

SHORTEST PATHS ON POLYMATROIDS AND HYPERGRAPHIC POLYTOPES

Jean Cardinal¹ and Raphael Steiner^{*2}

¹*Department of Computer Science, Université libre de Bruxelles (ULB), Brussels, Belgium
jean.cardinal@ulb.be*

²*Department of Mathematics, ETH Zürich, Zurich, Switzerland
raphaelmario.steiner@math.ethz.ch*

Submitted: Oct 25, 2024; Accepted: Mar 17, 2025; Published: Sep 15, 2025

© The authors. Released under the CC BY license (International 4.0).

Abstract. Base polytopes of polymatroids, also known as generalized permutohedra, are polytopes whose edges are parallel to a vector of the form $\mathbf{e}_i - \mathbf{e}_j$, where the $\{\mathbf{e}_i\}_{i \in [n]}$ are the canonical basis vectors of \mathbb{R}^n .

We consider the following computational problem: Given two vertices of a generalized permutohedron P , determine the length of a shortest path between them on the skeleton of P , where the length of a path is its number of edges. This captures many known flip distance problems, such as computing the minimum number of exchanges between two spanning trees of a graph, the rotation distance between binary search trees, the flip distance between acyclic orientations of a graph, or rectangulations of a square.

We prove that this general problem is NP-hard, even when restricted to very simple polymatroids in \mathbb{R}^n defined by $O(n)$ inequalities. Assuming $P \neq NP$, this shows that even when maximizing a linear functional over a polymatroid, we cannot hope for the existence of a computationally efficient simplex pivoting rule that performs a minimum number of nondegenerate pivoting steps to an optimal solution. Such results have previously been shown only for other, arguably more complicated classes of polytopes.

We also prove that the shortest path problem is inapproximable when the polymatroid is specified via an evaluation oracle for a corresponding submodular function, which strengthens a recent result by Ito, Kakimura, Kamiyama, Kobayashi, Maezawa, Nozaki, and Okamoto (ICALP'23). More precisely, we prove that it is NP-hard to approximate the length of a shortest path to within a factor $(1 + \varepsilon)$ for some absolute constant $\varepsilon > 0$, even when the polymatroid is a hypergraphic polytope, whose vertices are in bijection with acyclic orientations of a given hypergraph. The shortest path problem then amounts to computing the flip distance between two acyclic orientations of a hypergraph.

*Research supported by SNSF Ambizione grant No. 216071.

On the positive side, we provide a polynomial-time algorithm which, given any pair of acyclic orientations of a hypergraph, computes a connecting path whose length approximates the length of a shortest path to within a factor bounded by the maximum codegree of the hypergraph. Our result implies in particular an exact polynomial-time algorithm for computing shortest flip sequences between acyclic orientations of any linear hypergraph.

Keywords. Polymatroids, Generalized permutahedra, Hypergraphic polytopes, Simplex method, Shortest paths, Polytope diameter, Polytope skeleton, Flip distance, Combinatorial reconfiguration

Mathematics Subject Classifications. 90C05, 90C08, 90C27, 90C35, 90C49, 90C57, 90C60, 05C50, 05C65, 05B35, 52B40

1. Introduction

We consider the problem of computing (the length of) a shortest path between two vertices of a graph formed by the vertices and edges of an n -dimensional polytope, where the length of the path is its number of edges. Bounding the length of such paths has been the topic of intensive research since the birth of the theory of linear programming and the invention by Dantzig, around 1947, of the Simplex algorithm, that seeks solutions of a linear program by iteratively moving from one feasible solution to another, along the edges of the polytope. The performance of this method, parameterized by a so-called *pivoting rule*, is still the subject of many open questions. In particular, it is not known whether there exists a pivoting rule that makes the simplex algorithm run in strongly polynomial time. In a series of recent papers, the question of whether any efficient algorithm could find (approximately) shortest paths to a solution has been studied, and several negative results have been proved [ACH⁺21, DLKS22, IKK⁺23, IKK⁺22]. In particular, de Loera, Kafer and Sanitá [DLKS22] proved that finding (monotone) shortest paths to some optimal solution of an LP on the skeleton of a $\{0, 1\}$ -polytope is NP-hard. This result was later strengthened by Cardinal and Steiner [CS25], showing that computing paths between two perfect matchings of a bipartite graph on the perfect matching polytope whose length approximates the minimum length to within any constant factor is NP-hard. This implies that no simplex pivoting rule can perform an approximately optimal number of (nondegenerate) pivoting steps, even on linear programs defined by totally unimodular matrices.

Another motivation for studying the length of shortest paths on polytopes is the problem of computing the rotation distance between two binary search trees, or, equivalently, the flip distance between two triangulations of a convex polygon. The equivalence can be proved by a classical Catalan bijection that maps every triangulation to its dual tree. This problem, too, can be cast as computing the distance between two vertices of a polytope. The associated polytope is called the *associahedron*, and has been the topic of numerous investigations in various mathematical contexts, including data structures, topology, and algebra [STT22, Lod04, Pou14, CSZ15]. The complexity status of the rotation distance problem is another widely open question. It is not known to be NP-hard, but no polynomial-time algorithm is known either. The only known results are a simple polynomial-time 2-approximation algorithm and several polynomial-time fixed-parameter algorithms [CSJ09, LX23].

These two problems, the existence of efficient pivoting rules and computing the rotation dis-

tance between binary trees, are our main motivations for introducing the question of computing the distance between vertices of *base polytopes of polymatroids*, also known in the combinatorics literature as *generalized permutohedra*. In what follows, we will refer to them simply as *polymatroids*. From the optimization point of view, these are polytopes that generalize matroid base polytopes, while preserving the property that the vertices which are extremal with respect to a linear function can be found easily using a greedy algorithm [Sch03, Fuj05]. Hence just as in the case of perfect matching polytopes, polymatroids yield another class of linear programs that are solvable in strongly polynomial time. While polymatroids have been introduced and studied by Edmonds as early as 1970 [Edm03], they play a key role in more recent developments in algebraic combinatorics, where they are defined as polytopes whose normal fan is a coarsening of the normal fan of the permutohedron [PRW08, Pos09]. They are fundamental objects, among others, in the study of combinatorial Hopf algebras [AA23, LR12, Mee19]. Associahedra are classical examples of polymatroids, hence the rotation distance between binary search trees is a special case of a shortest path problem on a polymatroid.

Recently, Ito, Kakimura, Kamiyama, Kobayashi, Maezawa, Nozaki, and Okamoto [IKK⁺23] showed that unless $P = NP$, there is no polynomial-time algorithm for finding a shortest path between two vertices of a polymatroid, when it is given in the form of an oracle to its associated submodular rank function, a standard computation model for optimization on matroids and polymatroids. More precisely, they prove the hardness of computing the rotation distance between elimination trees on graphs, a problem that generalizes the rotation distance between binary search trees, and amounts to the lengths of shortest paths on a class of polymatroids known as *graph associahedra* [CD06, MP15].

1.1. Our results

We strengthen the result by Ito et al. [IKK⁺23] in two ways.

We first show that there exist simply-structured polymatroids in \mathbb{R}^n described by a set of $2n + 2$ inequalities, on which computing distances is strongly NP-hard. We show that shortest paths on those polytopes correspond to optimal solutions of the known *pure constant fixed charge transportation* problem, a special case of the classical fixed charge transportation problem.

We then consider the shortest paths in the setting where the polymatroid is specified by an oracle to a submodular function. Strengthening the aforementioned hardness result by Ito et al., we show that in this computational model, the length of a shortest path between two vertices of a polymatroid cannot even be efficiently approximated to within a factor $(1 + \varepsilon)$ for some absolute constant $\varepsilon > 0$, unless $P = NP$. To prove this result, we consider a special family of polymatroids defined by acyclic orientations of hypergraphs. These polytopes are known as *hypergraphical polytopes* [BBM19, CHM⁺23]. They form a natural generalization of polymatroids defined by acyclic orientations of graphs, and have recently been extensively studied both from an algebraic and combinatorial viewpoint [AA23, Reh22]. They also include the special cases of associahedra and graph associahedra mentioned above.

Finally, we provide a positive result about acyclic orientations of hypergraphs. We show that the shortest path problem on hypergraphical polytopes can be approximated within a factor equal to the maximum codegree of the hypergraph. This provides a new polynomial-time exact algorithm

for the shortest path problem on hypergraphic polytopes corresponding to *linear* hypergraphs, in which any two hyperedges intersect in at most one vertex.

Section 2 sets the stage for the problems we tackle and details a number of related computational questions. In Section 3 we prove the hardness of the shortest path problem on a polymatroid defined by linear inequalities. Section 4 describes the necessary definitions and background related to acyclic orientations of hypergraphs and the corresponding flip graphs. In Sections 5 and 6, we then present the proofs of the inapproximability result and the codegree approximation algorithm, respectively.

1.2. Other related works

Our contribution is also relevant in the context of *combinatorial reconfiguration* problems [IDH⁺11, vdH13, GIKO24]. In such problems, we are given two combinatorial objects together with local reconfiguration operations, and we are asked to decide whether it is possible to transform the first object into the other by iteratively applying a reconfiguration operation, and if it is, to find the minimum number of required steps. This problem is typically posed for structures defined as solutions of optimization problems, such as graph colorings or dominating sets [MN20], independent sets [IKO⁺20], cliques [IOO23], bounded-degree spanning trees [BIK⁺23], k -edge-connected orientations [IIK⁺23], among others. Our results can be seen as contributions in this vein, although in the cases we consider, the question of whether any two objects are connected is trivial, since the polytopal structure ensures connectivity (and in fact, d -connectivity, where d is the dimension of the polytope, from Balinski's Theorem [Bal61]).

2. Definitions and background

In what follows, we use the notation $\mathbf{p} = (p_1, p_2, \dots, p_n)$ for n -dimensional vectors, and let $[n] := \{1, 2, \dots, n\}$. We first consider the problem of finding the length of a shortest path between two vertices of a polytope specified by linear inequalities.

Problem 2.1 (Shortest paths on a polytope). Given a polytope P in \mathbb{R}^n , described by a system of linear inequalities, and two vertices \mathbf{a} and \mathbf{b} of P , each identified by n tight inequalities, find the length of a shortest path from \mathbf{a} to \mathbf{b} on the skeleton of P . Here the length of a path is defined as the number of edges.

This problem is sometimes referred to as *routing* on a polytope. It is closely related to the complexity of the Simplex method and the Hirsch conjecture. In its polynomial version, the latter states that the diameter of n -polytopes with m facets is bounded by a polynomial in n and m . If the polynomial Hirsch conjecture holds, then Problem 2.1 lies in NP. It is known that Problem 2.1 is NP-hard to approximate within any constant factor, even when P is restricted to be the perfect matching polytope of a bipartite graph [ACH⁺21, CS25]. We are interested in proving similar results in the case where the polytope is a polymatroid.

2.1. Polymatroids and generalized permutohedra

Base polytopes of polymatroids are also known as *generalized permutohedra*. In what follows, with a slight abuse of terminology, we will simply refer to them as *polymatroids*. They can be characterized in different ways [Sch03, Fuj05, DF10, ACEP20], generalizing the characterizations of matroid base polytopes by their edge lengths and the matroid rank function [GGMS87, Sch03].

Recall that a set function $f : 2^{[n]} \rightarrow \mathbb{R}$ is called *submodular* if the inequality

$$f(T) + f(U) \geq f(T \cap U) + f(T \cup U)$$

holds for all pairs U, T of subsets of $[n]$. We let \mathbf{e}_i denote the i^{th} canonical basis vector of \mathbb{R}^n .

Theorem 2.2 (Characterizations of polymatroids). *Given a polytope $P \subset \mathbb{R}^n$, the following statements are equivalent:*

1. Every edge of P is parallel to $\mathbf{e}_j - \mathbf{e}_i$ for a pair $i, j \in [n]$.
2. There exists a submodular function $f : 2^{[n]} \rightarrow \mathbb{R}$ such that

$$P = P_f := \left\{ \mathbf{x} \in \mathbb{R}^n : \sum_{i \in U} x_i \leq f(U) \forall U \subset [n], \text{ and } \sum_{i \in [n]} x_i = f([n]) \right\}.$$

The reader is referred to [Sch03, Fuj05, Pos09, AA23] for details on this classical statement. We focus on the case where f only takes on integral values. Then the right-hand sides of the inequalities defining P are integer, which implies that P is an integral polytope, and more precisely an *integer polymatroid*.

2.2. Shortest paths on polymatroids

We consider a variant of Problem 2.1 in which the input polytope P is encoded by a submodular function $f : 2^{[n]} \rightarrow \mathbb{R}$.

Problem 2.3 (Shortest paths on a polymatroid). *Given an oracle to a submodular function f , and two vertices \mathbf{a} and \mathbf{b} of the polytope P_f , compute the length of a shortest path from \mathbf{a} to \mathbf{b} on the skeleton of P_f .*

Several interesting computational problems can be reduced to Problem 2.3.

2.2.1 Matroids.

Clearly, when f is the rank function of a *matroid* M on the ground set $[n]$, then Problem 2.3 consists of computing the base exchange distance between two bases A and B of M , which is simply $|A \setminus B|$ (or, symmetrically, $|B \setminus A|$). In fact, a shortest path between two vertices \mathbf{x}, \mathbf{y} of P_f in this case can be found greedily using polynomially many oracle calls to the rank function f : We start by reading off the corresponding bases A and B as the supports of \mathbf{x} and \mathbf{y} . We then

greedily construct a sequence of bases $A = B_0, B_1, \dots, B_t = B$ inducing a shortest path from \mathbf{x} to \mathbf{y} on P_f as follows. For each $k \geq 0$ such that $B \setminus B_k \neq \emptyset$, we search through all potential pairs $(i, j) \in (B \setminus B_k) \times (B_k \setminus B)$ and test whether $(B_k \cup \{i\}) \setminus \{j\}$ is a basis (by using the oracle to decide if $f((B_k \cup \{i\}) \setminus \{j\}) = f([n])$). Such a pair always exists by the basis exchange axiom for matroids. Once a suitable pair (i, j) is found, we define $B_{k+1} := (B_k \cup \{i\}) \setminus \{j\}$ and proceed. Thus, Problem 2.3 is efficiently solvable in the case of matroids. In the case of graphic matroids, for instance, this amounts to computing the number of edge exchanges between two spanning trees of a graph. (See Bousquet et al. [BIK⁺23] for harder variants of this problem.)

2.2.2 Acyclic orientations of graphs and graphical zonotopes.

Given a simple, connected graph $G = ([n], E)$, consider the submodular function $f : 2^{[n]} \rightarrow \mathbb{N}$, where for a vertex subset $U \subseteq [n]$ the value of $f(U)$ is simply defined as the number of edges that have at least one endpoint in U . The polymatroid P_f defined by this choice of f is known as the *graphical zonotope* of G [Sta73, Gre77, GZ83]. The vertices of P_f are in one-to-one correspondence with the acyclic orientations of G , and its edges are in one-to-one correspondence with pairs of acyclic orientations that differ by a single edge flip. Given an acyclic orientation A of G , the entries of its corresponding vertex of P_f simply correspond to the in-degrees of the vertices in A .

In a way that is similar to the previous case, the flip distance problem for graphical zonotopes can be shown to be solvable easily: The flip distance between two acyclic orientations A and B of G is exactly equal to the number of edges that are oriented differently in A and B , and a sequence of $|A \Delta B|$ edge reversals that transforms A into B while keeping all intermediate orientations acyclic can be efficiently computed (see e.g. [FPS01]). We claim that in fact, a shortest path between two given vertices \mathbf{x}, \mathbf{y} of P_f can be found efficiently even when one is only given oracle access to the function f . To see this, note that given any pair $u, v \in [n]$, we can determine whether u and v are adjacent in G by computing the value of $f(u) + f(v) - f(\{u, v\})$, which is 1 if u and v are adjacent and 0 otherwise. Thus, using polynomially oracle calls to f we can determine G . Let A and B be the acyclic orientations of G whose in-degree sequences correspond to \mathbf{x} and \mathbf{y} . By repeatedly identifying a vertex of in-degree 0 and removing it, we can efficiently reconstruct the orientations A and B of G from \mathbf{x} and \mathbf{y} . Finally, knowing A, B and G we can efficiently find a shortest sequence of edge reversals that transforms A into B (see [FPS01]) and map this to a shortest path from \mathbf{x} to \mathbf{y} on the skeleton of P_f .

We refer the reader to Pilaud [Pil24], and Cardinal, Hoang, Merino, and Mütze [CHM⁺23] for recent developments on the combinatorics of acyclic orientations of graphs. One of the contributions of the present paper is the analysis of a hypergraph generalization of this problem.

2.2.3 Triangulations and associahedra.

When the submodular function is defined as

$$f(U) = |\{(i, j) \in [n]^2 : i < j \text{ and } U \cap \{i, i+1, \dots, j\} \neq \emptyset\}|, \quad (2.1)$$

then the polymatroids are Loday's associahedra [STT22, Pou14, Lod04, CSZ15, PSZ23]. Problem 2.3 in this special case is the problem of computing the flip distance between two triangulations of a convex polygon, where a flip consists of replacing the diagonal of a quadrilateral

formed by two adjacent triangles by its other diagonal. Equivalently, this amounts to computing the rotation distance between the dual binary trees. The more general problems of computing flip distances between triangulations of arbitrary point sets or polygons are known to be NP-hard [LP15, AMP15]. FPT algorithms have been designed [KSX17, CSJ09].

2.2.4 Elimination trees and graph associahedra.

An *elimination tree* of a graph G is obtained by selecting any vertex v of G as the root, and attaching elimination trees obtained by recursing on each of the connected components of $G - v$. A natural rotation operation connects some pairs of elimination trees, and the rotation graph on the elimination trees is the skeleton of the so-called *graph associahedron* of G .

The graph associahedron of G can be specified by a submodular function f defined as follows. Let V be the vertex set of G , and define $\mathcal{B}(G) \subset 2^V$ as the collection of nonempty subsets $S \subseteq V$ such that $G[S]$ (i.e., the induced subgraph of G with vertex set S) is connected. The collection $\mathcal{B}(G)$ is known as the *graphical building set* of G . The submodular function defining the associahedron of G is then defined as

$$f(U) = |\{S \in \mathcal{B}(G) : U \cap S \neq \emptyset\}|. \quad (2.2)$$

Graph associahedra were defined by Carr and Devadoss [CD06] and Postnikov [Pos09], and also considered by Postnikov, Rainer, and Williams [PRW08]. The structure of their skeleton was the subject of several recent works [MP15, CMM24, CPVP22]. The flip distance problem in this case is known to be NP-hard from the recent result of Ito et al. [IKK⁺23].

2.2.5 Partial permutations and stellohedra.

Stellohedra are associahedra of stars [MP15, CP16], and their vertices, the elimination trees of a star, can be shown to be one-to-one with *partial permutations* of the set of leaves [CMM24]. The rotations can be interpreted in that case as either adjacent transpositions or adding or removing an element in final position. The corresponding flip distance problem was recently shown to be solvable in polynomial time by Cardinal, Pournin, and Valencia-Pabon [CPVP24].

2.2.6 Rectangulations and quotientopes.

A *rectangulation* is a tessellation of the unit square by disjoint axis-aligned rectangles. Modulo natural equivalence classes, rectangulations can be shown to be in bijection with classes of pattern-avoiding permutations [MM23]. Similarly to triangulations, flip graphs on rectangulations can be defined and shown to be skeletons of polytopes [CP25, CSS18, Mee19]. These polytopes have been studied in [LR12, Rea12] and belong to the wider family of *quotientopes* [PS19, PPR23]. They are also polymatroids, and therefore the flip distance problems between rectangulations are also special cases of Problem 2.3. The complexity of these problems is open.

Note that graphical zonotopes, associahedra, graph associahedra, and stellohedra are all special cases of *hypergraphic polytopes*, that we consider extensively in Section 4.

3. NP-hardness of shortest paths on polymatroids defined by linear inequalities

We prove that Problem 2.1 is NP-hard even in the case where the polytope P is a polymatroid in the sense of Theorem 2.2. In fact, the polytope can be restricted to be a very special polymatroid.

Theorem 3.1. *Problem 2.1 is strongly NP-hard even when the input polytope is a polymatroid defined as the intersection of an axis-aligned box with a hyperplane.*

We now define precisely the family of polytopes we consider. Let $\mathbf{a}, \mathbf{b} \in \mathbb{N}^n$ such that $\sum_{i=1}^n a_i = \sum_{i=1}^n b_i = N$ for some $N \in \mathbb{N}$, and $a_i \neq b_i$ for all $i \in [n]$. In what follows, we let the vectors \mathbf{l} and \mathbf{u} be defined as $l_i := \min\{a_i, b_i\}$ and $u_i := \max\{a_i, b_i\}$, respectively. Note that $l_i < u_i$ for all $i \in [n]$. We consider the polytope

$$P := \left\{ \mathbf{x} \in \mathbb{R}^n : \mathbf{l} \leq \mathbf{x} \leq \mathbf{u} \text{ and } \sum_{i=1}^n x_i = N \right\}.$$

Problem 3.2 (Shortest path on a box). Given two vectors $\mathbf{a}, \mathbf{b} \in \mathbb{N}^n$ such that $\sum_{i=1}^n a_i = \sum_{i=1}^n b_i = N$ for some $N \in \mathbb{N}$, and P defined as above, determine the length of a shortest path between \mathbf{a} and \mathbf{b} on the skeleton of P .

Lemma 3.3. *Problem 3.2 is strongly NP-hard.*

We characterize the edges of P exactly and label them with pairs $(i, j) \in [n]^2$.

Lemma 3.4. *Two distinct vertices \mathbf{x} and \mathbf{y} of P are adjacent on the skeleton of P if and only if there exist two indices $i, j \in [n]$ such that:*

- $y_k = x_k \in \{l_k, u_k\}$ for all $k \in [n]$ with $k \neq i, j$,
- $y_i = x_i - \delta_{ij}$,
- $y_j = x_j + \delta_{ij}$,

where $\delta_{ij} = \min\{u_j - x_j, x_i - l_i\}$. In that case, we say that the oriented edge (\mathbf{x}, \mathbf{y}) is an (i, j) -flip.

Proof. For the first direction of the stated equivalence, suppose that \mathbf{x} and \mathbf{y} are adjacent on the skeleton of P . Then the segment $\text{conv}(\{\mathbf{x}, \mathbf{y}\})$ forms an edge of P (and thus, a face of P of dimension 1). Since P is the intersection of the axis-aligned box $Q := [a_1, b_1] \times \cdots \times [a_n, b_n]$ with the hyperplane H of equation $\sum_{i=1}^n x_i = N$, it follows that $\text{conv}(\{\mathbf{x}, \mathbf{y}\})$ is the intersection of a two-dimensional face of the box Q and the hyperplane H . Such a two-dimensional face is an axis-aligned rectangle

$$R = [l_i, u_i] \times [l_j, u_j]$$

lying in a plane orthogonal to all \mathbf{e}_k such that $k \neq i, j$. In this plane, the edge is simply the intersection of R with a line parallel to $\mathbf{e}_j - \mathbf{e}_i$. It is straightforward to check that the coordinates of its endpoints obey the equations above.

Vice-versa, suppose that two distinct vertices \mathbf{x}, \mathbf{y} of P satisfy the stated conditions. We claim that $\text{conv}(\{\mathbf{x}, \mathbf{y}\})$ is an edge of P . To see this, consider the polytope P' defined by the points $\mathbf{z} \in \mathbb{R}^n$ satisfying:

$$z_k \begin{cases} \geq l_k & \text{if } x_k = l_k \\ \leq u_k & \text{if } x_k = u_k \end{cases}$$

for all $k \in [n] \setminus \{i, j\}$, and $\sum_{i=1}^n z_i = N$. By definition, P' contains P . Furthermore, \mathbf{x} and \mathbf{y} lie on the one-dimensional intersection of the $n - 1$ hyperplanes defining P' . Therefore, they must also form an edge of P . \square

From Lemma 3.4, the edges of P are parallel to $\mathbf{e}_j - \mathbf{e}_i$ for a pair $i, j \in [n]$. Hence from Theorem 2.2, P is a polymatroid and Problem 3.2 can also be reduced to Problem 2.3. The submodular function $f : 2^{[n]} \rightarrow \mathbb{R}$ such that $P_f = P$ is

$$f(U) = \min \left\{ N - \sum_{i \in [n] \setminus U} l_i, \sum_{i \in U} u_i \right\}.$$

We say that an (i, j) -flip is a (U, V) -flip for some subsets $U \subseteq [n]$ and $V \subseteq [n]$ if $i \in U$ and $j \in V$. Let $S := \{i \in [n] : a_i = u_i\}$, and $T := \{i \in [n] : a_i = l_i\} = [n] \setminus S$. The set S is therefore the set of indices of components of \mathbf{a} that have to eventually decrease, while T denote those components that have to eventually increase.

We show that Problem 3.2 can be cast as a special case of the *fixed charge transportation* problem on the bipartite graph $S \times T$. This problem is a minimum cost maximum flow problem in which the cost of every edge is the sum of a variable cost proportional to the flow on the edge and a fixed charge for each edge with nonzero flow. It has a long history dating back to the pioneering work from Hirsch and Dantzig [HD68], and is known to be NP-hard. The special case where the cost is simply the number of edges with nonzero flow, and there is no variable cost, is known as the *pure constant fixed charge transportation* (PCFCT) problem. The PCFCT problem has been considered in particular by Hultberg and Cardoso [HC97] under the name *teacher assignment problem*, and by Schrenk, Finke, and Cung [SFC11]. In a sense, it is about finding the “most degenerate” solution to a set of flow and positivity constraints.

Problem 3.5 (Pure constant fixed charge transportation (PCFCT)). Given two sets S and T and positive integers $\{s_i\}_{i \in S}$ and $\{t_j\}_{j \in T}$ such that $\sum_{i \in S} s_i = \sum_{j \in T} t_j$, find the minimum possible number of nonzero values $|\{(i, j) \in S \times T : x_{ij} > 0\}|$ among reals $x_{ij} \geq 0$ for $(i, j) \in S \times T$ such that:

$$\begin{aligned} \sum_{j \in T} x_{ij} &= s_i, \forall i \in S \\ \sum_{i \in S} x_{ij} &= t_j, \forall j \in T. \end{aligned}$$

The PCFCT problem is known to be strongly NP-hard, see Hultberg and Cardoso [HC97]. Their proof is by reduction from the strongly NP-complete 3-partition problem. They also ob-




Figure 3.1: Illustration of Lemma 3.6. The top and bottom rows show the values of the s_i and t_j , respectively. There exists a same-sum partition of size $m = 3$, and a corresponding solution of value $|S| + |T| - m = 5 + 5 - 3 = 7$.

served that the problem can in fact be cast as a maximum partition problem. We define a *same-sum partition* of S and T as a collection of pairs (S_ℓ, T_ℓ) of nonempty subsets of S and T , such that $\{S_\ell\}$ is a partition of S , $\{T_\ell\}$ is a partition of T , and for each ℓ , we have $\sum_{i \in S_\ell} s_i = \sum_{j \in T_\ell} t_j$.

Lemma 3.6 (Hultberg and Cardoso [HC97]). *The optimal value of the PCFCT problem equals $|S| + |T| - m$, where m is the maximum size of a same-sum partition of S and T .*

The proof of Lemma 3.6 is simple. We give an example in Figure 3.1.

We now prove the following key lemma, establishing that the PCFCT problem can be reduced to Problem 3.2.

Lemma 3.7. *Any instance of the PCFCT problem can be converted (in polynomial time) into an equivalent instance (\mathbf{a}, \mathbf{b}) of Problem 3.2, so that the length of a shortest path is exactly equal to the optimal value of the PCFCT instance.*

Proof. Given an instance $(\{s_i\}_{i \in S}, \{t_j\}_{j \in T})$ of the PCFCT problem, we construct an equivalent instance (\mathbf{a}, \mathbf{b}) of Problem 3.2 as follows. First, we can assume without loss of generality that $S = [n_1]$ and $T = \{n_1 + 1, n_1 + 2, \dots, n_1 + n_2\}$, and let $n = n_1 + n_2$. We now let $\mathbf{a} = (s_1, s_2, \dots, s_{n_1}, 0, 0, \dots, 0)$, and $\mathbf{b} = (0, 0, \dots, 0, t_1, t_2, \dots, t_{n_2})$, and $u_i = s_i$ for each $i \in S$, $u_j = t_j$ for each $j \in T$, and finally $l_i = 0$ for all $i \in S \cup T$.

We first show that given an optimal solution to the PCFCT problem, we can find a path between \mathbf{a} and \mathbf{b} whose length is the value of the solution.

Recall from Lemma 3.6 that an optimal solution can be mapped to a same-sum partition $(S_\ell, T_\ell)_{\ell=1}^m$ of S and T of maximum size, such that the optimal value of the PCFCT equals $|S| + |T| - m$. To construct a path on the skeleton of P from \mathbf{a} to \mathbf{b} of length $|S| + |T| - m$, we consider the following simple algorithm.

The algorithm keeps track of two subsets X and Y of S and T , respectively, the current vertex \mathbf{x} of the polytope P , as well as a special “active” element v . Initially, the algorithm sets $X \leftarrow S, Y \leftarrow T, \mathbf{x} \leftarrow \mathbf{a}$ and picks $v \in S_1$ arbitrarily.

During the process, we maintain the following invariants.

1. \mathbf{x} is a vertex of P .
2. If ℓ is the index such that $v \in S_\ell \cup T_\ell$, then $X \subseteq S_\ell \cup \dots \cup S_m, Y \subseteq T_\ell \cup \dots \cup T_m$ and $v \in X \cap S_\ell$ or $v \in Y \cap T_\ell$.

3. $X = \{k \in S | x_k > 0\}$ and $Y = \{k \in T | x_k < t_k\}$.
4. $x_k = u_k = s_k$ for every $k \in X \setminus \{v\}$ and $x_k = l_k = 0$ for every $k \in Y \setminus \{v\}$.
5. For every $\ell = 1, \dots, m$, it holds that

$$\sum_{k \in X \cap S_\ell} x_k = \sum_{k \in Y \cap T_\ell} (t_k - x_k).$$

Note that these invariants hold after the initialization of (\mathbf{x}, v, X, Y) . Now, for $\ell = 1, \dots, m$, while $v \in S_\ell \cup T_\ell$, the algorithm updates (\mathbf{x}, v, X, Y) as follows.

It starts by picking a pair $(i, j) \in (X \cap S_\ell) \times (Y \cap T_\ell)$ such that $v \in \{i, j\}$. Note that this is always possible, since invariants (2), (3) and (5) together enforce that $X \cap S_\ell$ and $Y \cap T_\ell$ are nonempty.

The algorithm now performs an (i, j) -flip at the vertex \mathbf{x} of P (see the definition from Lemma 3.4). Concretely, we put $\delta \leftarrow \min\{x_i, t_j - x_j\}$ and update $x_i \leftarrow x_i - \delta, x_j \leftarrow x_j + \delta$. Since $x_k \in \{l_k, u_k\}$ for every $k \notin \{i, j\}$ by the invariants (3) and (4), Lemma 3.4 shows that this change of \mathbf{x} indeed corresponds to moving along an edge of the polytope P , in particular invariant (1) is preserved. Note that by definition of δ , we have $x_i = 0$ or $x_j = t_j$ after the flip. The algorithm now proceeds according to the following cases:

- If $x_i = 0$ and $x_j < t_j$ after the flip, then the algorithm updates $X \leftarrow X \setminus \{i\}$ and $v \leftarrow j$.
- Symmetrically, if $x_i > 0$ and $x_j = t_j$ after the flip, then the algorithm updates $Y \leftarrow Y \setminus \{j\}$ and $v \leftarrow i$.
- If $x_i = 0$ and $x_j = t_j$ after the flip, then we claim that we must have $X \cap S_\ell = \{i\}$ and $Y \cap T_\ell = \{j\}$. Indeed, if $x_i = 0$ and $x_j = t_j$ after the flip, then we must have had $x_i = t_j - x_j$ before the flip. Using invariant (5), this implies

$$\sum_{k \in (X \cap S_\ell) \setminus \{i\}} x_k = \sum_{k \in (Y \cap T_\ell) \setminus \{j\}} (t_k - x_k).$$

However, by invariant (4) we have $x_k = s_k$ for every $k \in (X \cap S_\ell) \setminus \{i\}$ and $x_k = 0$ for every $k \in (Y \cap T_\ell) \setminus \{j\}$. This implies that refining the same-sum partition of (S, T) by replacing the pair (S_ℓ, T_ℓ) with the two pairs $((X \cap S_\ell) \setminus \{i\}, (Y \cap T_\ell) \setminus \{j\})$ and $((S_\ell \setminus X) \cup \{i\}, (T_\ell \setminus Y) \cup \{j\})$ would yield a larger same-sum partition, a contradiction — unless $(X \cap S_\ell) \setminus \{i\}$ and $(Y \cap T_\ell) \setminus \{j\}$ are empty. This proves that $X \cap S_\ell = \{i\}$ and $Y \cap T_\ell = \{j\}$, as claimed above.

Finally, the algorithm proceeds to update $X \leftarrow X \setminus \{i\}, Y \leftarrow Y \setminus \{j\}$. If $\ell < m$, then it updates v to be any vertex of $S_{\ell+1}$, and the for loop now continues with the new index $\ell + 1$ instead of ℓ . If $\ell = m$, then the algorithm stops and returns the sequence of all vertices of P visited during the process.

Each of the three types of updates above preserve our invariants (1)–(5). It is clear from the above update rules that in every step of the algorithm, whenever we take a step along an edge of P to a new vertex, the value of $|X| + |Y|$ decreases by one, and in fact sometimes this value decreases by two, namely exactly when we depart the inner loop for a fixed $\ell = 1, \dots, m$ and make the intersection of X and Y with S_ℓ and T_ℓ empty.

The above implies that the algorithm finishes after a total of $|S| + |T| - m$ steps, and returns a sequence of $|S| + |T| - m + 1$ vertices of P , any two consecutive of which are adjacent on the skeleton. Note that at each point of the algorithm, by invariant (3) it holds that $|X| + |Y|$ is the Hamming distance between the current vertex \mathbf{x} of P and the target vertex \mathbf{b} . In particular, the last vertex of this sequence is equal to \mathbf{b} . It follows that the sequence of vertices of P that we output forms a path from \mathbf{a} to \mathbf{b} on the skeleton of P that has length $|S| + |T| - m$, as desired. This concludes the first part of the proof of the lemma.

We now turn to the second part of the proof, and show that given a path between \mathbf{a} and \mathbf{b} on the skeleton of P , we can find a solution of the PCFCT instance whose value is at most the length of the path. Let us consider all flips on this path, and construct the undirected graph G on $S \cup T$ in which two vertices i, j are adjacent whenever there is an (i, j) -flip in the path. Note that some flips might not be (S, T) -flips, so G is not necessarily bipartite. Each connected component $S_\ell \cup T_\ell$ of G must be such that $\sum_{i \in S_\ell} s_i = \sum_{j \in T_\ell} t_j$. Hence the connected components induce a same-sum partition of S and T , of some size m . So from Lemma 3.6, we can construct a solution of the PCFCT problem of value $|S| + |T| - m = n - m$. On the other hand, the number of flips in each connected component $S_\ell \cup T_\ell$ of G is at least $|S_\ell| + |T_\ell| - 1$, hence the total number of flips on the path is at least $\sum_{\ell=1}^m (|S_\ell| + |T_\ell| - 1) = n - m$. \square

Combining Lemma 3.7 with the fact that the PCFCT problem is strongly NP-hard proves Lemma 3.3, and therefore Theorem 3.1.

Observe that the shortest paths in the reduction above only involve (S, T) -flips, and therefore are monotone with respect to any objective function $\mathbf{w} \in \mathbb{R}^n$ such that $w_i < w_j$ for all pairs $(i, j) \in S \times T$. Hence it is also NP-hard to compute the length of a shortest *monotone* path with respect to such an objective function in Lemma 3.3. The result therefore rules out, unless $P = NP$, the existence of a polynomial-time computable simplex pivoting rule that would minimize the number of nondegenerate pivoting steps to an optimal solution of a linear program, even when it defines a polymatroid. Note that while it was previously shown (see [DLKS22]) that finding shortest monotone paths to *some* optimum of an LP is NP-hard, in our hardness reduction the optimal solution to the linear program (say, \mathbf{b}) can be known in advance, and yet Lemma 3.3 shows that finding a shortest path with this extra piece of information remains hard.

4. Flip distances between acyclic orientations of hypergraphs

We introduce hypergraphic polytopes, the associated flip distance problem, and our two results on the complexity of this problem: We prove that approximating the optimal value up to an error of ε is NP-hard for some explicit constant $\varepsilon > 0$, and on the other hand present an efficient algorithm that produces a connecting path whose length is within a factor of at most the so-called *maximum codegree* of the hypergraph from that of a shortest path.

4.1. Hypergraphic polytopes

Every hypergraph $H = (V, \mathcal{E})$, with edge-set $\mathcal{E} \subseteq 2^V \setminus \{\emptyset\}$, induces a corresponding submodular function f_H according to

$$f_H(U) := |\{e \in \mathcal{E} : U \cap e \neq \emptyset\}| \tag{4.1}$$

for every $U \subseteq V$. The base polytope induced by the polymatroid associated with f_H is then called the *hypergraphic polytope* of the hypergraph H , and denoted by P_H . Hypergraphic polytopes and the corresponding polymatroids have been studied in various guises in the past 50 years as covering hypermatroids [Hel74], Boolean polymatroids [VW93], or relational submodular functions [AA23]. Hypergraphic polytopes form an important family of generalized permutohedra [PRW08, Pos09, AA23, Reh22].

Interestingly, this definition encompasses most of the special cases we described in Section 2.2. When H is a graph, the polymatroids are the graphical zonotopes, whose vertices are in bijection with acyclic orientations of the graph. When H is defined as

$$H = \{\{i, i + 1, \dots, j\} : 1 \leq i < j \leq n\},$$

then one can check that Eqn. 4.1 matches Eqn. 2.1, hence the hypergraphic polytope is nothing but the associahedron. When the hyperedges correspond with building sets of some graph G , then the definition matches that of Eqn. 2.2, and the polymatroid is a graph associahedron. We refer to [PRW08, Pos09, BBM19, Reh22, CHM⁺23] for detailed treatments of the connections between the polytopes P_H and combinatorial structures.

4.2. Acyclic orientations of hypergraphs

It is known that the vertices of the hypergraphic polytope P_H are in one-to-one correspondence with so-called *acyclic orientations* of the hypergraph H [BBM19, Reh22, CHM⁺23], which we define now.

An *orientation* of a hypergraph $H = (V, \mathcal{E})$ is a mapping $h : \mathcal{E} \rightarrow V$ such that $h(e) \in e$ for every $e \in \mathcal{E}$. The element $h(e)$ is also called the *head* of an edge e . To every orientation h of H we can associate a directed graph D_h that has vertex set V and an arc from a vertex u to a vertex v if and only if there exists some $e \in \mathcal{E}$ such that $h(e) = v$ and $u \in e \setminus \{v\}$. Intuitively, D_h is the union, over all $e \in \mathcal{E}$, of the inwards-oriented stars centered at the head of e and whose leaves are all other elements of e . With this notion at hand, we call an orientation h of H *acyclic* if D_h is an acyclic digraph.

Adjacency of vertices on the skeleton of P_H has also been characterized, see Theorem 4.1 below. In order to state it, we need one more definition, which is that of a *flip*. Given two distinct acyclic orientations h_1, h_2 of a hypergraph $H = (V, \mathcal{E})$, we say that h_2 is obtained from h_1 by a *flip* at an ordered pair $(u, v) \in V^2$ of distinct vertices, if h_2 is obtained from h_1 by redefining the head of every hyperedge $e \in \mathcal{E}$ with $e \supseteq \{u, v\}$ that satisfies $h_1(e) = u$ as $h_2(e) = v$, while keeping all other head assignments the same as in h_1 . We note that the above definition of a flip is symmetric, in the following sense: If h_2 can be obtained from h_1 by flipping at (u, v) , then h_1 is obtained from h_2 by flipping at (v, u) . The reason for this is that since h_1 and h_2 are acyclic




Figure 4.1: An example of flip, showing all hyperedges $e \in \mathcal{E}$ such that $e \supseteq \{u, v\}$ and satisfying $h_1(e) = u$, and arrows pointing to their heads $h_1(e)$ (left) and $h_2(e)$ (right). The bottom figure shows the corresponding parts of the digraphs D_{h_1} and D_{h_2} .

orientations of H , it is not possible that there exist distinct edges $e, f \in \mathcal{E}$ with $\{u, v\} \subseteq e, f$ such that $h_i(e) = u$ and $h_i(f) = v$ for some $i \in \{1, 2\}$, for otherwise u, v would form a directed cycle of length two in D_{h_i} . In consequence, when h_2 can be obtained from h_1 by flipping at (u, v) , then $h_1(e) = u$ for every $e \in \mathcal{E}$ with $e \supseteq \{u, v\}$ and $h_1(e) \in \{u, v\}$, and $h_2(e) = v$ for every $e \in \mathcal{E}$ for every $e \in \mathcal{E}$ with $e \supseteq \{u, v\}$ and $h_2(e) \in \{u, v\}$. Therefore, h_1 can be obtained from h_2 by a flip at (v, u) . Also, note that as a consequence of the previous discussion, when h_2 is obtained from h_1 by flipping at (u, v) , then a flip at (v, u) for h_1 is impossible. Thus, in the following we can say without ambiguity that two acyclic orientations h_1 and h_2 are obtained from each other by flipping the unordered pair $\{u, v\}$ of distinct vertices if h_2 can be obtained from h_1 by flipping (u, v) or (v, u) . An illustration is given in Figure 4.1.

The following statement can be deduced as a special case of Theorem 2.18 in [BBM19].

Theorem 4.1. *Let $H = (V, \mathcal{E})$ be a hypergraph. Two vertices x, y of the polytope P_H corresponding to acyclic orientations h_1, h_2 of H are adjacent on the skeleton of P_H if and only if there exists a pair $\{u, v\} \subseteq V$ of distinct vertices such that h_1 and h_2 can be obtained from each other by flipping $\{u, v\}$.*

4.3. Flip distances between acyclic orientations of hypergraphs

Theorem 4.1 implies that the computation of shortest paths on the skeleton of a hypergraphic polytope P_H boils down to a flip distance problem between acyclic orientations of H , as follows:

Problem 4.2. Given a hypergraph $H = (V, \mathcal{E})$ and two acyclic orientations h_1, h_2 of H as input, find the minimum length of a sequence of vertex pairs $\{u_i, v_i\}, i = 1, \dots, \ell$ such that h_2 can be obtained from h_1 by flipping, in order, $\{u_1, v_1\}, \dots, \{u_\ell, v_\ell\}$.

Problem 4.2 generalizes several of the previous problems mentioned in Section 2. It clearly includes the special case of flip distances between acyclic orientations of graphs, when the hypergraph H is a graph. The flip distance between triangulations or binary search trees is the

special case of Problem 4.2 in which the underlying hypergraph is the set of intervals in $[n]$, as mentioned above.

In the following, we will prove several results regarding the computational complexity of Problem 4.2. On the negative side, we show that approximating the length of a shortest flip sequence between two acyclic orientations of a general hypergraph H is APX-hard, even restricted to hypergraphs H of bounded degree and bounded hyperedge size.

Theorem 4.3. *There exists an absolute constant $\Delta \in \mathbb{N}$ such that Problem 4.2 is APX-hard even when the input hypergraph $H = (V, \mathcal{E})$ is known to have maximum degree at most Δ and be such that $|e| \leq 3$ for every $e \in \mathcal{E}$.*

Since the input hypergraph H allows to compute the submodular function associated with H efficiently, Theorem 4.3 implies the following.

Corollary 4.4. *There exists a constant $\varepsilon > 0$ such that Problem 2.3 cannot be approximated to within a factor $(1 + \varepsilon)$ in polynomial time, unless $P = NP$.*

The proof of Theorem 4.3 is given in Section 5, and relies on a reduction from the vertex cover problem. Using, for instance, a result from Chlebík and Chlebíková [CC06] on the inapproximability of minimum vertex cover on graphs of degree at most 4, we can set $\varepsilon = 1/1300$ in the statement of Corollary 4.4.

On the positive side, we prove that shortest paths in the skeleton of P_H can be approximated efficiently up to a multiplicative factor that is bounded by the *maximum codegree* in H . Given a hypergraph $H = (V, \mathcal{E})$ and a pair $\{u, v\}$ of distinct vertices, the *codegree* of $\{u, v\}$ is defined as

$$\text{codeg}(\{u, v\}) := |\{e \in \mathcal{E} \mid \{u, v\} \subseteq e\}|,$$

and the maximum codegree of H is defined as

$$\Delta_2(H) := \max\{\text{codeg}(\{u, v\}) \mid u \neq v \in V\}.$$

Theorem 4.5. *There exists a polynomial-time algorithm that, given as input a hypergraph H and two acyclic orientations h_1, h_2 of H , outputs a flip sequence $\{u_1, v_1\}, \dots, \{u_\ell, v_\ell\}$ from h_1 to h_2 such that $\ell \leq \Delta_2(H) \cdot \text{dist}(h_1, h_2)$, where we denote by $\text{dist}(h_1, h_2)$ the length of a shortest flip sequence from h_1 to h_2 .*

In particular, for *linear hypergraphs*, these are the hypergraphs $H = (V, \mathcal{E})$ satisfying $|e_1 \cap e_2| \leq 1$ for every distinct $e_1, e_2 \in \mathcal{E}$, we obtain an efficient exact algorithm for computing a shortest flip sequence between two acyclic orientations.

Corollary 4.6. *There exists a polynomial-time algorithm that, given as input a linear hypergraph H and two acyclic orientations h_1, h_2 of H , outputs a shortest flip sequence from h_1 to h_2 .*

The proof of Theorem 4.5 is given in Section 6.




Figure 5.1: Twelve flips are sufficient to transform h_1 (left) into h_2 (right). The transitive reductions of the orientations are shown at each step, each of which is a broom, composed of a handle – a chain above vertex 0 – and a collection of independent vertices below 0.

5. Proof of Theorem 4.3

Our proof describes a reduction from the minimum vertex cover problem in graphs, which is well-known to be APX-hard. Recall that a vertex cover of a graph $G = (V, E)$ is a subset $X \subseteq V$ such that every edge $e \in E$ has at least one endpoint in X .

We use the following classical hardness result due to Håstad [Hs01]. While Håstad only stated hardness of approximation on general graphs, his result in fact applies to bounded-degree graphs, see for instance [DS05], page 449 for an explanation why we can restrict Theorem 5.1 to the bounded-degree case. See also Chlebík and Chlebíková [CC06] for inapproximability result for graphs of maximum degree 3 or 4.

Theorem 5.1 ([Hs01]). *There exist absolute constants $\Delta > 0$ and $\delta > 0$ such that the following holds. Unless $P=NP$, there exists no polynomial-time algorithm that achieves an approximation ratio of less than $1 + \delta$ for the following minimization problem.*

Input: A graph $G = (V, E)$ of maximum degree at most Δ .

Output: The size $|X|$ of a smallest vertex cover $X \subseteq V$ of G .

Before proceeding to the proof of Theorem 4.3, we describe a hypergraph that will be used as a gadget in our reduction. This gadget exploits the fact that stellohedra do not have the *non-leaving face property* [MP15, CP16] (a polytope P is said to have the non-leaving face property if every shortest path between two vertices x, y in the skeleton of P is fully contained in the minimal face of P containing x and y). The way this gadget is used in the reduction is inspired by previous reductions for the flip distance problem in geometric triangulations [AMP15, LP15].

The *star gadget* is the hypergraph $S = (V, \mathcal{E})$, where $V = \{0, 1, \dots, 6\}$ and $\mathcal{E} = \{\{0, i\} : i \in [6]\} \cup \{\{0, i, j\} : i, j \in [6], i \neq j\}$. Hence it is composed of a star graph, together with all triples of vertices containing exactly two leaves of the star. The vertex 0 will be called the *center* of the star gadget, and the vertices in $[6]$ will be referred to as the *leaves*.

We define two orientations h_1 and h_2 of the star gadget. In both orientations, all hyperedges of the form $e = \{0, i\}$ are such that $h_1(e) = h_2(e) = i$. On the other hand, for hyperedges of the form $e = \{0, i, j\}$, we let $h_1(e) = \min\{i, j\}$, and $h_2(e) = \max\{i, j\}$.




Figure 5.2: Illustration of the proof of Theorem 4.3. Given an instance G of the vertex cover problem on the left, we construct the hypergraph H on the right. We omit the depiction of the hyperedges of size three within each of the star gadgets.

Lemma 5.2. *The flip distance between the two orientations h_1 and h_2 of the star gadget is 12. However, any sequence of flips between h_1 and h_2 that does not involve the vertex 0 has length at least 15.*

Proof. We observe that the acyclic orientations of the star gadget S are one-to-one with brooms, as pictured in Figure 5.1. Given an acyclic orientation h , let $W \subseteq [6]$ be the set of vertices i whose edge $e = \{0, i\}$ is such that $h(e) = i$. Consider the transitive reduction D of the digraph D_h . Then D consists of two parts: A directed path, that we call the *handle*, originating from 0 and involving vertices of W , and a directed star of center 0 with leaves consisting in the vertices not in W , all pointing to 0. Indeed, the hyperedges of the form $e = \{0, i, j\}$ induce a total order on W , while the other vertices are incomparable in D .

The orientations h_1 and h_2 are both such that the set W contains all 6 vertices, but the order in the handle is reversed. There is a sequence of 12 flips that transforms h_1 into h_2 . It consists of iteratively removing each of the 6 vertices from the handle of the broom, by performing the 6 flips $\{0, 6\}, \{0, 5\}, \dots, \{0, 1\}$, then the same flips again, in the same order: $\{0, 6\}, \{0, 5\}, \dots, \{0, 1\}$ (see Figure 5.1).

On the other hand, if we are not allowed to flip any pair of the form $\{0, i\}$, we can only change the order of two adjacent vertices in the handle, by flipping a pair $\{i, j\}$ with $i, j > 0$. This amounts to transforming the identity permutation on 6 elements into the reverse identity permutation by adjacent transpositions. Since all pairs are reversed, this costs $\binom{6}{2} = 15$ flips. \square

Proof of Theorem 4.3. Consider an instance $G = (V, E)$ of the minimum vertex cover problem, with $|V| = n$ and $|E| = m$ and maximum degree bounded by an absolute constant Δ . We construct a hypergraph $H = (V', \mathcal{E})$ as follows. First, we choose and fix some linear ordering (V, \prec) of the vertex set of G . Next, for each vertex $v \in V$, we construct two vertices v^- and v^+ in V' , and connect them by an edge $\{v^-, v^+\}$ in \mathcal{E} . Then, for each edge $e = \{u, v\} \in E$ such that $u \prec v$, we include the following vertices and hyperedges in the hypergraph H :




Figure 5.3: Orientations of the edges of H in h_1 and h_2 . The only difference between h_1 and h_2 is the orientation of the hyperedges of size three of the star gadgets, not shown here.

1. A copy S_e of the star gadget,
2. an edge connecting the center of the star S_e to u^- ,
3. the edge $\{u^+, v^-\}$,
4. 6 edges connecting v^+ to each of the leaves of S_e .

The hypergraph H has $|V'| = 2n + 7m$ vertices and $|\mathcal{E}| = n + 29m$ hyperedges. An illustration is given in Figure 5.2.

We now consider two orientations h_1 and h_2 of H . In both orientations, for each vertex $v \in V$, the edge $\{v^-, v^+\}$ is oriented from v^- to v^+ , and for each edge $e = \{u, v\} \in E$, the edges of H of type 2 above are oriented from the center of S_e to u^- , the edges of type 3 are oriented from u^+ to v^- , and the edges of type 4 are oriented towards the leaves of S_e . Also, for each star gadget, all hyperedges of size two are oriented towards the leaves. These orientations are illustrated on Figure 5.3. The only difference between h_1 and h_2 are the orientations of the hyperedges of size three of the star gadgets. For each such hyperedge of the form $e = \{0, i, j\}$, we set $h_1(e) = \min\{i, j\}$, and $h_2(e) = \max\{i, j\}$.

Claim 5.3. *The orientations h_1 and h_2 are acyclic.*

Proof. To certify that the digraphs D_{h_1} and D_{h_2} corresponding to the orientations h_1, h_2 of H are indeed acyclic, we will describe two linear orderings of V' that are *topological orderings*, i.e., all directed edges in D_{h_1} (or, D_{h_2} , respectively) have the property that their starting point precedes their endpoint in the ordering.

In both orderings, for D_{h_1} and D_{h_2} , we start by listing all the 0-vertices of all the star gadgets S_e for all $e \in E$. Note that these vertices form an independent set in both D_{h_1} and D_{h_2} . Then, we start by listing the pairs of vertices v^-, v^+ (in this order) for all $v \in V$, where we order the pairs according to the linear order \prec that we initially fixed on V . The latter makes sure that all the edges directed from u^+ to v^- for some edge $\{u, v\} \in E$ with $u \prec v$ are indeed going forward in the linear ordering we define on V' . At the end of the ordering, we list all the vertices

that are leaves of some star $S_e, e \in E$. While it does not matter in which order we put the leaves of different stars, it is crucial that within each of the stars, we order the leaves in decreasing order (i.e. $6, 5, \dots, 1$) for the orientation D_{h_1} and in the opposite order (i.e. $1, 2, \dots, 6$) for the orientation D_{h_2} . It is not hard to check by going through all directed edges of D_{h_1} and D_{h_2} that both of these orderings for D_{h_1} and D_{h_2} are indeed topological, concluding the proof. ■

This finishes the description of our instance of the flip distance problem, composed of H together with its two acyclic orientations h_1 and h_2 .

Let us denote by x the size of a minimum vertex cover of G . We now prove the following:

Claim 5.4. *The length of a shortest flip sequence between h_1 and h_2 is equal to*

$$2x + 12m.$$

Proof. We first show that there exists a flip sequence transforming h_1 into h_2 , of length at most $2x + 12m$. First observe that all edges of H of the form $\{v^-, v^+\}$ can be flipped independently, without affecting acyclicity (indeed, in the topological ordering of V' for D_{h_1} described in the proof of the previous claim, one may just flip the order in which v^- and v^+ occur for every $v \in X$, and the arising ordering is still topological). Now pick a minimum vertex cover $X \subseteq V$ of G , and flip all edges of the form $\{v^+, v^-\}, v \in X$. This needs x flips.

We now observe that after these flips, no directed path in the associated digraph leaves and goes back to a star gadget. Indeed, consider any star gadget S_e corresponding to an edge $e = uw$ of G with $u \prec w$. Then $u \in X$ or $w \in X$. This means that in the orientation of H obtained after flipping all pairs $\{v^+, v^-\}, v \in X$ we have that u^- is a sink (if $u \in X$) or w^+ is a source (if $w \in X$). In the first case, every directed path leaving S_e at its center must end already at u^- and thus cannot return to S_e , and in the second case, every directed path entering S_e at one of its leaves must begin already at w^+ and thus could not have started in S_e .

Hence all star gadgets can now be flipped independently from the rest of the hypergraph, without creating new directed cycles in the digraphs corresponding to the orientations. We next flip all star gadgets using exactly 12 flips per gadget, as described in Lemma 5.2. Finally, we can flip back all edges $\{v^-, v^+\}$ for $v \in X$. By the same arguments as above, these flips do not affect acyclicity. The total length of the described flip sequence is $2x + 12m$.

We now show that any flip sequence transforming h_1 into h_2 must have length at least $2x + 12m$. Consider such a flip sequence of length ℓ , and let R be the set of edges ab of size two of H , excluding the edges of the star gadgets, such that the pair $\{a, b\}$ occurs in the flip sequence from h_1 and h_2 . Since the orientations h_1 and h_2 agree on these edges, and since they are not included in any larger hyperedges of H , the flip sequence must flip each such pair at least twice. So in total the flip sequence involves at least $2|R|$ such flips. We let $r = |R|$.

We now partition the set E of edges of G into two subsets. For each edge $e = \{u, v\} \in E$ with $u \prec v$, consider the (six) paths from the center of S_e to a leaf of S_e , and going through the vertices $u^-, u^+, v^-,$ and v^+ . We now denote by $E_1 \subseteq E$ the set of edges $e \in E$ such that no edge of any of these six paths is included in R , and let $E_2 = E \setminus E_1$.

From Lemma 5.2, each star gadget S_e requires 15 flips if $e \in E_1$, since in that case any flip involving the center of the star would create a directed cycle with one of the paths associated

with e , and thus the edges of size 2 in S_e can never be flipped in this case. On the other hand, if $e \in E_2$, we know that the star gadget requires at least 12 flips.

Therefore the total length of the flip sequence satisfies

$$\ell \geq 15|E_1| + 12|E_2| + 2r.$$

We now observe that

$$x \leq r + |E_1|. \quad (5.1)$$

Indeed, define a set $W \subseteq V$ of vertices as follows: We put $w \in W$ if and only if one of the edges that connect a center of a star gadget to w^- is in R , or if w^+ is incident to an edge in R . It is straightforward to check that for every edge $e \in E_2$, at least one of its endpoints is in W . Pause to note that $|W| \leq |R| = r$. Thus, adding to W for each edge $e \in E_1$ one of its endpoints, we obtain a vertex cover of G , which is of size at most $r + |E_1|$.

It remains to bound ℓ as follows:

$$\begin{aligned} \ell &\geq 15|E_1| + 12|E_2| + 2r \\ &= 15|E_1| + 12(m - |E_1|) + 2r \\ &\geq 2|E_1| + 12m + 2r \\ &\geq 2(x - r) + 12m + 2r \text{ (from Eq. (5.1))} \\ &= 2x + 12m. \end{aligned}$$

This concludes the proof of the claim. ■

To finish the proof, we need to show that if there exists a $(1 + \varepsilon)$ -approximation of the flip distance between h_1 and h_2 , then we can obtain a $(1 + \delta)$ -approximation of the size of a minimum vertex cover of G , where $\delta = \delta(\varepsilon)$ tends to 0 with ε . Let d be a $(1 + \varepsilon)$ -approximation of the flip distance. Note that by the definition of a vertex cover, it holds that $m = |E| \leq \Delta x$. We have

$$2x + 12m \leq d \leq (1 + \varepsilon)(2x + 12m).$$

Hence

$$\begin{aligned} x &\leq \frac{1}{2}(d - 12m) \\ &\leq \frac{1}{2}((1 + \varepsilon)(2x + 12m) - 12m) \\ &= (1 + \varepsilon)x + 6\varepsilon m \\ &\leq (1 + \varepsilon + 6\varepsilon\Delta)x, \end{aligned}$$

so $\frac{d-12m}{2}$ is a $(1 + \delta)$ -approximation for x where $\delta(\varepsilon) = (1 + 6\Delta)\varepsilon = O(\varepsilon)$. □

For a specific value of ε in Corollary 4.4, one can for instance use the result of Chlebík and Chlebíková [CC06] on the inapproximability of vertex cover with $\Delta = 4$ and $\delta = 1/52$, yielding $\varepsilon = 1/1300$.




Figure 6.1: Illustration of the first claim in the proof of Lemma 6.1. The figure on the right illustrates the case where $h_2(e) = h_2(e')$ and a is the first arc of Q .

6. Proof of Theorem 4.5

We now prove Theorem 4.5. At the heart of our proof is the following key lemma, which we show first.

Lemma 6.1. *Let $H = (V, \mathcal{E})$ be a hypergraph and let h_1, h_2 be distinct acyclic orientations of H . Then there exists $e \in \mathcal{E}$ such that $h_1(e) \neq h_2(e)$ and $(h_1(e), h_2(e))$ is flippable in h_1 .*

Proof. Let $\mathcal{E}' := \{e \in \mathcal{E} \mid h_1(e) \neq h_2(e)\}$. Towards a contradiction, suppose that for every edge $e \in \mathcal{E}'$, the pair $(h_1(e), h_2(e))$ is not flippable in h_1 . By definition, this means that for every $e \in \mathcal{E}'$ there exists a directed cycle in the digraph D_{h^e} , where $h^e : \mathcal{E} \rightarrow V$ denotes the orientation of H defined by

$$h^e(f) := \begin{cases} h_2(e), & \text{if } h_1(f) = h_1(e) \text{ and } \{h_1(e), h_2(e)\} \subseteq f \\ h_1(f), & \text{otherwise} \end{cases}$$

for every $f \in \mathcal{E}$. In the following, we denote by $D \subseteq D_{h_1}$ the *transitive reduction* of D_{h_1} (this is the subdigraph of D_{h_1} containing all those arcs that do not have an edge-disjoint parallel directed path of length greater than one). Our first claim below is illustrated on Figure 6.1.

Claim 6.2. *For every $e \in \mathcal{E}'$, there exists some $e' \in \mathcal{E}'$ and a directed path Q in D starting at $h_2(e)$ and ending at $h_1(e)$, such that the following statements hold:*

1. Q has length at least two.
2. There exists an arc $a = (x, y)$ on Q , such that $x, y \in e'$ and $h_1(e') = y$.
3. The path W composed of the prefix of Q from $h_2(e)$ to x followed by the arc $(x, h_2(e'))$ (if $x \neq h_2(e')$) is a directed path in D_{h_2} . It is of length at least one unless $h_2(e) = h_2(e')$ and a is the first arc of Q .

Proof. Let C be a shortest directed cycle in D_{h^e} . Then C must use at least one arc of the form $(z, h_2(e))$ for some vertex $z \in f$ and $f \in \mathcal{E}$ with $\{h_1(e), h_2(e)\} \subseteq f$, $h_1(f) = h_1(e)$

and $h^e(f) = h_2(e)$. Also, note that if $z = h_1(e)$, then C must have length at least three, since $(h_2(e), h_1(e))$ cannot be an arc of D_{h^e} . Observe that the directed path $P := C - (z, h_2(e))$ in D_{h^e} also forms a directed path from $h_2(e)$ to z in the digraph D_{h_1} . Now consider the path P' in D_{h_1} , defined as $P' := P$ if $z = h_1(e)$, and as $P' := P + (z, h_1(e))$ if $z \neq h_1(e)$. Then P' is a directed path in D_{h_1} from $h_2(e)$ to $h_1(e)$. Since C has length at least three if $z = h_1(e)$, P' must have length at least two.

This proves that there is a directed path of length at least two in D_{h_1} from $h_2(e)$ to $h_1(e)$. Let us now define Q as a longest directed path in D_{h_1} from $h_2(e)$ to $h_1(e)$. Then Q also has length at least two, and the first item of the claim holds. Furthermore, the maximality of Q implies that Q is fully contained in the transitive reduction D of D_{h_1} .

Let $h_2(e) = q_0, q_1, \dots, q_\ell = h_1(e)$ be the sequence of vertices of Q in order. Then by definition of D_{h_1} , there exist edges $f_1, \dots, f_\ell \in \mathcal{E}$ such that $q_{i-1} \in f_i$ and $h_1(f_i) = q_i$, for every $i = 1, \dots, \ell$. Suppose first that $h_1(f_i) = h_2(f_i)$ for every $i = 1, \dots, \ell$. Then $(q_{i-1}, q_i) \in A(D_{h_2})$ for every $i = 1, \dots, \ell$, and clearly we also have $(h_1(e), h_2(e)) = (q_\ell, q_0) \in A(D_{h_2})$. This is a contradiction to h_2 being an acyclic orientation of H . Consequently, at least one of f_1, \dots, f_ℓ must be contained in \mathcal{E}' . Let j be the smallest index such that $f_j \in \mathcal{E}'$. We now define $e' := f_j$ and $a = (x, y) := (q_{j-1}, q_j)$. It follows directly from these definitions that the second item of our claim holds.

Thus, only the third item remains to be verified, which states that W , defined as the prefix of Q up to x , together with the arc $(x, h_2(e'))$ if $x \neq h_2(e')$, is a directed path in D_{h_2} . To see this, note that by minimality of the choice of j , we have that $h_2(f_i) = h_1(f_i) = q_i$ and thus $(q_{i-1}, q_i) \in A(D_{h_2})$ for every $0 \leq i \leq j-1$. Furthermore, since $q_{j-1} \in e' = f_j$, we have $(q_{j-1}, h_2(e')) = (q_{j-1}, h_2(f_j)) \in A(D_{h_2})$ or $q_{j-1} = h_2(e')$. Thus the sequence W is a directed walk from $h_2(e) = q_0$ to $h_2(e')$ in D_{h_2} . It is easy to see that W is of positive length unless $h_2(e) = h_2(e')$ and $j = 1$ (hence a is the first arc on Q). This concludes the proof of the claim. \blacksquare

We next define an infinite sequence $(e_i)_{i=1}^\infty$ of hyperedges $e_i \in \mathcal{E}'$ of H , as follows. First we pick an element $e_1 \in \mathcal{E}'$ arbitrarily (note that $\mathcal{E}' \neq \emptyset$ as $h_1 \neq h_2$). Then, for every $i \geq 1$, having defined e_i , we define $e_{i+1} \in \mathcal{E}'$ as the hyperedge e' given by the previous claim, applied with e_i in place of e . Concretely, we pick e_{i+1} as an element of \mathcal{E}' such that there exists

- a directed path Q_i of length at least two in D from $h_2(e_i)$ to $h_1(e_i)$,
- an arc $a_i = (x_i, y_i)$ on Q_i such that $x_i \in e_{i+1}$ and $h_1(e_{i+1}) = y_i$,
- and a directed walk W_i in D_{h_2} from $h_2(e_i)$ to $h_2(e_{i+1})$ which is of length at least 1, unless $h_2(e_i) = h_2(e_{i+1})$ and a_i is the first arc of Q_i .

Claim 6.3. *There exists $i_0 \in \mathbb{N}$ such that for every $i \geq i_0$ we have that $h_2(e_i) = h_2(e_{i_0})$ and a_i is the first arc of Q_i .*

Proof. Towards a contradiction, suppose there exist infinitely many distinct indices $i_1 < i_2 < i_3 < \dots$ such that for every $k \geq 1$, we either have $h_2(e_{i_k}) \neq h_2(e_{i_{k+1}})$ or a_{i_k} is not the first arc on Q_{i_k} . By the third item above, this means that W_{i_k} is of length at least 1 for

every $k \geq 1$. Since H is finite, there must exist distinct values $r < s$ such that $h_2(e_{i_r}) = h_2(e_{i_s})$. But then the concatenation of the directed walks $W_{i_r}, W_{i_r+1}, \dots, W_{i_s-1}$ forms a closed directed walk in D_{h_2} which is of length at least 1. This is a contradiction to the fact that D_{h_2} is acyclic, and we thereby conclude the proof of our claim. ■

In the following, let us denote $v := h_2(e_{i_0})$. By the previous claim, we have that $h_2(e_i) = v$ and $a_i = (x_i, y_i)$ is the first arc on Q_i for every $i \geq i_0$. In particular, $(h_2(e_{i_0+1}), h_1(e_{i_0+1})) = (v, y_{i_0}) = (x_{i_0}, y_{i_0}) = a_{i_0} \in A(Q_{i_0}) \subseteq A(D)$ is contained in D . However, we also have that Q_{i_0+1} is a directed path in $D \subseteq D_{h_1}$ starting in $h_2(e_{i_0+1})$ and ending in $h_1(e_{i_0+1})$. Since Q_{i_0+1} has length at least two by the first item above, it is a directed path also in $D_{h_1} - (h_2(e_{i_0+1}), h_1(e_{i_0+1}))$ parallel to the arc $(h_2(e_{i_0+1}), h_1(e_{i_0+1}))$. This is a contradiction, since by the above the arc $(h_2(e_{i_0+1}), h_1(e_{i_0+1}))$ is contained in the transitive reduction D of D_{h_2} . This final contradiction shows that our initial assumption was incorrect, concluding the proof of the lemma by the principle of contradiction. □

Using Lemma 6.1, we are now ready to give the proof of Theorem 4.5.

Proof of Theorem 4.5. We describe an algorithm that takes as input a hypergraph $H = (V, \mathcal{E})$ and two distinct acyclic orientations h_1, h_2 of H , and returns a flip sequence $\{u_1, v_1\}, \dots, \{u_\ell, v_\ell\}$ transforming h_1 into h_2 , as follows.

The algorithm starts by going through the list of hyperedges e in \mathcal{E} and checks whether $h_1(e) \neq h_2(e)$ and whether the pair $(h_1(e), h_2(e))$ is flippable in h_1 . The latter can be done by checking whether the digraph D_{h^e} defined as in the proof of Lemma 6.1 is acyclic (which can be done in polynomial time). Once the algorithm has found an edge $e \in \mathcal{E}$ with $h_1(e) \neq h_2(e)$ such that $(h_1(e), h_2(e))$ is flippable in h_1 (such an edge e exists by Lemma 6.1), it proceeds to define $\{u_1, v_1\} := \{h_1(e), h_2(e)\}$. If $h^e = h_2$, then the algorithm returns the flip sequence $\{u_1, v_1\}$ of length 1 from h_1 to $h^e = h_2$. Otherwise, if $h^e \neq h_2$, it performs a recursive call on the input pair h^e, h_2 of distinct acyclic orientations, which gives back a flip sequence $\{u'_1, v'_1\}, \dots, \{u'_{\ell'}, v'_{\ell'}\}$ from h^e to h_2 . The algorithm now returns the concatenated flip sequence $\{u_1, v_1\}, \{u'_1, v'_1\}, \dots, \{u'_{\ell'}, v'_{\ell'}\}$ from h_1 to h_2 .

Note that with every recursive call of the above algorithm, the number of hyperedges f for which the two considered orientations of H assign different heads strictly decreases.

Hence, the above algorithm returns in polynomial time a flip sequence of length at most $|\{e \in \mathcal{E} | h_1(e) \neq h_2(e)\}|$ from h_1 to h_2 .

We next show that

$$|\{e \in \mathcal{E} | h_1(e) \neq h_2(e)\}| \leq \Delta_2(H) \cdot \text{dist}(h_1, h_2),$$

where $\text{dist}(h_1, h_2)$ denotes the length of a shortest flip sequence transforming h_1 into h_2 . To see this, note that if we flip at a pair $\{u, v\}$ of vertices in an orientation h of H , by definition of a flip this can only change the value of $h(e)$ for those $e \in E$ with $\{u, v\} \subseteq e$. However, by definition of the codegree there are at most $\Delta_2(H)$ such hyperedges e for every fixed pair $\{u, v\}$. Consequently, in an optimal flip sequence of length $\text{dist}(h_1, h_2)$ transforming h_1 to h_2 , every single flip changes the heads of at most $\Delta_2(H)$ hyperedges, and in total at least $|\{e \in \mathcal{E} | h_1(e) \neq h_2(e)\}|$ head assignments have to be changed. This shows

that $|\{e \in \mathcal{E} | h_1(e) \neq h_2(e)\}| \leq \Delta_2(H) \cdot \text{dist}(h_1, h_2)$, as desired. It follows that the above polynomial-time algorithm indeed returns a $\Delta_2(H)$ -approximation of a shortest flip sequence from h_1 to h_2 in H . This concludes the proof. \square

As a final remark, we observe that the proof of Theorem 4.5 as given above yields an upper bound $\text{dist}(h_1, h_2) \leq |\mathcal{E}|$ for the flip distance between any pair of acyclic orientations h_1, h_2 of the hypergraph $H = (V, \mathcal{E})$. This directly implies the following bound on the diameter of hypergraphic polytopes, which appears to be novel. In particular, this bound is better than what follows from the general upper bounds for diameters of polymatroids, compare e.g. [Top92].

Corollary 6.4. *For every hypergraph H , the diameter of the associated hypergraphic polytope P_H is bounded from above by the number of hyperedges in H .*

References

- [AA23] Marcelo Aguiar and Federico Ardila. Hopf monoids and generalized permutahedra. *Mem. Amer. Math. Soc.*, 289(1437):vi+119, 2023. doi:10.1090/memo/1437.
- [ACEP20] Federico Ardila, Federico Castillo, Christopher Eur, and Alexander Postnikov. Coxeter submodular functions and deformations of Coxeter permutahedra. *Adv. Math.*, 365:107039, 36, 2020. doi:10.1016/j.aim.2020.107039.
- [ACH⁺21] Oswin Aichholzer, Jean Cardinal, Tony Huynh, Kolja Knauer, Torsten Mütze, Raphael Steiner, and Birgit Vogtenhuber. Flip distances between graph orientations. *Algorithmica*, 83(1):116–143, 2021. doi:10.1007/s00453-020-00751-1.
- [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. doi:10.1007/s00454-015-9709-7.
- [Bal61] M. L. Balinski. On the graph structure of convex polyhedra in n -space. *Pacific J. Math.*, 11(2):431–434, 1961. URL: <http://projecteuclid.org/euclid.pjm/1103037323>, doi:10.2140/pjm.1961.11.431.
- [BBM19] Carolina Benedetti, Nantel Bergeron, and John Machacek. Hypergraphic polytopes: combinatorial properties and antipode. *J. Comb.*, 10(3):515–544, 2019. doi:10.4310/JOC.2019.v10.n3.a4.
- [BIK⁺23] Nicolas Bousquet, Takehiro Ito, Yusuke Kobayashi, Haruka Mizuta, Paul Ouvrard, Akira Suzuki, and Kunihiro Wasa. Reconfiguration of spanning trees with degree constraints or diameter constraints. *Algorithmica*, 85(9):2779–2816, 2023. doi:10.1007/s00453-023-01117-z.
- [CC06] Miroslav Chlebík and Janka Chlebíková. Complexity of approximating bounded variants of optimization problems. *Theoret. Comput. Sci.*, 354(3):320–338, 2006. doi:10.1016/j.tcs.2005.11.029.
- [CD06] Michael P. Carr and Satyan L. Devadoss. Coxeter complexes and graph-associahedra. *Topology Appl.*, 153(12):2155–2168, 2006. doi:10.1016/j.topol.2005.08.010.

- [CHM⁺23] Jean Cardinal, Hung P. Hoang, Arturo Merino, Ondřej Mička, and Torsten Mütze. Combinatorial generation via permutation languages. V. Acyclic orientations. *SIAM J. Discrete Math.*, 37(3):1509–1547, 2023. doi:10.1137/23M1546567.
- [CMM24] Jean Cardinal, Arturo I. Merino, and Torsten Mütze. Combinatorial generation via permutation languages. IV. elimination trees. *ACM Trans. Algorithms*, 21(1):13:1–13:41, 2024. doi:10.1145/3689633.
- [CP16] Cesar Ceballos and Vincent Pilaud. The diameter of type D associahedra and the non-leaving-face property. *European J. Combin.*, 51:109–124, 2016. doi:10.1016/j.ejc.2015.04.006.
- [CP25] Jean Cardinal and Vincent Pilaud. Rectangulotopes. *European J. Combin.*, 125:Paper No. 104090, 25, 2025. doi:10.1016/j.ejc.2024.104090.
- [CPVP22] Jean Cardinal, Lionel Pournin, and Mario Valencia-Pabon. Diameter estimates for graph associahedra. *Ann. Comb.*, 26(4):873–902, 2022. doi:10.1007/s00026-022-00598-z.
- [CPVP24] Jean Cardinal, Lionel Pournin, and Mario Valencia-Pabon. The rotation distance of brooms. *European J. Combin.*, 118:Paper No. 103877, 16, 2024. doi:10.1016/j.ejc.2023.103877.
- [CS25] Jean Cardinal and Raphael Steiner. Inapproximability of shortest paths on perfect matching polytopes. *Math. Program.*, 210(1-2):147–163, 2025. doi:10.1007/s10107-023-02025-4.
- [CSJ09] Sean Cleary and Katherine St. John. Rotation distance is fixed-parameter tractable. *Inform. Process. Lett.*, 109(16):918–922, 2009. doi:10.1016/j.ipl.2009.04.023.
- [CSS18] Jean Cardinal, Vera Sacristán, and Rodrigo I. Silveira. A note on flips in diagonal rectangulations. *Discrete Math. Theor. Comput. Sci.*, 20(2):Paper No. 14, 22, 2018. doi:10.23638/dmtcs-20-2-14.
- [CSZ15] Cesar Ceballos, Francisco Santos, and Günter M. Ziegler. Many non-equivalent realizations of the associahedron. *Combinatorica*, 35(5):513–551, 2015. doi:10.1007/s00493-014-2959-9.
- [DF10] Harm Derksen and Alex Fink. Valuative invariants for polymatroids. *Adv. Math.*, 225(4):1840–1892, 2010. doi:10.1016/j.aim.2010.04.016.
- [DLKS22] Jesús A. De Loera, Sean Kafer, and Laura Sanità. Pivot rules for circuit-augmentation algorithms in linear optimization. *SIAM J. Optim.*, 32(3):2156–2179, 2022. doi:10.1137/21M1419994.
- [DS05] Irit Dinur and Samuel Safra. On the hardness of approximating vertex cover. *Ann. of Math. (2)*, 162(1):439–485, 2005. doi:10.4007/annals.2005.162.439.
- [Edm03] Jack Edmonds. Submodular functions, matroids, and certain polyhedra. In *Combinatorial optimization—Eureka, you shrink!*, volume 2570 of *Lecture Notes in Comput. Sci.*, pages 11–26. Springer, Berlin, 2003. doi:10.1007/3-540-36478-1_2.

- [FPS01] Komei Fukuda, Alain Prodon, and Tadashi Sakuma. Notes on acyclic orientations and the shelling lemma. *Theoret. Comput. Sci.*, 263(1-2):9–16, 2001. *Combinatorics and computer science (Palaiseau, 1997)*. doi:10.1016/S0304-3975(00)00226-7.
- [Fuj05] Satoru Fujishige. *Submodular functions and optimization*, volume 58 of *Annals of Discrete Mathematics*. Elsevier B. V., Amsterdam, second edition, 2005.
- [GGMS87] I. M. Gelfand, R. M. Goresky, R. D. MacPherson, and V. V. Serganova. Combinatorial geometries, convex polyhedra, and Schubert cells. *Adv. in Math.*, 63(3):301–316, 1987. doi:10.1016/0001-8708(87)90059-4.
- [GIKO24] Tatsuya Gima, Takehiro Ito, Yasuaki Kobayashi, and Yota Otachi. Algorithmic meta-theorems for combinatorial reconfiguration revisited. *Algorithmica*, 86(11):3395–3424, 2024. doi:10.1007/s00453-024-01261-0.
- [Gre77] Curtis Greene. Acyclic orientations. In *Proceedings of the NATO Advanced Study Institute held in Berlin (West Germany), September 1-10, 1976*, 1977.
- [GZ83] Curtis Greene and Thomas Zaslavsky. On the interpretation of Whitney numbers through arrangements of hyperplanes, zonotopes, non-Radon partitions, and orientations of graphs. *Trans. Amer. Math. Soc.*, 280(1):97–126, 1983. doi:10.2307/1999604.
- [HC97] Tim H. Hultberg and Domingos M. Cardoso. The teacher assignment problem: A special case of the fixed charge transportation problem. *Eur. J. Oper. Res.*, 101(3):463–473, 1997. doi:10.1016/S0377-2217(96)00082-3.
- [HD68] Warren M. Hirsch and George B. Dantzig. The fixed charge problem. *Naval Res. Logist. Quart.*, 15(3):413–424, 1968. URL: <https://doi.org/10.1002/cpa.3160140313>, doi:10.1002/nav.3800150306.
- [Hel74] Thorkell Helgason. Aspects of the theory of hypermatroids. In *Hypergraph Seminar (Proc. First Working Sem., Ohio State Univ., Columbus, Ohio, 1972; dedicated to Arnold Ross)*, volume Vol. 411 of *Lecture Notes in Math.*, pages 191–213. Springer, Berlin-New York, 1974.
- [Hs01] Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001. doi:10.1145/502090.502098.
- [IDH⁺11] Takehiro Ito, Erik D. Demaine, Nicholas J. A. Harvey, Christos H. Papadimitriou, Martha Sideri, Ryuhei Uehara, and Yushi Uno. On the complexity of reconfiguration problems. *Theoret. Comput. Sci.*, 412(12-14):1054–1065, 2011. doi:10.1016/j.tcs.2010.12.005.
- [IIK⁺23] Takehiro Ito, Yuni Iwamasa, Naonori Kakimura, Naoyuki Kamiyama, Yusuke Kobayashi, Shun-ichi Maezawa, Yuta Nozaki, Yoshio Okamoto, and Kenta Ozeki. Monotone edge flips to an orientation of maximum edge-connectivity à la Nash-Williams. *ACM Trans. Algorithms*, 19(1):Art. 6, 22, 2023. doi:10.1145/3561302.

- [IKK⁺22] Takehiro Ito, Naonori Kakimura, Naoyuki Kamiyama, Yusuke Kobayashi, and Yoshio Okamoto. Shortest reconfiguration of perfect matchings via alternating cycles. *SIAM J. Discrete Math.*, 36(2):1102–1123, 2022. doi:10.1137/20M1364370.
- [IKK⁺23] Takehiro Ito, Naonori Kakimura, Naoyuki Kamiyama, Yusuke Kobayashi, Shunichi Maezawa, Yuta Nozaki, and Yoshio Okamoto. Hardness of finding combinatorial shortest paths on graph associahedra. In *50th International Colloquium on Automata, Languages, and Programming*, volume 261 of *LIPICs. Leibniz Int. Proc. Inform.*, pages Art. No. 82, 17. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2023. doi:10.4230/lipics.icalp.2023.82.
- [IKO⁺20] Takehiro Ito, Marcin Kamiński, Hirotaka Ono, Akira Suzuki, Ryuhei Uehara, and Katsuhisa Yamanaka. Parameterized complexity of independent set reconfiguration problems. *Discrete Appl. Math.*, 283:336–345, 2020. doi:10.1016/j.dam.2020.01.022.
- [IOO23] Takehiro Ito, Hirotaka Ono, and Yota Otachi. Reconfiguration of cliques in a graph. *Discrete Appl. Math.*, 333:43–58, 2023. doi:10.1016/j.dam.2023.01.026.
- [KSX17] Iyad Kanj, Eric Sedgwick, and Ge Xia. Computing the flip distance between triangulations. *Discrete Comput. Geom.*, 58(2):313–344, 2017. doi:10.1007/s00454-017-9867-x.
- [Lod04] Jean-Louis Loday. Realization of the Stasheff polytope. *Arch. Math. (Basel)*, 83(3):267–278, 2004. doi:10.1007/s00013-004-1026-y.
- [LP15] Anna Lubiw and Vinayak Pathak. Flip distance between two triangulations of a point set is NP-complete. *Comput. Geom.*, 49:17–23, 2015. doi:10.1016/j.comgeo.2014.11.001.
- [LR12] Shirley Law and Nathan Reading. The Hopf algebra of diagonal rectangulations. *J. Combin. Theory Ser. A*, 119(3):788–824, 2012. doi:10.1016/j.jcta.2011.09.006.
- [LX23] Haohong Li and Ge Xia. An $O(3.82^k)$ time FPT algorithm for convex flip distance. In *40th International Symposium on Theoretical Aspects of Computer Science*, volume 254 of *LIPICs. Leibniz Int. Proc. Inform.*, pages Art. No. 44, 14. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2023. doi:10.4230/lipics.stacs.2023.44.
- [Mee19] Emily Meehan. The Hopf algebra of generic rectangulations. arXiv:1903.09874, 2019. arXiv:1903.09874.
- [MM23] Arturo Merino and Torsten Mütze. Combinatorial generation via permutation languages. III. Rectangulations. *Discrete Comput. Geom.*, 70(1):51–122, 2023. doi:10.1007/s00454-022-00393-w.
- [MN20] C. M. Mynhardt and S. Nasserars. Reconfiguration of colourings and dominating sets in graphs. In *50 years of combinatorics, graph theory, and computing*, Discrete Math. Appl. (Boca Raton), pages 171–191. CRC Press, Boca Raton, FL, [2020] ©2020.

- [MP15] Thibault Manneville and Vincent Pilaud. Graph properties of graph associahedra. *Sém. Lothar. Combin.*, 73:Art. B73d, 31, 2015.
- [Pil24] Vincent Pilaud. Acyclic reorientation lattices and their lattice quotients. *Ann. Comb.*, 28(4):1035–1092, 2024. doi:10.1007/s00026-024-00697-z.
- [Pos09] Alexander Postnikov. Permutohedra, associahedra, and beyond. *Int. Math. Res. Not. IMRN*, 2009(6):1026–1106, 2009. doi:10.1093/imrn/rnn153.
- [Pou14] Lionel Pournin. The diameter of associahedra. *Adv. Math.*, 259:13–42, 2014. doi:10.1016/j.aim.2014.02.035.
- [PPR23] Arnau Padrol, Vincent Pilaud, and Julian Ritter. Shard polytopes. *Int. Math. Res. Not. IMRN*, 2023(9):7686–7796, 2023. doi:10.1093/imrn/rnac042.
- [PRW08] Alex Postnikov, Victor Reiner, and Lauren Williams. Faces of generalized permutohedra. *Doc. Math.*, 13:207–273, 2008. doi:10.4171/dm/248.
- [PS19] Vincent Pilaud and Francisco Santos. Quotientopes. *Bull. Lond. Math. Soc.*, 51(3):406–420, 2019. doi:10.1112/blms.12231.
- [PSZ23] Vincent Pilaud, Francisco Santos, and Günter M. Ziegler. Celebrating Loday’s associahedron. *Arch. Math. (Basel)*, 121(5-6):559–601, 2023. doi:10.1007/s00013-023-01895-6.
- [Rea12] Nathan Reading. Generic rectangulations. *European J. Combin.*, 33(4):610–623, 2012. doi:10.1016/j.ejc.2011.11.004.
- [Reh22] Sophie Rehberg. Pruned inside-out polytopes, combinatorial reciprocity theorems and generalized permutahedra. *Electron. J. Combin.*, 29(4):Paper No. 4.36, 31, 2022. doi:10.37236/10371.
- [Sch03] Alexander Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency*, volume 24 of *Algorithms and Combinatorics*. Springer, 2003.
- [SFC11] Susann Schrenk, Gerd Finke, and Van-Dat Cung. Two classical transportation problems revisited: pure constant fixed charges and the paradox. *Math. Comput. Modelling*, 54(9-10):2306–2315, 2011. doi:10.1016/j.mcm.2011.05.039.
- [Sta73] Richard P. Stanley. Acyclic orientations of graphs. *Discrete Math.*, 5(2):171–178, 1973. doi:10.1016/0012-365X(73)90108-8.
- [STT22] Daniel D. Sleator, Robert E. Tarjan, and William P. Thurston. Rotation distance, triangulations, and hyperbolic geometry. In *Collected works of William P. Thurston with commentary. Vol. III. Dynamics, computer science and general interest*, pages 343–377. Amer. Math. Soc., Providence, RI, [2022] ©2022. Reprint of [0928904].
- [Top92] Donald M. Topkis. Paths on polymatroids. *Math. Programming*, 54(3):335–351, 1992. doi:10.1007/BF01586058.
- [vdH13] Jan van den Heuvel. The complexity of change. In *Surveys in combinatorics 2013*, volume 409 of *London Math. Soc. Lecture Note Ser.*, pages 127–160. Cambridge Univ. Press, Cambridge, 2013.

- [VW93] Dirk Vertigan and Geoff Whittle. Recognizing polymatroids associated with hypergraphs. *Combin. Probab. Comput.*, 2(4):519–530, 1993. doi:10.1017/S0963548300000882.