

Flip distances between graph orientations

Oswin Aichholzer¹ Jean Cardinal² Tony Huynh^{2*} Kolja Knauer^{3†}
Torsten Mütze⁴ Raphael Steiner⁴ Birgit Vogtenhuber^{1‡}

¹TU Graz, Austria, {oaich, bvogt}@ist.tugraz.at

²Université libre de Bruxelles (ULB), Belgium, jcardin@ulb.ac.be,
tony.bourbaki@gmail.com

³Université Aix-Marseille, France, kolja.knauer@lis-lab.fr

⁴TU Berlin, Germany, {steiner, muetze}@math.tu-berlin.de

February 17, 2019

Abstract

Flip graphs are a ubiquitous class of graphs, which encode relations induced on a set of combinatorial objects by elementary, local changes. Skeletons of associahedra, for instance, are the graphs induced by quadrilateral flips in triangulations of a convex polygon. For some definition of a flip graph, a natural computational problem to consider is the flip distance: Given two objects, what is the minimum number of flips needed to transform one into the other?

We consider flip graphs on orientations of simple graphs, where flips consist of reversing the direction of some edges. More precisely, we consider so-called α -orientations of a graph G , in which every vertex v has a specified outdegree $\alpha(v)$, and a flip consists of reversing all edges of a directed cycle. We prove that deciding whether the flip distance between two α -orientations of a planar graph G is at most two is NP-complete. This also holds in the special case of perfect matchings, where flips involve alternating cycles. This problem amounts to finding geodesics on the common base polytope of two partition matroids, or, alternatively, on an alcoved polytope. It therefore provides an interesting example of a flip distance question that is computationally intractable despite having a natural interpretation as a geodesic on a nicely structured combinatorial polytope.

We also consider the dual question of the flip distance between graph orientations in which every cycle has a specified number of forward edges, and a flip is the reversal of all edges in a minimal directed cut. In general, the problem remains hard. However, if we restrict to flips that only change sinks into sources, or vice-versa, then the problem can be solved in polynomial time. Here we exploit the fact that the flip graph is the cover graph of a distributive lattice. This generalizes a recent result from Zhang, Qian, and Zhang (Acta. Math. Sin.-English Ser., 2019).

1 Introduction

The term [flip](#) is commonly used in combinatorics to refer to an elementary, local, reversible operation that transforms one combinatorial object into another. Such flip operations naturally yield a [flip graph](#), whose vertices are the considered combinatorial objects, and two of them are adjacent if they differ by a single flip. A classical example is the flip graph of triangulations of a convex polygon [[STT88](#), [Pou14](#)]; see [Figure 1](#). The vertex set of this graph are all triangulations of the polygon, and two triangulations are adjacent if one can be obtained from the other by replacing the diagonal of a quadrilateral

*Supported by ERC Consolidator Grant 615640-ForEFront

†Partially supported by ANR grants GATO: ANR-16-CE40-0009-01, DISTANCIA: ANR-17-CE40-0015, and CAPPS: ANR-17-CE40-0018

‡Partially supported by the Austrian Science Fund (FWF): I 3340-N35

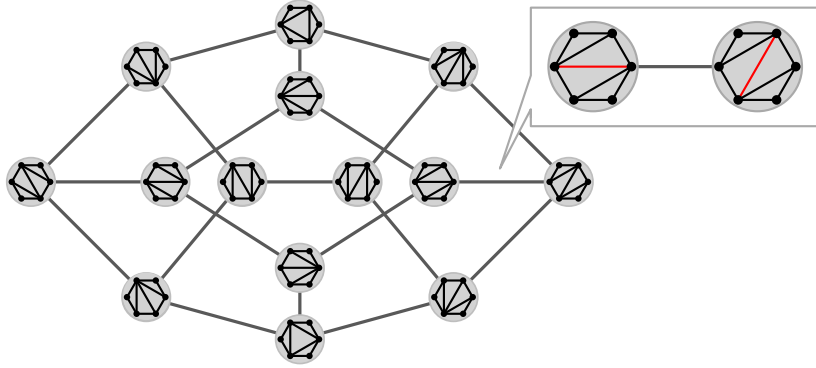


Figure 1: The flip graph of triangulations of a convex polygon.

formed by two triangles by the other diagonal. Similar flip graphs have also been investigated for triangulations of general point sets in the plane [LRS10], triangulations of topological surfaces [Neg94], and planar graphs [BH09, BV11]. The **flip distance** between two combinatorial objects is the minimum number of flips needed to transform one into the other. It is known that computing the flip distance between two triangulations of a simple polygon [AMP15] or of a point set [LP15] is NP-hard. The latter is known to be fixed-parameter tractable [KSX17]. On the other hand, the NP-hardness of computing the flip distance between two triangulations of a convex polygon is a well-known open question [LZ98, Rog99, DP02, CSJ09, CSJ10, LEP10]. Flip graphs involving other geometric configurations have also been studied, such as flip graphs of non-crossing perfect matchings of a point set in the plane, where flips are with respect to alternating 4-cycles [HHN02], or alternating cycles of arbitrary length [HHNRC05]. Other flip graphs include the flip graph on plane spanning trees [AAHV07], the flip graph of non-crossing partitions of a point set or dissections of a polygon [HHNOP09], the mutation graph of simple pseudoline arrangements [Rin57], the Eulerian tour graph of a Eulerian graph [ZG87], and many others. There is also a vast collection of interesting flip graphs for non-geometric objects, such as bitstrings, permutations, combinations, and partitions [FKMS18].

In essence, a flip graph provides the considered family of combinatorial objects with an underlying structure that reveals interesting properties about the objects. It can also be a useful tool for proving that a property holds for all objects, by proving that one particularly nice object has the property, and that the property is preserved under flips. Flip graphs are also an essential tool for solving fundamental algorithmic tasks such as random and exhaustive generation, see e.g. [AF96] and [PW98].

The focus of the present paper is on flip graphs for orientations of graphs satisfying some constraints. First, we consider so-called α -orientations, in which the outdegree of every vertex is specified by a function α , and the flip operation consists of reversing the orientation of all edges in a directed cycle. We study the complexity of computing the flip distance between two such orientations. An interesting special case of α -orientations corresponds to perfect matchings in bipartite graphs, where flips involve alternating cycles. We also consider the dual notion of c -orientations, in which the number of forward edges along each cycle is specified by a function c . Here a flip consists of reversing all edges in a directed cut. We also analyze the computational complexity of the flip distance problem in c -orientations.

There are several deep connections between flip graphs and polytopes. Specifically, many interesting flip graphs arise as the (1-)skeleton of a polytope. For instance, flip graphs of triangulations of a convex polygon are skeletons of associahedra [CSZ15], and flip graphs of regular triangulations of a point set in the plane are skeletons of secondary polytopes (see [LRS10, Chapter 5]). Associahedra are generalized by **quotientopes** [PS18], whose skeletons yield flip graphs on rectangulations [CSS18], bitstrings, permutations, and other combinatorial objects. Moreover, flip graphs of acyclic orientations




Figure 2: Two α -orientations of a graph and a flip between them, where the values of α are depicted on the vertices.

or strongly connected orientations of a graph are skeletons of graphical and co-graphical zonotopes, respectively (see [Pos09, Section 2]). Similarly, as we show below, flip graphs on α -orientations are skeletons of matroid intersection polytopes. We also consider vertex flips in c -orientations, inducing flip graphs that are distributive lattices and in particular subgraphs of skeletons of certain distributive polytopes. These polytopes specialize to flip polytopes of planar α -orientations, are generalized by the polytope of tensions of a digraph, and form part of the family of [alcoved polytopes](#) (see [FK11]).

In the next section, we give the precise statements of the computational problems we consider, connections with previous work, and the statements of our main results. In Section 4, we give the proof of our first main result, showing that computing the flip distance between α -orientations and between perfect matchings is NP-hard even for planar graphs. Section 5 presents the proof of our second main result, where we give a polynomial time algorithm to compute the vertex flip distance between c -orientations. Finally, in Section 6 we show that computing the distance between c -orientations, when double vertex flips are also allowed, is NP-hard.

2 Problems and main results

Flip distance between α -orientations. Given a graph G and some $\alpha : V(G) \rightarrow \mathbb{N}_0$, an α -orientation of G is an orientation of the edges of G in which every vertex v has outdegree $\alpha(v)$. An example for a graph and two α -orientations for this graph is given in Figure 2. A flip of a directed cycle C in some α -orientation X consists of the reversal of the orientation of all edges of C , as shown in the figure. Edges with distinct orientations in two given α -orientations X and Y induce a Eulerian subdigraph of both X and Y . They can therefore be partitioned into an edge-disjoint union of cycles in G which are directed in both X and Y . Hence the reversal of each such cycle in X gives rise to a flip sequence transforming X into Y and vice versa. We may thus define the [flip distance](#) between two α -orientations X and Y to be the minimum number of cycles in a flip sequence transforming X into Y . We are interested in the computational complexity of determining the flip distance between two given α -orientations.

Problem 1. *Given a graph G , some $\alpha : V(G) \rightarrow \mathbb{N}_0$, a pair X, Y of α -orientations of G and an integer $k \geq 0$, decide whether the flip distance between X and Y is at most k .*

The crucial difficulty of this problem is that a shortest flip sequence transforming X into Y may flip edges that are oriented the same in X and Y an even number of times, to reach Y with fewer flips compared to only flipping edges that are oriented differently in X and Y ; see the example in Figure 3. This motivates the following variant of the previous problem:

Problem 2. *Given G, α, X, Y, k as in Problem 1, decide whether the flip distance between X and Y is at most k , where we only allow flipping edges that are oriented differently in X and Y .*




Figure 3: An α -orientation X of a graph. The α -orientation Y obtained by flipping the four directed facial cycles C_1, \dots, C_4 can be reached with fewer flips by flipping only the three directed facial cycles D_1, D_2, D_3 in this order.

From α -orientations to perfect matchings. The flexibility in choosing a function α for a set of α -orientations on a graph allows us to capture numerous relevant combinatorial structures, some of which are listed below:

- domino and lozenge tilings of a plane region [Rém04, Thu90],
- planar spanning trees [GL86],
- (planar) bipartite perfect matchings [LZ03],
- (planar) bipartite d -factors [Pro02, Fel04],
- Schnyder woods of a planar triangulation [Bre00],
- Eulerian orientations of a (planar) graph [Fel04],
- k -fractional orientations of a planar graph with specified outdegrees [BF12],
- contact representations of planar graphs with homothetic triangles, rectangles, and k -gons [Fel13, GLP12, FSS18a, FSS18b].

In the following, we focus on perfect matchings of bipartite graphs. Consider any bipartite graph G with bipartition (V_1, V_2) equipped with

$$\alpha : V(G) \rightarrow \mathbb{N}_0, \quad \alpha(x) := \begin{cases} 1 & \text{if } x \in V_1, \\ d_G(x) - 1 & \text{if } x \in V_2. \end{cases}$$

With this definition, in each α -orientation of G , the edges directed from V_1 to V_2 form a perfect matching. This is illustrated in Figure 4. Conversely, given a perfect matching M of G , orienting all edges of M from V_1 to V_2 and all the other edges from V_2 to V_1 yields an α -orientation of the above type. Furthermore, the directed cycles in any α -orientation of G correspond to the alternating cycles in the associated perfect matching. Flipping an alternating cycle in a perfect matching corresponds to exchanging matching-edges and non-matching-edges. An example of the flip graph of perfect matchings of a graph is given in Figure 5. In this special case, Problem 1 boils down to:

Problem 3. Given a bipartite graph G , a pair X, Y of perfect matchings in G and an integer $k \geq 0$, decide whether the flip distance between X and Y is at most k .

The example from Figure 3 can be easily modified to show that when transforming X into Y using the fewest number of flips, we may have to flip alternating cycles that are not in the symmetric difference of X and Y ; see the example in Figure 6. If we restrict the flips to only use cycles in the symmetric difference of X and Y , then the problem of finding the flip distance becomes trivial, as the symmetric difference is a collection of disjoint cycles, and each of them has to be flipped, so Problem 2 is trivial for perfect matchings.




Figure 4: An α -orientation of a bipartite graph and the corresponding perfect matching.




Figure 5: The flip graph of perfect matchings of a graph. The solid edges indicate flips along facial cycles, and the dashed edges indicate flips along non-facial cycles.




Figure 6: A perfect matching X in a graph. The perfect matching Y obtained by flipping the four alternating facial cycles C_1, \dots, C_4 can be reached with fewer flips by flipping only the three alternating facial cycles D_1, D_2, D_3 in this order.

Flip graphs and matroid intersection polytopes. We proceed to give a geometric interpretation of the flip distance between α -orientations as the distance in the skeleton of a 0/1-polytope.

Recall that a **matroid** is an abstract simplicial complex (E, \mathcal{I}) , where $\mathcal{I} \subseteq 2^E$ satisfies the independent set augmentation property. The elements of \mathcal{I} are called **independent sets**. A **base** of the matroid is an inclusionwise maximal independent set.

It is well-known that perfect matchings in a bipartite graph $G = (V_1 \cup V_2, E)$ are common bases of two partition matroids (E, \mathcal{I}_1) and (E, \mathcal{I}_2) , in which a set of edges is independent if no two share an endpoint in V_1 , or, respectively, in V_2 .

Similarly, α -orientations can be defined as common bases of two partition matroids. In this case, every edge of the graph G is replaced by a pair of parallel arcs, one for each possible orientation of the edge. One matroid encodes the constraint that for every edge exactly one orientation is chosen. The second matroid encodes the constraint that each vertex v has exactly $\alpha(v)$ outgoing arcs.

The **common base polytope** of two matroids is a 0/1-polytope obtained as the convex hull of the characteristic vectors of the common bases. Adjacency of two vertices of this polytope has been characterized by Frank and Tardos [FT88]. A shorter proof was given by Iwata [Iwa02]. We briefly recall their result in the next theorem. To state the theorem, consider a matroid $M = (E, \mathcal{I})$, a base $B \in \mathcal{I}$, and a subset $F \subseteq E$. The **exchangeability graph** $G(B, F)$ of M is a bipartite graph with $B \setminus F$ and $F \setminus B$ as vertex bipartition, and edge set $\{ij \mid B \setminus \{i\} \cup \{j\} \text{ is a basis}\}$. This definition and the theorem are illustrated in Figure 7 for the two partition matroids whose common bases are perfect matchings of a graph.

Theorem 1 ([FT88, Iwa02]). *For two matroids $M^+ = (E, \mathcal{I}^+)$ and $M^- = (E, \mathcal{I}^-)$, two common bases $A, B \in \mathcal{I}^+ \cap \mathcal{I}^-$ are adjacent on the common base polytope if and only if all the following conditions hold:*

- (i) *the exchangeability graph $G(A, B)$ of M^+ has a unique perfect matching P^+ ,*
- (ii) *the exchangeability graph $G(B, A)$ of M^- has a unique perfect matching P^- ,*
- (iii) *$P^+ \cup P^-$ is a single cycle.*

From this theorem we can conclude that the flip graphs we consider on perfect matchings and α -orientations are precisely the skeletons of the corresponding polytopes of common bases.

It is interesting to compare Problems 1 and 3 with the analogous problems for other families of matroid polytopes. For instance, it is known that for two bases A, B of a matroid, the exchangeability graph $G(A, B)$ has a perfect matching [Bru69]. Hence A can be transformed into B by performing $|A \Delta B|/2$ exchanges of elements (where $A \Delta B$ is the symmetric difference of A and B), which is also the distance in the skeleton of the base polytope of the matroid. On the other hand, the problem of computing the flip distance between two triangulations of a convex polygon amounts to computing distances in skeletons of associahedra, which are known to be polymatroids (see [AA17] and references therein). This problem is neither known to be in P nor known to be NP-hard. Also note that for other families of combinatorial polytopes, testing adjacency is already intractable. This is the case for instance for the polytope of the Traveling Salesman Problem (TSP) [Pap78], whose skeleton is known to have diameter at most 4 [RC98]. On the other hand, the corresponding polytope is known to be the common base polytope of three matroids.

Another important class of combinatorial polytopes are **alcoved polytopes**, see [LP07]. It is known that the flip graphs of planar α -orientations are skeletons of alcoved polytopes, see [FK11]. Thus, by our results below, flip distances in this class are also NP-hard to compute.

Hardness of flip distance between perfect matchings and α -orientations. We prove that Problem 3 is NP-complete, even for 2-connected bipartite subcubic planar graphs and $k = 2$. This clearly implies that Problem 1 is NP-complete as well.

Theorem 2. *Given a 2-connected bipartite subcubic planar graph G and a pair X, Y of perfect matchings in G , deciding whether the flip distance between X and Y is at most two is NP-complete.*




Figure 7: Two common bases A and B (left and middle) of the matroids M^+ and M^- , where M^+ and M^- have as independent sets all subsets of edges of the graph where no two share an endpoint in the set of circled vertices, or the set of squared vertices, respectively. The right hand side shows the exchangeability graphs $G(A, B)$ of M^+ (solid edges) and $G(B, A)$ of M^- (dashed edges). As the conditions of Theorem 1 are met, the two bases are adjacent in the common base polytope, and adjacent in the flip graph shown in Figure 5.

As direct consequences of the proof of Theorem 2 we get:

Corollary 1. *Unless $P = NP$, deciding whether the flip distance between two perfect matchings is at most k is not fixed-parameter tractable with respect to parameter k .*

Corollary 2. *Unless $P = NP$, the flip distance between two perfect matchings is not approximable within a multiplicative factor $3/2 - \epsilon$ in polynomial time, for any $\epsilon > 0$.*

We also prove that Problem 2 is NP-complete, even for 4-regular graphs and $k = 2$.

Theorem 3. *Given a 4-regular graph G and a pair X, Y of α -orientations of G , deciding whether the flip distance between X and Y is at most two is NP-complete. Moreover, the problem remains NP-complete if we only allow flipping edges that are oriented differently in X and Y .*

The proofs of Theorem 2 and 3 are presented in Section 4.

From α -orientations in planar graphs to c -orientations. In what follows, we generalize the problem, via planar duality, to flip distances in so-called c -orientations.

Consider an arbitrary 2-connected plane graph G and its planar dual G^* . Then for any orientation D of the edges of G , the directed dual D^* of D is obtained by orienting any dual edge forward if it crosses a left-to-right arc in D in a simultaneous plane embedding of G and G^* , and backward otherwise; see Figure 8. Edge sets of directed cycles in D correspond to edge sets of minimal **directed cuts** in D^* and vice-versa. Hence D is acyclic (respectively, strongly connected) if and only if D^* is strongly connected (respectively, acyclic). A **directed vertex cut** is a cut consisting of all edges incident to a sink or a source vertex. Directed facial cycles in D are in bijection with the directed vertex cuts in D^* , and vice versa. The unbounded face in the plane embedding of D can be chosen such that it corresponds to a fixed vertex \top in D^* .

Let D be an α -orientation of G . Given a minimal cut in D separating $U \subseteq V(D)$ from $\bar{U} := V(D) \setminus U$, we denote by $\delta^+(U)$ the edges pointing from U to \bar{U} in D . We also let $d_D^+(v)$ denote the outdegree of vertex v in D . We have

$$|\delta^+(U)| = \sum_{v \in U} d_D^+(v) - |E(G[U])| = \sum_{v \in U} \alpha(v) - |E(G[U])|,$$




Figure 8: Duality between flips in α -orientations (solid edges) and in c -orientations (dashed edges).

which only depends on α and G . Consequently, the set of orientations of G^* which are directed duals of α -orientations of G can be characterized by the property that for every cycle C in G^* , the number of edges in clockwise direction is fixed by a certain value $c(C)$ independent of the orientation. The flip operation between α -orientations of D consists of the reversal of a directed cycle. In the corresponding set of dual orientations of D^* , this translates to the reversal of the orientations of the edges in a minimal directed cut, as shown on Figure 8.

The same notion has been investigated more generally without planarity conditions under the name of c -orientations by Propp [Pro02] and Knauer [Kna07]. Given a graph G , we can fix an arbitrary direction of traversal for each cycle C . Given a graph and an assignment $c(C) \in \mathbb{N}_0$ to each cycle in G , one may define a c -orientation of G to be an orientation having exactly $c(C)$ edges in forward direction for every cycle C in G . Note that it is sufficient to define the function c on a cycle basis of G , which consists of linearly many cycles. The flip operation on the set \mathcal{R}_c of such c -orientations of a graph is defined as the reversal of all edges in a minimal directed cut. It is not difficult to see that flips make the set of c -orientations of a graph connected (this will be noted in Section 5).

From the duality between planar α -orientations and planar c -orientations, determining flip distances between α -orientations of 2-connected planar graphs reduces to determining flip distances between the dual c -orientations. Note that planar duals of bipartite graphs are exactly the Eulerian planar graphs. Theorem 2 therefore directly yields:

Corollary 3. *Given a Eulerian planar multigraph G and a pair X, Y of c -orientations of G , deciding whether the flip distance between X and Y is at most two is NP-complete.*

c -orientations and distributive lattices. A more local operation consists of flipping only directed vertex cuts, induced by sources and sinks, excluding a fixed vertex \top . We will refer to this special case as a **vertex flip**. Specifically, given a pair of c -orientations X, Y of a graph G with a fixed vertex \top , we aim to transform X into Y using only vertex flips at vertices distinct from \top .

A c -orientation X of G may not be acyclic, that is, there is a cycle C in G which is directed in X . According to the definition of a c -orientation, this means that C keeps the same orientation in every c -orientation of G . Consequently, any (minimal) directed cut in a c -orientation of G is disjoint from $E(C)$. Contracting the cycle C in G , we end up with a smaller graph G' containing the same (minimal) directed cuts, such that the c -orientations of G are determined by their corresponding orientations on G' . We can therefore safely assume that the c -orientations that we consider are all acyclic. Similarly, G will be assumed to be connected.

Problem 4. *Given a connected graph G with a fixed vertex \top and a pair X, Y of acyclic c -orientations, what is the length of a shortest vertex flip sequence transforming X into Y ?*

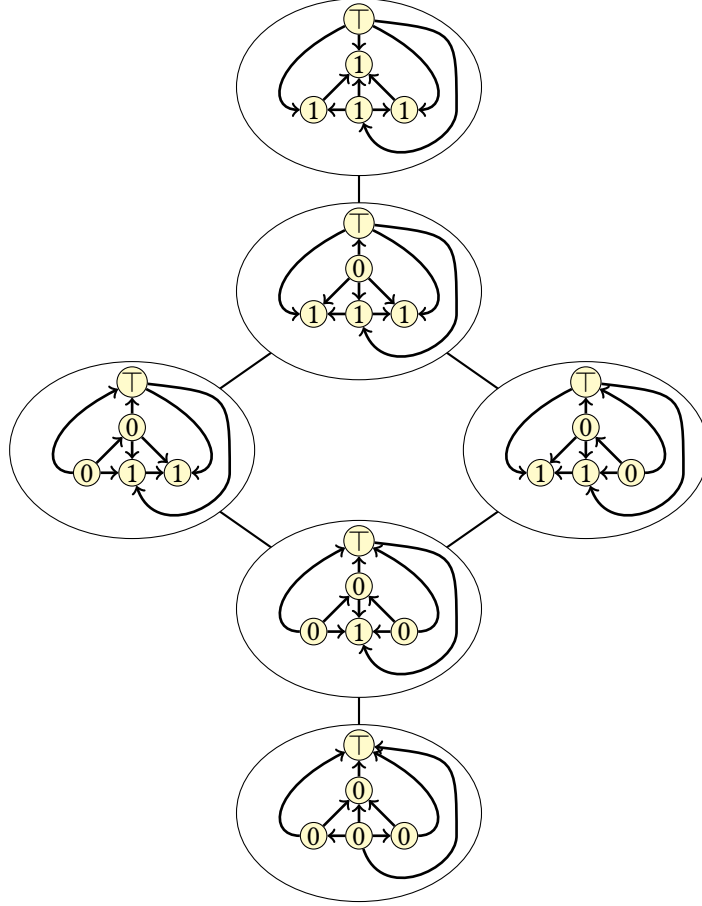


Figure 9: The distributive lattice induced by vertex flips in c -orientations. The reference orientation at the bottom is the directed dual D^* of the orientation D of the graph G used in Figures 4 and 5, where some parallel arcs incident with \top are grouped together for simplicity. The numbers depicted at the vertices indicate the number of times that each vertex is flipped with respect to the reference orientation.

We should convince ourselves that under the assumptions made above, every pair of c -orientations is reachable from each other by vertex flips. This property is provided in a much stronger way by a distributive lattice structure on the set \mathcal{R}_c ; see Figure 9. The next theorem is a special case of Theorem 1 in Propp [Pro02] where the c -orientations are acyclic.

Theorem 4 ([Pro02, Kna07]). *Let G be a graph with fixed vertex \top and \mathcal{R}_c a set of acyclic c -orientations of G . Then the partial order \leq_c on \mathcal{R}_c in which Y covers X if and only if Y can be obtained from X by flipping a source defines a distributive lattice on \mathcal{R}_c .*

Hence Problem 4 consists of finding shortest paths in the cover graph of a distributive lattice, where the size of the lattice can be exponential in the size of the input G .

Every distributive lattice is a lattice of c -orientations. We next point out that every distributive lattice is isomorphic to the distributive lattice induced by a set of c -orientations of a graph. This relationship was described by Knauer [Kna08].

In order to represent a given distributive lattice L by an isomorphic lattice of c -orientations, we need to construct a corresponding digraph $D(L)$. For this purpose, we shortly recall a classical result

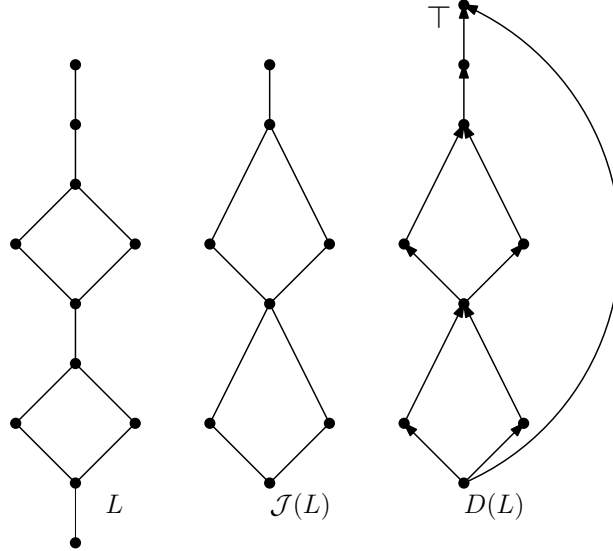


Figure 10: A distributive lattice L represented by its Hasse diagram (left), the corresponding subposet of join-irreducible elements $\mathcal{J}(L)$ (middle), and the digraph $D(L)$ associated with the lattice (right).

from lattice theory, [Birkhoff's Theorem](#) (see [DP02]).

For any distributive lattice L , $\mathcal{J}(L)$ is the subposet of L induced by the set of [join-irreducible](#) elements, these are the elements of L covering exactly one element. On the other hand, given any poset P we may look at the distributive lattice $\mathcal{O}(P)$ formed by the downsets of P ordered by inclusion. Birkhoff's Theorem in our setting asserts that those two operations are inverse in the sense that $P \cong \mathcal{J}(\mathcal{O}(P))$ for any finite poset P and $\mathcal{O}(\mathcal{J}(L)) \cong L$ for any finite distributive lattice.

The idea is to define a digraph $D(L)$ whose vertex set consists of the elements of $\mathcal{J}(L)$ with an additional vertex \top . The digraph is obtained from the natural upward-orientation of $\mathcal{J}(L)$ plus additional arcs from all the sinks and sources to \top . Let $G(L)$ be the underlying graph of $D(L)$ and let $L_c(D(L))$ denote the distributive lattice obtained from the set of c -orientations of $G(L)$ containing $D(L)$ with fixed (non-flippable) vertex \top . An example of this construction is provided in [Figure 10](#).

The following theorem is an easy consequence of Birkhoff's Theorem.

Theorem 5 ([Kna08]). *Let L be any distributive lattice and $D(L)$ be the corresponding digraph as defined above. Then $L \cong L_c(D(L))$.*

[Theorem 9](#) below gives a natural geometric embedding of the lattice L depending on the digraph $D(L)$. This embedding is such that all values $z_X(x)$ are 0 or 1, and the vectors $z_X(\cdot)$ are exactly the characteristic vectors of the downsets of $\mathcal{J}(L)$. The convex hull of those vectors is known as the [order polytope](#) of $\mathcal{J}(L)$ [Sta86], which is a particular case of the above-mentioned alcoved polytopes. The problem of computing vertex flip distances between elements of L encoded by c -orientations of $G(L)$ therefore boils down to computing the distance between two downsets of $\mathcal{J}(L)$ in their inclusion lattice, which is a simple special case of [Problem 4](#).

Facial flips in planar graphs. When we consider [Problem 4](#) on planar graphs, restricting to vertex flips and considering the dual plane graph amounts to considering only flips of directed facial cycles, excluding the outer face whose dual vertex is \top . We refer to these as [facial flips](#). Felsner [Fel04] considered distributive lattices induced by facial flips. The following computational problem is a special case of [Problem 4](#).

Problem 5. *Given a 2-connected plane graph G and a pair X, Y of strongly connected α -orientations, what is the length of a shortest facial flip sequence transforming X into Y ?*

Zhang, Qian, and Zhang [ZQZ19] recently provided a closed formula for this flip distance, which can be turned into a polynomial-time algorithm. We prove the analogous stronger statement for Problem 4.

Theorem 6. *There is a polynomial-time algorithm that, given a graph G with a fixed vertex \top and a pair X, Y of c -orientations of G , outputs a shortest vertex flip sequence between X and Y .*

In the planar case, this directly translates to a polynomial-time algorithm for Problem 5. The proof of Theorem 6 is presented in Section 5. In [FK09], the distributive lattice structure on c -orientations is generalized to so-called Δ -bonds, also known as **tensions**. We believe that our proof of Theorem 6 can be generalized to these objects.

Flip distance with larger cut sets. While computing the cut flip distance between c -orientations is an NP-hard problem in general (Theorem 2), there is a polynomial-time-algorithm for computing the distance when only using vertex flips (Theorem 6). It is natural to ask for a threshold between the hard and easy cases of flip distance problems. Our hardness reduction in Section 4 involves very long directed cycles, which correspond to flips of directed cuts in the dual c -orientations with cut sets of large size. Consequently, one may hope that the problem gets easier when restricting the sizes of the cut sets involved in a flip sequence. Our last result destroys this hope:

Theorem 7. *Let X, Y be c -orientations of a connected graph G with fixed vertex \top . It is NP-hard to determine the length of a shortest cut flip sequence transforming X into Y , which consists only of minimal directed cuts with interiors of size at most two.*

We will present the proof of Theorem 7 in Section 6.

3 Preliminaries

In this section we recall some standard terminology concerning digraphs and posets that will be used repeatedly in the paper.

Cuts and cut sets. Given a directed graph D and a subset $U \subseteq V(D)$ of vertices, we denote by $\delta(U)$ the set of all arcs in $E(D)$ having one end in U and the other in $\bar{U} := V(D) \setminus U$. By $\delta^+(U)$, we denote the set of all arcs in $E(D)$ directed from U to \bar{U} . If $\delta(U) = \delta^+(U)$, we call $S = \delta(U)$ a **directed cut** or **dicut** induced by U , and U is referred to as a **cut set** of S . In the case that D is weakly connected, the cut set is uniquely determined by the dicut.

Posets and lattices. A **partially ordered set**, or **poset** for short, is a pair (P, \prec) , where P is a set and \prec is a reflexive, antisymmetric and transitive binary relation on P . Posets can be represented more compactly by their minimal comparabilities: We say that $x \prec y$ is a **cover relation**, or y **covers** x , if there is no z in the poset with $x \prec z \prec y$. This defines the **cover graph** of P , which has the elements of P as vertices, and an edge for every cover relation. The **Hasse diagram** of P is a drawing of the cover graph in the plane, where vertices are represented by distinct points and for every cover relation $x \prec y$, the edge between x and y is drawn as a straight line going upwards from x to y .

The **downset** of an element y in P is the set of all x with $x \prec y$. A poset P is called a **lattice**, if for any two elements x and y in P there is a unique smallest element z such that $x \prec z$ and $y \prec z$, and a unique largest element z such that $z \prec x$ and $z \prec y$. These elements are called the **join** and the **meet** of x and y , respectively. A lattice is called **distributive**, if the join and meet operations distribute over each other.

4 Flip distance between perfect matchings and between α -orientations

The proof of Theorem 2 is by reduction from the following NP-complete problem.

Theorem 8 (Plesniak [Ple79]). *Deciding directed Hamiltonicity of orientations of cubic planar graphs without sinks and sources is NP-complete.*

It is clear that the above problem remains NP-complete if we additionally assume 2-connectivity of the cubic graph.

Proof of Theorem 2. As each flip sequence of length at most two can be used as a polynomially verifiable certificate, the problem is clearly in NP.

We now provide a reduction of the decision problem in Theorem 8 to Problem 3. So suppose we are given an orientation D of a 2-connected cubic planar graph without sinks and sources, and assume without loss of generality that $|V(D)| \geq 3$. Given D , we define an undirected graph $G = G(D)$ as follows; see Figure 11: For each vertex $v \in V(D)$ we create a vertex x_v in G , and for each arc $e \in E(D)$ we create a pair of vertices x_e^+, x_e^- in G . The edges of G are defined as follows: For each arc $e \in E(D)$, we connect x_e^+ and x_e^- with an edge in G . Furthermore, we denote by V_1 and V_2 the vertices of D with outdegree 1 or 2, respectively. For each $v \in V_1$, if $e, f \in E(D)$ are the two incoming arcs at v and g is the outgoing arc, then we add the edges $x_e^+ x_v$, $x_f^+ x_v$, $x_e^+ x_g^-$, and $x_f^+ x_g^-$ to G . Similarly, for each $v \in V_2$, if $e, f \in E(D)$ are the two outgoing arcs at v and g is the incoming arc, then we add the edges $x_e^- x_v$, $x_f^- x_v$, $x_e^- x_g^+$, and $x_f^- x_g^+$ to G . We refer to the 4-cycles in G formed by these edges as C_4 -gadgets. Note that G is subcubic, planar, 2-connected, and bipartite. Specifically, the bipartition is given by $\{x_v \mid v \in V_1\} \cup \{x_e^- \mid e \in E(D)\}$ and $\{x_v \mid v \in V_2\} \cup \{x_e^+ \mid e \in E(D)\}$.




Figure 11: C_4 -gadgets to construct the undirected graph $G = G(D)$ (right) from the digraph D (left) in the proof of Theorem 2. The edges of the matchings X and Y in G are indicated by bold solid and dashed lines, respectively.

We construct a pair of perfect matchings X, Y on G as follows. The first matching X is defined by fixing a particular perfect matching on each C_4 -gadget, and the second matching Y is obtained from

X by flipping all cycles formed by the C_4 -gadgets; see Figure 11. We claim that X and Y have flip distance at most two in G if and only if D has a directed Hamiltonian cycle. From this the theorem follows.

First, assume there is a directed Hamiltonian cycle H in D . We define a pair of cycles C_1, C_2 in G , where C_1 and C_2 both contain all the edges $\{x_e^+ x_e^- \mid e \in E(H)\}$, plus additional edges defined as follows. For each vertex $v \in V(D)$, consider the corresponding C_4 -gadget in G with the two incident edges corresponding to the edges incident with v on H . The endpoints of those edges on the gadget divide it into two alternating paths in X , one with matching edges at both ends and one with non-matching edges at both ends. We add the edges on those two types of paths to C_1 or C_2 , respectively. Note that C_1 is an alternating cycle in X . Moreover, after flipping C_1 , the cycle C_2 is alternating, and flipping C_2 yields Y , as each edge in a C_4 -gadget gets flipped once while the remaining edges are flipped an even number of times and thus remain unchanged.

For the reverse implication, assume that X and Y are connected by flipping at most two alternating cycles. As the symmetric difference $X \Delta Y$ contains at least three disjoint cycles (recall the assumption $|V(D)| \geq 3$), exactly two cycles C_1, C_2 are flipped to transform X into Y , and neither C_1 nor C_2 is one of the 4-cycles formed by the C_4 -gadgets. As the edges outside the gadgets remain unchanged, they are covered by both C_1, C_2 or by neither of them. We claim that $H := \{e \in E(D) \mid x_e^+ x_e^- \in E(C_1)\}$ is the arc set of a directed Hamiltonian cycle in D . Since up to isomorphism, H is obtained from $E(C_1)$ by contraction of the C_4 -gadgets, H forms a cycle in D (here we need that D is cubic). If the cycle H is not a directed cycle in D , there would be some $v \in V_1$ with two incoming incident edges from H . However, in this case, the path in the corresponding C_4 -gadget contained in C_1 consists of two edges, one of which is not in X , contradicting that C_1 is an alternating cycle in X . Finally, the directed cycle H has to be spanning. Indeed, if there is a C_4 -gadget not covered by C_1 , then C_2 would have to cover exactly this gadget, a contradiction. \square

Proof of Theorem 3. We use the following hardness result of Peroche [Pér84]. Given a digraph D , where each vertex has indegree and outdegree equal to 2, it is NP-complete to decide if $E(D)$ is the union of two directed Hamiltonian cycles. Given such a digraph D , let \overleftarrow{D} be the digraph obtained from D by reversing the direction of every arc. We regard D and \overleftarrow{D} as α -orientations X and Y of the same underlying graph G , where $\alpha(v) = 2$ for all $v \in V(G)$. The theorem follows by observing that the flip distance between X and Y is at most 2 if and only if $E(D)$ is the union of two directed Hamiltonian cycles. Moreover, the same statement holds when we only allow flipping edges that are oriented differently in X and Y . \square

5 Vertex flip distance between c -orientations

In this section we prove Theorem 6.

Recall that, given a graph G with a fixed vertex \top , we only allow vertex flips at vertices distinct from \top . In the case that G is connected, we distinguish between two types of dicuts as follows: we say that a dicut S in an orientation of G is **positive** with respect to \top if and only if the uniquely determined cut set U of S does not contain \top . Otherwise the dicut is called **negative**. We also define the **interior** of S , denoted $\text{Int}(S)$, as the cut set U of S if S is positive and as its complement \overline{U} if S is negative. That is, $\text{Int}(S)$ is the set of vertices on the side of the cut opposite to \top .

The following lemma is needed to decompose the edges of certain digraphs into dicuts with nested cut sets. Formally, for a digraph D and dicuts $S_1 = \delta(U_1)$ and $S_2 = \delta(U_2)$, the pair S_1, S_2 is called **laminar** if either $U_1 \subseteq U_2$, $U_2 \subseteq U_1$, $U_1 \cap U_2 = \emptyset$, or $U_1 \cup U_2 = V(D)$; see Figure 12. A family of dicuts in D is called **laminar** if all of its pairs are laminar. A **balanced digraph** is a digraph in which every cycle of the underlying graph has the same number of forward and backward edges.

Lemma 1. *Let D be a balanced digraph. Then $E(D)$ can be decomposed into a laminar family of disjoint minimal dicuts.*

Proof. We prove the statement by induction on $|E(D)|$. It is clearly true if $E(D) = \emptyset$. Assume for the induction step that $|E(D)| = k \geq 1$ and the statement holds for all digraphs with less than k arcs.

As D is balanced, it is obviously acyclic and therefore contains a source $s \in V(D)$. Each cycle in the underlying graph of D that contains s has exactly one forward and one backward edge incident to s . Therefore, the digraph D' obtained from D by contracting the set $\delta^+(\{s\})$ of arcs incident to s is still balanced and has less than k edges. By the induction hypothesis, there exists a laminar decomposition of $E(D') = E(D) \setminus \delta^+(\{s\})$ into disjoint dicuts in D' and thus in D . Note that the dicuts in D' are exactly those of D disjoint from $\delta^+(\{s\})$, and laminarity is preserved. Hence adding a decomposition of the directed vertex cut $\delta^+(\{s\})$ into minimal dicuts to the collection gives rise to a decomposition of $E(D)$ into disjoint minimal dicuts. This resulting decomposition is also laminar, as the cut set of the new minimal cuts is $\{s\}$, which is either contained in or disjoint from each cut set of the original decomposition of $E(D')$. \square

For every pair X, Y of c -orientations on a graph G , the **difference** $X \setminus Y$ denotes the set of arcs in X whose orientation is reversed in Y . The difference forms a balanced subdigraph of X (and reversing them yields a balanced subdigraph of Y). Consequently, Lemma 1 provides another proof that c -orientations can be reached from one another by flipping minimal dicuts.

We now consider the partial order defined on acyclic c -orientations of an n -vertex graph such that the covering relation corresponds to flipping a source vertex. Recall that by Theorem 4, this partial order is a distributive lattice. We now reuse a result from Propp [Pro02] and Felsner and Knauer [FK09] that gives an embedding of this distributive lattice into \mathbb{N}^{n-1} , which led to the introduction of distributive polytopes by Felsner and Knauer [FK11]. This theorem is illustrated in Figure 9, where the values of the functions z_X are depicted in the vertices of the graph G .

Theorem 9 ([Pro02, FK09]). *Let G be a graph on n vertices with a fixed vertex \top , X an acyclic c -orientation of G , and denote by X_{\min} the minimal element of the associated distributive lattice. Then the number of times $z_X(x)$ a vertex $x \in V(G) \setminus \{\top\}$ is flipped in an upward lattice path from X_{\min} to X is independent of the sequence. The resulting function $z_X : \mathcal{R}_c \rightarrow \mathbb{N}^{n-1}$ is a lattice embedding. That is, for every $x, y \in \mathbb{N}^{n-1}$ corresponding to c -orientations of G , the join and meet correspond to $\min(x, y)$ and $\max(x, y)$, respectively.*

In other words, the distributive lattice on \mathcal{R}_c is isomorphic to an induced sublattice of the componentwise dominance order on \mathbb{N}^{n-1} . We call a vertex flip sequence **monotone** if every flipped vertex is either only flipped as a source or only as a sink. With this definition, Theorem 9 yields the following:

Corollary 4. *Let G be a graph with fixed vertex \top and X, Y a pair of acyclic c -orientations on G . Then every monotone vertex flip sequence transforming X into Y has minimal length.*

Consider two c -orientations X, Y of G . Our goal is to construct a monotone flip sequence from X to Y . By Lemma 1, there is a laminar decomposition of the difference $X \setminus Y$ into minimal dicuts whose reversal yields Y . Denote by $\mathcal{S} = \mathcal{S}(X, Y)$ the dicuts of one such decomposition. We construct a poset P on \mathcal{S} by the inclusion order of the interiors of the minimal dicuts. That is, for $S, T \in \mathcal{S}$, S is ordered before T in P if and only if $\text{Int}(S) \subseteq \text{Int}(T)$; see Figure 12. Since \mathcal{S} is laminar, the cover graph of P is a forest, with the additional property that every non-maximal element S is covered by a unique other element, which we denote by $\text{cov}(S)$. Moreover, for each vertex $x \in V(G)$ in the interior of at least one of the cuts in \mathcal{S} , we let S_x be the (unique) minimal element of the poset P such that $x \in \text{Int}(S_x)$. Also, for each $S \in \mathcal{S}$ we denote by $\underline{\text{Int}}(S) := \{x \in V(G) \mid S_x = S\} \subseteq \text{Int}(S)$ the set of vertices in the interior of S but in none of the interiors of the cuts covered by S in the poset.




Figure 12: A laminar collection \mathcal{S} of disjoint minimal dicuts of $X \setminus Y$ (left), where the positive ones are dashed, and the negative ones are dotted, and the corresponding poset P of dicuts in \mathcal{S} ordered by inclusion (right) with its associated signed weights $\text{sgn}(S) \cdot w(S)$, $S \in \mathcal{S}$.

For each dicut $S \in \mathcal{S}$ we define an integer **weight** $w(S)$ and a **sign** $\text{sgn}(S) \in \{+, 0, -\}$ as follows; see Figure 12. If $S \in \mathcal{S}$ is a maximal element in P then we define $w(S) := 1$, and $\text{sgn}(S) := +$ if S is positive and $\text{sgn}(S) := -$ otherwise. For every sign $s \in \{+, 0, -\}$ and dicut $S \in \mathcal{S}$, we say that S **agrees** with s , if either $s = 0$, or if $s = +$ and S is positive, or $s = -$ and S is negative. For every non-maximal $S \in \mathcal{S}$, we inductively define

$$w(S) := \begin{cases} w(\text{cov}(S)) + 1 & \text{if } S \text{ agrees with } \text{sgn}(\text{cov}(S)), \\ w(\text{cov}(S)) - 1 & \text{otherwise,} \end{cases}$$

and

$$\text{sgn}(S) := \begin{cases} \text{sgn}(\text{cov}(S)) & \text{if } \text{sgn}(\text{cov}(S)) \neq 0 \text{ and } w(S) \neq 0, \\ + & \text{if } \text{sgn}(\text{cov}(S)) = 0, w(S) \neq 0 \text{ and } S \text{ is positive,} \\ - & \text{if } \text{sgn}(\text{cov}(S)) = 0, w(S) \neq 0 \text{ and } S \text{ is negative,} \\ 0 & \text{if } w(S) = 0. \end{cases}$$

It follows from this definition that the weights are non-negative and that $\text{sgn}(S) = 0$ if and only if $w(S) = 0$ for every $S \in \mathcal{S}$. We will see that given a minimal dicut S in \mathcal{S} , the weight $w(S)$ describes the number of times each vertex which lies in $\text{Int}(S)$ will be flipped, whereas $\text{sgn}(S)$ captures the direction in which (all) these vertices are flipped. That is, a positive sign means that vertices are flipped from sources to sinks, while a negative sign means that vertices are flipped from sinks to sources. We will need the following auxiliary statement.

Lemma 2. *Let X, Y be acyclic c -orientations of a connected graph G with fixed vertex \top . If $S = X \setminus Y$ is a positive dicut, then there is a vertex flip sequence transforming X into Y such that only vertices in $\text{Int}(S)$ are flipped, each exactly once from source to sink.*

The analogous statement for negative dicuts holds with sources and sinks exchanged.

Proof. We prove the statement by induction on $|\text{Int}(S)|$. If $\text{Int}(S)$ is a single vertex, then S corresponds to the arcs incident to a source, and the statement holds.

For the induction step assume $|\text{Int}(S)| = k \geq 2$ and that the claim holds for all positive cuts in c -orientations of G whose interiors have size less than k . Since X is acyclic, the induced subdigraph $X[\text{Int}(S)]$ is also acyclic and thus contains a source $x \in \text{Int}(S)$. Since $\text{Int}(S)$ is the cut set of S , x is a source in X as well. Let Z be the c -orientation obtained from X by a vertex flip at x . It follows that the cut $\delta(\text{Int}(S) \setminus \{x\})$ in Z is positive with interior of size $k - 1$. By induction, there is a vertex flip sequence from Z to Y such that only vertices in $\text{Int}(S) \setminus \{x\}$ are flipped, each exactly once from source to sink. Starting with a vertex flip at x and continuing with this flip sequence yields a flip sequence from X via Z to Y with the desired properties. \square

We are now in position to prove the main result of this section.

Theorem 10. *Let X, Y be acyclic c -orientations of a connected graph G . There is a monotone vertex flip sequence transforming X into Y which can be computed in time polynomial in $|E(G)|$.*

Proof. Consider the following strengthening of the theorem:

Claim: *Let X, Y be acyclic c -orientations of a connected graph G and $\mathcal{S} = \mathcal{S}(X, Y)$ a laminar decomposition of $X \setminus Y$ into disjoint minimal dicuts. Then there is a monotone vertex flip sequence from X to Y , such that every flipped vertex x is contained in the interior of some dicut of \mathcal{S} , and x is flipped $w(S_x)$ times from source to sink if $\text{sgn}(S_x) = +$, and from sink to source otherwise.*

We prove this claim by induction on the size of \mathcal{S} . The statement is clearly true if $X = Y$ (which means that $|\mathcal{S}| = 0$), settling the base case of the induction. Assume for the induction step that we are given a pair $X \neq Y$ of c -orientations and a laminar decomposition \mathcal{S} of $X \setminus Y$ of size $k \geq 1$. Assume that the claim holds for all pairs of c -orientations with a laminar decomposition of size less than k .

In the poset P on \mathcal{S} we consider a minimal element corresponding to a cut $S \in \mathcal{S}$, i.e., we have $\underline{\text{Int}}(S) = \text{Int}(S)$ and all vertices $x \in \text{Int}(S)$ satisfy $S_x = S$. Lemma 2 gives a vertex flip sequence F_1 that flips only vertices in $\text{Int}(S)$, each exactly once from source to sink if S is positive and from sink to source if S is negative. Applying this flip sequence to X , we obtain an intermediate c -orientation Z that differs from X only by the reversal of all edges in S . Consequently, $\mathcal{S} \setminus \{S\}$ is a laminar decomposition of $Z \setminus Y$ into minimal dicuts in Z of size $k - 1$. By induction, we also have a vertex flip sequence F_2 transforming Z into Y with the aforementioned properties.

Note that the weights and signs of all dicuts $T \in \mathcal{S} \setminus \{S\}$ defined with respect to \mathcal{S} or $\mathcal{S} \setminus \{S\}$ are the same, so we may simply write $w(T)$ and $\text{sgn}(T)$. Furthermore, the set $\underline{\text{Int}}(T)$ defined with respect to \mathcal{S} is a subset of the same set defined with respect to $\mathcal{S} \setminus \{S\}$. To complete the induction step, we distinguish two cases.

The first case is that S is a maximal element in P , or that S agrees with $\text{sgn}(\text{cov}(S))$. In this case, we claim that the concatenation F of F_1 and F_2 is a flip sequence transforming X via Z into Y with the desired properties. It suffices to check this for the vertices in $\underline{\text{Int}}(S) = \text{Int}(S)$, since for all other vertices, the claimed properties follow inductively (they are never flipped in F_1 , so their behavior in F will be the same as in F_2). If S is a maximal element in P , then $w(S) = 1$ and every vertex $x \in \text{Int}(S)$ will be flipped exactly once. Moreover, according to Lemma 2, if S is positive, i.e., $\text{sgn}(S) = +$, then x is flipped from source to sink, and if S is negative and $\text{sgn}(S) = -$, then x is flipped sink to source. It remains to consider the subcase that S is not maximal, i.e., $\text{cov}(S)$ exists. Consider any vertex $x \in \text{Int}(S)$. During the flip sequence F , the vertex x is flipped once in F_1 and $w(\text{cov}(S))$ times in F_2 , so $w(S) = w(\text{cov}(S)) + 1$ times in total, as required. Moreover, the assumption that S agrees with $\text{sgn}(\text{cov}(S))$ means that either $\text{sgn}(\text{cov}(S)) = +$ and S is positive or $\text{sgn}(\text{cov}(S)) = -$ and S is negative, or $w(\text{cov}(S)) = \text{sgn}(\text{cov}(S)) = 0$. We conclude from the inductive assumption that in those

three cases, x is only flipped from source to sink in both F_1 and F_2 , or only sink to source in both, or only once in F_1 but not in F_2 , respectively. Consequently, x satisfies the inductive claim in all cases.

The second case is that S is not maximal, i.e., $\text{cov}(S)$ exists, and that S does not agree with $\text{sgn}(\text{cov}(S))$. This means that $w(S) = w(\text{cov}(S)) - 1$. Without loss of generality, assume that S is positive and consequently $\text{sgn}(\text{cov}(S)) = -$ (the other case is symmetric). Consider again the vertex flip sequence F obtained by concatenating F_1 and F_2 . This flip sequence would transform X via Z into Y , however, we will not actually apply F , but modify the sequence as follows. By Lemma 2, F_1 flips each vertex in $\text{Int}(S)$ exactly once from source to sink. By induction, in F_2 , each vertex in $\text{Int}(S)$ contained in the interior of $\text{cov}(S)$ (defined with respect to $\mathcal{S} \setminus \{S\}$) is flipped from sink to source. Let x be the last element of F_1 and consider the subsequence x, x_1, \dots, x_k, x of F starting with x and ending with the first occurrence of x in F_2 . None of the vertices x_1, \dots, x_k is adjacent to x in G , because after the first vertex flip at x (from source to sink) all edges incident with x are incoming, and in F_2 we only flip sinks to sources. This shows that deleting the first two occurrences of x from F preserves the number of and direction of all flips at vertices distinct from x , and still transforms X into Y . Repeated application of this argument produces a reduced vertex flip sequence F' transforming X into Y such that each vertex $x \in V(G) \setminus \text{Int}(S)$ is flipped the same number of times and in the same direction as in F_2 . That is, x is flipped $w(S_x)$ times from source to sink if $\text{sgn}(S_x) = +$, and $w(S_x)$ times from sink to source if $\text{sgn}(S_x) = -$. On the other hand, every $x \in \text{Int}(S)$ is missing its first occurrence but is flipped in the same way from sink to source for all remaining occurrences. This implies that x is flipped $w(S) = w(\text{cov}(S)) - 1$ times from sink to source, as it should. This proves that F' is a vertex flip sequence from X to Y satisfying the conditions in our claim, completing its proof.

It remains to verify that the recursive algorithm obtained from this inductive argument runs in time polynomial in $m := |E(G)|$. First of all, the number of dicuts in any laminar decomposition, which corresponds to the number of induction steps, is bounded by the number of edges m . Consequently, it suffices to show that the number of operations needed in one induction step is bounded by a polynomial in m . Specifically, we need to compute the cover relations of P , the weights and signs of the dicuts, find a minimal element of the poset P , test its properties for the case distinction and construct the resulting flip sequence by concatenation and possibly deletion of double occurrences, all of which can be done in time $\mathcal{O}(m^2)$. This proves an upper bound of $\mathcal{O}(m^3)$ for the total number of steps performed for the construction of the monotone flip sequence to transform X into Y . Finally, a laminar decomposition \mathcal{S} of $X \setminus Y$ as guaranteed by Lemma 1 can be computed in polynomial time, by following the recursive strategy explained in the proof of the lemma. This completes the proof. \square

Combining Corollary 4 and Theorem 10 yields Theorem 6.

6 Flip distance with larger cut sets

In this section we prove Theorem 7 by reduction from the following NP-hard problem.

Given a (finite) poset (P, \prec) , its **height** is the maximum size k of a chain $x_1 \prec x_2 \prec \dots \prec x_k$ in P . A **linear extension** of P is a sequence (x_1, \dots, x_n) of all elements of P such that $x_i \prec x_j$ implies that $i < j$. Given a linear extension $L = (x_1, \dots, x_n)$ of P , a **jump** is a pair x_i, x_{i+1} in L for which $x_i \not\prec x_{i+1}$ in P . Conversely, a **bump** is a pair x_i, x_{i+1} such that $x_i \prec x_{i+1}$. The **jump number** $s(P)$ of P is the minimum number of jumps among all linear extensions of P . The **Jump Number Problem** is the algorithmic problem of computing the jump number of a poset given by its comparabilities.

Theorem 11 ([Pul75, Mül90]). *Determining the jump number of a poset of height two is NP-hard.*

Proof of Theorem 7. We provide a Turing-reduction of the Jump Number Problem for posets of height two to the problem stated in the theorem. For this purpose, assume we are given a poset (P, \prec) of height two with bipartite Hasse diagram $G = (P_1 \cup P_2, E)$ as an instance for the Jump Number

Problem. We may assume that P has no isolated elements and that P_1 contains all minimal elements and P_2 all maximal elements of the poset. We construct an auxiliary Hasse diagram G' from G by adding an additional unique maximal element \top , and connecting it with edges to all vertices of G . We construct two orientations X, Y of G' as follows: In both orientations all edges are oriented from P_1 to P_2 . Moreover, in X all edges incident with \top are oriented towards \top , while in Y all these edges are oriented away from \top . As X, Y are obtained from each other by flipping all edges incident with \top (this flip is not allowed, though, as \top is the fixed vertex), they are c -orientations with respect to the same c .

Let d denote the minimal flip distance between these c -orientations according to the conditions of the theorem. We will complete the proof by showing that $s(P) = d - 1$.

We first show that $s(P) \geq d - 1$. For this argument, let $L = (x_1, \dots, x_n)$ be an arbitrary linear extension of P . As P has height two, the elements $\{x_1, \dots, x_n\}$ of P are partitioned into subsets B_1, \dots, B_m of size one or two, such that for all B_i, B_j with $i < j$, the elements from B_i appear before the elements from B_j in L , and such that the two-element sets B_i contain exactly all bump pairs. We define a flip sequence that starts with the orientation X and consecutively flips the cuts induced by B_1, \dots, B_m . Since for all $1 \leq i \leq m$, B_i and $\overline{B_i} := (P \cup \{\top\}) \setminus B_i$ induce connected subgraphs of G' , these are indeed minimal cuts. Moreover, each of these cuts is flippable. This is obviously true for B_1 , as B_1 induces a dicut in X . Now assume inductively that the cuts induced by B_1, \dots, B_{k-1} for some $k \geq 2$ have been flipped. As L is a linear extension of P , all elements in the downset of B_k in P but not in B_k are in one of the B_i with $i < k$. This implies that every arc between some $x \in B_k$ and $y \notin B_k$ is oriented from x to y in the current orientation, and thus B_k is indeed flippable.

In this flip sequence, every arc in X not incident to \top will be flipped zero or two times and thus maintains its original orientation, while all the edges incident to \top get reversed, as they are incident to exactly one set B_i . Consequently, the flip sequence transforms X into Y , proving that $d \leq m$. As m equals the number of jumps in L plus 1 (every non-jump is a bump within one of the B_i), this yields $d - 1 \leq s(P)$.

We now show that $s(P) \leq d - 1$. Assume that $B_1, \dots, B_m \subseteq P$ are the cut sets of size one or two appearing (in this order) in a shortest flip sequence transforming X into Y . We may assume that among all shortest flip sequences, this sequence also minimizes $|B_1| + |B_2| + \dots + |B_m|$. Since each vertex $x \in P$ has an outgoing arc to \top in X which must be reversed during the flip sequence, x must be contained in at least one of the B_i . We claim that x is contained in at most one of the B_i . That is, the B_i are pairwise disjoint. Assume to the contrary that $x \in B_i \cap B_j$ for some $i < j$ and that B_i, B_j is the only intersecting pair among B_i, B_{i+1}, \dots, B_j (by minimizing $j - i$). In particular, none of the cut sets B_{i+1}, \dots, B_{j-1} contains x , and x is the only vertex flipped multiple times in this subsequence. We are then in one of the four cases $B_i = B_j = \{x\}$, or $B_i = \{x, y\}$ and $B_j = \{x\}$, or $B_i = \{x\}$ and $B_j = \{x, z\}$, or $B_i = \{x, y\}$ and $B_j = \{x, z\}$ for some elements $y, z \in P$ distinct from x . Since no vertex adjacent to x in G' can be flipped by B_{i+1}, \dots, B_{j-1} , it follows that in each of these cases, the sequence

$$B_1, \dots, B_{i-1}, B_i \setminus \{x\}, B_{i+1}, \dots, B_{j-1}, B_j \setminus \{x\}, B_{j+1}, \dots, B_m$$

is a valid flip sequence from X to Y of length at most m and with decreased sum $|B_1| + |B_2| + \dots + |B_m|$, a contradiction. This proves that the cut sets B_i are pairwise disjoint.

The B_i are flipped one after the other and by definition of X , the dicut induced by B_i is flippable if and only if all the elements in the downset of B_i with respect to P but not in B_i were flipped before. Therefore, by listing the elements in the sets B_1, \dots, B_m in this relative order, and ordering the elements within each B_i according to their order in P , we obtain a linear extension L of P whose jumps are exactly those pairs having elements in two consecutive sets B_i . It follows that there are $m - 1 = d - 1$ jumps in L , proving that $s(P) \leq d - 1$.

Combining these arguments shows that $s(P) = d - 1$, and using Theorem 11 we obtain the claimed hardness result. \square

7 Open problems

Recall that Problem 2 asks for a shortest flip sequence of directed cycles transforming one α -orientation X into another one Y , where we only allow flipping edges that are oriented differently in X and Y . Since the set of edges that are oriented differently in X and Y form a Eulerian subdigraph D of both X and Y , we have the following natural question:

Question 1. *What is the smallest number of directed cycles into which a Eulerian digraph can be decomposed?*

We have seen in Theorem 3 that from a computational point of view, this problem is hard for general digraphs, but we wonder what happens when adding planarity constraints. The aforementioned question can also be studied in terms of upper bounds as a function of the number of vertices, which is related to the famous Hajós conjecture on undirected Eulerian graphs, see [Lov68]. Another interesting undirected variant of Question 1 is the following:

Question 2. *Given a graph G with a Eulerian subgraph H . What is the smallest number of cycles of G such that their symmetric difference is H ?*

Concerning our proof of Theorem 7, we believe that for any bound on the size of the cuts, the corresponding flip distance will be NP-hard to compute. On the other hand, we use very particular graphs as gadgets, and we do not know the complexity of the corresponding problem for planar α -orientations. We think the following is an interesting special case:

Question 3. *Let X, Y be perfect matchings of a planar bipartite 3-connected graph G . What is the complexity of determining the distance of X and Y with respect to alternating cycles that are either a face or the symmetric difference of two incident faces?*

The feeling that this problem might be tractable is supported by the following observation. It is not difficult to show that every height two poset with bipartite planar Hasse diagram has dimension at most two. It then follows from [SS87] that the restriction of the Jump Number Problem to such posets is solvable in polynomial time, and thus, the hardness reduction presented in the previous section fails.

Acknowledgements

This work was initiated during the workshop “Order & Geometry” 2018 in Ciężen Palace. We thank the organizers and participants of this workshop for the stimulating atmosphere.

References

- [AA17] Marcelo Aguiar and Federico Ardila. Hopf monoids and generalized permutahedra. ArXiv: 1709.07504, September 2017.
- [AAHV07] Oswin Aichholzer, Franz Aurenhammer, Clemens Huemer, and Birgit Vogtenhuber. Gray code enumeration of plane straight-line graphs. *Graphs Combin.*, 23(5):467–479, 2007.
- [AF96] David Avis and Komei Fukuda. Reverse search for enumeration. *Discrete Appl. Math.*, 65(1-3):21–46, 1996. First International Colloquium on Graphs and Optimization (GOI), 1992 (Grimentz).
- [AMP15] Oswin Aichholzer, Wolfgang Mulzer, and Alexander Pilz. Flip distance between triangulations of a simple polygon is NP-complete. *Discrete Comput. Geom.*, 54(2):368–389, 2015.

- [BF12] Olivier Bernardi and Éric Fusy. A bijection for triangulations, quadrangulations, pentagulations, etc. *J. Combin. Theory Ser. A*, 119(1):218–244, 2012.
- [BH09] Prosenjit Bose and Ferran Hurtado. Flips in planar graphs. *Comput. Geom.*, 42(1):60–80, 2009.
- [Bre00] Enno Brehm. 3-orientations and schnyder-3-tree-decompositions, 2000. Diploma Thesis, Freie Universität Berlin.
- [Bru69] Richard A. Brualdi. Comments on bases in dependence structures. *Bull. Austral. Math. Soc.*, 1:161–167, 1969.
- [BV11] Prosenjit Bose and Sander Verdonschot. A history of flips in combinatorial triangulations. In *Computational Geometry - XIV Spanish Meeting on Computational Geometry, EGC 2011, Dedicated to Ferran Hurtado on the Occasion of His 60th Birthday, Alcalá de Henares, Spain, June 27-30, 2011, Revised Selected Papers*, pages 29–44, 2011.
- [CSJ09] Sean Cleary and Katherine St. John. Rotation distance is fixed-parameter tractable. *Inform. Process. Lett.*, 109(16):918–922, 2009.
- [CSJ10] Sean Cleary and Katherine St. John. A linear-time approximation for rotation distance. *J. Graph Algorithms Appl.*, 14(2):385–390, 2010.
- [CSS18] Jean Cardinal, Vera Sacristán, and Rodrigo I. Silveira. A note on flips in diagonal rectangulations. *Discrete Mathematics & Theoretical Computer Science*, 20(2), 2018.
- [CSZ15] Cesar Ceballos, Francisco Santos, and Günter M. Ziegler. Many non-equivalent realizations of the associahedron. *Combinatorica*, 35(5):513–551, 2015.
- [DP02] Brian A. Davey and Hilary A. Priestley. *Introduction to lattices and order*. Cambridge University Press, New York, second edition, 2002.
- [Fel04] Stefan Felsner. Lattice structures from planar graphs. *Electron. J. Combin.*, 11(1):Research Paper 15, 24, 2004.
- [Fel13] Stefan Felsner. Rectangle and square representations of planar graphs. In *Thirty essays on geometric graph theory*, pages 213–248. Springer, New York, 2013.
- [FK09] Stefan Felsner and Kolja Knauer. ULD-lattices and Δ -bonds. *Combin. Probab. Comput.*, 18(5):707–724, 2009.
- [FK11] Stefan Felsner and Kolja Knauer. Distributive lattices, polyhedra, and generalized flows. *European J. Combin.*, 32(1):45–59, 2011.
- [FKMS18] Stefan Felsner, Linda Kleist, Torsten Mütze, and Leon Sering. Rainbow cycles in flip graphs. In *34th International Symposium on Computational Geometry, SoCG 2018, June 11-14, 2018, Budapest, Hungary*, pages 38:1–38:14, 2018.
- [FSS18a] Stefan Felsner, Hendrik Schrezenmaier, and Raphael Steiner. Equiangular polygon contact representations. In *Graph-Theoretic Concepts in Computer Science - 44th International Workshop, WG 2018, Cottbus, Germany, June 27-29, 2018, Proceedings*, pages 203–215, 2018.
- [FSS18b] Stefan Felsner, Hendrik Schrezenmaier, and Raphael Steiner. Pentagon contact representations. *Electron. J. Combin.*, 25(3):Paper 3.39, 38, 2018.

- [FT88] András Frank and Éva Tardos. Generalized polymatroids and submodular flows. *Math. Programming*, 42(3, (Ser. B)):489–563, 1988. Submodular optimization.
- [GL86] Patrick M. Gilmer and Richard A. Litherland. The duality conjecture in formal knot theory. *Osaka J. Math.*, 23(1):229–247, 1986.
- [GLP12] Daniel Gonçalves, Benjamin Lévêque, and Alexandre Pinlou. Triangle contact representations and duality. *Discrete Comput. Geom.*, 48(1):239–254, 2012.
- [HHN02] M. Carmen Hernando, Ferran Hurtado, and Marc Noy. Graphs of non-crossing perfect matchings. *Graphs Combin.*, 18(3):517–532, 2002.
- [HHNOP09] Clemens Huemer, Ferran Hurtado, Marc Noy, and Elsa Omaña-Pulido. Gray codes for non-crossing partitions and dissections of a convex polygon. *Discrete Appl. Math.*, 157(7):1509–1520, 2009.
- [HHNRC05] Michael E. Houle, Ferran Hurtado, Marc Noy, and Eduardo Rivera-Campo. Graphs of triangulations and perfect matchings. *Graphs Combin.*, 21(3):325–331, 2005.
- [Iwa02] Satoru Iwata. On matroid intersection adjacency. *Discrete Math.*, 242(1-3):277–281, 2002.
- [Kna07] Kolja Knauer. Partial orders on orientations via cycle flips. Master’s thesis, Faculty of Mathematics, TU Berlin, 2007.
- [Kna08] Kolja Knauer. Distributive lattices on graph orientations. In *Semigroups, acts and categories with applications to graphs*, volume 3 of *Math. Stud. (Tartu)*, pages 79–91. Est. Math. Soc., Tartu, 2008.
- [KSX17] Iyad Kanj, Eric Sedgwick, and Ge Xia. Computing the flip distance between triangulations. *Discrete Comput. Geom.*, 58(2):313–344, 2017.
- [LEP10] Fabrizio Luccio, Antonio Mesa Enriquez, and Linda Pagli. Lower bounds on the rotation distance of binary trees. *Inform. Process. Lett.*, 110(21):934–938, 2010.
- [Lov68] László Lovász. On covering of graphs. *Theory of Graphs, Proc. Colloq. Tihany, Hungary 1966*, 231-236 (1968), 1968.
- [LP07] Thomas Lam and Alexander Postnikov. Alcoved polytopes. I. *Discrete Comput. Geom.*, 38(3):453–478, 2007.
- [LP15] Anna Lubiw and Vinayak Pathak. Flip distance between two triangulations of a point set is NP-complete. *Comput. Geom.*, 49:17–23, 2015.
- [LRS10] Jesús A. De Loera, Jörg Rambau, and Francisco Santos. *Triangulations: Structures for Algorithms and Applications*. Number 25 in Algorithms and Computation in Mathematics. Springer, 2010.
- [LZ98] Ming Li and Louxin Zhang. Better approximation of diagonal-flip transformation and rotation transformation. In *Computing and Combinatorics, 4th Annual International Conference, COCOON ’98, Taipei, Taiwan, R.o.C., August 12-14, 1998, Proceedings*, pages 85–94, 1998.
- [LZ03] Peter Che Bor Lam and Heping Zhang. A distributive lattice on the set of perfect matchings of a plane bipartite graph. *Order*, 20(1):13–29, 2003.

- [Mül90] Haiko Müller. Alternating cycle-free matchings. *Order*, 7(1):11–21, 1990.
- [Neg94] Seiya Negami. Diagonal flips in triangulations of surfaces. *Discrete Math.*, 135(1-3):225–232, 1994.
- [Pap78] Christos H. Papadimitriou. The adjacency relation on the traveling salesman polytope is NP-complete. *Math. Programming*, 14(3):312–324, 1978.
- [Pér84] Bernard Péroche. NP-completeness of some problems of partitioning and covering in graphs. *Discrete Appl. Math.*, 8(2):195–208, 1984.
- [Ple79] Ján Plesník. The NP-completeness of the Hamiltonian cycle problem in planar digraphs with degree bound two. *Inform. Process. Lett.*, 8(4):199–201, 1979.
- [Pos09] Alexander Postnikov. Permutohedra, associahedra, and beyond. *Int. Math. Res. Not. IMRN*, (6):1026–1106, 2009.
- [Pou14] Lionel Pournin. The diameter of associahedra. *Adv. Math.*, 259:13–42, 2014.
- [Pro02] James Propp. Lattice structure for orientations of graphs. ArXiv: math/0209005, September 2002.
- [PS18] Vincent Pilaud and Francisco Santos. Quotientopes. ArXiv: 1711.05353, August 2018.
- [Pul75] William R. Pulleyblank. On minimizing setups in precedence constrained scheduling. Technical report, University of Bonn, 1975. Technical Report 81105-OR.
- [PW98] James Propp and David Wilson. Coupling from the past: a user’s guide. In *Microsurveys in discrete probability (Princeton, NJ, 1997)*, volume 41 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 181–192. Amer. Math. Soc., Providence, RI, 1998.
- [RC98] Fred J. Rispoli and Steven Cosares. A bound of 4 for the diameter of the symmetric traveling salesman polytope. *SIAM J. Discrete Math.*, 11(3):373–380, 1998.
- [Rém04] Eric Rémila. The lattice structure of the set of domino tilings of a polygon. *Theoret. Comput. Sci.*, 322(2):409–422, 2004.
- [Rin57] Gerhard Ringel. Über Geraden in allgemeiner Lage. *Elem. Math.*, 12:75–82, 1957.
- [Rog99] Rodney O. Rogers. On finding shortest paths in the rotation graph of binary trees. In *Proceedings of the Thirtieth Southeastern International Conference on Combinatorics, Graph Theory, and Computing (Boca Raton, FL, 1999)*, volume 137, pages 77–95, 1999.
- [SS87] George Steiner and Lorna K. Stewart. A linear time algorithm to find the jump number of 2-dimensional bipartite partial orders. *Order*, 3(4):359–367, 1987.
- [Sta86] Richard P. Stanley. Two poset polytopes. *Discrete Comput. Geom.*, 1(1):9–23, 1986.
- [STT88] Daniel D. Sleator, Robert E. Tarjan, and William P. Thurston. Rotation distance, triangulations, and hyperbolic geometry. *J. Amer. Math. Soc.*, 1(3):647–681, 1988.
- [Thu90] William P. Thurston. Conway’s tiling groups. *Amer. Math. Monthly*, 97(8):757–773, 1990.
- [ZG87] Fu Ji Zhang and Xiao Feng Guo. Hamilton cycles in directed Euler tour graphs. *Discrete Math.*, 64(2-3):289–298, 1987.
- [ZQZ19] Wei Juan Zhang, Jian Guo Qian, and Fu Ji Zhang. Distance between α -orientations of plane graphs by facial cycle reversals. *Acta Mathematica Sinica, English Series*, Jan 2019.