Flip distance on convex polygons and interval graphs on the circle

Jad Issa
American University of Beirut

May, 2021

Abstract

For a given polygon, it is possible to transform one triangulation into another by flipping one diagonal for another intersecting diagonal. The minimum number of flips required is an `NP`-complete problem for simple polygons and an open problem is whether it is in `P` or `NP`-complete for convex polygons. Among others, an application of a result about convex polygons is in counting the number of rotations required to transform one binary tree into another.

We show that flip distance is equivalent to token sliding and examine the resulting class of 'cyclical interval graphs' to which convex polygons correspond and build notions of the theory of cyclical interval graphs. We finally use these graphs to explain attempts to a solution of flip distance using this equivalence.

Contents

1 Introduction

1.1 Definitions

Definition 1.1 (Convex polygon)**.** A *convex polygon* is a polygon such that its interior is convex; i.e. any line segment between any two points x, y in the interior of the polygon lies completely inside the polygon.

Definition 1.2 (Triangulation)**.** For a polygon, a maximal set of pairwise non-intersecting diagonals is called a *triangulation* of that polygon.

Definition 1.3 (Flip)**.** For a given polygon P and two triangulations T, T' such that $T - T' = \{\delta\}$ and $T' - T = \{\delta'\}$ and such that δ, δ' are intersecting diagonals, the operation (δ, δ') is a *flip* transforming T into T'. We will write $(\delta, \delta') : T \mapsto T'$ to indicate the operation and its action on T.

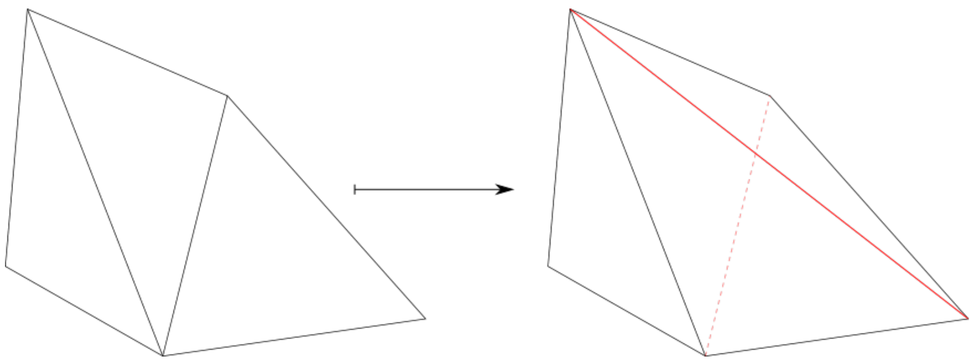

Figure 1: A flip (δ, δ'). δ is the dotted red line (removed), and δ' is the solid red line (added).

Theorem 1.1 (Flip graph)**.** *Triangulations of a polygon form an undirected graph according to the flip operation.*

Proof. If $(\delta, \delta') : T \mapsto T'$, then $(\delta', \delta) : T' \mapsto T$. □

Problem 1.1 (Flip distance)**.** The flip distance between triangulations T, T' of a polygon P is the distance between T and T' in the flip graph of P; i.e., the length of the shortest sequence of flips required to transform T into T', provided that is possible. The flip distance problem asks whether the distance between two triangulations T, T' is at most k.

1.2 Application

The flip distance problem is particularity interesting for convex polygons because of a correspondence found in [STT86] to rotation distance, which is another reconfiguration problem concerned with binary trees.

Definition 1.4 (Rotation)**.** For a rooted tree T, consider a structure where x is a vertex with y as a parent and α, β as child subtrees, and y has γ a subtree disjoint from x. The rotation of T at x replaces the subtree of y with one where x is the root, y is one of is children while the other is α, and β, γ are children of y. A pictorial representation can be found in figure 2.

Problem 1.2 (Rotation distance)**.** The rotation distance problem asks whether, given two rooted binary trees T, T', it's possible to transform T into T' in at most k rotations.

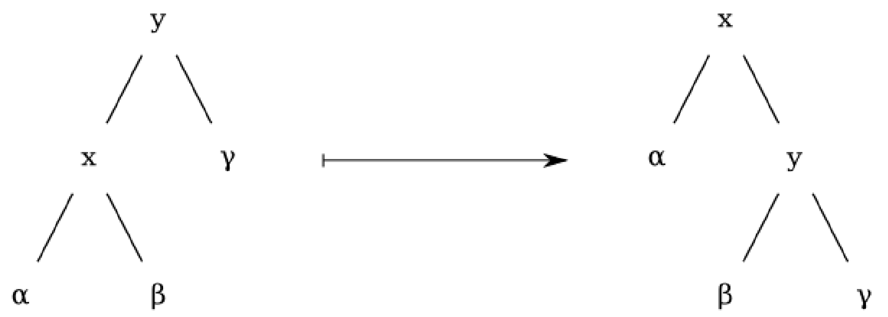

Figure 2: A tree rotation at x.

1.3 Current status

Currently, it is known that, on simple polygons, which are polygons that do not have self-intersections and split the plane into exactly two regions: an interior and an exterior, the problem is NP-complete [AMP15]. However, it is an open problem whether NP-hardness is maintained for convex polygons. As explained above, the answer to the convex case matters for more than classification, and work in this paper shows that the convex case is in fact a **very** special case (corollary of theorem 4.1).

2 Token sliding

Flip distance can be reduced (polynomially) to token sliding, which is a problem about independent set reconfiguration. The advantage of this reduction is that more results are known about token sliding and it is possible to solve our problem without reference to triangulations.

Problem 2.1 (Token sliding)**.** Given a graph G and two independent sets I, I', a token slide consists of replacing $u \in I$ by $v \notin I$ such that $I - u + v$ is an independent set and $uv \in e(G)$. The token sliding problem asks whether I' is reachable from I using at most k token slides.

Definition 2.1 (Diagonal graph)**.** For a polygon P, consider the undirected graph where vertices are diagonals, including the edges, and they are adjacent whenever the diagonals intersect. We call this graph the *diagonal graph* of P and will denote it G_P.

Theorem 2.1 (The reduction)**.** *Flip distance between T, T' is equivalent to token sliding between T, T' as independent sets of the diagonal graph.*

Proof. For a graph G and a triangulation T is a set of diagonals of P; therefore, it is a set of vertices of G_P. Furthermore, the diagonals from T are pairwise non-intersecting; therefore, they are pairwise non-adjacent vertices of G_P. A flip picks a pair of intersecting diagonals of which exactly one is in T, removes it, and picks the other one. This corresponds to a token slide in G_P. The converse holds by the same reasoning. $\square$

3 Cyclical interval graphs

3.1 Introduction

By reducing flip distance to token sliding, we now examine the structure of the diagonal graphs of convex polygons. It turns out, these graphs are similar to interval graphs, but not precisely so. We will call them cyclical interval graphs. The reason for the name will soon become apparent once we examine regular n-gons.

Definition 3.1 (Cyclical interval graph)**.** A cyclical interval graph is a graph whose vertices correspond to closed intervals of the real line and such that two vertices are adjacent whenever the corresponding intervals intersect non-trivially; precisely, if X, Y are two such intervals, then they are adjacent as vertices iff $X \cap Y$ is none of $\emptyset, X, Y$. We also require that no intervals be repeated.

Remark. We require non-repetition for convenience in the later definitions; however, there may be uses for the structure without this condition.

Remark. If the graph is finite, the intervals can be assumed to live entirely in a finite, closed, interval of the real line. After normalization, this interval can be even be assumed to be $[0, 1]$.

Cyclical interval graphs are the equivalent of intervals graphs when the real line is quotiented out into a circle; hence the naming. In particular, if we add to the number line a point $\infty = +\infty = -\infty$ (equivalent to 'sticking' the line together, and can be done with the arctan function), and allow intervals to use it to bridge $+\infty$ and $-\infty$; i.e., we allow $[a, b]$ where $a \geq b$ to be an interval, then, cyclical interval graphs behave like interval graphs in the precise way formalized below

If $[a, b] \cap [a', b'] = \emptyset$, then there is no edge in interval graphs, nor in cyclical interval graphs. If $[a, b] \cap [a', b'] = [a', b']$, then $[b, a] \cap [a', b'] = \emptyset$ and there is no edge just as in interval graphs. Here, $[b, a]$ is one of those intervals that contain ∞. Symmetry insures the other case.

The same trick of adding the point ∞ can be used when the graph is finite, and infinities are not introduced except symbolically; i.e. it has nothing to do with the halting of a computation.

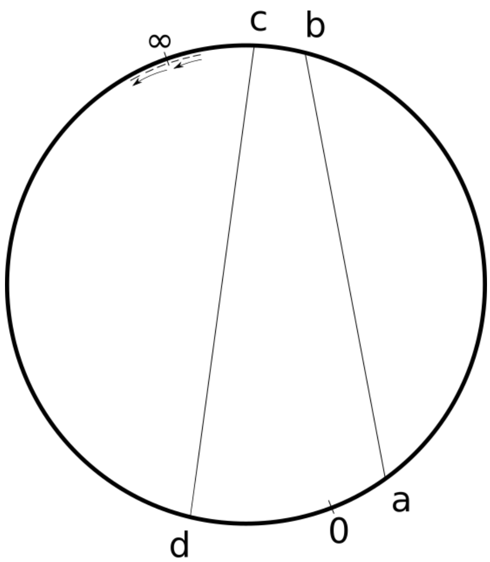

Figure 3: A visual representation of why $[a, b]$ and $[d, c]$ do not intersect even though $[a, b] \subseteq [d, c]$. Consider $[d, c] = [c, (\infty), d]$, then $[a, b] \cap [c, d] = \emptyset$.

3.2 Support and precision

We now provide some measures to classify cyclical interval graphs in ways that are relevant to our study of convex polygons.

Definition 3.2 (Support (general))**.** If the cyclical interval representation of a cyclical interval graph G is $\{[a_i, b_i]\}_{i \in I}$, then let the *support* S_G of G in this representation be the smallest connected closed set containing $\bigcup_{i \in I} [a_i, b_i]$.

Corollary (Support (finite))**.** *For a finite cyclical interval graph, the support is a closed interval* $[a, b]$

Proof. We have finitely many intervals $[a_i, b_i]$; each of which having a minimum a_i and a maximum b_i. Therefore, the sets $\{a_i\}_{i\in I}, \{b_i\}_{i\in I}$ have a minimum and a maximum respectively, say a, b respectively.

Then, $a, b \in S_G$ since $a, b \in \bigcup_{i\in I}[a_i, b_i]$, and since S_G is connected, then $[a, b] \subseteq S_G$. Furthermore, since $[a, b]$ is a closed connected set containing $\bigcup_{i\in I}[a_i, b_i]$, then $S_G \subseteq [a, b]$ because S_G is the smallest such set. Therefore, $S_G = [a, b]$. □

Definition 3.3 (Precision)**.** When possible, a cyclical interval graph is of precision $k \in \mathbb{N}$ if k is the least number such that G can be represented using integer intervals and support of length $k - 1$. If no such k is found, the precision is defined to be $+\infty$.

Theorem 3.1. *Any cyclical interval graph of precision $k \in \mathbb{N}$ can be represented using integers on $[1, k]$.*

Proof. If the support found is $[\eta, \eta+k-1]$, then subtract $\eta-1$ from each end of each interval. □

Theorem 3.2. *Any finite cyclical interval graph with intervals with rational endpoints has a finite precision.*

Proof. Out of all finitely many interval endpoints, pick the one with the largest denominator. Multiply all endpoints by the given denominator. This homothety preserves intersections; however, now we have integer endpoints. Since G is finite, the support is a finite closed interval with integer endpoints; therefore, of integer length. □

Remark. This means that any cyclical interval graph whose cyclical interval representation is possible in standard representation in finitely many bits has a finite precision. This characterizes classes of cyclical interval graphs in a similar way to other measures such as the clique number ω, and the chromatic number χ for general graphs.

Definition 3.4 (Complete cyclical interval graph of precision k; $\mathcal{CIK}_k$)**.** For any natural k, the complete cyclical interval graph of precision k, denoted $\mathcal{CIK}_k$ is the graph with includes all integer intervals from $[1, k]$ exactly once.

Theorem 3.3. *Any cyclical interval of precision k which also has a maximal representation in a support of length $k - 1$ is isomorphic to the complete cyclical interval graph with precision k.*

Proof. Using theorem 3.1, we can get a representation on $[1, k]$ which must also be maximal. Therefore it has the same representation as the complete cyclical interval graph; thus, it is isomorphic to the $\mathcal{CIK}_k$. □

Theorem 3.4. *Any finite cyclical interval graph with intervals with rational endpoints is an induced subgraph of some complete cyclical interval graph.*

Proof. Using 3.2 and 3.3, we get that the graph is k-precise and therefore, that its vertices are a subset of the vertices of $\mathcal{CIK}_k$. As for the edges, they are not determined by the graph itself, rather by the intersection of the intervals which is completely inflexible and predetermined after picking the intervals. Therefore, we have an induced subgraph of $\mathcal{CIK}_k$. □

The importance of theorem 3.4 is that it ensures that precision is a stable concept for which completeness makes sense. Finally, we define an open version of cyclical interval graphs which we will use later.

Definition 3.5 (Open cyclical interval graph)**.** An *open* cyclical interval graph is a cyclical interval graph with open intervals instead of closed intervals.

Remark. Any open cyclical interval graph of precision at most k can be represented as a cyclical interval graph of precision at most $3k-2$ by replacing (a, b) with $[3a+1, 3a-1]$ for a, b integers.

4 Connections to flip distance

Theorem 4.1. *The diagonal graph of any simple n-gon is an at most n-precise open cyclical interval graph.*

Proof. Number the vertices starting from some arbitrary vertex and moving in the anti-clockwise direction from 1 to n. Any diagonal can then be represented as (i, j).

Consider all possible diagonals (i', j') where $(i', j') \subseteq (i, j)$; i.e. $i \leq i' \leq j' \leq j$. If i', j' appear both on the same side of the line (ij), then (i', j') cannot intersect (i, j).

Otherwise, suppose (i', j') intersects (i, j), then, since we have a simple polygon, the edges going from i' to j' can only cross the (ij) line strictly outside the $[ij]$ line segment at some point a. Now, consider the continuous deformation of $[i'a]$ into $[i'j']$ by moving a to j' across the edges of the polygon. We know that $[ja]$ is not inside the polygon because otherwise j would be an internal point, which is impossible in simple polygons; therefore, at some point across $[i'a]$, there must have been a last point y_0 in the polygon such that $y_0 \neq a$ and $[y_0 a]$ is not in the polygon. However, for the point j', we assumed the diagonal $[i'j']$ exists; therefore, y_1, the point that was continuously transformed from y_0 must have become after a. Because the transformation is continuous, there must have been some point where $y = x$ (where x is the point that is moving from a to j'). However, this is a self-intersection which is impossible in simple polygons.

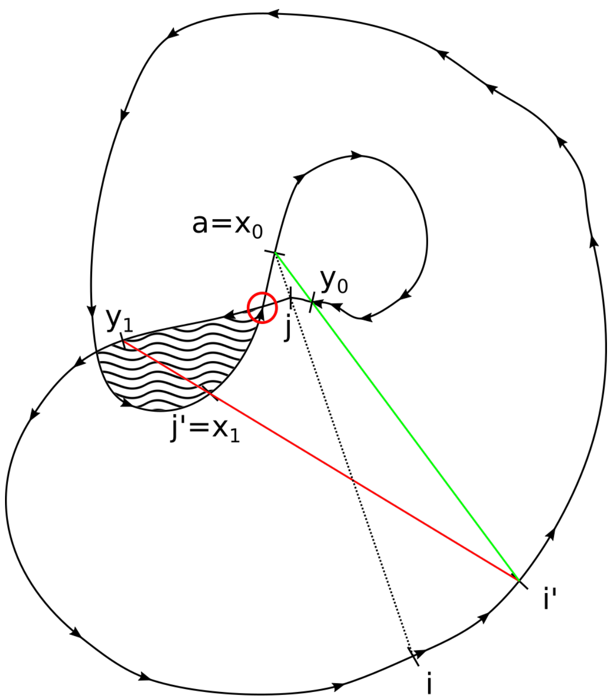

Figure 4: Visualization of how the continuous transformation must lead to a self-intersection point (red circle). There is also possibly a complete area which has intersection (wavy area).

We conclude that such a pathological diagonal does not exist. Therefore, anytime $(i', j') \subseteq (i, j)$, we do not have any intersection.

The remaining case is where $i < i' < j < j'$. In this case, if the diagonals cross as lines and they cross outside at least one of the segments, then by the argument above, either one of them must be an invalid diagonal or the polygon cannot be simple. On the other hand, if they don't intersect as lines and instead both i', j' lie on one side of (ij), then, consider the diagonal $(i'j')$. There must be precisely 2 edges that contain i', one entering i' and one leaving i' according to the order given to the circumference. Given the strict inequalities, it is impossible that $(i'j')$ be an edge; therefore, it is a proper diagonal completely contained strictly inside the polygon. As such, in must lie strictly between the 2 edges that contain i' in terms of angles (otherwise there would be a third edge containing i'). The same applies to j'. Furthermore, the path after i' and until j cannot pass by the segment $[i'j']$ otherwise we would have a self-intersecting polygon or (i', j') would be invalid. Therefore it must cross $(i'j')$ from outside $[i'j']$. Moving across the edges from i' to j then to j', at some point, we must close the loop $i' \cdots j \cdots j'i'$ with the added diagonal (i', j'). Since the diagonal (i', j') must be between the two edges on j', one of the edges of the polygon must be inside the loop. However, this edge cannot leave the loop by the edges of the polygon because that would create self-intersections, and it cannot leave through the diagonal (i', j') as that would invalidate it as a diagonal. Therefore, this case is impossible and $(i, j), (i', j')$ must always intersect in this case.

What we have is a cyclical interval graph realized by a representation using integer intervals from $[1, n]$; therefore, at has precision at most n.

□

Corollary. *The diagonal graph of a convex n-gon is the complete open cyclical interval graph with precision n.*

Proof. Every pair of vertices makes up a line that is completely inside the polygon due to convexity. Therefore, every interval in $[1, n]$ represents a unique diagonal or edge; thus, the graph is complete. □

Remark. Notice that this corollary is an equivalence, meaning that if the diagonal graph of an n-gon is the open $\mathcal{CIK}_k$, then this polygon must be convex because it has all of the possible diagonals. This means that we were able to **completely** capture what convexity means in very different, and possibly more familiar terms as the completeness of a graph.

One related problem that seems promising is token sliding on proper interval graphs. Proper interval graphs are interval graphs where no interval is properly contained in another one. The difference between these graphs and cyclical interval graphs is that in cyclical interval graphs, we only require the non-existence of an edge between intervals that contain one another, while in proper interval graphs, it is required that the intervals never exist simultaneously.

Incidentally, token sliding on proper interval graphs has been shown to be in `P` in [YU21]. However, the slight relaxation of the condition for proper interval graphs to get cyclical interval graphs makes the problem `NP`-hard

Theorem 4.2. *Token sliding on cyclical interval graphs is `NP`-hard.*

Proof. Triangulation on simple n-gons is `NP`-complete[AMP15], yet can be poly-reduced to token sliding on open cyclical interval graphs. Therefore, the latter must be `NP`-hard. (Closed) cyclical interval graphs include also include polynomially larger representations for open cyclical interval graphs. Therefore, token sliding on cyclical interval graphs is `NP`-hard in general. □

5 Conclusion

The results from the given work are many-fold. Primarily, the focus is on using the knowledge about the diagonal graph of convex polygons to solve the problem of the complexity of flip distance on such polygons. The main contributions provided by this paper can be summarized as the following.

1. The introduction of a very basic framework for studying cyclical interval graphs (section 3).

2. The characterization of the diagonal graphs of simple n-gons as open cyclical interval graphs of precision at most n (theorem 4.1).

3. The faithful specialization and reformulation of the convexity condition as the completeness condition on open cyclical interval graphs (corollary of theorem 4.1).

4. One complexity-theoretic result on the `NP`-hardness of token sliding on general cyclical interval graphs.

However this paper raises other questions, the most important of which is how to proceed given this structure to solve flip distance on convex polygons. More questions that might be of interest are listed below.

1. Is token sliding on cyclical interval graphs `NP`-complete or harder (`PSPACE`-complete)? What about open cyclical interval graphs? How many, and which conditions do we have to add to get inside `NP` and/or inside `P`?

2. If cyclical interval graphs are analogous to interval graphs when the real line is quotiented into a circle, then are what would be similar constructions for more complex topological spaces and which are useful and where?

3. Is the implication of 4.1 an equivalence, and what would an answer to that imply? One result is that if it is an equivalence, we can reduce token sliding on open cyclical interval graphs to flip distance on simple polygons, which would make token sliding on open cyclical interval graphs not only `NP`-hard, but also `NP`-complete. My conjecture is that it's not an equivalence; however, I could not prove that.

References

[AMP15] Oswin Aichholzer, Wolfgang Mulzer, and Alexander Pilz. Flip distance between triangulations of a simple polygon is np-complete. *Discrete & Computational Geometry*, 54(2):368–389, 2015.

[STT86] Daniel D Sleator, Robert E Tarjan, and William P Thurston. Rotation distance, triangulations, and hyperbolic geometry. *Proceedings of the eighteenth annual ACM symposium on Theory of computing - STOC '86*, page 122–135, Nov 1986.

[YU21] Takeshi Yamada and Ryuhei Uehara. Shortest reconfiguration of sliding tokens on subclasses of interval graphs. *Theoretical Computer Science*, 863:53–68, 2021.