Quadratic Reconfiguration of Plane Spanning Paths on Two-Layer Point Sets

Samuel Avril[1][0009-0007-2333-6825], Yan Gerard[2][0000-0002-2664-0650], and Guilherme D. da Fonseca[3][0000-0002-9807-028X]

[1] ENS de Lyon, Lyon, France
samuel.avril@ens-lyon.fr
[2] LIMOS, Université Clermont Auvergne, Clermont-Ferrand, France
yan.gerard@uca.fr
[3] LIS, Aix–Marseille Université, Marseille, France
guilherme.fonseca@lis-lab.fr

Abstract. Consider a crossing-free polygonal path with n vertices and a flip operation that removes an edge and adds another one, without changing the vertex set and keeping the path crossing free. It is conjectured that any two such paths can be transformed into each other by a sequence of flips. Recently, Kleist et al. (2024) proved the conjecture for the special cases of vertices arranged in two convex layers. However, the number of flips in their proof may be exponential. In this paper, we show that $O(n^2)$ flips suffice. More precisely, for a two-layer point set $S = L_0 \cup L_1$, where L_0 and L_1 denote the outer and inner layers, respectively, our method uses $O(|L_0| + |L_1|^2)$ flips.

Keywords: Flip graphs · Geometric reconfiguration · Computational geometry

1 Introduction

Reconfiguration studies a family of configurations that are changed into each other using a flip operation. In our case, the *configurations* are crossing-free spanning paths of a set S of n points and the *flip* operation removes an edge and adds another one while preserving a valid configuration. Given a set S, there is an associated *flip graph* whose vertices are configurations and edges are defined by flips. The *diameter* of the flip graph is the number of flips needed, in the worst case, to transform one path into another. We refer to paths in the flip graph as *flip sequences* to avoid confusion. An example of flip sequence is drawn Fig. 1.

The problem was introduced by [7] in 2001, together with an implicit *conjecture* that the flip graph is connected and a proof of it for the special cases of points S in *convex* position (S is the set of convex hull vertices). The first polynomial bound on the diameter for the convex case appears in [3], at most $2n - 5$. Later a tight bound of $2n - 5$ for $n \leq 4$ and $2n - 6$ for $n \geq 5$ has been shown [5]. In the convex case, the shortest flip sequence can be computed in linear time [1].

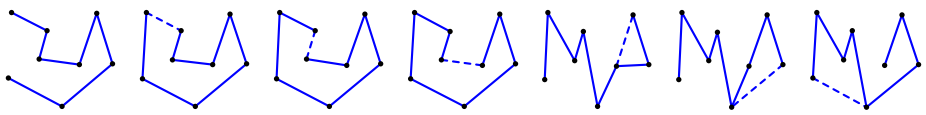


Fig. 1: A flip sequence connecting a pair of paths

Proving the conjecture for arbitrary point sets remains challenging and many authors considered several other special cases; we cite the most relevant for this paper. For convex position plus one *interior* point, the diameter is at most $2n - 4$ [2]. For convex position plus one *exterior* point, the diameter is at most $2n$ [4]. Both cases are special cases of a more general *two-layer* point set $S = L_0 \cup L_1$ where the *outer layer* L_0 is the set of vertices in the convex hull of S and the *inner layer* L_1 is also in convex position. In [6], they show that the flip graph for two layers is connected, but the diameter implied by their proof is exponential. In this paper, we prove the following quadratic bound.

Theorem 1. *For a two-layer point set $S = L_0 \cup L_1$, where L_0 and L_1 denote the outer and inner layers, the diameter of the flip graph is $O(|L_0| + |L_1|^2)$ and computing a flip sequence connecting a pair of paths takes $O(|L_0| + |L_1|^2)$ time.*

2 Canonical Paths and Potential Reduction

This section assembles the main elements allowing us to prove Theorem 1. Let $S = L_0 \cup L_1$ be a two-layer point set with L_0 and L_1 the outer and inner layer respectively. Let $n_0 = |L_0|$ and $n_1 = |L_1|$.

We first prove the connectivity of a particular family of paths that we will call *canonical paths* (Subsection 2.2). The second step is to show that any arbitrary plane spanning path on a two-layer point set can be transformed into a canonical path. We achieve this goal by defining a potential function Φ on paths (subsection 2.3). The potential has three properties (subsection 2.4):

- the potential of a path is 1 if and only if that path is canonical
- the potential is strictly positive, so that the set of the canonical paths is exactly the class of paths with minimum potential
- if the potential of a path is not minimal i.e if the path is not canonical, we provide an algorithm for getting a sequence of flips decreasing the potential by at least 1. The number of flips required for a potential decreasing is precisely the key explaining the bound on the diameter of the flip graph given in Theorem 1.

In other words, a sequence of flips decreasing the potential of a path until the minimum value 1 transforms the initial path into a canonical path. As all the canonical paths are connected, the connection of any path to a canonical path proves the connectivity of the whole flip graph. It yields Theorem 1.

2.1 Terminology

Before starting to define the canonical paths and the potential, we introduce some notations. For two vertices $u, v \in S$ in a path P, we say that u *sees* v if the interior of the segment uv does not intersect any other edge of P.

To distinguish the neighbors of the vertices in a path, we may orient the path. When a path is oriented, its first and last vertices are called the *start* and *end*, respectively. An *st-path* is an oriented path with start s and end t. For a vertex u, we denote by $\text{in}(u)$ and $\text{out}(u)$ its incoming and outgoing edges, when they exist. We write $\text{pred}(u)$ and $\text{succ}(u)$ for the predecessor and successor of u, respectively. Whenever an edge is written as uv in an oriented path, we implicitly assume that $u = \text{pred}(v)$.

For an st-path P, the subpath from $u \in S$ to $v \in S$ is denoted by $P[u : v]$. A prefix of the path is any subpath of the form $P[s : u]$, $u \in S$, and a suffix of the path is any path of the form $P[u : t]$, $u \in S$.

2.2 Reconfiguration of Canonical Paths

A path is *canonical* if it first visits all vertices of one layer in cyclic order and then all vertices of the other layer in cyclic order. The directions chosen on the two layers do not need to be the same. For instance, all paths in Figure 2 are canonical.

The unique edge of a canonical path connecting an outer point to an inner point is called a *bridge*. A canonical path is completely determined by its bridge, and, for each layer, the direction (clockwise or counterclockwise) in which the path traverses that layer from the corresponding extremity of the bridge.

To reconfigure canonical paths, two useful observation can be made:

- The traversal direction of either layer can be reversed in 1 flip.
- Either extremity of the bridge can be shifted to their layer neighbor they are not connected with in 1 flip, if it does not introduce crossing.

Using these observations, we can reconfigure canonical paths in linear flips while preserving their structure.

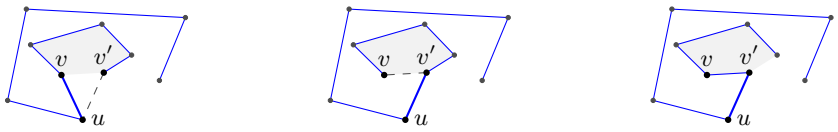


Fig. 2: Counterclockwise bridge shift on the inner layer

Lemma 2. *The subgraph of the flip graph of canonical paths on a two-layer point set is connected with diameter $O(n_0 + n_1)$ flips.*

Proof. Let P and P' be two canonical paths with bridges uv and $u'v'$, where $u, u' \in L_0$ and $v, v' \in L_1$.

Consider a triangulation of the ring between L_0 and L_1 that contains both uv and $u'v'$ and such that and no triangle connects three points of L_0. In such a triangulation, the edges joining L_0 and L_1 form a cyclic sequence T where each consecutive edge differ only by one endpoint.

Fix a direction along this sequence, say counterclockwise. First, if necessary, flip the edges adjacent to the bridge to connect u and v to $\mathrm{ccw}(u)$ and $\mathrm{ccw}(v)$ respectively.

Starting from uv, repeatedly flip the bridge to the next one of T in counterclockwise direction, and do another flip to connect the extremity of the bridge that shifted to its clockwise neighbor again. An Illustration of these two flips is given in Figure 2 This second flip preserves the loop invariant that the path is always a canonical path in which both extremities of the bridge are connected to their clockwise neighbor.

Repeating this process eventually transforms the bridge uv into $u'v'$. Finally, at most two additional flips are needed to adjust the traversal directions on the two layers so that the resulting canonical path is exactly P'

At the end of that loop we can correct the orientation of the path at the extremity of the bridge to match P' in 2 flips. Therefore P can be transformed into P'.

Moreover, the sequence T contains exactly $n_0 + n_1$ edges. So the whole process uses at most $2(n_0 + n_1) + 4$ flips, proving the claim. □

2.3 Potential

We now introduce the potential used for connecting any path to a canonical path. Its value is based on a classification of edges according to their position with respect to the two layers. We distinguish five types of edges:

- *level edges*, joining two consecutive vertices on the same layer;
- *bridges*, joining an outer and an inner vertex without crossing the convex hull of L_1;
- *small diagonals*, joining two non-consecutive vertices of L_1;
- *medium diagonals*, joining an outer and an inner vertex through the interior of the convex hull of L_1;
- *big diagonals*, joining two non-consecutive vertices of L_0.

These edge types are represented in Figure 3.

We assign a weight to each edge type and define the potential $\Phi(P)$ of a path P as

$$\Phi(P) = 0L(P) + 1B(P) + 1S(P) + 3M(P) + 4G(P) \tag{1}$$

where $L(P)$, $B(P)$, $S(P)$, $M(P)$ and $G(P)$ denote the numbers of level edges, bridges, small diagonals, medium diagonals and big diagonals in P, respectively. The coefficients in (1) are chosen so that canonical paths are exactly the paths of minimum potential.

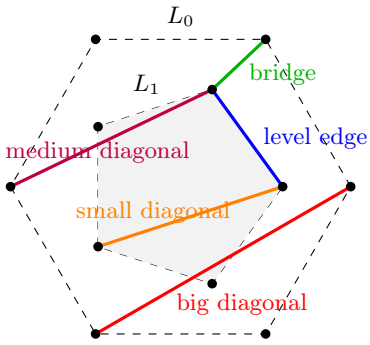


Fig. 3: Illustration of the five edge types in a two-layer point set.

2.4 Potential Properties and Proof of Theorem 1

The first properties are straightforward. First, by definition, the potential is a strictly positive integer. Secondly, paths of minimal potential are canonical:

Lemma 3. *The minimum value of Φ is* 1*, and a path has potential* 1 *if and only if it is canonical.*

Proof. Every spanning path must contain an edge joining the two layers. Such an edge is either a bridge or a medium diagonal, both of which contribute at least 1 to the potential. Therefore $\Phi(P) \geq 1$ for every path P.

Now let P be a path with $\Phi(P) = 1$. Since every edge type except level edges has weight at least 1, the path must contain exactly one bridge and all its remaining $n_0 + n_1 - 2$ edges must be level edges.

Consequently, the restriction of P to L_0 contains exactly $n_0 - 1$ level edges and the restriction of P to L_1 contains exactly $n_1 - 1$ level edges. Hence the vertices of each layer are traversed consecutively in cyclic order. The unique bridge connects the two layer-traversals, so P is a canonical path.

Conversely, every canonical path consists of a single bridge and $n_0 + n_1 - 2$ level edges, and therefore has potential 1. □

Lemma 4. *If a path is not canonical, we can reduce its potential by at least* 1 *with a sequence of at most $O(n_1)$ flips.*

Lemma 5. *Throughout the process of diminishing the potential of the path, at most $O(n_1)$ potential reduction steps require more than 1 flip.*

These last two lemmas enable the proof of Theorem 1.

Proof (Theorem 1). Let P be a path. By Lemma 4, we can diminish its potential in at most $O(n_1)$ flips, until it has least potential and becomes canonical (Lemma 3). Among each of these reduction steps, at most $O(n_1)$ of them may require more than 1 flip to be performed, according to Lemma 5. Therefore, since $\Phi(P) =$

$O(n_0 + n_1)$, this means that the reconfiguration of P into a canonical path requires at most $O(n_0 + n_1^2)$.

Moreover, using Lemma 2, we can reconfigure any canonical paths into one another in $O(n_0 + n_1)$ flip, meaning we can reconfigure any path into one another in $O(n_0 + n_1^2)$ flips. □

The later material in Sections 3, 7 and 5 are respectively devoted to the proofs of Lemma 4 and Lemma 5. An organization of the dependencies between the previous and next claims is illustrated in Fig. 4

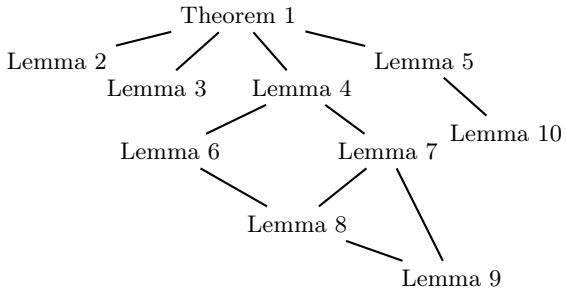


Fig. 4: The tree of dependencies of the different lemmas allowing us to prove Theorem 1

3 Potential Reduction

In this section, we provide a sketch of proof of Lemma 4. The different procedures constituting the algorithm leading to this lemma and their respective proof of correction are given in Section 4.

To this end, we describe an algorithm that, given a path P with $\Phi(P) > 1$, performs a sequence of flips that decreases the potential. The algorithm uses two different procedures, depending on whether an endpoint of P is on the outer layer or if both are in the inner layer. We give a high level description of these procedures.

The procedure OuterReduction takes as input an oriented path whose start vertex lies on the outer layer. Its goal is to create a new level edge incident to the start. When such an edge cannot be created immediately, it tries the same thing with the end. If this fails again, the procedure uses its end vertex to change the orientation of the path next to the start vertex, so that a new level edge incident to the start can be created. This procedure performs a constant flip sequence.

Lemma 6. *Let P be an st-path where the start $s \in L_0$ is an outerpoint. Then the procedure OuterReduction reduces the potential of the path in $O(1)$ flips.*

The procedure InnerReduction handles the case where both endpoints lie on the inner layer. As before, the primary objective is to create a new level edge incident to the start vertex. If this fails, depending on the type of the incoming edge of the layer neighbor of the start, the procedure either performs a direct reduction or transforms the path into a simpler configuration that admits further reductions. This procedure may perform $O(n_1)$ flips in certain cases.

Lemma 7. *Let P be an st-path where both endpoints are in L_1. Then the procedure InnerReduction reduces the potential of the path in $O(n_1)$ flips.*

The two lemmas 6 and 7 directly imply Lemma 4. Their two procedures OuterReduction and InnerReduction may invoke some common subprocedures ObstrReduction and DiagObstrReduction when common cases arise. These two subprocedures handle cases where some visibility between some points are obstructed by some edges of the path. These cases always involve edges with heavy weight, so the goal of these subprocedures is to find a way to flip them.

Lemma 8. *Let P be a path starting at $s \in L_1$ and let q be a neighbor not seen by s because some edge of P blocks the visibility between them. Then the procedure ObstrReduction reduces the potential of the path in $O(1)$ flips.*

Lemma 9. *Let P be a path starting at $s \in L_1$. Suppose that P contains a diagonal uv, where either uv is a big diagonal or uv is a medium diagonal with u a cyclic neighbor of s in L_1. Assume that s sees u but not v, and that there exists an inner vertex $p \in \triangle suv$ that sees both s and v. Moreover, assume that the concave quadrilateral $suvp$ is empty.*

Let u' be the last outer vertex on $P[s:u]$ and let v' be the first neighbor of v visited by P. Assume further that

1. *$P[s:p]$ lies entirely in L_1 and visits its vertices in cyclic order from s to p;*
2. *and $P[p:u']$ visits the remaining inner vertices between s and p, together with the outer vertices from v' to u', preserving their cyclic order on L_0.*

Then DiagObstrReduction reduces the potential of P using $O(1)$ flips.

We detail the procedures and prove Lemmas 6 to 9 in the following section.

4 Procedures reducing the potential of a path

In this section, we provide the different procedures reducing the potential of a path, and prove their correctness and their complexity. We start by presenting the two main procedures, OuterReduction (Section 4.1) and InnerReduction (Section 4.2), and finish by the last two sub procedures ObstrReduction and DiagObstrReduction (Section 4.3).

4.1 Procedure OuterReduction and proof of Lemma 6

Procedure OuterReduction.
Input: a path P whose start $s \in L_0$ lies on the outer layer,
and a neighbor s' of s that does not share an edge with it.
Assume $s' = \text{ccw}(s)$ by symmetry.
Let t be the end of the path.

1. If $s' = t$, pick any non-level edge and flip it to st.
2. Else if the incoming edge of s' is not a level edge, flip it to ss'.
3. Otherwise: Let t' be a neighbor of t that does not share an edge with it.
   (a) If t' does not see t, reverse the order of the path and apply the ObstrReduction procedure on that obstruction.
   (b) Else if the outgoing edge of t' is not a level edge, flip it to tt'.
   (c) Else if $t' \in P[s : s']$, perform

$$\text{out}(t') \longrightarrow t't, \qquad \text{in}(s') \longrightarrow ss'$$

   (d) Otherwise,
   Let uv, with $u = \text{pred}(v)$, be the last small diagonal in $P[\text{succ}(s') : t]$, and let v' be neighbor of v not seen by t'. Then perform $uv \longrightarrow ut$, and apply recursively OuterReduction using v' for the non connected neighbor of the new start v.

Proof (Lemma 6). Let P be an st-path starting at $s \in L_0$, and assume $|L_0| \geq 3$, otherwise the lemma is trivial. Let $s' \in L_0$ be one of the cyclic neighbor of s and assume $s' = \text{ccw}(s)$ by symmetry. Since these points are neighbors of the outer layer, they necessarily see each other.

Therefore the flips in Cases 1 and 2 are legal and decrease the potential.

So now suppose the premise of Case 3 is verified: the incoming edge of $s' \neq t$ is a level edge. Then the path has a particular structure. Indeed, as illustrated in Figure 5, the prefix $P[s : s']$ wraps around the suffix $P[\text{succ}(s') : t]$ so that the latter is forced to stay confined in L_1. In particular, the end $t \in L_1$ is an inner point. This restriction on the suffix means that $P[s : s']$ must visit all outer points, and we can prove that all the points from L_0 must appear in a cyclic order in that prefix.

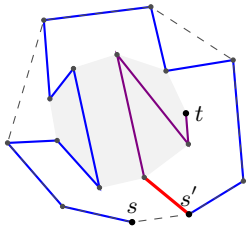


Fig. 5: Structure of the path in Case 3 of procedure OuterReduction

Also, since the incoming edge of s' is a level edge, its outgoing edge is not as s' does not share an edge with s. So the remaining cases try to reverse the orientation of the path in s' if it cannot be immediately be reduced near t.

Let $t' \in L_1$ be a neighbor of t such that t' does not share an edge with t. Such a t' always exists unless $|L_1| \leq 2$ and the path already has potential 1. We distinguish cases depending on visibility between t and t' and the type of the outgoing edge of t'.

- **Case 3.(a)**: If t does not see t', then applying ObstrReduction reduces the potential in $O(1)$ according to Lemma 8.
- **Case 3.(b)**: If t sees t' and the outgoing edge of t' is not a level edge, then $\text{out}(t') \to t't$ creates a new level edge, diminishing the potential.

Now let's suppose that t sees t', and $\text{out}(t')$ is a level edge.

- **Case 3.(c)**: If $t' \in P[s : s']$ or equivalently $s' \in P[t' : t]$, the first flip $(\text{out}(t') \to t't)$ is legal and does not change the potential. The second flip $\text{in}(s') \to ss'$ is also legal as s' sees s. Moreover, after the first flip the orientation of the path in $s' \in P[t' : t]$ is reversed, exchanging $\text{in}(s')$ with the non-level edge $\text{out}(s')$. Therefore the second flip flips a non-level edge to a level edge, which diminishes the potential. The flip sequence produced is illustrated in Figure 6

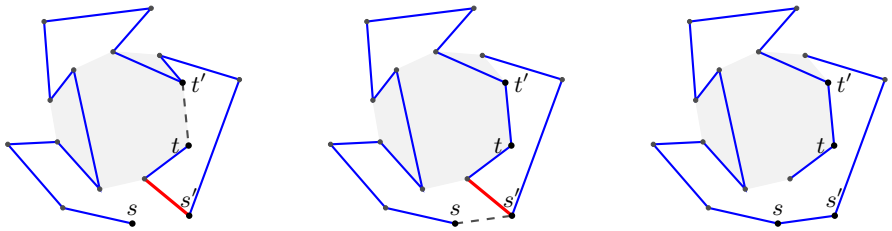


Fig. 6: Reduction of potential in Case 3.(c) of Procedure OuterReduction.

- **Case 3.(d)**: Suppose $t' \in P[\text{succ}(s') : t]$. In this case, we can show that $t' = \text{succ}(s')$. Indeed, if that was not the case, then as $P[\text{succ}(s') : t] \subset L_1$ and $\text{out}(t')$ is a level edge, $\text{in}(t')$ must be a small diagonal. This diagonal would cut L_1 in two parts, separating $\text{succ}(t')$ and t the two neighbors of t' and introducing a crossing with the suffix $P[\text{succ}(t') : t]$. Therefore in this case $t' = \text{succ}(s')$.
  From this observation we can deduce that $P[t' : t]$ necessarily contains a small diagonal, as if it was not the case, $P[t' : t]$ would visit all inner points in cyclic order, which would imply that the path is already canonical.
  Now, as in the procedure, let uv be the last diagonal of $P[t' : t]$. Then it splits L_1 in half and we can define v' to be the neighbor of v not seen by t'. The flip

$uv \to ut$ does not necessarily decrease the potential, but it changes the end of the path to v, so that applying again the procedure with v' for the choice of the neighbor of v will not lead to that case again, since $v' \neq t$ so $v' \neq \text{succ}(s')$. And since all the other cases decrease the potential in $O(1)$ flips, this one does as well. The flip sequence produced in this case is illustrated in Figure 7



Fig. 7: Reduction Case 3.(d) to other cases of procedure OuterReduction

This proves by case analysis that the procedure OuterReduction can always reduce the potential of a non-canonical path with an outer endpoint in $O(1)$ flips. □

4.2 Procedure InnerReduction and proof of Lemma 7

Procedure InnerEndpointsReduce. *Input:* a path P whose start $s \in L_1$ and end $t \in L_1$ lie on the inner layer,
and a neighbor s' of s that does not share an edge with it.
Assume $s' = \text{ccw}(s)$ by symmetry.

1. If s does not see s', apply the ObstrReduction procedure on that obstruction.
2. Else if the incoming edge of s' is not a level edge, flip it to ss'.
3. Else if $s' = t$, pick any non-level edge and flip it to ss'.
4. Else if the outgoing edge of s' is a small diagonal: Let $x = \text{succ}(s')$. Repeat until the potential decreases: let s be the start of the path, and $s' = \text{ccw}(s)$, and perform
$$s'x \to sx, \qquad \text{in}(s) \to s's.$$
5. Else if the outgoing edge of s' is a medium diagonal: Let $x = \text{succ}(s')$. If s and t are on the same side of $s'x$ in L_1, flip $\text{ccw}(s')s' \to ss'$, rename $\text{ccw}(s')$ as s, and take the reflect the path so that $s' = \text{ccw}(s)$ again. Then
   (a) If x sees s, repeat until the potential decreases or s no longer sees x: let s be the start of the path, and $s' = \text{ccw}(s)$, and perform
   $$s'x \to sx, \qquad \text{in}(s) \to s's.$$
   If the potential has not decreased, then follow the procedure DiagObstrReduction with the start s and diagonal $s'x$.

(b) If x does not see s, follow the procedure DiagObstrReduction with the start s and diagonal $s'x$.

6. Otherwise, the outgoing edge of s' is a bridge.
   (a) If $P[s : s'] \not\subset L_1$, Let $u = \text{succ}(s')$.
   If s', u, succ(u) are not in clockwise order, flip ccw(s')s' to ss', rename ccw(s') as s, and reflect the path so that s' becomes ccw(s) again, and the order of s', u, succ(u) becomes clockwise.
   Let $v = \text{ccw}(u)$ and $v' = \text{succ}(v)$.
      i. If s sees v', perform $vv' \to sv'$, then $s'u \to vu$.
      ii. Else let pq be the last small diagonal of $P[v' : s']$, with $p = \text{pred}(q)$, and let $p' = \text{ccw}(p)$. Then flip $pq \to sq$ and recursively apply InnerEndpointsReduce using p' as the neighbor of the new start p.
   (b) Otherwise $P[s : s'] \subset L_1$. Let uv be the first small diagonal of $P[s : s']$, with $u = \text{pred}(v)$. let u' be the neighbor of u not seen by s'. Then flip uv to sv, and recursively apply InnerEndpointsReduce using u' as the neighbor of the new start u.

Proof (Lemma 7). Let P be an st-path with both endpoints on the inner layer. Let $s' \in L_1$ be a cyclic neighbor of s that does not share an edge with s, and assume $s' = \text{ccw}(s)$ by symmetry.

In the first three cases we try to create a new level edge between s and s'.

- **Case 1:** If s does not see s', then ObstrReduction reduces the potential in $O(1)$ according to Lemma 8.
- **Case 2:** If s sees s' and the incoming edge of s' is a non level edge, then flipping it to ss' instantly decreases the potential
- **Case 3:** If s sees s' and $s' = t$ is the end, then we can pick any non level edge and flip it to ss' to decrease the potential.

So now, suppose that s sees s', and the incoming edge of $s' \neq t$ is a level edge. Cases 4-6 follow different strategies depending on the type of out(s').

- **Case 4:** If out(s') is a small diagonal, the idea is that we can rotate that diagonal, with $x = \text{succ}(s')$ as an anchor, via the flips $s'x \to sx$ followed by in(s) $\to s's$. We repeat applying this case until a level edge is created as illustrated in Figure 8. The validity of the sequence comes from two loop invariants:
  - $s, s' \in L_1$ are neighbors and see each other;
  - $s'x$ is a small diagonal with $x = \text{succ}(s')$,

  and the termination is implied by the fact that after an iteration of the loop, the couple (s, s') is shifted clockwise, so if the loop does not stop earlier, the diagonal $s'x$ will rotate until it becomes a level edge, when s' becomes ccw(x). This loop can rotate that diagonal at most $O(n_1)$ times, meaning that the reduction step in this case takes $O(n_1)$ flips.

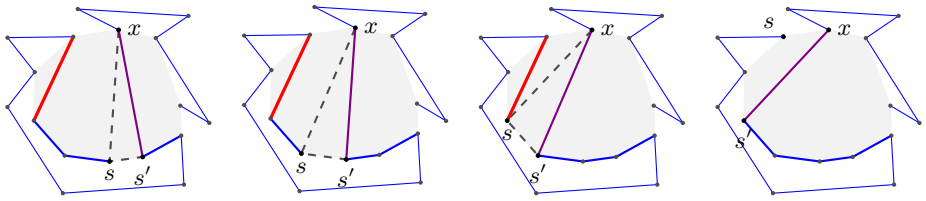


Fig. 8: Rotation of a small diagonal in $O(n_1)$ flips until the potential decreases.

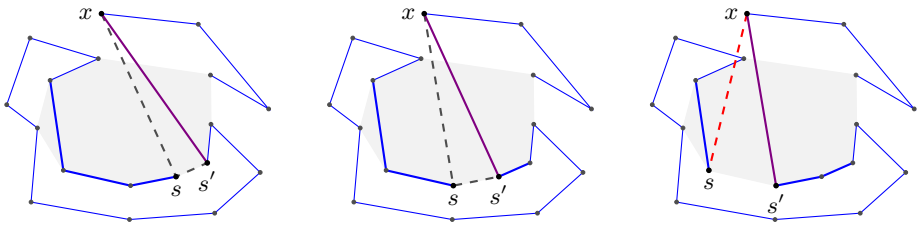


Fig. 9: Rotation of a medium diagonal in Case 5.(a) of InnerReduction in $O(n_1)$ flips.

– **Case 5:** If $\text{out}(s')$ is a medium diagonal, the idea is to use first rotate the diagonal until the potential diminishes as illustrated in Figure 9, or $x = \text{succ}(s')$ does not see s (Case 5.(a)), and then if the potential has not already be diminished, use the procedure DiagObstrReduction on the obstruction between s and x (Case 5.(b)).

  The proof for the validity of Case 5.(a) is similar to Case 4, so let's focus on Case 5.(b). If s does not see x, and t is on the other side of the medium diagonal, we can show that the path verifies the premises of Lemma 9. Let $u = s'$ and $v = x$. Then since v does not see s, the path must enter the triangle $\triangle suv$ without crossing uv (planarity) or su (assumption). Thus the existance of an inner point $p \in \triangle suv$ the closest to uv. Then we can prove that p sees s and v.

  Also, we can observe that the prefix $P[s : u]$ must contain an outer point. Let $u' \in L_0$ be the last one of that prefix. Then the idea is that $P[v : u']$ separates the set of point in 2 parts, and the prefix must visit all the points of it's left side. With this observation and the fact that s sees u, we can deduce that the path cannot visit any outer point before reaching p without crossings. And similar planarity restrictions enforce the prefix $P[s : p]$ to visit a subset of points of L_1 in clockwise order, and the subpath $P[p : u']$ to contain at least all the points of L_0 between $v' = \text{ccw}(v)$ and u' in counterclockwise order. Therefore the premises of DiagObstrReduction are met and it can reduce the potential in $O(1)$ flips. An example of a path in this case is given in Figure 10.

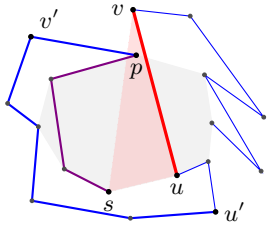


Fig. 10: Prefix structure of the path in Case 5.(b) of InnerReduction

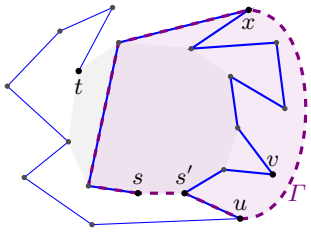


Fig. 11: Illustration of the dashed polygonal line separating the points in two parts in Case 6.(a) of InnerReduction

– **Case 6:** If out(s') is a bridge, the natural goal is to remove it. We subdivise that case since the path can have different properties depending on the shape of the prefix.
  - **Case 6.(a):** If $P[s : s'] \not\subset L_1$ To enforce stong properties on the path, suppose s', $u = \text{succ}(s')$ and succ(u) are in clockwise order (otherwise we can come back to this case with a single flip). Then let x be the first outer point visited by the path. Then we can show that the polygonal line formed by $us's$ and $P[s : x]$ splits the points in half, with the suffix $P[u : t]$ on the left side, and the subpath $P[x : s']$ on the left as illustrated in Figure 11. We can deduce from this structure that the prefix $P[s : u]$ must visit every outer point from x to u in clockwise order. So the last outer point visited before u was $v = \text{ccw}(u)$. Let $v' = \text{succ}(v)$.
    If s sees v' (Case 6.(a).i), then the flip $vv' \to sv'$ transforms a bridge or medium diagonal into a small diagonal, and $s'u \to vu$ creates a level edge on the L_0 as illustrated in Figure 12

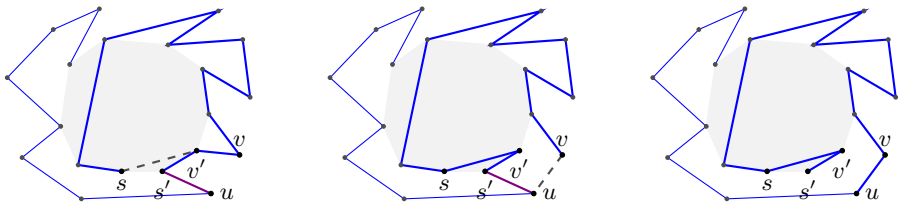


Fig. 12: Reduction of the potential in Case 6.(a).ii of InnerReduction

Otherwise (Case 6.(a).ii), there must be a small diagonal in $P[v' : s']$ blocking the vision between s and v. Let pq be the last small diagonal of that subpath. Then $P[q : s']$ only contains level edges and s sees s' so the flip $pq \to sq$ is valid and does not change the potential.

Now, let's show that applying the same procedure on the new path does not make a loop. Let $p' = \text{ccw}(p)$. That point does not share an edge with p because otherwise $P[v' : p']$ would have had to cross pq. Therefore when applying this procedure, the choice of p' for the neighbor of the start p of the new path P' is valid, and we can show that in P', if the incoming edge of p' is a level edge, its outgoing edge is a bridge and it sees p, then that bridge must be crossing the subpath $P'[p : p']$. Therefore Case 6 cannot be applied again in P', so one of the other ones will. So in the worst case, the potential will be diminished in $O(n_1)$ flips. An illustration of that case is presented in Figure 13.

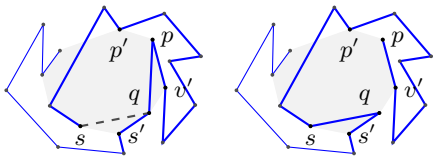


Fig. 13: Reduction from Case 6.(a).ii to other cases of InnerReduction.

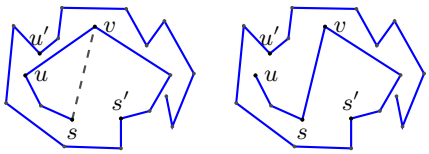


Fig. 14: Reduction from Case 6.(b) to other cases of InnerReduction.

- **Case 6.(b)** Otherwise $P[s : s'] \subset L_1$. In this last case, there is no immediate reduction possible in the general case, so we perform a flip to reduce that case to the other ones. Let uv be the first small diagonal of $P[s : s']$. It must exist as $t \in L_1$ is not in $P[s : s']$. Then u has neighbor u' not seen by s, and u' cannot be connected to u because otherwise $P[u' : s']$ would need to cross uv.
  So after the flip $uv \to sv$, u becomes the start of the new path P' and u' is a valid choice for the neighbor of u in the recursive call of the procedure for P'. Also, we can be sure that P' will not fall again into Case 6.(b) since $u \notin P[s : s']$, so $u \notin P'[u : s']$ and therefore $P'[u : u'] \not\subset L_1$. So in the worst case, the potential will be diminished in $O(n_1)$ flips. This reduction of case is represented in Figure 14.

Therefore each case uses at most $O(n_1)$ flips to reduce the potential.

4.3 Obstruction procedures and proof of Lemmas 8 and 9

Procedure ObstrReduction.
Input: a path P whose start $s \in L_1$ does not see one of its neighbor $q \in L_1$.
Assume $q = \text{ccw}(s)$ by symmetry.
Let uv be the closest edge to s intersecting sq, with $u = \text{pred}(v)$.

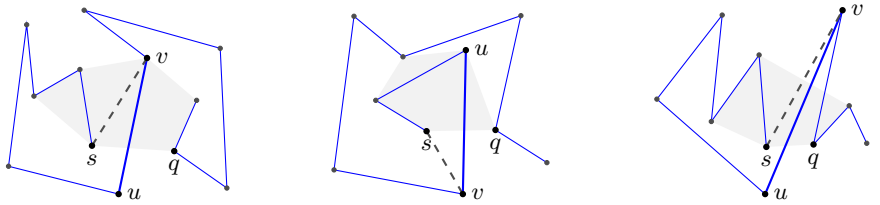
1. If v sees s, perform $uv \longrightarrow sv$.
2. Else follow the procedure DiagObstrReduction.

Procedure DiagObstrReduction. *Input:* A path P starting at $s \in L_1$ such that there is a medium or big diagonal uv, with $v \in L_0$, intersecting the convex hull of the inner layer next to s and another vertex p. Suppose also that p is inside the triangle $\triangle suv$, and that the prefix $P[s:p]$ visits a subset of L_1 in circular order from s to p, and the subpath $P[p:u']$ visits all the points between s and p skipped by the prefix (without particular order), and all the outer points in order from v' (the first neighbor of v visited by after p) to u' (the last outer point visited before u).

Suppose without loss of generality that the crossing of the diagonal uv is counterclockwise to s. Let $s' = \text{pred}(p)$, $v' = \text{ccw}(v)$ and $p' = \text{ccw}(p)$.

1. If $\text{succ}(p) \in L_1$, perform
$$s'p \to sp, \qquad p\,\text{succ}(p) \to s'\,\text{succ}(p), \qquad uv \to pv.$$
2. Else if $p' \in P[s:p]$, perform
$$s'p \to sp, \qquad pv' \to s'v', \qquad uv \to pv.$$
3. Else if p' does not see p:
Let ab, with $b = \text{succ}(a)$, be the closest edge intersecting pp' and perform
$$s'p \to sp, \qquad ab \to s'b.$$
4. Otherwise, perform
$$s'p \to sp.$$
   (a) If $\text{succ}(p') \in L_1$, let $y = \text{succ}(p')$ and perform
$$p'y \to s'y, \qquad v'p \to p'p, \qquad uv \to vv'.$$
   (b) Else if $\text{pred}(p'), \text{succ}(p') \in L_0$,
   Let $x = \text{pred}(p')$ and $y = \text{succ}(p')$ and perform
$$xp' \to s'p', \qquad p'y \to xy.$$
   (c) Otherwise, $\text{succ}(p') \in L_0$ and $\text{pred}(p') \in L_1$.
      i. if s' does not see $\text{cw}(s')$, recursively apply ObstrReduction on the obstruction between s' and $\text{cw}(s')$.
      ii. Else if the incoming edge of $\text{cw}(s')$ is not level, perform
$$\text{in}(\text{cw}(s')) \to s'\,\text{cw}(s').$$
      iii. Otherwise, perform
$$\text{in}(\text{cw}(s')) \to s'\,\text{cw}(s'), \qquad \text{in}(p') \to p'\,\text{cw}(\text{cw}(s')).$$

Proof (Lemma 8). Let P be an st-path whose start $s \in L_1$ does not see one of its neighbor q, and assume $q = \text{ccw}(s)$ by symmetry. Let uv be the closest edge to s intersecting sq. Then uv is either a medium diagonal or a big diagonal, so it has a lot of weight. Therefore the goal is to remove it. We can observe that because of the definition of uv, s sees either u or v.



(a) v is an inner point (b) v is an outer point next to s (c) v is an outer point far from s

Fig. 15: Immediate potential reduction in case 1 of ObstrReduction

- **Case 1.** If s sees v, then the flip $uv \to sv$ is legal and we can show by case exhaustion that no matter the type of uv initially, that flip decreases the potential. The different subcases are illustrated in Figure 15.
- **Case 2.** Otherwise, s does not see v but s sees u, so we can deduce that the diagonal uv must be a big diagonal, otherwise, s would have been able to see v. We can now show that the path has the strong structural properties that enable the use of the procedure DiagObstrReduction.
  Since s does not see v, the path must enter the triangle $\triangle suv$ without crossing uv (planarity) or su (assumption). so that triangle contains at least 1 point of L_1 inside of it. Let p be the closest point to uv inside that triangle. Then we can prove that there is no other point inside the concave quadrilateral $suvp$ or that would contradict the definition of p, therefore p sees both s and v.
  Now, since uv is a big diagonal, it splits the point set in two parts, and it particular, it must separate the start s from the end t because otherwise an entire side could not be visited without crossings. This means that the prefix $P[s:u]$ must visit the entire left side of uv. Moreover, since the quadrilateral $suvp$ cannot be crossed by the path, we can deduce that the prefix $P[s:p]$ cannot visit any point of L_0 without introducing crossings or violating the definition of uv or p. Therefore $P[p:u]$ must visit at least all the outer points from $\text{ccw}(v)$ to u. With the same argument, we can observe that $P[s:p]$ must visit that subset of L_1 in clockwise order and all the outer points of the left side are visited in counterclockwise order in $P[p:u]$. The structure of the prefix of that path is illustrated in Figure 16
  Therefore the premises of the procedure DiagObstrReduction are met and it can reduce the potential in $O(1)$ flips.

Therefore in both cases, the potential of the path can be reduced in $O(1)$ flips. □

Proof (Lemma 9). The main goal of this procedure is to remove the diagonal uv by flipping it to pv, or to remove any other heavy weight edge encountered in the prefix. This procedure does it by successively eliminating simple cases

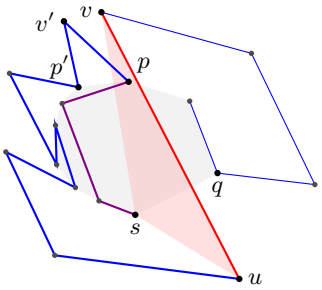


Fig. 16: Illustration of the particular prefix the path has in Case 2 of ObstrReduction

to constrain the shape of the path until it forces the appearance of a medium diagonal in the prefix that is easy to remove.

Let $s' = \text{pred}(p)$. A first easy case is when the successor of p is an inner point.

– **Case 1.** If $\text{succ}(p) \in L_1$, we can make p the new starting point with 2 flips: $s'p \to sp$ followed by $p\,\text{succ}(p) \to s'\,\text{succ}(p)$. These flips are legal since p sees s and $s'p\,\text{succ}(p)$ is an empty triangle of L_1 where two of the borders cannot be crossed by planarity, so s' sees $\text{succ}(p)$. After these two flips, the potential may have increased by 1 at most, However, the last flip $uv \to pv$ transform a medium or big diagonal into a bridge, diminishing the potential by at least 2. So at the end of that sequence, the potential has decreased by at least 1. This flip sequence is illustrated in Figure 17

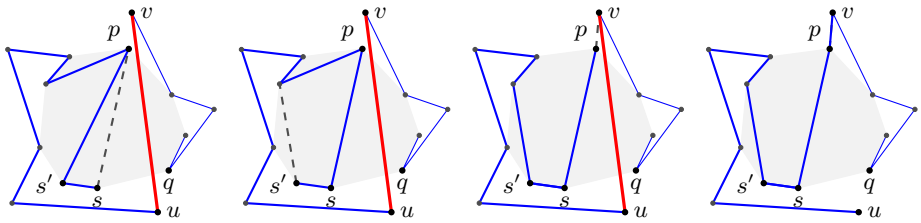


Fig. 17: Reduction of potential in Case 1 of DiagObstrReduce

Suppose now that $\text{succ}(p) \in L_0$. Since by assumption the subpath $P[p:u]$ visits all the outer points from $v' = \text{ccw}(v)$ to u in counterclockwise order, we can deduce that $\text{succ}(p) = v'$. Let $p' = \text{ccw}(p)$. The case analysis is now refined based on the position of p' in the path

– **Case 2.** If $p' \in P[s:p]$, then $p' = s'$ since by assumption on $P[s:p]$, that prefix visit points of L_1 in clockwise order. The sequence of flips we perform

in this case is the same as in Case 1, except this time the second flip is legal because v' and p' see each other. Indeed, if that was not the case, the convexity of L_1 would prevent the quadrilateral $suvp$ from being concave, which would mean that s would see v, contradicting the assumption made on the initial path.
So at the end of that sequence (illustrated in Figure 18, the potential has decreased by at least 1 in $O(1)$ flips.

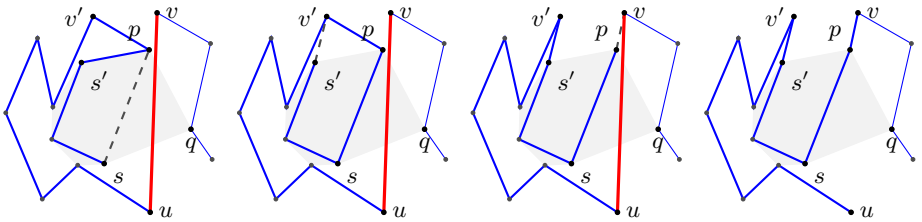


Fig. 18: Reduction of potential in Case 2 of DiagObstrReduction

– **Case 3.** If $p' \notin P[s : p]$ and p' does not see p, then there is a medium diagonal blocking the view between those points. Let ab be the closest edge to p intersecting pp'. Then we can show that necessarily $a \in L_0$ and $b \in L_1$, otherwise there would need to be some crossings in the path. We can also show that s' sees b since the triangle $\triangle s'pb$ is empty and the path cannot cross $s'p$ (planarity) nor can it cross bp (definition of ab). So after the first neutral flip $s'p \to sp$, the flip $ab \to s'b$ reduced the potential by at least 2 as illustrated in Figure 19.

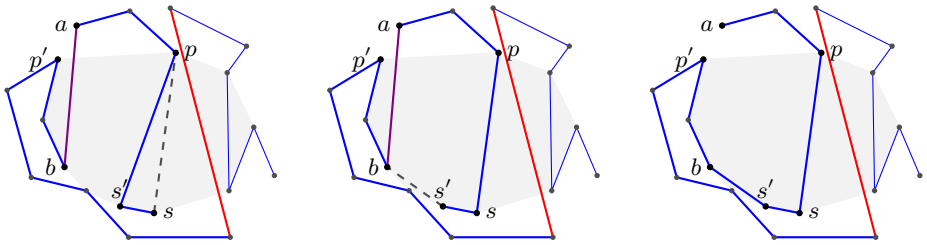
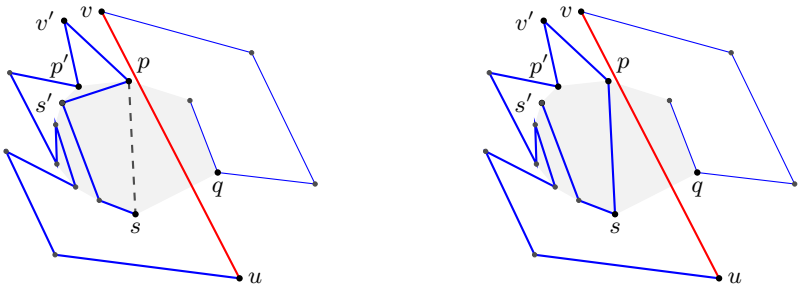


Fig. 19: Reduction of potential in Case 3 of DiagObstrReduction.

Now suppose p' sees p and $p' \notin P[s : p]$ (Case 4.), i.e. $p' \in P[v' : u]$. Then like all the previous cases, we start flipping $s'p \to sp$ to form a new path P' as

illustrated in Figure 20. In this case, we know that s' is not a cyclic neighbor of p in the path, so $s'p$ is a small diagonal, just like sp. Therefore $\Phi(P) = \Phi(P')$.

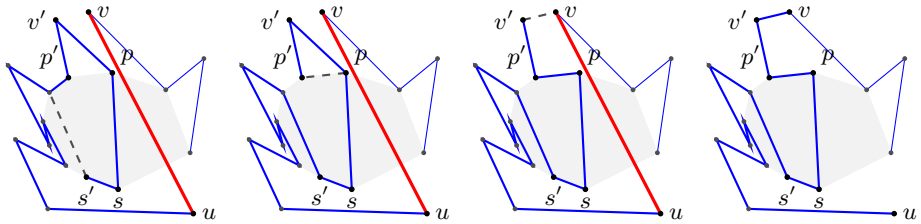


(a) The original path P. (b) The new path P'.

Fig. 20: First flip of Case 4 of DiagObstrReduction.

We refine again the case analysis based on the position of the predecessor and the successor of p'. We can observe that 2 cases can be solved easily.

– **Case 4.(a)**: If $\text{succ}(p') \in L_1$, This case resembles the flip sequence from Case 1: with the same empty triangles arguments we can deduce the required visibility to prove that all the flips performed in this case are legal. The potential could stay the same after the first 2 flips but decreases by at least 3 with the last one as illustrated in Figure 21.

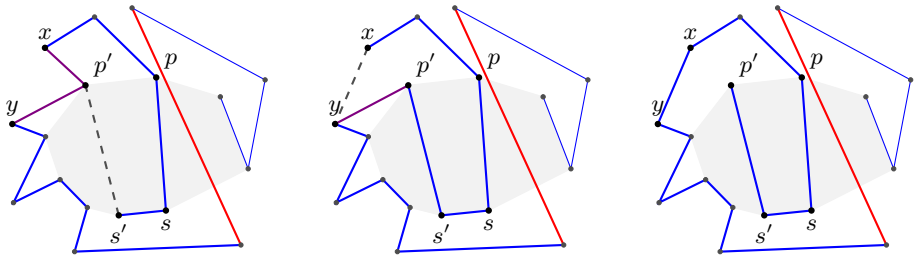


(a) (b) (c) (d)

Fig. 21: Reduction of potential in Case 4.(a) of DiagObstrReduction.

– **Case 4.(b)** If $\text{succ}(p') \in L_0$ and $\text{pred}(p') \in L_0$, then by the assumption on the path, we know that $x = \text{pred}(p')$ is the clockwise neighbor of $y = \text{succ}(p')$. In this case we can just connect them in 2 flips. By the same empty triangle argument we know that s' sees p', and since they are outer layer neighbor, we know that x sees y, thus the validity of both flips. Moreover, that sequence

of flips transforms two bridges into a level edge and a small diagonal, so the potential decreases by 1 as illustrated in Figure 22.



(a) (b) (c)

Fig. 22: Reduction of potential in Case 4.(b) of DiagObstrReduction.

Now suppose $\text{succ}(p') \in L_0$ and $\text{pred}(p') \in L_1$. In this case, the outgoing edge of p' is so constrained that it is forced to be a medium diagonal as we will show in the relevant subcase. Therefore, we shift the objective from flipping uv to flipping $\text{out}(p')$. However, the edge is in the wrong direction to be flipped directly. Therefore we use $\text{cw}(s')$ to change the orientation of that edge to flip it (Case 4.(c).iii), if the potential cannot be reduced immediately (Case 4.(c).i-ii).

- **Case 4.(c).i** If s' does not see $\text{cw}(s')$, because of the structure of the prefix assumed in the hypotheses of the lemma, that blocking edge cannot be a big diagonal, therefore it is a medium one, so applying procedure ObstructionReduce will reduce the potential with its case 1 in 1 flip.
- **Case 4.(c).ii** If s' sees $\text{cw}(s')$ and $\text{in}(\text{cw}(s'))$ is not a level edge, then the flip $\text{in}(\text{cw}(s')) \to s'\,\text{cw}(s)$ creates a level edge, reducing the potential.
- **Case 4.(c).iii** if s' sees $\text{cw}(s')$ and $\text{in}(\text{cw}(s'))$ is a level edge, then the flip $\text{in}(\text{cw}(s')) \to s'\,\text{cw}(s)$ makes $\text{cw}(\text{cw}(s'))$ the new start, and reverses the order of the path in $P'[s' : \text{cw}(s')]$ as displayed in Figure 23.
  The second flip is valid since after the first flip, p' sees s' which is connected to $\text{cw}(s')$ which sees $\text{cw}(\text{cw}(s'))$, and all these points form an empty quadrilateral, so p sees $\text{cw}(\text{cw}(s'))$.
  Now to prove that the potential reduces, we must prove that in P', the outgoing edge of p' is a medium diagonal, and that after the first flip, the order of the path in p' is reversed, meaning that that medium diagonal is flipped for a small one in the second flip.
  Since $p' \in L_1$ and $\text{succ}(p') \in L_0$, the outgoing edge of p' is either a bridge or a medium diagonal. Suppose for contradiction that it is a bridge, and consider the simple loop $\Gamma = vp \cup pp' \cup p'\,\text{succ}(p') \cup \gamma_{v,\text{succ}(p')}$ where $\gamma_{v,\text{succ}(p')}$ lies outside $\text{conv}(S)$ except at its endpoints v and $\text{succ}(p')$. By the Jordan curve theorem, Γ separates the plane into two regions. Since $p'\,\text{succ}(p')$ is a

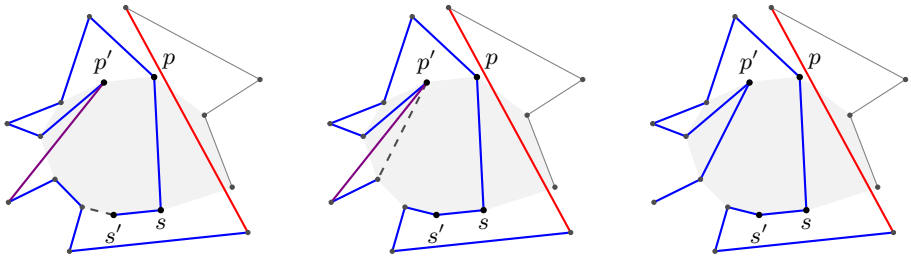


Fig. 23: Reduction of potential in Case 4.(c).iii of DiagObstrReduction.

bridge and p' is a neighbor of p, the region containing $v' = \mathrm{ccw}(v)$ contains no point of L_1 other than p and p' at its border, so in particular, it does not contain $\mathrm{pred}(p')$. Therefore $P'[v' : \mathrm{pred}(p')]$ must cross Γ, but it cannot cross vp and pp' (assumption), it cannot cross $p'\,\mathrm{succ}(p')$ (planarity) and it cannot cross $\gamma_{v,\mathrm{succ}(p')}$ (outside the hull). This contradiction shows that $\mathrm{out}(p')$ is a medium diagonal.
We can also show with a different loop that p' must be visited before $\mathrm{cw}(s')$ in P', therefore after the first flip, the medium diagonal of p' becomes its incoming edge and the potential diminishes by at least 2 in the second flip.

So in any case, the potential diminishes by at least 1 in $O(1)$ flips. □

5 Proof of Lemma 5

Using Lemma 4 alone, we would get a flip complexity of $O(n_1(n_0 + n_1))$ in the worst case. In this section we prove that the expensive cases where we may need $\Omega(n_1)$ flips in a reduction step does not happen too often. In fact, we show that at most $O(n_1)$ reduction steps may need strictly more than 1 flip to be performed.

First, we can observe that the number of bridges $B(P)$, the number of small diagonals $S(P)$, and the number of medium diagonals $M(P)$ are all bounded by $2n_1$ as these edge types all have an extremity in L_1. Therefore, we just need to show that most of the big diagonals can be removed in 1 flip throughout the reduction process.

Let $\mathcal{I}$ be the set of boundary segments of $\mathrm{conv}(L_1)$ induced by the vertices of L_1.

Let P be a path. For $A, B \in \mathcal{I}$, let $G_P(A, B)$ denote the set of big diagonals whose two intersections with $\mathrm{conv}(L_1)$ belong respectively to the intervals A and B. We also denote by $G_P(\emptyset)$ the set of big diagonals that do not intersect $\mathrm{conv}(L_1)$, so that $G_P(\emptyset) \cup \bigcup_{A,B\in\mathcal{I}} G_P(A, B)$ forms a partition of the set of big diagonals of P. Because of the planarity of the path, the support of that partition is $O(n_1)$ (Lemma 10).

Lemma 10. *For any given path P, the number of non-empty sets $G_P(A, B)$ is $O(n_1)$.*

Proof. Associate to each non-empty set a representative big diagonal. Since big diagonals of a plane path are pairwise non-crossing, these representatives form a family of non-crossing chords of dual of the convex polygon $\text{conv}(L_1)$.

The graph whose vertices are the segments of $\text{conv}(L_1)$ (i.e. $\mathcal{I}$) and whose edges correspond to non-empty sets is therefore outerplanar. Since $|\mathcal{I}| = n_1$, it contains only $O(n_1)$ edges. Hence only $O(n_1)$ sets are non-empty.

This partition of the big diagonals enables the following observation.

Observation 11. *Let P be a path. As soon as one big diagonal of $G \in \{G_P(\emptyset)\} \cup \{G_P(A, B), A, B \in \mathcal{I}\}$ as been flipped, we can transform every remaining diagonal of G to level edges using one flip per diagonal.*

Proof. Suppose $G = G_P(A, B) \neq \emptyset$, for $A, B \in \mathcal{I}$, and suppose we are about to flip one diagonal of G. We can order them by the order in which they appear in the path (that would correspond to left-to-right or right-to-left if they were vertical). Let uv be the first and xy be last diagonal of G.

Since each big diagonal separates the point set in two parts, we can deduce that they also separate the endpoints (planarity and Hamiltonicity). Applying this principle on uv and xy, we can deduce that there is no endpoint in the region of the plane between the lines uv and xy. So no other diagonal from G can be visible by the endpoints, therefore, by symmetry we can suppose that uv is the diagonal we are about to flip, and that u will become the next starting point.

By definition of $G_P(A, B)$, there is no inner point in the region between the lines uv and xy, so the vertices of the subpath $P[u : y]$ are in convex position. This forces a particular structure where each successive diagonal are connected by level edges, and as soon as uv has been flipped, so that u becomes an endpoint, we can successively flip all the diagonals in order to level edges like in the convex case.

The proof is the similar when $G = G_P(\emptyset)$. □

Therefore, according to Observation 11, at most $1 + |G_P(A, B) \neq \emptyset| = O(n_1)$ big diagonals may need more than 1 flip in the reconfiguration process. We can now prove Lemma 5.

Proof (Lemma 5). According to Observation 11, at most $B(P) + S(P) + 3M(P) + 4(1 + |G_P(A, B) \neq \emptyset|) = O(n_1)$ reduction steps may require more than 1 flip. □

Acknowledgments and Disclosure of Interests. Use of AI: Generative AI tools were used for literature search, language polishing, manuscript formatting, and improving the exposition of the paper. They were also used as an auxiliary aid to review the presentation of proofs and identify possible inconsistencies. Generative AI tools were not used to generate the scientific results, proofs, algorithms, or conclusions of this work. All AI-assisted content was reviewed and verified by the authors, who take full responsibility for the final manuscript. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Aichholzer, O., Dorfer, J.: A linear time algorithm for the flip distance between plane spanning paths in convex point sets. Comput. Geom. Topol. **5**, 2:1–2:17 (2026)
2. Aichholzer, O., Knorr, K., Mulzer, W., Obenaus, J., Paul, R., Vogtenhuber, B.: Flipping plane spanning paths. In: Proc. 17th Int. Conf. Workshops Algorithms Comput. (WALCOM). LNCS, vol. 13973, pp. 49–60 (2023)
3. Akl, S.G., Islam, M.K., Meijer, H.: On planar path transformation. Inf. Process. Lett. **104**, 59–64 (2007)
4. Boucard, V., da Fonseca, G.D., Rivier, B.: Further connectivity results on plane spanning path reconfiguration. arXiv:2407.00244 (2024)
5. Chang, J.M., Wu, R.Y.: On the diameter of geometric path graphs of points in convex position. Inf. Process. Lett. **109**, 409–413 (2009)
6. Kleist, L., Kramer, P., Rieck, C.: On the connectivity of the flip graph of plane spanning paths. In: Proc. 50th Int. Workshop Graph-Theoretic Concepts Comput. Sci. (WG). LNCS, vol. 14760, pp. 327–342 (2025)
7. Rivera-Campo, E., Urrutia-Galicia, V.: Hamilton cycles in the path graph of a set of points in convex position. Comput. Geom. **18**, 65–72 (2001)