Lines: 355

Flipping Plane Spanning Paths ⋆

Oswin Aichholzer[1], Kristin Knorr[2], Wolfgang Mulzer[2], Johannes Obenaus[2], Rosna Paul[1], and Birgit Vogtenhuber[1]

[1] Institute of Software Technology, Graz University of Technology, Austria
{oaich,bvogt,ropaul}@ist.tugraz.at

[2] Institut für Informatik, Freie Universität Berlin, Germany
{kristin.knorr,wolfgang.mulzer,johannes.obenaus}@fu-berlin.de

Abstract. Let S be a planar point set in general position, and let $\mathcal{P}(S)$ be the set of all plane straight-line paths with vertex set S. A flip on a path $P \in \mathcal{P}(S)$ is the operation of replacing an edge e of P with another edge f on S to obtain a new valid path from $\mathcal{P}(S)$. It is a long-standing open question whether for every given point set S, every path from $\mathcal{P}(S)$ can be transformed into any other path from $\mathcal{P}(S)$ by a sequence of flips. To achieve a better understanding of this question, we show that it is sufficient to prove the statement for plane spanning paths whose first edge is fixed. Furthermore, we provide positive answers for special classes of point sets, namely, for wheel sets and generalized double circles (which include, e.g., double chains and double circles).

Keywords: flips · plane spanning paths · generalized double circles

1 Introduction

Reconfiguration is a classical and widely studied topic with various applications in multiple areas. A natural way to provide structure for a reconfiguration problem is by studying the so-called *flip graph*. For a class of objects, the flip graph has a vertex for each element and adjacencies are determined by a local flip operation (we will give the precise definition shortly). In this paper we are concerned with transforming plane spanning paths via edge flips.

Let S be a set of n points in the plane in general position (i.e., no three points are collinear), and let $\mathcal{P}(S)$ be the set of all plane straight-line spanning paths for S, i.e., the set of all paths with vertex set S whose straight-line embedding

⋆ This work was initiated at the 2nd Austrian Computational Geometry Reunion Workshop in Strobl, June 2021. We thank all participants for fruitful discussions. J.O. is supported by ERC StG 757609. O.A. and R.P. are supported by FWF grant W1230. B.V. is supported by FWF Project I 3340-N35. K.K. is supported by the German Science Foundation (DFG) within the research training group 'Facets of Complexity' (GRK 2434). W.M. is partially supported by the German Research Foundation within the collaborative DACH project *Arrangements and Drawings* as DFG Project MU 3501/3-1, and by ERC StG 757609.

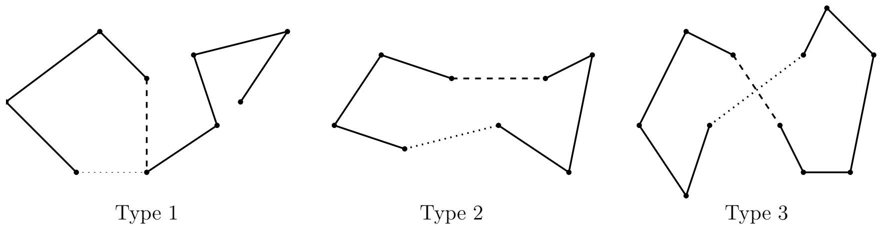

Figure 1. The three types of flips in plane spanning paths.

on S is crossing-free. A *flip* on a path $P \in \mathcal{P}(S)$ is the operation of removing an edge e from P and replacing it by another edge f on S such that the graph $(P \setminus e) \cup f$ is again a path from $\mathcal{P}(S)$. Note that the edges e and f might cross. The *flip graph* on $\mathcal{P}(S)$ has vertex set $\mathcal{P}(S)$ and two vertices are adjacent if and only if the corresponding paths differ by a single flip. The following conjecture will be the focus of this paper:

Conjecture 1 (Akl et al. [3]). For every point set S in general position, the flip graph on $\mathcal{P}(S)$ is connected.

Related work. For further details on reconfiguration problems in general we refer the reader to the surveys of Nishimura [10] and Bose and Hurtado [4]. Connectivity properties of flip graphs have been studied extensively in a huge variety of settings, see, e.g., [6,7,8,9,11] for results on triangulations, matchings and trees.

In our setting of plane spanning paths, flips are much more restricted, making it more difficult to prove a positive answer. Prior to our work only results for point sets in convex position and very small point sets were known. Akl et al. [3], who initiated the study of flip connectivity on plane spanning paths, showed connectedness of the flip graph on $\mathcal{P}(S)$ if S is in convex position or $|S| \leq 8$. In the convex setting, Chang and Wu [5] derived tight bounds concerning the diameter of the flip graph, namely, $2n-5$ for $n = 3, 4$, and $2n-6$ for $n \geq 5$.

For the remainder of this paper, we consider the flip graph on $\mathcal{P}(S)$ (or a subset of $\mathcal{P}(S)$). Moreover, unless stated otherwise, the word *path* always refers to a path from $\mathcal{P}(S)$ for an underlying point set S that is clear from the context.

Flips in plane spanning paths. Let us have a closer look at the different types of possible flips for a path $P = v_1, \ldots, v_n \in \mathcal{P}(S)$ (see also Figure 1). When removing an edge $v_{i-1}v_i$ from P with $2 \leq i \leq n$, there are three possible new edges that can be added in order to obtain a path (where, of course, not all three choices will necessarily lead to a plane path in $\mathcal{P}(S)$): v_1v_i, $v_{i-1}v_n$, and v_1v_n. A flip of *Type 1* is a valid flip that adds the edge v_1v_i (if $i > 2$) or the edge $v_{i-1}v_n$ (if $i < n$). It results in the path $v_{i-1}, \ldots, v_1, v_i, \ldots, v_n$, or the path $v_1, \ldots, v_{i-1}, v_n, \ldots, v_i$. That is, a Type 1 flip inverts a contiguous chunk from one of the two ends of the path. A flip of *Type 2* adds the edge v_1v_n and has the

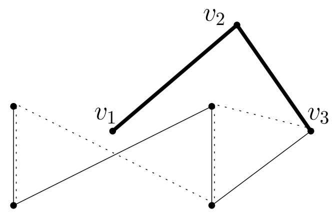

Figure 2. Example where the flip graph is disconnected if the first three vertices of the paths are fixed. No edge of the solid path can be flipped, but there is at least one other path (dotted) with the same three starting vertices.

additional property that the edges $v_{i-1}v_i$ and v_1v_n do not cross. In this case, the path P together with the edge v_1v_n forms a plane cycle. If a Type 2 flip is possible for one edge $v_{i-1}v_i$ of P, then it is possible for all edges of P. A Type 2 flip can be simulated by a sequence of Type 1 flips, e.g., flip v_1v_2 to v_1v_n, then flip v_2v_3 to v_1v_2, then v_3v_4 to v_2v_3, etc., until flipping $v_{i-1}v_i$ to $v_{i-2}v_{i-1}$. A flip of *Type 3* also adds the edge v_1v_n, but now the edges v_1v_n and $v_{i-1}v_i$ cross. Note that a Type 3 flip is only possible if the edge v_1v_n crosses exactly one edge of P, and then the flip is possible only for the edge $v_{i-1}v_i$ that is crossed.

Contribution. We approach Conjecture 1 from two directions. First, we show that it is sufficient to prove flip connectivity for paths with a fixed starting edge. Second, we verify Conjecture 1 for several classes of point sets, namely wheel sets and generalized double circles (which include, e.g., double chains and double circles).

Towards the first part, we define, for two distinct points $p, q \in S$, the following subsets of $\mathcal{P}(S)$: let $\mathcal{P}(S,p)$ be the set of all plane spanning paths for S that start at p, and let $\mathcal{P}(S,p,q)$ be the set of all plane spanning paths for S that start at p and continue with q. Then for any S, the flip graph on $\mathcal{P}(S,p,q)$ is a subgraph of the flip graph on $\mathcal{P}(S,p)$, which in turn is a subgraph of the flip graph on $\mathcal{P}(S)$. We conjecture that all these flip graphs are connected:

Conjecture 2. For every point set S in general position and every $p \in S$, the flip graph on $\mathcal{P}(S,p)$ is connected.

Conjecture 3. For every point set S in general position and every $p, q \in S$, the flip graph on $\mathcal{P}(S,p,q)$ is connected.

Towards Conjecture 1, we show that it suffices to prove Conjecture 3:

Theorem 1. *Conjecture 2 implies Conjecture 1.*

Theorem 2. *Conjecture 3 implies Conjecture 2.*

Note that the analogue of Conjecture 3 for paths where the first $k \geq 3$ vertices are fixed, does not hold: Figure 2 shows a counterexample with 7 points and $k = 3$.

Towards the flip connectivity for special classes of point sets, we consider wheel sets and generalized double circles. A point set is in *wheel configuration* if it has exactly one point inside the convex hull. For generalized double circles we defer the precise definition to Section 4, however, intuitively speaking a generalized double circle is obtained by replacing each edge of the convex hull by a flat enough concave chain of arbitrary size (as depicted on the right). We show that the flip graph is connected in both cases:

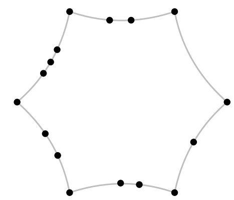

Theorem 3. *(⋆) Let S be a set of n points in wheel configuration. Then the flip graph (on $\mathcal{P}(S)$) is connected with diameter at most $2n-4$.*

Theorem 4. *(⋆) Let S be a set of n points in generalized double circle configuration. Then the flip graph (on $\mathcal{P}(S)$) is connected with diameter $O(n^2)$.*

Finally, we remark that using the order type database [1], we are able to computationally verify Conjecture 1 for every set of $n \leq 10$ points in general position (even when using only Type 1 flips).[3]

Notation. We denote the convex hull of a point set S by $\mathrm{CH}(S)$. All points $p \in S$ on the boundary of $\mathrm{CH}(S)$ are called *extreme points* and the remaining points are called *interior* points. All results marked by a (⋆) have a full proof in the full version of this paper [2].

2 A Sufficient Condition

In this section we prove Theorem 1 and Theorem 2.

Lemma 1. *(⋆) Let S be a point set in general position and $p, q \in S$. Then there exists a path $P \in \mathcal{P}(S)$ which has p and q as its end vertices.*

Theorem 1. *Conjecture 2 implies Conjecture 1.*

Proof. Let S be a point set and $P_s, P_t \in \mathcal{P}(S)$. If P_s and P_t have a common endpoint, we can directly apply Conjecture 2 and the statement follows. So assume that P_s has the endpoints v_a and v_b, and P_t has the endpoints v_c and v_d, which are all distinct. By Lemma 1 there exists a path P_m having the two endpoints v_a and v_c. By Conjecture 2 there is a flip sequence from P_s to P_m with the common endpoint v_a, and again by Conjecture 2 there is a further flip sequence from P_m to P_t with the common endpoint v_c. This concludes the proof. □

[3] The source code is available at `https://github.com/jogo23/flipping_plane_spanning_paths`.

Towards Theorem 2, we first have a closer look at what edges form *viable* starting edges. For a given point set S and points $p, q \in S$, we say that pq forms a *viable* starting edge if there exists a path $P \in \mathcal{P}(S)$ that starts with pq. For instance, an edge connecting two extreme points that are not consecutive along $\text{CH}(S)$ is not a viable starting edge. The following lemma shows that these are the only non-viable starting edges.

Lemma 2. *($\star$) Let S be a point set in general position and $u, v \in S$. The edge uv is a viable starting edge if and only if one of the following is fulfilled: (i) u or v lie in the interior of* $\text{CH}(S)$*, or (ii) u and v are consecutive along* $\text{CH}(S)$*.*

The following lemma is the analogue of Lemma 1:

Lemma 3. *($\star$) Let S be a point set in general position and $v_1 \in S$. Further let $S' \subset S$ be the set of all points $p \in S$ such that $v_1 p$ forms a viable starting edge. Then for two points $q, r \in S'$ that are consecutive in the circular order around v_1, there exists a plane spanning cycle containing the edges $v_1 q$ and $v_1 r$.*

Theorem 2. *Conjecture 3 implies Conjecture 2.*

Proof. Let S be a point set and $v_1 \in S$. Further let $P, P' \in \mathcal{P}(S, v_1)$. If P and P' have the starting edge in common, then we directly apply Conjecture 3 and are done. So let us assume that the starting edge of P is $v_1 v_2$ and the starting edge of P' is $v_1 v_2'$. Clearly $v_2, v_2' \in S'$ holds. Sort the points in S' in radial order around v_1. Further let $v_x \in S'$ be the next vertex after v_2 in this radial order and C be the plane spanning cycle with edges $v_1 v_2$ and $v_1 v_x$, as guaranteed by Lemma 3.

By Conjecture 3, we can flip P to $C \setminus v_1 v_x$. Then, flipping $v_1 v_2$ to $v_1 v_x$ we get to the path $C \setminus v_1 v_2$, which now has $v_1 v_x$ as starting edge. We iteratively continue this process of "rotating" the starting edge until reaching $v_1 v_2'$. □

Theorems 1 and 2 imply that it suffices to show connectedness of certain subgraphs of the flip graph. A priori it is not clear whether this is an easier or a more difficult task – on the one hand we have smaller graphs, making it easier to handle. On the other hand, we may be more restricted concerning which flips we can perform, or exclude certain "nice" paths.

3 Flip Connectivity for Wheel Sets

Akl et al. [3] proved connectedness of the flip graph if the underlying point set S is in convex position. They showed that every path in $\mathcal{P}(S)$ can be flipped to a *canonical path* that uses only edges on the convex hull of S. To generalize this approach to other classes of point sets, we need two ingredients: (i) a set of *canonical paths* that serve as the target of the flip operations and that have the property that any canonical path can be transformed into any other canonical path by a simple sequence of flips, usually of constant length; and (ii) a strategy to flip any given path to some canonical path.

Recall that a set S of $n \geq 4$ points in the plane is a *wheel set* if there is exactly one interior point $c_0 \in S$. We call c_0 the *center* of S and classify the edges on S as follows: an edge incident to the center c_0 is called a *radial* edge, and an edge along $\mathrm{CH}(S)$ is called *spine* edge (the set of spine edges forms the *spine*, which is just the boundary of the convex hull here). All other edges are called *inner* edges. The *canonical paths* are those that consist only of spine edges and one or two radial edges.

We need one observation that will also be useful later. Let S be a point set and $P = v_1, \ldots, v_n \in \mathcal{P}(S)$. Further, let v_i $(i \geq 3)$ be a vertex such that no edge on S crosses $v_1 v_i$. We denote the face bounded by $v_1, \ldots, v_i, v_1$ by $\Phi(v_i)$.

Observation 5. *Let S be a point set, $P = v_1, \ldots, v_n \in \mathcal{P}(S)$, and v_i $(i \geq 3)$ be a vertex such that no edge on S crosses $v_1 v_i$. Then all vertices after v_i (i.e., $\{v_{i+1}, \ldots, v_n\}$) must entirely be contained in either the interior or the exterior of $\Phi(v_i)$.*

Theorem 3. *($\star$) Let S be a set of n points in wheel configuration. Then the flip graph (on $\mathcal{P}(S)$) is connected with diameter at most $2n - 4$.*

Proof (Sketch). Let $P = v_1, \ldots, v_n \in \mathcal{P}(S)$ be a non-canonical path and w.l.o.g., let $v_1 \neq c_0$. We show how to apply suitable flips to increase the number of spine edges of P. By Lemma 2, $v_1 v_2$ can only be radial or a spine edge. In the former case we can flip the necessarily radial edge $v_2 v_3$ to the spine edge $v_1 v_3$. In the latter case, let v_a with $a \neq 2$ be a neighbor of v_1 along the convex hull. Then, either $v_{a-1} v_a$ is not a spine edge and hence, we can flip it to $v_1 v_a$, or otherwise we show, using Observation 5, that P actually already is a canonical path. □

4 Flip Connectivity for Generalized Double Circles

The proof for generalized double circles is in principle similar to the one for wheel sets but much more involved. For a point set S and two extreme points $p, q \in S$, we call a subset $CC(p, q) \subset S$ *concave chain* (chain for short) for S, if (i) $p, q \in CC(p, q)$; (ii) $CC(p, q)$ is in convex position; (iii) $CC(p, q)$ contains no other extreme points of S; and (iv) every line ℓ_{xy} through any two points $x, y \in CC(p, q)$ has the property that all points of $S \setminus CC(p, q)$ are contained in the open halfplane bounded by ℓ_{xy} that contains neither p nor q. Note that the extreme points p and q must necessarily be consecutive along $\mathrm{CH}(S)$. If there is no danger of confusion, we also refer to the spanning path from p to q along the convex hull of $CC(p, q)$ as the concave chain.

A point set S is in *generalized double circle* position if there exists a family of concave chains such that every inner point of S is contained in exactly one chain and every extreme point of S is contained in exactly two chains. We denote the class of generalized double circles by GDC. For $S \in \mathsf{GDC}$, it is not hard to see that the union of the concave chains forms an uncrossed spanning cycle (cf. the full version [2]). Figure 3 gives an illustration of generalized double circles.

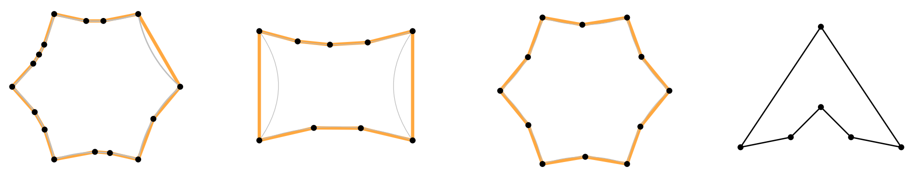

(a) (b) double chain (c) double circle (d)

Figure 3. (a-c) Examples of generalized double circles (the uncrossed spanning cycle is depicted in orange). (d) A point set that is *not* a generalized double circle, but admits an uncrossed spanning cycle.

Before diving into the details of the proof of Theorem 4, we start by collecting preliminary results in a slightly more general setting, namely for point sets S fulfilling the following property:

(P1) there is an *uncrossed* spanning cycle C on S, i.e., no edge joining two points of S crosses any edge of C.

A point set fulfilling (P1) is called *spinal* point set. When considering a spinal point set S, we first fix an uncrossed spanning cycle C, which we call *spine* and all edges in C *spine edges*. For instance, generalized double circles are spinal point sets and the spine is precisely the uncrossed spanning cycle formed by the concave chains as described above. Whenever speaking of the spine or spine edges for some point set without further specification, the underlying uncrossed cycle is either clear from the context, or the statement holds for any choice of such a cycle. Furthermore, we call all edges in the exterior/interior of the spine *outer/inner edges*.

We define the *canonical paths* to be those that consist only of spine edges. Note that this definition also captures the canonical paths used by Akl et al. [3], and that any canonical path can be transformed into any other by a single flip (of Type 2). Two vertices incident to a common spine edge are called *neighbors*.

Valid flips. We collect a few observations which will be useful to confirm the validity of a flip. Whenever we apply more than one flip, the notation in subsequent flips refers to the original path and not the current (usually we apply one or two flips in a certain step). Figure 4 gives an illustration of Observation 6.

Observation 6. *Let S be a spinal point set, $P = v_1, \ldots, v_n \in \mathcal{P}(S)$, and v_1, v_a ($a \neq 2$) be neighbors. Then the following flips are valid (under the specified additional assumptions):*

(*a*) *flip $v_{a-1}v_a$ to v_1v_a*

(*b*) *flip v_av_{a+1} to $v_{a-1}v_{a+1}$* (*if the triangle $\Delta v_{a-1}v_av_{a+1}$ is empty and (b) is performed subsequently after the flip in (a)*)

(*c*) *flip v_av_{a+1} to v_1v_{a+1}* (*if the triangle $\Delta v_1v_av_{a+1}$ is empty and $v_{a-1}v_a$ is a spine edge*)

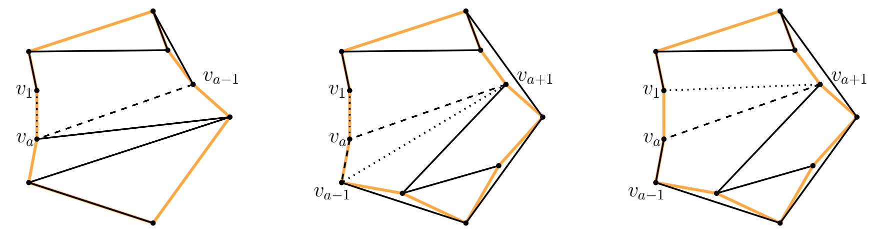

Figure 4. *Left to right:* Illustration of the three flips in Observation 6. The spine is depicted in orange and edge flips are indicated by replacing dashed edges for dotted (in the middle, the two flips must of course be executed one after the other).

Strictly speaking, in Observation 6(c) we do not require $v_{a-1}v_a$ to be a spine edge, but merely to be an edge not crossing v_1v_{a+1}. The following lemma provides structural properties for generalized double circles, if the triangles in Observation 6(b,c) are non-empty, i.e., contain points from S (see also Figure 5 (left)):

Lemma 4. *($\star$) Let $S \in \mathsf{GDC}$ and $p, q, x \in S$ such that p and q are neighbors. Further, let the triangle Δpqx be non-empty. Then the following holds:*

(i) At least one of the two points p, q is an extreme point (say p),
(ii) x does not lie on a common chain with p and *q, but shares a common chain with either p or q (the latter may only happen if q is also an extreme point).*

Combinatorial distance measure. In contrast to the proof for wheel sets, it may now not be possible anymore to directly increase the number of spine edges and hence, we need a more sophisticated measure. Let C be the spine of a spinal point set S and $p, q \in S$. Further let $o \in \{\text{cw}, \text{ccw}\}$ be an orientation. We define the *distance* between p, q *in direction* o, denoted by $d^o(p, q)$, as the number of spine edges along C that lie between p and q in direction o. Furthermore, we define the *distance* between p and q to be

$$d(p, q) = \min\{d^{\text{cw}}(p, q), d^{\text{ccw}}(p, q)\}.$$

Note that neighboring points along the spine have distance 1. Using this notion, we define the *weight* of an edge to be the distance between its endpoints and the (overall) weight of a path on S to be the sum of its edge weights.

Our goal is to perform weight-decreasing flips. To this end, we state two more preliminary results (see also Figure 5 (middle) and (right)):

Observation 7. *Let S be a spinal point set, p, q, r be three neighboring points in this order (i.e., q lies between p and r), and $s \in S \setminus \{p, q, r\}$ be another point. Then $d(p, s) < d(q, s)$ or $d(r, s) < d(q, s)$ holds.*

Combining Observation 6 and Observation 7, it is apparent that we can perform weight-decreasing flips whenever $\Delta v_{a-1}v_av_{a+1}$ and $\Delta v_1v_av_{a+1}$ are empty.

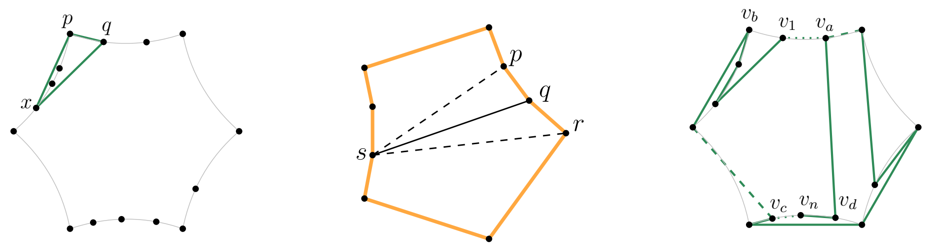

Figure 5. *Left:* Illustration of Lemma 4. If p and q are neighbors, x has to lie on the depicted chain in order to obtain a non-empty triangle Δpqx. *Middle:* Illustration of Observation 7. One of the dashed edges has smaller weight than the solid: $d(s,q) = 4$; $d(s,p) = 4$; $d(s,r) = 3$. *Right:* Illustration of Lemma 5. The initial path is depicted by solid and dashed edges. Flipping the dashed edges to the dotted edges increases the number of spine edges.

Lemma 5. *($\star$) Let S be a spinal point set, $P = v_1, \ldots, v_n \in \mathcal{P}(S)$, and v_a, v_b $(a, b \neq 2)$ be neighbors of v_1 as well as v_c, v_d $(c, d \neq n-1)$ be neighbors of v_n. If* $\max(a,b) > \min(c,d)$*, then the number of spine edges in P can be increased by performing at most two flips, which also decrease the overall weight of P.*

Note that v_b or v_d in Lemma 5 may not exist, if the first or last edge of P is a spine edge. Lemma 5 essentially enables us to perform weight decreasing flips whenever the path traverses a neighbor of v_n before it reached both neighbors of v_1. We are now ready to prove Theorem 4, but briefly summarize the proof strategy from a high-level perspective beforehand:

High level proof strategy. To flip an arbitrary path $P \in \mathcal{P}(S)$ to a canonical path, we perform iterations of suitable flips such that in each iteration we either

(i) increase the number of spine edges along P, while not increasing the overall weight of P, or
(ii) decrease the overall weight of P, while not decreasing the number of spine edges along P.

Note that for the connectivity of the flip graph it is not necessary to guarantee the non increasing overall weight in the first part. However, this will provide us with a better bound on the diameter of the flip graph.

Theorem 4. *($\star$) Let S be a set of n points in generalized double circle configuration. Then the flip graph (on $\mathcal{P}(S)$) is connected with diameter $O(n^2)$.*

Proof (Sketch). Let $P = v_1, \ldots, v_n \in \mathcal{P}(S)$ be a non-canonical path. We show how to iteratively transform P to a canonical path by increasing the number of spine edges or decreasing its overall weight. Let v_a $(a \neq 2)$ be a neighbor of v_1.

We can assume, w.l.o.g, that v_1 and v_n are not neighbors (i.e., $a < n$), since otherwise we can flip an arbitrary (non-spine) edge of P to the spine edge

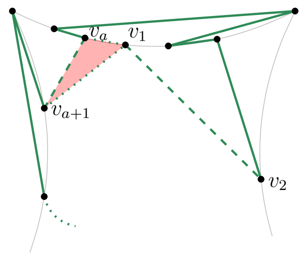

Figure 6. Illustration of Case 1. If v_1v_2 is not a spine edge and $\Delta v_1v_av_{a+1}$ is empty, we make progress by flipping the dashed edges to the dotted.

v_1v_n (performing a Type 2 flip). Furthermore, we can also assume w.l.o.g., that $v_{a-1}v_a$ is a spine edge, since otherwise we can flip $v_{a-1}v_a$ to the spine edge v_1v_a (Observation 6(a)). This also implies that the edge v_av_{a+1}, which exists because $a < n$, is not a spine edge, since v_a already has the two neighbors v_{a-1} and v_1.

We distinguish two cases – v_1v_2 being a spine edge or not:

Case 1: v_1v_2 is not a spine edge.

This case is easier to handle, since we are guaranteed that both neighbors of v_1 are potential candidates to flip to. In order to apply Observation 6, we require $\Delta v_1v_av_{a+1}$ to be empty. If that is the case we apply the following flips (see also Figure 6):

$$\text{flip } v_av_{a+1} \text{ to } v_1v_{a+1} \qquad \text{and} \qquad \text{flip } v_1v_2 \text{ to } v_1v_a,$$

where the first flip results in the path $v_a, \ldots, v_1, v_{a+1}, \ldots, v_n$ (and is valid by Observation 6(c)) and the second flip results in the path $v_2, \ldots, v_a, v_1, v_{a+1}, \ldots, v_n$ (valid due to Observation 6(a)). Together, the number of spine edges increases, while the overall weight does not increase.

If $\Delta v_1v_av_{a+1}$ is not empty we need to be more careful, using Lemma 4 (details can be found in the full version [2]).

Case 2: v_1v_2 is a spine edge.

In this case we will consider P from both ends v_1 and v_n. Our general strategy here is to first rule out some easier cases and collect all those cases where we cannot immediately make progress. For these remaining "bad" cases we consider the setting from both ends of the path.

Again, we skip the analysis of the easier cases and just summarize the six "bad" cases. These "bad" cases always involve v_1, v_a, or v_{a-1} being an extreme point. Instead of spelling all these cases out, we give an illustration in Figure 7.

In the remainder of the proof we settle these "bad" cases by arguing about both ends of the path, i.e, we consider all $\binom{6}{2} + 6 = 21$ combinations of "bad" cases.

We exclude several combinations as follows. By Lemma 5, we can assume that $a < c$ holds (otherwise there are weight decreasing flips) and hence, no

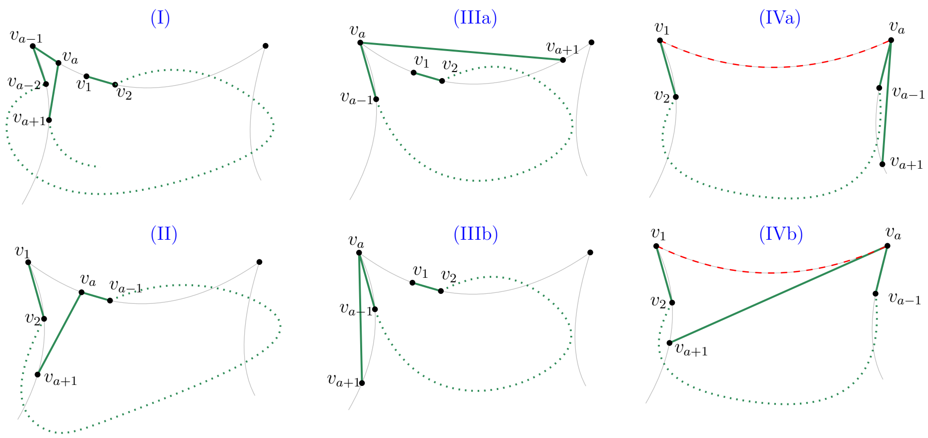

Figure 7. The six “bad” cases. The solid edges depict the fixed edges of the corresponding “bad” case and the red arcs (here and in the following) indicate that there is no vertex other than the two extreme points lying on this chain.

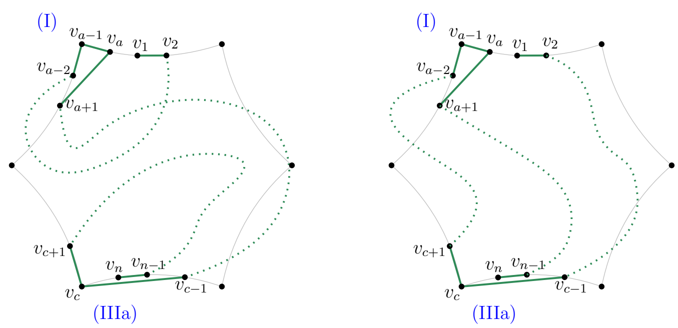

Figure 8. (I) and (IIIa) cannot be combined in a plane manner (left), except if the path traverses a neighbor of v_n before those of v_1, i.e., $c < a$ holds (right).

“bad” case where v_{a+1} is in the interior of $\Phi(v_a)$ can be combined with a “bad” case having v_n or v_c as extreme point (Observation 5). This excludes (almost) all combinations involving (I), (II), or (IVb); see Figure 8 for an example.

For the remaining cases, we try to decrease the weight of P by flipping $v_a v_{a+1}$ either to $v_1 v_{a+1}$ or $v_{a-1} v_{a+1}$ (see Observation 7). If these flips are valid they are either weight-decreasing or we can identify disjoint regions that must each contain at least $n/2$ vertices, which will result in a contradiction. Again, we skip the details of this analysis.

Iteratively applying the above process transforms P to a canonical path and the $O(n^2)$ bound for the required number of flips also follows straightforwardly. □

5 Conclusion

In this paper, we made progress towards a positive answer of Conjecture 1, though it still remains open in general. We approached Conjecture 1 from two directions and believe that Conjecture 3 might be easier to tackle, e.g. for an inductive approach. For all our results we used only Type 1 and Type 2 flips (which can be simulated by Type 1 flips). It is an intriguing question whether Type 3 flips are necessary at all.

Concerning the approach of special classes of point sets, of course one can try to further adapt the ideas to other classes. Most of our results hold for the setting of spinal point sets; the main obstacle that remains in order to show flip connectivity for the point sets satisfying condition (P1) would be to adapt Lemma 4. A proof for general point sets, however, seems elusive at the moment.

Lastly, there are several other directions for further research conceivable, e.g. non-straight-line drawings.

References

1. Aichholzer, O., Aurenhammer, F., Krasser, H.: Enumerating order types for small point sets with applications. Order **19**, 265–281 (2002). https://doi.org/10.1023/A:1021231927255
2. Aichholzer, O., Knorr, K., Mulzer, W., Obenaus, J., Paul, R., Vogtenhuber, B.: Flipping plane spanning paths (2022). https://doi.org/10.48550/ARXIV.2202.10831
3. Akl, S.G., Islam, M.K., Meijer, H.: On planar path transformation. Information Processing Letters **104**(2), 59–64 (2007). https://doi.org/10.1016/j.ipl.2007.05.009
4. Bose, P., Hurtado, F.: Flips in planar graphs. Computational Geometry **42**(1), 60–80 (2009). https://doi.org/10.1016/j.comgeo.2008.04.001
5. Chang, J.M., Wu, R.Y.: On the diameter of geometric path graphs of points in convex position. Information Processing Letters **109**(8), 409–413 (2009). https://doi.org/10.1016/j.ipl.2008.12.017
6. Hernando, C., Hurtado, F., Noy, M.: Graphs of non-crossing perfect matchings. Graphs and Combinatorics **18**(3), 517–532 (2002)
7. Houle, M., Hurtado, F., Noy, M., Rivera-Campo, E.: Graphs of triangulations and perfect matchings. Graphs and Combinatorics **21**, 325–331 (09 2005). https://doi.org/10.1007/s00373-005-0615-2
8. Lawson, C.L.: Transforming triangulations. Discrete mathematics **3**(4), 365–372 (1972)
9. Nichols, T.L., Pilz, A., Tóth, C.D., Zehmakan, A.N.: Transition operations over plane trees. Discrete Mathematics **343**(8), 111929 (2020). https://doi.org/10.1016/j.disc.2020.111929
10. Nishimura, N.: Introduction to reconfiguration. Algorithms **11**(4) (2018). https://doi.org/10.3390/a11040052
11. Wagner, K.: Bemerkungen zum Vierfarbenproblem. Jahresbericht der Deutschen Mathematiker-Vereinigung **46**, 26–32 (1936), http://eudml.org/doc/146109