, the second (square) particle in Figure 2 must move away from its destination to let the first particle pass.

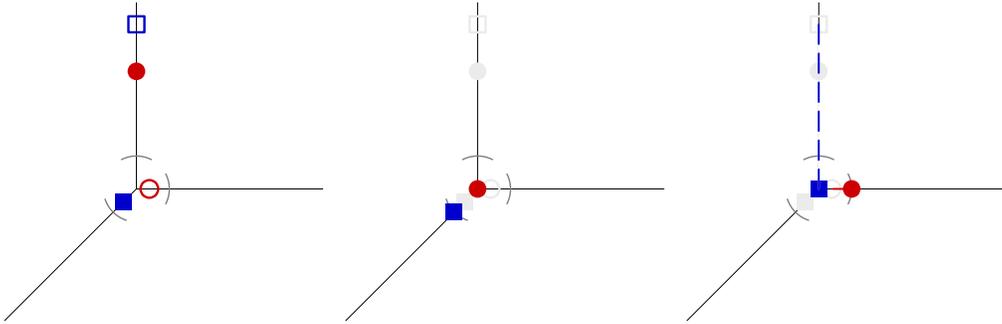


Figure 2: With a and b given as in the left image, the geodesic will travel the path which has intermediate stages depicted in the middle and right images. After the particles arrive at the configuration in the right image, they move directly to their destinations.

Observe that there is a second feasible path from a to b , in which the second particle moves into the right arm of the Y-graph and the first particle moves into the bottom; next the second particle moves through the vertex until both particles can travel directly to their destinations. Although this path appears almost obviously longer, note that in the limit, taken as the points a_2 and b_1 approach the vertex, the two paths have equal length. Thus it is important to carefully determine the geodesics based on the relative locations of a and b . In particular, we will see that most points of the total cut locus require an *orientation switch*, in which an empty arm is used to allow the particles to reconfigure (see Figure 4, right). When an orientation switch is necessary, there are two possible considerations: which particle will allow the other to pass, and which arms the particles will use to execute the orientation switch. These considerations are discussed in greater depth in Section 2.2.

Remark It is interesting to note, in the case such that G is homeomorphic to an interval and C is considered with the ℓ_1 metric, that both the total cut locus and its complement are nonempty, yet there exists a single GMPR on the entirety of $C \times C$. This serves as a counterexample to the somewhat intuitive notion that, by definition of \mathcal{C} , a GMPR on \mathcal{C}^c should not be compatible with one on \mathcal{C} .

Acknowledgments. We are grateful to an anonymous referee for suggestions which corrected some inconsistencies and simplified several proofs.

2 The proof of Theorem 1: Geodesic motion planning on (F_ε, ℓ_i)

To begin, we exclusively consider the ℓ_2 metric on F_ε . In Section 2.4 we make appropriate modifications to the ℓ_2 case to establish the ℓ_1 case. The main difficulty in the proof of the ℓ_2 case of Theorem 1 is to determine the total cut locus of F_ε , so that geodesic motion planning rules (GMPRs) can be constructed and the upper bound on GC can be established. Before launching into this analysis, we formalize the lower bound in terms of $\text{TC}(F(G, 2))$.

Lemma 3 *Let G be a tree and suppose that $\varepsilon > 0$ is not larger than the length of any edge. Then there is a $(\mathbb{Z}_2$ -equivariant) deformation retraction from $F(G, 2)$ to F_ε .*

Proof Let $(a, b) \in F(G, 2)$ with $\mathbf{d}(a, b) \leq \varepsilon$. If a or b is a vertex v , move the other one along its edge to distance ε from v . If a and b lie on the same edge, move them apart uniformly until they are at distance ε from one another. If this motion causes one of them to reach a vertex v , stop it at the vertex and move the other one to distance ε from v . If a vertex v lies between a and b , move them apart with speed proportional to their distances to v , until they are at distance ε from one another. \square

The remainder of this section is dedicated to establishing the upper bounds on GC. We begin by introducing a planar representation of the configuration space F_ε which is convenient for depicting certain paths in F_ε . Consider $(a, b) \in F_\varepsilon \times F_\varepsilon$, so that a_1, a_2, b_1 , and b_2 each lie on some open arm of the Y-graph. We define

$$Z := Z(a, b) = \{c = (c_1, c_2) \in F_\varepsilon \mid c_i \text{ is in the same arm as } a_i \text{ or } b_i\}.$$

We emphasize that the set Z depends on the points a and b , hence will change based on the locations of a_1, a_2, b_1 , and b_2 .

Consider particles a_1 and a_2 in the top arm of the Y-graph, b_1 in the right arm, and b_2 in the bottom arm (see Figure 3). In this example, Z consists of points $c \in F_\varepsilon$ such that c_1 is in the top or right arm, and c_2 is in the top or bottom arm.

The right side of Figure 3 depicts a *representation* of the set Z . The choice of representation is indicated by the directed arcs labeled x and y in the left image. These arcs define the meaning and orientation of the x and y axes in the right image. In this case, the x -axis *always indicates the position of the first particle*, as follows: the negative x -axis represents the negative distance from the vertex to the first particle, assuming that the first particle is in the top arm; the positive x -axis represents the

distance from the vertex to the first particle, assuming that the first particle is in the right arm. The negative (positive) y -axis is similar; it always indicates distance for the second particle, with respect to the top (bottom) arm. The interior of the rectangular strip is *forbidden*; for any point inside, the distance between the particles is less than ε , hence the point is not an element of F_ε .

If the first particle enters the bottom arm, or if the second particle enters the right arm, the configuration of the particles ceases to be an element of Z , hence is not representable on the axes shown. Thus we sometimes refer to elements/subsets of Z as *representable*.

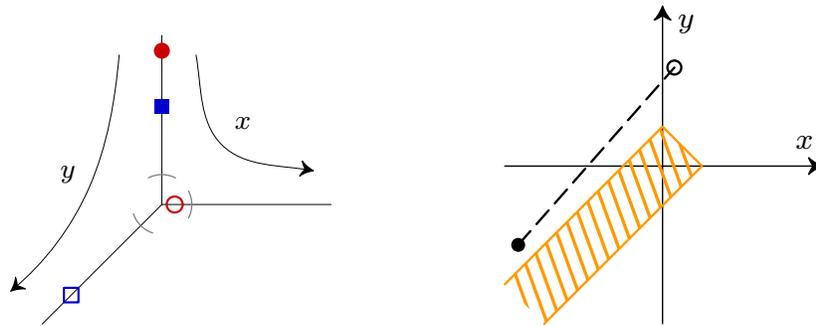


Figure 3: Left: An element of $F_\varepsilon \times F_\varepsilon$ such that a_1 and a_2 share the top arm; the directed arcs indicate the meaning and orientation of the axes on the right side. Right: Representations of the initial point a (solid) and destination point b (hollow), along with the unique geodesic connecting them.

In the next sections, we rely primarily on geometric intuition to present the GMPRs, and we use the representation only as a visual aid to depict the geodesic between a and b . We show in Section 2.5 that the notion of representability can be used to formalize these intuitive statements. For example, by noting that the map taking (Z, ℓ_2) to (\mathbb{R}^2, ℓ_2) is an isometry, it is often convenient to argue the uniqueness of geodesics in F_ε by using the uniqueness of geodesics in the corresponding representation. As a related example, in the left side of Figure 1, it is clear that a geodesic path from a to b has the property that if the particles follow the geodesic, the first particle will not leave the top arm, and the second particle will not enter the top arm. We tacitly assume such statements in the following sections before verifying them in Section 2.5.

2.1 Geodesic motion planning on (F_ε, ℓ_2) for the Y-graph

We begin our investigation in the case of $k = 3$ arms, though the methods here generalize to higher values of k . In this subsection, G always refers to the Y-graph, and F_ε always

refers to the 2-point ordered ε -configuration space of G , considered with the ℓ_2 metric. We define the following partition of $F_\varepsilon \times F_\varepsilon$.

Definition 3 We partition $F_\varepsilon \times F_\varepsilon$ into the following subsets. In the notation below, a subscript indicates the exact number of arms which the points occupy; a superscript indicates whether the orientation of the initial configuration a agrees with (“+”) or disagrees with (“-”) the orientation of the target configuration b . In particular:

- (a) X_1^+ (resp. X_1^-) consists of points $(a, b) \in X$ such that a_1, a_2, b_1, b_2 all lie on a single arm of Y (vertex included), and such that the relative orientation of the starting points a_1 and a_2 agrees with (resp. disagrees with) that of b_1 and b_2 .
- (b) X_2^+ (resp. X_2^-) consists of points $(a, b) \in X$ such that a_1, a_2, b_1, b_2 all lie on exactly two arms of Y (vertex included), and such that the relative orientation of the starting points a_1 and a_2 agrees with (resp. disagrees with) that of b_1 and b_2 ; see Figure 4, left (resp. right).
- (c) X_3 consists of points $(a, b) \in X$ such that each of the three arms of Y (vertex excluded) contains at least one of a_1, a_2, b_1, b_2 (see Figures 3, 5, and 6).

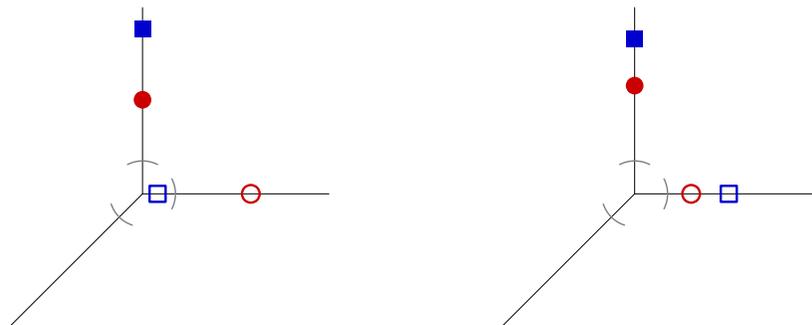


Figure 4: Left: An example configuration in X_2^+ ; the particles can travel directly to their destination points. Right: An example configuration in X_2^- ; the particles must undergo an orientation switch before traveling to their destination points.

Proposition 4 The sets X_1^+, X_2^+ , and X_3 are disjoint from the total cut locus \mathcal{C} .

Proof of Proposition 4 For points in X_1^+ and X_2^+ , there is a unique geodesic taking the direct path from a to b .

By the definition of X_3 , each of the three arms contains a point, hence two points must share one arm. By utilizing symmetries, we may assume that a_1 shares an arm with

another point. Indeed, in the case that a_2 shares an arm with b_1 (resp. b_2), the argument follows from the case in which a_1 and b_2 (resp. b_1) share an arm, by the \mathbb{Z}_2 -symmetry swapping a and b . In case the two destination points b_1 and b_2 share an arm, the argument follows by reversing the geodesic path in the case that a_1 and a_2 share an arm. Thus we have three cases to consider: that a_1 shares an arm with a_2 , b_2 , or b_1 .

These three cases are depicted with their representations in Figures 3, 5, and 6, respectively. In each case the particles can move directly to their destinations and do not enter unused arms. We reiterate that this intuitive proof can be formalized using the notion of representability, first by showing that any geodesic from a to b must be representable (see Lemma 11), and then by observing that the representation map $(Z, \ell_2) \rightarrow (\mathbb{R}^2, \ell_2)$ is an isometry, so that the unique geodesic between the representations of a and b corresponds to a unique geodesic between a and b . \square

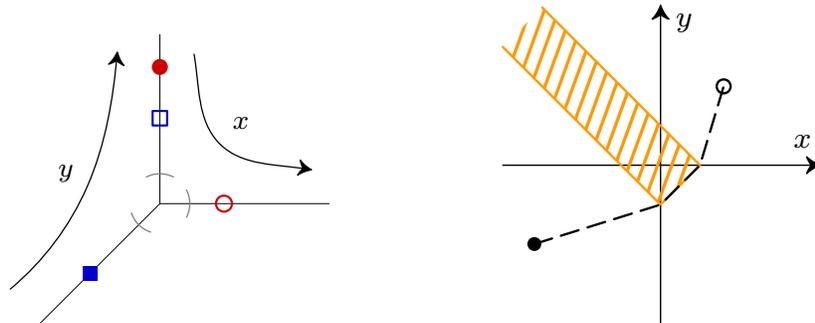


Figure 5: Left: An element of X_3 such that a_1 and b_2 share the top arm. Right: Representations of a , b , and the unique geodesic connecting them.

2.2 The total cut locus of F_ϵ

We turn our attention to the total cut locus \mathcal{C} of F_ϵ . By Proposition 4, the total cut locus is contained in $X_1^- \cup X_2^-$. If $(a, b) \in X_i^-$, $i = 1, 2$, any path from a to b must undergo an *orientation switch*, in which an empty arm is used to allow the particles to reconfigure; for example, an orientation switch is necessary for the configuration depicted in the right side of Figure 4. If G is a star graph with $k > 3$ arms, then any (a, b) which requires an orientation switch is an element of the total cut locus, because one always needs to designate which empty arm is used for the orientation switch. In the case of the Y graph with $k = 3$ arms, not all of X_2^- is contained in \mathcal{C} . In particular, the arm used for the orientation switch is pre-determined by the relative locations of a

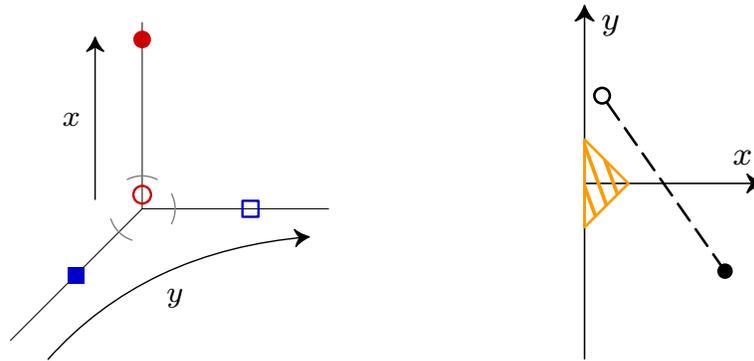


Figure 6: Left: An element of X_3 such that a_1 and b_1 share the top arm. Right: Representations of a , b , and the unique geodesic connecting them.

and b , so (a, b) is only in \mathcal{C} when there is not a preferred particle which moves onto the free arm to let the other pass. We will formalize these ideas now.

As there exists a GMPR on the complement of the total cut locus, the ℓ_2 case of Theorem 1(a) is a consequence of the following:

Proposition 5 *There exists a GMPR on the total cut locus \mathcal{C} .*

As $\mathcal{C} = (X_1^- \cup X_2^-) \cap \mathcal{C}$, we first consider the spaces X_1^- and X_2^- separately.

Lemma 6 *The set X_1^- is a subset of the total cut locus \mathcal{C} , and there exists a GMPR on X_1^- .*

Proof For points in X_1^- , an orientation switch is necessary, and there exist exactly two geodesics from a to b . In particular, order the arms clockwise, with arm 1 at the top, and consider the ordering mod 3. If the points all lie on arm i , there is a unique geodesic for which the second particle uses arm j for the orientation switch, for each $j \neq i$.

To see that there exist no other geodesics, consider the configuration depicted in Figure 7. According to the definition of Z , the only representable points are those for which both particles are in the top arm. There are two natural ways to extend the representable region, depicted either in the left or the center of Figure 7. In either case, the unique geodesic has the representation depicted on the right (although the positive axes have different meaning depending on the choice). If a path is not representable in either of these two representations, then at some time during the trajectory, both particles lie

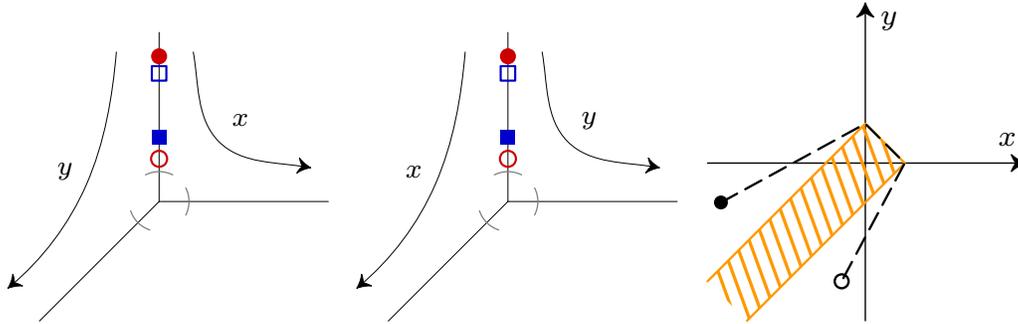


Figure 7: Left and Center: An element of X_1^- with two possible choices of representation. Right: The unique geodesic in either representation.

either in the open right arm or the open bottom arm. Such a path is non-minimizing (see the displayed statement in the proof of Lemma 11).

Now a GMPR on X_1^- can be defined: choose the geodesic for which the particle which starts off closer to the vertex uses arm $i + 1$ and the other particle uses arm $i + 2$ for the orientation switch. \square

All points of X_2^- require orientation switches, but not all points are in the total cut locus. Assume momentarily that a_1 lies on the top arm, that none of a_2 , b_1 , or b_2 lies above a_1 , and that points not on the top arm lie on the bottom arm. Then there are three possible permutations of the a_i and b_i , from top to bottom: $a_1b_2b_1a_2$, $a_1b_2a_2b_1$, and $a_1a_2b_2b_1$ (it is possible that $a_1 = b_2$ or $a_2 = b_1$ or $a_2 = b_2$). Other permutations beginning with a_1 have agreeing orientation and are not elements of X_2^- .

Considering cases dictated by the three possible permutations, as well as the number of a_i and b_i which lie on each arm, we see that there is a unique geodesic in the following cases:

- if the closed top arm (i.e. including the vertex) contains exactly three of the a_i and b_i , or if the closed bottom arm contains exactly three of the a_i and b_i , then a geodesic is uniquely determined. See Figure 8 (left and center) for two possible configurations – other possibilities permute the a_i and b_i but have similar behavior.
- if a_1 and a_2 share the open top arm and b_1 and b_2 share the open bottom arm, there exists a unique geodesic (see Figure 8, right).

Note that in the case of $k > 3$ arms, all such points lie in the total cut locus, because an arm choice is necessary – see Section 2.3 for further discussion.

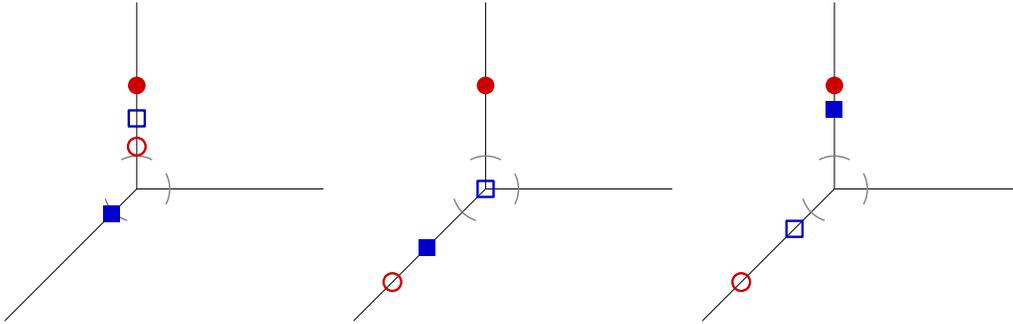


Figure 8: Three elements of X_2^- in the complement of the total cut locus. Left: the circle moves into the right arm first. Center: the square moves into the right arm first. Right: The square moves into the right arm first.

Similar arguments may be made in the cases which do not adhere to our specific assumptions. For example, if b_1 lies above a_1 , there is a symmetry swapping the particles. If either a_2 or b_2 lie above a_1 , there is a symmetry swapping the initial configuration with the final configuration, and the unique geodesic is the reversed geodesic from the above case. Finally, the same arguments can be made with respect to any of the three arms.

Thus it remains to consider the case in which a_1 and b_2 share an open arm, and a_2 and b_1 share a different open arm. As such, we define

$$X_2^{opp} = \{(a, b) \in X_2^- \mid a_i \text{ shares an open arm with } b_{i+1}\}.$$

A point of X_2^{opp} is depicted in Figure 9, left. For $(a, b) \in X_2^{opp}$, every path from a to b must undergo an orientation switch, and paths from a to b may be distinguished based on which particle moves through the vertex first for the orientation switch (i.e. to enter the empty arm so that the other particle may pass). We refer to a path such that the i th particle passes through the vertex first as a *path of type i* .

We claim that for each particle i , there exists a unique length-minimizing type i path. Indeed, any length-minimizing type 1 path must pass through the point $p \in X$ depicted in Figure 9, right. Now (a, p) and (p, b) are in the complement of the total cut locus by Proposition 4, so there are unique geodesics connecting a to p and p to b . Their concatenation is the unique length-minimizing type 1 path from a to b . Similarly, there is a unique length-minimizing type 2 path from a to b .

For at least one value of i , the unique length-minimizing type i path is the minimal geodesic from a to b . Occasionally these two paths have equal length, in which case

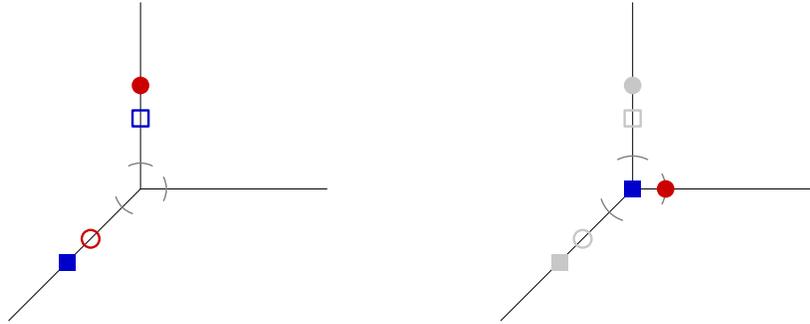


Figure 9: Left: An element of X_2^{opp} . Right: A length-minimizing type 1 path from a to b must pass through the depicted point p .

the point (a, b) is in the total cut locus. Such points are characterized by a certain algebraic condition relating the distances from $a_1, a_2, b_1,$ and b_2 to the vertex, but it is not necessary to determine this condition explicitly. Instead, we simply define

$X_2^{eq} = \{(a, b) \in X_2^{opp} \mid \text{the unique length-minimizing type } i \text{ paths have equal length}\}$, and $X_2^n = X_2^- - X_2^{eq}$. We have shown that X_2^n is disjoint from the total cut locus \mathcal{C} , despite the fact that orientation switches are needed for all points of X_2^- . However, with $k > 3$ arms, every orientation switch requires a choice of empty arm to use for the switch, so X_2^n is part of the total cut locus for star graphs (see Section 2.3).

Lemma 7 *There exists a GMPR on $X_2^{eq} = X_2^- \cap \mathcal{C}$.*

Proof We define the GMPR on X_2^{eq} so that the particle which enters arm i for the orientation switch is the particle which begins on arm $i + 1$, i.e. the arm adjacent to i in the clockwise direction. □

Finally, we prove Proposition 5, that there exists a GMPR on the total cut locus \mathcal{C} .

Proof of Proposition 5 We have exhibited GMPRs on X_1^- and on X_2^{eq} , so it remains to show that these are compatible on the union \mathcal{C} . In particular, we must check that the GMPRs agree at the points of X_1^- which are limit points of X_2^{eq} : these are points such that $a_1 = b_2$ lie at the vertex, or $a_2 = b_1$ lie at the vertex.

We consider a fixed limit point $(a, b) \in X_2^{eq}$; without loss of generality, we assume that a_1 and b_2 lie on the top arm and that a_2 and b_1 lie on the vertex; see Figure 10. Then the GMPR would choose the geodesic for which the second particle moves into the

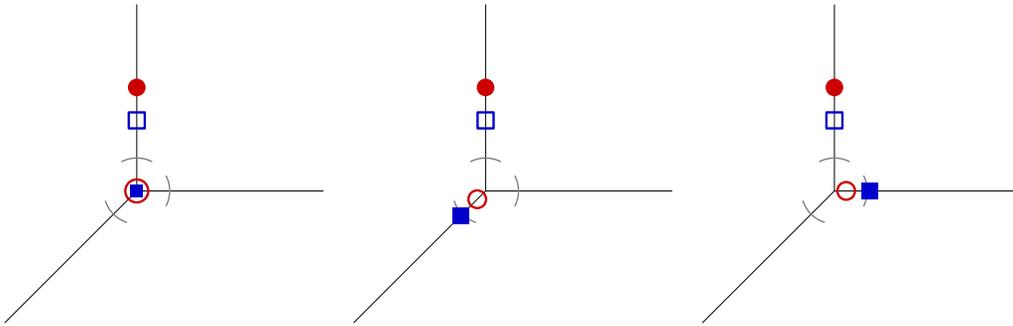


Figure 10: Left: A limit point (a, b) of X_2^{eq} contained in X_1^- . The GMPR indicates that the second particle enters the right arm, the first enters the bottom, and then both move to their destination. Center: A possible point of a sequence converging to (a, b) . The GMPR indicates that the second particle enters the right arm, the first enters the bottom, and then both move to their destination. Right: A possible point of a sequence converging to (a, b) . The GMPR indicates that the second particle stays in the right arm (moving back if necessary), the first particle enters the bottom arm, and then both move to their destination.

right arm and the first particle moves into the bottom arm to let the second particle pass into the top arm. Consider a sequence of points $(a^n, b^n) \in X_2^{eq}$ approaching (a, b) . If for some n , a_2^n and b_1^n lie in the bottom arm, then the GMPR chooses the geodesic for which the second particle moves into the empty (right) arm and the first moves into the bottom arm. If for some n , a_2^n and b_1^n lie in the right arm, then the GMPR chooses the geodesic for which the first particle enters the empty (bottom) arm (and if necessary, the second particle moves within the right arm to the boundary of the ε -ball around the vertex) to let the second particle pass. Thus in the limit (a, b) , the limiting path is exactly the chosen geodesic at (a, b) . \square

Remark To prove Theorem 1(a) in the ℓ_2 case we exhibited two GMPRs: one on the total cut locus \mathcal{C} and one on its complement \mathcal{C}^c . Note that no GMPR on either X_1^- or X_2^{eq} is compatible with a GMPR on \mathcal{C}^c .

2.3 Geodesic motion planning on (F_ε, ℓ_2) for star graphs

We are now equipped to study the geodesic motion planning problem on a star graph G , with $k > 3$ arms emanating from a single vertex. The techniques are similar to those for the case $k = 3$, and many of the results from the previous section will be used.

Let F_ε be the 2-point ordered ε -configuration space of G . In addition to the subsets of $F_\varepsilon \times F_\varepsilon$ defined in Definition 3, we consider the set X_4 consisting of points $(a, b) \in$

$F_\varepsilon \times F_\varepsilon$ such that there exist four arms of G occupied. Given $(a, b) \in X_4$, we say that a path from (a, b) is *type i* if the i th particle passes through the vertex first. Recall that type i paths were defined in Section 2.2 for elements of X_2^{opp} ; however, for $(a, b) \in X_2^{opp}$, the particle which moves through the vertex first does so to allow the other to pass, whereas for $(a, b) \in X_4$, the particle which moves through the vertex first may travel directly to its destination.

For each particle i , there is a unique length-minimizing path among those of type i . At least one of these is a geodesic. As in the case of X_2 , we define:

$X_4^{eq} = \{(a, b) \in X_4 \mid \text{the unique length-minimizing type } i \text{ paths have equal length}\}$,
 and $X_4^n = X_4 - X_4^{eq}$ (see Figure 11). Thus by definition, $X_4^{eq} = X_4 \cap \mathcal{C}$.

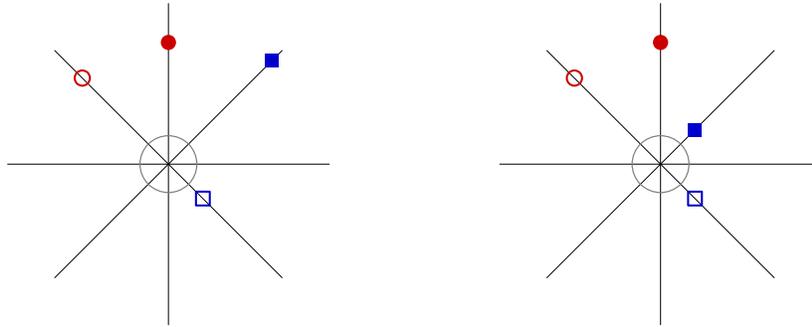


Figure 11: Left: An element of X_4^n . Right: An element of X_4^{eq} .

Recall the discussion prior to Lemma 7: with $k > 3$ arms, X_2^n belongs to the total cut locus, since any orientation switch requires a choice of empty arm, and, unlike the situation on the Y-graph, there are now multiple such choices.

Now the sets X_4^n and X_4^{eq} , together with $X_3 \cup X_2^+ \cup X_2^{eq} \cup X_2^- \cup X_1^+ \cup X_1^-$, partition $F_\varepsilon \times F_\varepsilon$. The following result establishes three GMPRs on $F_\varepsilon \times F_\varepsilon$.

Proposition 8 *Let G be a star graph with k arms, $k > 3$. The complement of the total cut locus is given by $X_4^n \cup X_3 \cup X_2^+ \cup X_1^+$, and there exist GMPRs on*

- (a) $X_4^n \cup X_3 \cup X_2^+ \cup X_1^+ \cup X_2^{eq}$
- (b) $X_1^- \cup X_4^{eq}$
- (c) X_2^n ,

hence $\text{GC}(F_\varepsilon) \leq 2$.

We compare to the situation of the Y-graph. There, the sets X_4^n and X_4^{eq} do not exist, the complement of the total cut locus is $X_3 \cup X_2^+ \cup X_1^+ \cup X_2^n$, and we exhibited a GMPR on $X_1^- \cup X_2^{eq}$. With more than three arms, one cannot define compatible GMPRs on X_1^- and X_2^{eq} . So we instead show that the GMPR on X_2^{eq} is compatible with that on the complement of the total cut locus, that X_4^{eq} admits a GMPR compatible with that on X_1^- , and that X_2^n admits a GMPR.

Proof We first remark that $X_4^n \cup X_2^+ \cup X_1^+$ are in the complement of the total cut locus by definition, and X_3 is in the complement of the total cut locus by Proposition 4 and because geodesics do not enter unused arms unless an orientation switch is necessary (this can be formalized in the same manner as Lemma 11).

Rule 1. As the complement of the total cut locus, there is a GMPR on $X_4^n \cup X_3 \cup X_2^+ \cup X_1^+$. The GMPR on X_2^{eq} can be defined in the case of $k > 3$ arms: designate arm i for the orientation switch, where i is the smallest index such that arm i is empty but arm $i + 1$ is occupied, and use the particle beginning on arm $i + 1$ for the switch.

To show that the GMPR on X_2^{eq} is compatible with that of the complement of the total cut locus, we can show that the closures of the sets are disjoint.

The closure of X_2^{eq} is contained in $X_2^{eq} \cup X_1^-$. If a limit point (a, b) of $X_4^n \cup X_3$ uses only two arms, then one of a_i or b_i must lie at the vertex, hence $(a, b) \notin X_2^{eq}$. Finally, $X_2^+ \cup X_1^+$ is closed.

Rule 2. To motion plan on X_4^{eq} , observe that X_4^{eq} is a union of finitely many disjoint open sets, defined by indicating which arms contain which particles. It is enough to motion plan on one of these sets, and the GMPR may be defined simply by always letting the first particle through the vertex first.

The GMPR on X_1^- is defined in Lemma 6; the same argument holds verbatim for star graphs with $k > 3$ arms.

Now X_1^- is closed, and limit points (a, b) of X_4^{eq} occupy at least two arms, so $(a, b) \notin X_1^-$. Thus the GMPRs are compatible on the union.

Rule 3. Finally, for points $(a, b) \in X_2^n$, the particle which must move first for the orientation switch is pre-determined, and so we only must choose which arm is used for the switch. We define the GMPR to use the empty arm of lowest index. \square

With the ℓ_2 analysis complete, we now turn our attention to the ℓ_1 metric on F_ε .

2.4 The ℓ_1 metric on F_ε and the proof of Theorem 1

The ℓ_2 -case of Theorem 1 follows from Propositions 5 and 8, and so it remains to consider the ℓ_1 -case. To distinguish between geodesics in various metrics, we will use the terminology “ ℓ_i -geodesic” to refer to a geodesic in (F_ε, ℓ_i) .

Remark Because the ℓ_1 and ℓ_2 metrics induce the same topology on F_ε , the induced compact-open topologies on the path space PF_ε are equivalent. In particular, if $E \subset F_\varepsilon \times F_\varepsilon$, and if $s : E \rightarrow PF_\varepsilon$, then s is continuous in the ℓ_1 -induced compact-open topology on PF_ε if and only if s is continuous in the ℓ_2 -induced compact-open topology on PF_ε . Thus it is not necessary to distinguish notions of continuity when considering geodesics in different metrics.

To adapt the ℓ_2 argument to ℓ_1 , the general idea is as follows: we will use essentially the same partition as in the ℓ_2 case, apart from small modifications to the decompositions of X_2 and X_4 . In particular, the sets X_2^{eq} , X_2^n , X_4^{eq} , and X_4^n , which were previously defined in terms of the ℓ_2 metric, will be redefined in terms of the ℓ_1 metric. The actual GMPRs, which by definition map a point (a, b) to an ℓ_1 -geodesic from a to b , will always follow *local* ℓ_2 -geodesics; that is, paths which locally, but perhaps not globally, minimize ℓ_2 -distance.

In particular, we recall that most points $(a, b) \in F_\varepsilon$ lie in the ℓ_1 total cut locus: unless one particle stays fixed throughout a minimizing trajectory, one can perturb any ℓ_1 -length-minimizing path, by changing the speed of one of the particles, without changing the ℓ_1 -length. Even pausing one particle does not penalize the ℓ_1 -length; as long as a motion from a to b does not involve unnecessary backtracking, the motion corresponds to an ℓ_1 -geodesic. It follows that if an ℓ_2 -geodesic γ follows a direct path from a to b , as is the case when $(a, b) \in X_1^+ \cup X_2^+$, then γ is also an ℓ_1 -geodesic. More generally, we have the following.

Lemma 9 *Let $(a, b) \in X_1 \cup X_2^+ \cup X_3 \cup (X_2^- - X_2^{opp})$ and let γ be an ℓ_2 -geodesic from a to b . Then γ is an ℓ_1 -geodesic.*

Proof For points $(a, b) \in X_1^+ \cup X_2^+$, the uniquely-determined ℓ_2 -geodesic γ moves a directly to b , hence is an ℓ_1 -geodesic. For points $(a, b) \in X_3$, γ is also uniquely determined, and although such geodesics may contain a brief motion which moves a particle farther from its destination, such motions are strictly necessary. To see this, observe in Figure 2 that any ℓ_1 -geodesic must pass through the points p and q depicted, respectively, in the center and right images, and so the length of any ℓ_1 -geodesic

from a to b is equal to the sum $\ell_1(a,p) + \ell_1(p,q) + \ell_1(q,b)$. The ℓ_2 -geodesic γ also passes through p and q . Moreover, the three points (a,p) , (p,q) , and (q,b) all lie in X_2^+ , so γ is the concatenation of the three uniquely-determined ℓ_2 -geodesics connecting these points, each of which is also an ℓ_1 -geodesic. Thus γ also has ℓ_1 -length $\ell_1(a,p) + \ell_1(p,q) + \ell_1(q,b)$ and hence is an ℓ_1 -geodesic.

Similar arguments apply to points $(a,b) \in X_1^- \cup (X_2^- - X_2^{opp})$. For such (a,b) , there is an ℓ_2 -geodesic from a to b corresponding to each choice of arm(s) used for the orientation switch. However, once the arm choices are fixed, one may again determine point(s) through which any ℓ_1 -geodesic must pass, and similar arguments may be made. \square

Lemma 9 establishes that all ℓ_2 -geodesics are ℓ_1 -geodesics, except perhaps ℓ_2 -geodesics from a to b for points $(a,b) \in X_2^{opp} \cup X_4$. In the next example we give explicit points of X_2^{opp} and of X_4 such that Lemma 9 fails. We recall that for $(a,b) \in X_2^{opp} \cup X_4$, a path from a to b is *type i* if the i th particle passes through the vertex first. For $(a,b) \in X_2^{opp}$, and for each particle i and empty arm j , there is a unique ℓ_2 -length minimizing path among those of type i using arm j for the orientation switch. For $(a,b) \in X_4$, no orientation switch is necessary, but for each particle i , there is a unique type i ℓ_2 -length minimizing path.

Example 1 Let G be a star graph with $k > 3$ arms with $\varepsilon = 2$. Consider $(a,b) \in X_4$, such that $|a_1| = 1$, $|a_2| = 2$, $|b_1| = 2$, and $|b_2| = 5$ (see Figure 12, left); here $|\cdot|$ represents distance from the vertex. Let $\gamma_{i,j}$ represent any path which achieves the minimum ℓ_j -length among those of type i . Let $d_{i,j}$ represent the length of $\gamma_{i,j}$. Then

$$d_{1,1} = 10, \quad d_{2,1} = 12, \quad d_{1,2} = 1 + \sqrt{8} + 5 \approx 8.82, \quad \text{and} \quad d_{2,2} = \sqrt{5} + \sqrt{8} + \sqrt{13} \approx 8.67.$$

In particular, the unique ℓ_2 -geodesic $\gamma_{2,2}$ is not an ℓ_1 -geodesic, but the path $\gamma_{1,2}$, which is not an ℓ_2 -geodesic but does minimize ℓ_2 -length among type 1 paths, is an ℓ_1 -geodesic. A similar phenomenon occurs for $(a,b) \in X_2^{opp}$ with $|a_1| = 1$, $|a_2| = 2$, $|b_1| = 5$, and $|b_2| = 2$ (see Figure 12, right).

We reiterate the intuitive motion planning strategy in the ℓ_1 -case: if $(a,b) \in X_2^{opp} \cup X_4$ has the property that all ℓ_1 -minimizing geodesics are of type i , then choose an ℓ_1 -geodesic $\gamma_{i,2}$ (which may not be an ℓ_2 -geodesic). Now we will formalize these rules, defined on the following sets:

$$(X_2^q)_1 = \{(a,b) \in X_2^{opp} \mid \text{there exist } \ell_1\text{-geodesics of types 1 and 2}\},$$

$$(X_4^q)_1 = \{(a,b) \in X_4 \mid \text{there exist } \ell_1\text{-geodesics of types 1 and 2}\}.$$

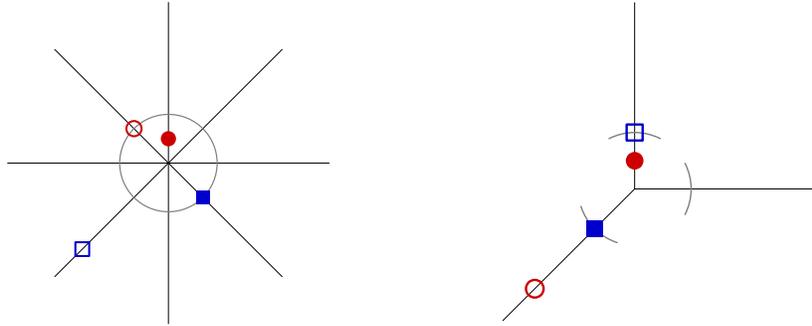


Figure 12: Elements of X_4 (left) and X_2^{opp} (right) such that the unique ℓ_2 -geodesic exhibits qualitative behavior distinct from any ℓ_1 -geodesic.

Note that $(X_2^{eq})_1$ and $(X_4^{eq})_1$ are defined in the same manner as X_2^{eq} and X_4^{eq} , with “ ℓ_1 -geodesic” in place of “ ℓ_2 -geodesic”. Also analogous to the ℓ_2 case, we define $(X_2^n)_1 = X_2^- - (X_2^{eq})_1$ and $(X_4^n)_1 = X_4 - (X_4^{eq})_1$.

We define the following GMPRs on these new sets. We note that, qualitatively, all rules exactly match those in the ℓ_2 case, as presented in Sections 2.2 and 2.3; the only difference is that the actual paths are only locally ℓ_2 -length minimizing and may not be ℓ_2 -geodesics.

Lemma 10 *There exist GMPRs on the sets $(X_2^{eq})_1$, $(X_2^n)_1$, $(X_4^{eq})_1$, and $(X_4^n)_1$.*

Proof Recall that $\gamma_{i,j}$ represents any path which achieves the minimum ℓ_j -length among those of type i . For $(a, b) \in X_4$, $\gamma_{i,2}$ is unique. For $(a, b) \in X_2^-$, we also use the notation $\gamma_{i,2,k}$ to refer to the unique ℓ_2 -length minimizing path from a to b among those type i paths which use leg k for the orientation switch.

On $(X_4^{eq})_1$, there are ℓ_1 -geodesics of type 1 and 2; among them we choose $\gamma_{1,2}$.

On $(X_4^n)_1$, there are only ℓ_1 -geodesics of one type i ; among them we choose $\gamma_{i,2}$.

On $(X_2^{eq})_1$, there are ℓ_1 -geodesics of type 1 and 2; among them we choose $\gamma_{i,2,k}$, where k is the smallest index such that arm k is empty but arm $k + 1$ is occupied, and i is the particle on arm $k + 1$.

On $(X_2^n)_1$, there are only ℓ_1 -geodesics of one type i ; among them we choose $\gamma_{i,2,k}$, where k is the unused arm of least index. □

When studying compatibility issues in the following proof of Theorem 1, it is important to note that for $G = Y$ and $(a, b) \in (X_2^n)_1 - X_2^{opp}$, the path chosen by the above rule is

the unique ℓ_2 -geodesic, which is an ℓ_1 -geodesic by Lemma 9. Also by Lemma 9, all ℓ_2 geodesics on $X_1 \cup X_2^+ \cup X_3$ are ℓ_1 -geodesics, so we define the GMPRs on these sets to use the same geodesics as used in the ℓ_2 case.

We are now prepared to formalize the proof of Theorem 1 in both the ℓ_1 and ℓ_2 cases.

Proof of Theorem 1 All lower bounds on GC follow from the known values of TC, together with Lemma 3. The upper bounds for the ℓ_2 case for parts (a) and (b) follow, respectively, from Propositions 5 and 8. In the ℓ_1 case, it remains to study the compatibility of the GMPRs above. We emphasize that the rules are qualitatively identical to those in the ℓ_2 case, so there are only a few compatibility issues to check.

(a) Let $G = Y$. We claim that there exist two GMPRs, using the same partition as in the ℓ_2 case, except for the above replacements. In particular, there are rules on the sets:

- $X_3 \cup X_2^+ \cup X_1^+ \cup (X_2^n)_1$
- $X_1^- \cup (X_2^{eq})_1$.

The rules on X_1^- and $(X_2^{eq})_1$ have the same qualitative behavior as in the ℓ_2 case; the verification that these are compatible is the same as in the proof of Proposition 5.

The GMPR on $X_3 \cup X_2^+ \cup X_1^+$ uses the unique ℓ_2 -geodesic and is well-defined by Lemma 9. We only must show that the GMPR on $(X_2^n)_1$ is compatible with that on $X_3 \cup X_2^+ \cup X_1^+$. Note that this was not an issue in the ℓ_2 case as all sets were in the complement of the ℓ_2 total cut locus.

To show compatibility, first observe that $X_2^+ \cup X_1^+$ is closed, and limit points of $(X_2^n)_1$ are contained in $(X_2^n)_1 \cup X_1^- \cup (X_2^{eq})_1$. A limit point (a, b) of X_3 may be an element of $(X_2^n)_1$, but it cannot be an element of X_2^{opp} since at least one of $a_1, a_2, b_1, \text{ or } b_2$ lie at the vertex. Since the GMPR on $(X_2^n)_1 - X_2^{opp}$ uses the unique ℓ_2 -geodesic from a to b (see the comment below the proof of Lemma 10), as does the rule on X_3 , the rules are compatible.

(b) Let G be a star graph with $k > 3$ arms. We claim that there are rules on the sets:

- $(X_4^n)_1 \cup X_3 \cup X_2^+ \cup X_1^+ \cup (X_2^{eq})_1$
- $X_1^- \cup (X_4^{eq})_1$
- $(X_2^n)_1$.

Once again, the second and third rules have the same qualitative behavior as in the ℓ_2 case; the verification that these are compatible is the same as in the proof of Proposition 8.

For the first rule, the GMPRs on $(X_4^n)_1$ and $(X_2^{eq})_1$ are defined above, and the GMPR on $X_3 \cup X_2^+ \cup X_1^+$ uses the unique ℓ_2 -geodesic. Furthermore, $(X_2^{eq})_1$ and $(X_4^n)_1 \cup X_3 \cup X_2^+ \cup X_1^+$ do not share any limit points (analogous to the proof of Proposition 8). Thus we only must show that the GMPR on $(X_4^n)_1$ is compatible with that on $X_3 \cup X_2^+ \cup X_1^+$. Note that this was not an issue in the ℓ_2 case as all sets were in the complement of the ℓ_2 total cut locus.

Suppose that $(a, b) \in X_3 \cup X_2^+$ is a limit point of $(X_4^n)_1$. We claim that if a point $(a', b') \in (X_4^n)_1$ is sufficiently close to (a, b) , then (a', b') is in the complement of the ℓ_2 total cut locus. This would establish that the GMPRs at (a', b') and (a, b) both use unique ℓ_2 -geodesics, hence the rules are compatible.

Observe that at least one (and at most two) of $a_1, a_2, b_1,$ and b_2 lie at the vertex; we assume a_1 is at the vertex, possibly shared with b_2 . The GMPR indicates that (a, b) should use the unique ℓ_2 -geodesic γ from a to b ; this geodesic γ has the property that the first particle will immediately move down the arm which contains its destination point b_1 (there is nothing to be gained by waiting at the vertex nor by moving onto another arm). Thus γ may be considered a type 1 path, since the first particle moves through the vertex before the second particle. Any type 2 path (i.e. a path such that the first particle moves onto an empty arm to allow the second particle to go through the vertex “first”) is strictly longer. Therefore a sufficiently nearby point $(a', b') \in (X_4^n)_1$ will also have the property that any ℓ_1 - or ℓ_2 -geodesic is type 1, hence the ℓ_2 -geodesic is unique. \square

2.5 Representations of configurations in F_ε

In previous sections we regularly appealed to intuition regarding the existence and uniqueness of geodesics in various cases. Here we formalize one statement using the notion of representability; similar arguments may be made for all such statements.

Lemma 11 *If $(a, b) \in X_3$, then every geodesic from a to b is representable.*

Proof We reference Figure 3 in the following argument.

Consider any path $\gamma : [0, 1] \rightarrow F_\varepsilon$ from a to b , and suppose some portion of γ is not representable. In particular, let $(t_1, t_2) \subset [0, 1]$ be some interval such that each $\gamma(t_i)$ is representable, but no $\gamma(t), t \in (t_1, t_2)$, is representable. Thus at each time t_i , one of the two particles is at the vertex of the Y-graph, and so $\gamma(t_i)$ is represented on one of the four branches of the axes in the right image of Figure 3. We will first show:

If $\gamma(t_1)$ and $\gamma(t_2)$ lie on the same axis branch, then $\gamma|_{(t_1, t_2)}$ is not minimizing.

Indeed, the geodesic in F_ε between two points $\gamma(t_1)$ and $\gamma(t_2)$, whose representations lie on the same axis branch, is representable: the representation is an isometry and the geodesic in the representation is the straight line segment along that axis. Since $\gamma|_{(t_1, t_2)}$ is not representable, it is not minimizing.

Now observe that if $\gamma(t_1)$ is on the positive y -axis, then $\gamma(t_2)$ also must be; since if the second particle is on the bottom arm when the first particle enters the bottom axis, it must stay there until the first particle leaves. The same argument applies to the positive x -axis.

The remaining option is that one point (we assume it is $\gamma(t_1)$) is represented on the negative x -axis, and $\gamma(t_2)$ is represented on the negative y -axis. In this situation, the path $\gamma|_{(t_1, t_2)}$ undergoes an orientation switch. In this case, the second particle enters the right arm while the first particle is on the top arm, then the first particle enters the bottom arm, then the second particle re-enters the top arm, and the first particle moves back to the vertex. However, this path is exactly the same length as the corresponding path in which the roles of the bottom and right arms are swapped in the previous sentence. This latter path *is* representable, so we may redefine γ on (t_1, t_2) with this representable path, without changing its length. In this way, every non-representable path is the same length as some representable path. However, there is a unique representable geodesic connecting a and b , and it does not involve any orientation switches, so it is not the same length as any non-representable path. \square

Together with the fact that the representation map $(Z, \ell_2) \rightarrow (\mathbb{R}^2, \ell_2)$ is an isometry, Lemma 11 is used to verify that elements of X_3 are not in the total cut locus of F_ε .

When $(a, b) \notin X_3$, it is possible that an orientation switch is necessary, and it is possible that some non-representable path and its representable replacement are actually geodesics. This was the case for cut locus points in X_1^- and X_2^- . Nevertheless, in these cases the same argument may be used to show that these geodesics are unique among those of certain types; e.g., among geodesics of type i using arm k for an orientation switch.

3 The proof of Theorem 2: Geodesic motion planning on (C, ℓ_1)

We now turn our attention to the proof of Theorem 2. In case G is homeomorphic to an interval, a single geodesic motion planning rule (GMPR) may be defined explicitly on

C. This establishes part (a) of Theorem 2. For the remainder of this section, G refers to a tree which is not homeomorphic to an interval. We assume that G has no degree-2 vertices.

The main accomplishment is to partition $C \times C$ into three ENRs (two for the Y graph), on each of which there is a continuous choice of a geodesic from the initial configuration to the ending configuration.

We begin with a definition and elementary proposition.

Definition 4 Let Q be a subset of a tree G . The *convex hull* $S(Q)$ of Q is the minimal connected subtree of G containing Q ; equivalently, it is the union of (the images of) all minimizing geodesics connecting points of Q .

If x_1, x_2, x_3 are any three points on a tree G , then the image of the geodesic path from x_2 to x_3 is contained in the union of the paths from x_1 to x_2 and x_1 to x_3 . Therefore $S(Q)$ may equivalently be defined as the union of the images of all geodesic paths from one fixed point of Q to all other points of Q . When Q is a set of four points on G , $S(Q)$ is the union of three paths emanating from a single point, and so we have the following:

Proposition 12 *If Q is a set of four points on a tree G , then $S(Q)$ must be one of the following five types, pictured in Figure 13.*

- Y_1 . *One point of Q is at a vertex, from which paths to the other three are disjoint.*
- Y_2 . *There is a vertex from which there are disjoint paths to three of the points of Q , and the fourth point of Q lies in the interior of one of these paths.*
- X . *There is a vertex from which paths to all four points are disjoint.*
- H . *There are two vertices, from each of which there are disjoint paths to two points of Q , and these four paths are also disjoint from the path between the two vertices.*
- I . *Two points of Q lie inside the path connecting the other two.*

We number the degree-1 vertices of G from 1 to k in clockwise order around the tree. In Figure 14, we depict a tree with numbered vertices, and a selection of four points on it of each of our five types.

If v is a vertex and $w \in G$ is distinct from v , let \overline{vw} denote the path component of $G - \{v\}$ containing w . For example, if v is the degree-4 vertex in the graph in Figure 14, and if w lies on the edge going down and to the left from it, then \overline{vw} contains all points between v and the degree-1 vertices labeled 1, 2, and 3. Another example is

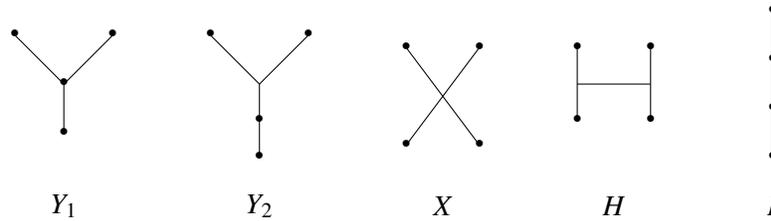


Figure 13: Types of $S(Q)$

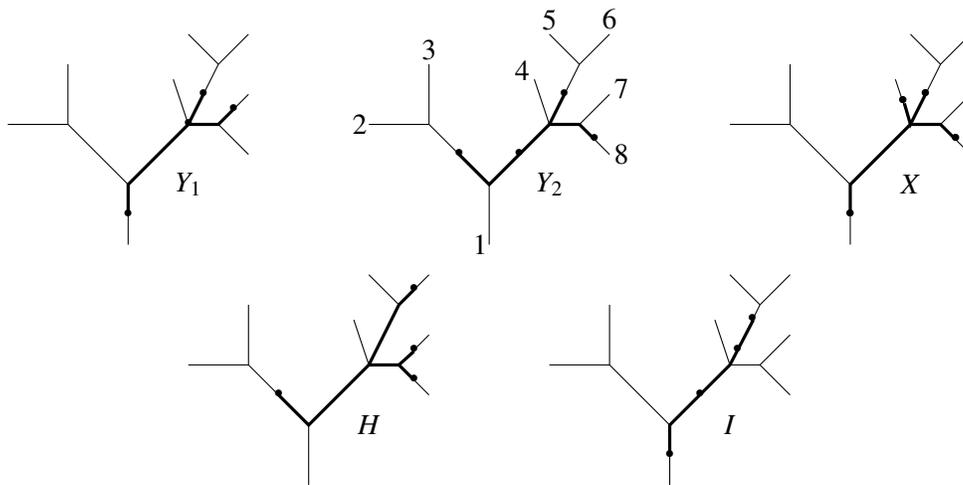


Figure 14: Various $S(Q)$ on a graph.

when v is the vertex where arms 5 and 6 meet, and w lies on the edge coming down from it; in this case \overline{vw} contains all points between v and the degree-1 vertices labeled 1, 2, 3, 4, 7, and 8.

Definition 5 Let Q be a set of 4 points on G such that $S(Q)$ is not of type I . Let w be a point in Q which is not in the interior of $S(Q)$. If $S(Q)$ is of type Y or X , let v be the unique vertex in the interior of $S(Q)$ which has degree at least 3 inside $S(Q)$. If $S(Q)$ is of type H , there are two such vertices; let v be the one nearer to w . The *arm number* of w (considered with respect to Q), is the smallest of the numbers of the degree-1 vertices in \overline{vw} .

For example, if Q consists of points on the arms going out to 5 and 6, a point at the vertex where those arms meet, and a point w on the arm going out to 7 and 8, then the

vertex v and the component \overline{vw} are the same as in the previous example, so the arm number for w is 1.

We will be considering moving from $\{a_1, a_2\}$ to $\{b_1, b_2\}$, with $Q = \{a_1, a_2, b_1, b_2\}$. We will depict a_1 and a_2 by black dots (B), and b_1 and b_2 by white (open) dots (W). It is possible that b_1 or b_2 might coincide with a_1 or a_2 , and so we also consider the possibility that in Y_2 the two dots on an edge coincide (and have different colors), and similarly for two adjacent dots on an I diagram.

We subdivide the I diagrams into two types:

- I_1 . The endpoints of I contain dots of different colors.
- I_2 . The endpoints of I only contain dots of the same color.

Proposition 13 *Let G be a tree. Partition $C \times C$ into three sets as follows. We use the dot color conventions as described above.*

- E_1 . $S(Q)$ is of type Y_1 or Y_2 such that the arm with two dots on it contains dots of the same color.
- E_2 . $S(Q)$ is of type I_1 .
- E_3 . Everything else. Thus $S(Q)$ is of type X , H , I_2 , or Y_2 such that the arm with two dots on it contains dots of each color.

There is a GMPR on each of E_1 , E_2 , and E_3 .

Remark This implies that $\text{GC}(C, \ell_1) \leq 2$, and hence implies Theorem 2 when G is not the Y graph, since $\text{TC}(C) \geq 2$ when $G \neq Y$.

Proof of Proposition 13 Figure 15 depicts the Y_2 cases in E_1 . By placing the dot in the inside of an arm at the vertex, we obtain the Y_1 cases of E_1 . The numbers 1, 2, and 3 indicate the relative order (smallest to largest) of arm numbers associated to the three non-interior points of $S(Q)$. The path that we select is the one that first does uniform motion from the black dot to the white dot indicated by the arrows in the diagram. It is followed immediately by uniform motion from the other black dot to the other white dot. Our uniform motions are always parametrized proportionally to arc length. One can see from the diagrams that collision is avoided. The analogous motion is performed for the associated Y_1 diagram. These geodesic paths clearly vary continuously with $(\{a_1, a_2\}, \{b_1, b_2\})$. If a dot on a single-dot edge moves to the vertex, an I -diagram is obtained, but, as these are in E_2 , there is no compatibility concern.

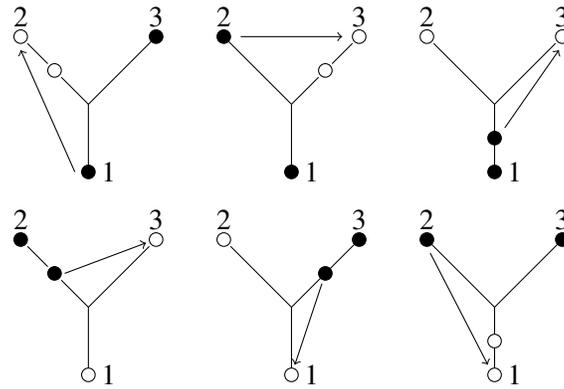


Figure 15: Y_2 graphs in E_1 .

The I_1 diagrams are in the order either BBWW or BWBW, with adjacent B and W possibly at the same position. The numbering of arms is not relevant. We choose the path which first moves uniformly from the B closest to the outer W, followed immediately by uniform motion from the other B to the other W.

This completes the proof that we have a GMPR on E_1 and E_2 . The argument for E_3 which follows is rather similar.

For the I_2 diagrams of the form BWWB and WBBW, we use simultaneous uniform motion from the B to the adjacent W.

For X diagrams, we first move uniformly from the black dot with smallest arm number to the white dot with smallest arm number, and then uniformly from the other black dot to the other white dot. As one or two (differently colored) dots approach the vertex, a Y_1 or I_1 diagram is obtained. Since these are not in E_3 , we need not worry about compatibility. A similar rule is used for an H diagram in which the arms emanating from each of the vertices v_0 and v_1 contain dots of the same color; we first move uniformly from the black dot with smallest arm number to the white dot with smallest arm number, and then uniformly between the other two dots. The limit as one or two (differently colored) dots approach a vertex is either a Y_2 diagram in E_1 or an I_1 diagram, and so again compatibility is not an issue.

For the Y_2 diagrams in E_3 , we can use simultaneous uniform motion between the two dots on one arm, and between the dots on the other two arms, always from black to white, of course. There is no risk of collision. If one of the dots which are unaccompanied on their arm approaches the vertex, we obtain either an I_2 diagram with compatible

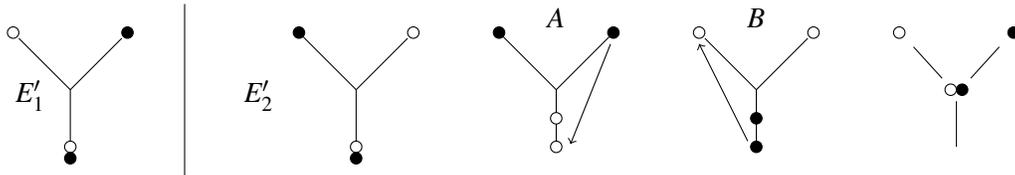


Figure 16: Some Y_2 diagrams and an I diagram.

simultaneous motion or an I_1 diagram. If one or both of the dots on the doubly-occupied arm approach the vertex, the limiting diagram is not in E_3 , so compatibility is not an issue. Similarly for H diagrams with dots of different colors on each of the two arms emanating from each vertex, move simultaneously and uniformly between the dots on the two arms emanating from each vertex. The limit as one or two of the dots approach their vertex is either a Y_2 or I diagram in E_3 with compatible motion or else an I diagram not in E_3 .

This completes the GMPR on E_3 and thus completes the proof of the theorem. \square

The following result implies Theorem 2 when G is the Y graph, since $TC(C) \geq 1$.

Proposition 14 *If G is the Y graph, then $C \times C$ can be partitioned into two ENRs with a GMPR on each.*

Proof There are no X or H diagrams. We can describe all the diagrams in a rotation-invariant way. The Y_2 diagrams and one I diagram are placed into the two sets E'_1 and E'_2 as suggested in Figure 16. The depiction of contiguous white and black dots means that they can appear in either order or at the same point, and in the last diagram they may appear anywhere between the two end points.

All Y_1 diagrams are placed in E'_2 , and all I diagrams except the ones of the type illustrated in Figure 16 are placed in E'_1 . The GMPR is simultaneous uniform linear motion on all the I diagrams and the first two diagrams in Figure 16. For the I diagrams, there is only one way that this can be done without collision, while for the first two in Figure 16, it is between the two bottom dots and between the two top dots. For diagrams A and B in Figure 16 and the Y_1 diagrams obtained from them by moving the inner point to the vertex, call the point moving with the arrow “point 1” and the point moving between the other two dots “point 2.” We use uniform linear motion for point 1, while point 2 moves uniformly from the black dot to the vertex, and then uniformly (at

a usually different rate) from the vertex to the white dot in such a way that for Diagram A (resp. B) it arrives at the vertex after (resp. before) the first dot.

Let d_1 (resp. d_3) denote the distance from the vertex of the black (resp. white) dot involved in the uniform motion with the arrow, and d_2 (resp. d_4) the distance from the vertex of the other black (resp. white) dot. Noting that point 1 hits the vertex when $t = d_1/(d_1 + d_3)$, we choose to have point 2 hit the vertex when $t_0 = \max(2d_1/(2d_1 + d_3), d_2/(d_2 + d_4))$ in Diagram A, and when $t_0 = \min(d_1/(d_1 + 2d_3), d_2/(d_2 + d_4))$ in Diagram B. This gives the appropriate values of t_0 if $d_4 = 0$ or $d_2 = 0$, implying uniform linear motion on the Y_1 diagrams. It is important to note, too, that the limiting motion as d_1 (resp. d_3) approaches 0 in Diagram A (resp. B) is the simultaneous uniform linear motion of the I diagram in E'_2 .

Since most of the I diagrams are in E'_1 , while most of the Y diagrams are in E'_2 , only a small amount of checking of compatibility of the motion in an I diagram that is a limit of Y diagrams in the same E'_i is required, and this is easily done. The only I diagram that is a limit of Y diagrams in E'_1 is in E'_2 . Each of the first three E'_2 diagrams in Figure 16 can approach an I diagram with an impossible motion of moving from an outer black dot to an outer white dot, but the arrangement of black and white in these is opposite to the I diagram in Figure 16, so that limiting diagram is in E'_1 . \square

Proof of Theorem 2 The lower bounds on GC follow from the known values of TC. The upper bounds follow from Propositions 13 and 14. \square

References

- [1] M.R.Bridson and A.Haefliger, *Metric Spaces of Non-Positive Curvature*, Grundlehren der mathematischen Wissenschaften Vol. 319, Springer-Verlag (1999).
- [2] D.M.Davis, *An n -dimensional Klein bottle*, Proc. Roy. Soc. Edinburgh Sect. A **149** (2019) 1207–1221.
- [3] D.M.Davis, *Geodesics in the configuration spaces of two points in \mathbb{R}^n* , arXiv:2001.00850.
- [4] D.M.Davis and D.Recio-Mitter, *Geodesic complexity of n -dimensional Klein bottles*, arXiv:1912.07411.
- [5] M.Farber, *Topological complexity of motion planning*, Discr. Comp. Geom. **29** (2003) 211–221.
- [6] M.Farber, *Collision-free motion planning on graphs*, in *Algorithmic Foundations of Robotics VI*, Springer (2005) 123–128.

- [7] R.Ghrist, *Configuration spaces and braid groups on graphs in robotics*, AMS/IP Studies Advanced Math **24** (2001) 29–40.
- [8] D.Recio-Mitter, *Geodesic complexity of motion planning*, arXiv:2002.07693.

D.M. Davis, Department of Mathematics, Lehigh University, Bethlehem, PA 18015, USA

M. Harrison, Dept. Math. Sciences, Carnegie Mellon University, Pittsburgh, PA 15213, USA

D. Recio-Mitter, Department of Mathematics, Lehigh University, Bethlehem, PA 18015, USA

dmd1@lehigh.edu, mah5044@gmail.com, d.reciomitter@yahoo.com