

LISTING FACES OF POLYTOPES

NASTARAN BEHROOZNIA, SOFIA BRENNER, ARTURO MERINO, TORSTEN MÜTZE,
CHRISTIAN RIECK, AND FRANCESCO VERCIANI

ABSTRACT. This paper investigates the problem of listing faces of polytopes that represent combinatorial objects, such as hypercubes, permutahedra, associahedra, and their generalizations. Firstly, we consider the face lattice, which is the inclusion order of all faces of a polytope, and we seek a Hamiltonian cycle in its cover graph, i.e., for any two consecutive faces, one must be a subface of the other, and their dimensions differ by 1. We construct such Hamiltonian cycles for hypercubes, permutahedra, B -permutahedra, associahedra, cyclic polytopes, 3-dimensional polytopes, graph associahedra of chordal graphs, and quotientopes. Secondly, we consider facet-Hamiltonian cycles, which are cycles on the (1-)skeleton of a polytope that enter and leave every facet exactly once. This notion was recently introduced by Akitaya, Cardinal, Felsner, Kleist, and Lauff [SODA 2025], where the authors conjectured that B -permutahedra admit a facet-Hamiltonian cycle for all dimensions. We construct such facet-Hamiltonian cycles in this paper, thus establishing their conjecture as a theorem. A key tool we use are so-called rhombic strips, which are planar spanning subgraphs of the cover graph of the face lattice in which every face is a 4-cycle. Specifically, we construct a rhombic strip in the face lattice of the hypercube of any dimension, and characterize the existence of rhombic strips in the face lattice of 3-dimensional polytopes. Our constructions yield time- and space-efficient algorithms for computing the aforementioned cycles and thus for listing the corresponding combinatorial objects, including ordered set partitions and dissections of a convex polygon.

1. INTRODUCTION

1.1. Combinatorial polytopes. Convex polytopes are objects of fundamental interest that tie together geometric, combinatorial, algebraic and algorithmic concepts and problems. In this paper, we are particularly interested in polytopes whose vertices represent combinatorial objects, such as the set of all bitstrings of length n , the set of all permutations of $[n] := \{1, \dots, n\}$, or the set of all binary trees with n vertices. The corresponding polytopes are the well-known *hypercube*, whose edges connect pairs of bitstrings that differ in a single bit, the *permutahedron*, whose edges connect pairs of permutations that differ in an adjacent transposition, and the *associahedron*, respectively, whose edges connect pairs of binary trees that differ in a tree rotation; see Figure 1. An equivalent model of the associahedron is to consider the dual graphs of the binary trees,

(Nastaran Behrooznia) DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF WARWICK, UNITED KINGDOM

(Sofia Brenner) INSTITUT FÜR MATHEMATIK, UNIVERSITÄT KASSEL, GERMANY

(Arturo Merino) UNIVERSIDAD DE O'HIGGINS, RANCAGUA, CHILE

(Torsten Mütze) INSTITUT FÜR MATHEMATIK, UNIVERSITÄT KASSEL, GERMANY

(Christian Rieck) INSTITUT FÜR MATHEMATIK, UNIVERSITÄT KASSEL, GERMANY

(Francesco Verciani) INSTITUT FÜR MATHEMATIK, UNIVERSITÄT KASSEL, GERMANY

E-mail addresses: `nastaran.behrooznia@warwick.ac.uk`, `sofia.brenner@mathematik.uni-kassel.de`,
`arturo.merino@uoh.cl`, `tmuetze@mathematik.uni-kassel.de`, `christian.rieck@mathematik.uni-kassel.de`,
`francesco.verciani@mathematik.uni-kassel.de`.

This project was supported by German Science Foundation grant 522790373. Arturo Merino was supported by ANID FONDECYT Iniciación No. 11251528.

An extended abstract of this paper appeared in SODA 2026 [BBM⁺26].

namely triangulations of a convex polygon, and each tree rotation translates to a flip operation that changes exactly one diagonal in the triangulation.

Recently, there has been an extensive line of work to define and analyze vast classes of polytopes that generalize the aforementioned three special polytopes. One of these generalizations are *graph associahedra* [CD06, Pos09, Dev09], which are parameterized by an underlying graph H and have as vertices all elimination trees of H , with edges connecting pairs of elimination trees that differ in a tree rotation (if H is a perfect matching, a complete graph, or a path, respectively, we obtain the hypercube, permutahedron, and associahedron as special cases). Another important generalization are *quotientopes* [PS19, PPR23], which arise from lattice congruences of the weak order on permutations [Rea05].




FIGURE 1. Examples of 3-dimensional combinatorial polytopes: (a) the hypercube; (b) the permutahedron; (c) the associahedron.

1.2. Shortest paths on the skeleton. The *(1-)skeleton* of a polytope P is the graph $G(P)$ formed by the vertices and edges of P . If P is d -dimensional, the graph $G(P)$ is known to be d -connected by Balinski's theorem. The skeleton captures much information relevant for various applications.

A fundamental algorithmic problem is to compute a shortest path on the skeleton of a polytope, for a given pair of input vertices. For the associahedron, we are given two binary trees and ask for a shortest sequence of rotations to transform one tree into the other (this is relevant in the context of balancing binary search trees). This long-standing problem was shown to be NP-hard in a recent breakthrough by Dorfer [Dor26]. Nonetheless, a 2-approximation algorithm [CSJ10] and various fixed-parameter algorithms are known [CSJ09, LX25]. Recently, Cunha, Sau, Souza and Valencia-Pabon [CSSV25] proved that the shortest path problem is fixed-parameter tractable even on graph associahedra, a problem that generalizes the rotation distance between binary trees. Complementing this result, Cardinal, Pournin and Valencia-Pabon [CPVP24] found a polynomial-time algorithm for computing rotation distance in graph associahedra of complete split graph, and this is the only known non-trivial family of graphs for which such a positive result is known. Furthermore, Cardinal and Steiner [CS25b] showed that the shortest path problem is NP-hard on base polytopes of polymatroids, even if the polytope is defined by few inequalities. They also proved inapproximability when the polymatroid is a hypergraphic polytope, whose vertices are in bijection with acyclic orientations of a given hypergraph; see also [CSSV26].

Ito, Kakimura, Kamiyama, Kobayashi, and Okamoto [IKK⁺22] considered the perfect matching polytope, i.e., the polytope defined as the convex hull of the characteristic vectors of all perfect matchings of a graph. They proved that shortest paths are NP-hard to compute, even if the underlying graph is planar or bipartite, but the problem can be solved efficiently if the graph is outerplanar. Cardinal and Steiner [CS25a] showed that shortest paths on the perfect matching polytope are also hard to approximate.

The *diameter* of $G(P)$ is the shortest distance, maximized over all pairs of vertices of P . This parameter has been heavily studied for various classes of combinatorial polytopes [STT88, Pou14, Pou17, CLPL18, CPVP21, Ber22], as it provides a lower bound for the running time of the simplex algorithm when optimizing a linear objective function over P .

1.3. Hamiltonicity of the skeleton. Another fundamental property with algorithmic applications is to search for a longest path instead of a shortest path on the skeleton. Formally, a *Hamiltonian path/cycle* is a path/cycle on the skeleton that visits every vertex exactly once. The hypercube, permutahedron and associahedron are known to have Hamiltonian cycles by classical algorithms. Specifically, the well-known *binary reflected Gray code*, first described in patents by the Bell Labs researchers George R. Stibitz and Frank Gray from 1943 and 1953, respectively, computes a Hamiltonian cycle in the hypercube. The *Steinhaus-Johnson-Trotter* algorithm [Ste64, Joh63, Tro62] computes a Hamiltonian cycle in the permutahedron. Similar constructions of Hamiltonian cycles are also known for the associahedron [LRR93, HN99]. More generally, all graph associahedra admit a Hamiltonian cycle if the underlying graph H has at least two edges [MP15]. If H is chordal, then a Hamiltonian path can be computed efficiently by a simple greedy algorithm [CMM25]. Furthermore, a variant of that algorithm can be used to compute a Hamiltonian path in all quotientopes [HM21].

The aforementioned Hamiltonian paths and cycles on combinatorial polytopes are instances of so-called *combinatorial Gray codes* [Sav97, Müt23]. This term refers to a listing of combinatorial objects such that any two consecutive objects differ in a ‘small change’. In the case of hypercubes, permutahedra, and associahedra, respectively, this is a single bit being flipped, an adjacent transposition or tree rotation being applied. The purpose of such ‘small change’ listings is to enable fast generation algorithms for the combinatorial objects, ideally in time $\mathcal{O}(1)$ per generated object. Such algorithms are often referred to as *loopless*.

1.4. Hamiltonicity of the face lattice. In this work, we aim to list not only vertices and edges of a polytope P , but also the faces of all other dimensions of P . For this we consider the *face lattice* $L(P)$, i.e., the inclusion order of all faces of P , which captures the complete combinatorial structure of P ; see Figure 2 (a)–(c). The cover relations in the face lattice are pairs of faces whose dimension differs by 1 where one is a subface of the other. The face lattice includes two special faces, called *trivial* faces, namely the empty set $\emptyset$ that constitutes the unique minimum of the face lattice $L(P)$, and the entire polytope P that constitutes the unique maximum of $L(P)$. The dimension of the trivial face $\emptyset$ is defined to be -1 .

We seek a Hamiltonian cycle in the cover graph of the face lattice, denoted $G(L(P))$, i.e., a cyclic listing of all faces of P , in which every face appears exactly once, such that the dimension of any two cyclically consecutive faces differs by 1, and one is a subface of the other; see Figure 2 (d1)+(d2). In other words, this is a Gray code listing of all the faces of P with respect to the inclusion order. While obviously there is a large variety of computational problems related to face lattices of polytopes (see, e.g., [FR94, RG96, KP02]), the question for Hamiltonicity of its cover graph appears to be novel to the best of our knowledge.




FIGURE 2. (a) A 3-polytope P ; (b) a Schlegel diagram of P , a 3-connected plane graph; (c) the face lattice of P ; (d1)+(d2) a Hamiltonian cycle in $G(L(P))$; (e1)+(e2) two distinct facet-Hamiltonian cycles of P (of different lengths); (f) a rhombic strip of $G(L(P))$. The rhombi are colored according to ranks for clarity. In this and all of the following pictures of rhombic strips in our paper, we display a grayed-out copy of the leftmost vertices at the right-hand side, in order to depict the ‘wrap-around’ edges on the sphere.

1.5. A brave conjecture and some evidence. We raise the following brave conjecture.

Conjecture 1. *For any polytope P of dimension $d \geq 1$, the graph $G(L(P))$ has a Hamiltonian cycle.*

For polytopes of dimension 1 and 2, Conjecture 1 is trivially true. For a simplex P , the face lattice $L(P)$ is the hypercube, and therefore a Hamiltonian cycle in $G(L(P))$ is given by the aforementioned binary reflected Gray code. In this paper, we verify Conjecture 1 for a number of further interesting polytopes:

- hypercubes (**Theorem 6**; see Figure 3 (a1)–(a4));
- permutahedra and B -permutahedra (**Theorems 8 and 9**; see Figure 3 (b2)–(b4));
- associahedra (**Theorem 12**; see Figure 4);
- cyclic polytopes (**Theorem 14**);
- 3-dimensional polytopes (**Theorem 18**);
- graph associahedra of chordal graphs (**Theorem 22**);
- quotientopes (**Theorem 24**).

In this list and in the rest of the introduction, the new theorems established in this work are highlighted in magenta. For a polytope P , we write P^* for the *polar* polytope of P . The face lattice $L(P^*)$ is obtained by turning the face lattice $L(P)$ upside down, an operation that clearly preserves Hamiltonian cycles. Consequently, the polars of all the polytopes mentioned before also satisfy Conjecture 1, including for example all cross-polytopes, which are the polars of hypercubes. As mentioned before, while hypercubes, permutahedra and associahedra are special cases of graph associahedra and quotientopes, it is still worth to discuss them separately, as this gives new explicit listings and algorithms.

Strikingly, we did not find any counterexamples to Conjecture 1. Additional hope for the conjecture is fueled by the following observation: The cover graph of the face lattice $G(L(P))$ is bipartite, with the partition classes given by the parity of the dimension of the faces; see Figure 5.

We write $d := \dim(P)$ for the dimension of P , and f_i , $i = -1, 0, 1, \dots, d$ for the number of i -dimensional faces of P , i.e., $f = (f_{-1}, f_0, f_1, \dots, f_{d-1}, f_d) = (1, f_0, f_1, \dots, f_{d-1}, 1)$ is the *f-vector* of P . A necessary condition for $G(L(P))$ to have a Hamiltonian cycle is that its partition classes have the same size, i.e.,

$$\sum_{\substack{i=-1, \dots, d \\ i \text{ odd}}} f_i = \sum_{\substack{i=-1, \dots, d \\ i \text{ even}}} f_i$$

which is equivalent to Euler-Poincaré's famous formula

$$\sum_{i=-1}^d (-1)^i f_i = 0, \tag{1}$$

valid for any polytope P . In fact, our construction of a Hamiltonian cycle in $G(L(P))$ for 3-dimensional polytopes yields another proof of (1) for the case $d = 3$, i.e., for the formula

$$-1 + f_0 - f_1 + f_2 - 1 = 0,$$

which can be rewritten in the more familiar form

$$|V(P)| - |E(P)| + |F(P)| = 2,$$

where $V(P)$, $E(P)$, and $F(P)$ are the sets of vertices, edges, and faces of P , respectively. Specifically, a bijection between the two partition classes of $G(L(P))$ is obtained by taking every second edge from our Hamiltonian cycle.




FIGURE 3. Hamiltonian cycles in the cover graph of the face lattice of hypercubes Q_n for $n = 1, 2, 3, 4$ (a1)–(a4) and permutahedra Π_n for $n = 2, 3, 4$ (b2)–(b4). The encoding for hypercubes is 0=white, 1=black, --=gray and for permutahedra it is 1=blue, 2=green, 3=yellow, 4=red, where for values in the same block of an ordered set partition the corresponding rectangle is striped horizontally with the colors of values in that block.




FIGURE 4. Hamiltonian cycles in the cover graph of the face lattice of the associahedron A_n for $n = 4, 5, 6$. The vertical bars delimit groups of dissections obtained from the same parent dissection in the previous sequence.




FIGURE 5. The two partition classes of the bipartite graph $G(L(P))$.

Similarly, proving Conjecture 1 for larger values of d (maybe $d = 4$ as a next step), would provide another proof of (1) in the more general d -dimensional case.

1.6. Facet-Hamiltonian cycles. Akitaya, Cardinal, Felsner, Kleist and Lauff [ACF⁺25] investigated a new type of cycles in polytopes. Given a polytope P of dimension d , each $(d - 1)$ -dimensional face is called a *facet*. A *facet-Hamiltonian cycle* in P is a cycle C in the skeleton of P that enters and leaves every facet of P exactly once; see Figure 2 (e1)+(e2). Formally, for every facet F of P , the intersection $C \cap F$ is connected and nonempty. The notion of a *facet-Hamiltonian path* is defined analogously. The authors show that permutahedra and A -, B -, C -, D -associahedra admit facet-Hamiltonian cycles, and the same is true for graph associahedra of wheels, fans and complete split graphs. For graph associahedra of complete bipartite graphs and caterpillars, they construct facet-Hamiltonian paths.

One of the conjectures from [ACF⁺25] concerns the *B-permutahedron*, which has as vertices all signed permutations of $[n]$, i.e., permutations of $[n]$ in which every entry has a positive/negative sign; see Figure 6. In the figure, entries of the permutation with a negative sign are overlined. The edges of the B -permutahedron connect pairs of signed permutations that either differ in an adjacent transposition, preserving all signs, or in a complementation of the sign of the first entry.




FIGURE 6. The 2- and 3-dimensional B -permutahedra. A facet-Hamiltonian cycle in the 3-dimensional B -permutahedron is highlighted.

Conjecture 2 ([ACF⁺25, Conj. 1]). *The B -permutahedron of any dimension $d \geq 2$ has a facet-Hamiltonian cycle.*

In this work, we resolve Conjecture 2 affirmatively and establish it as a theorem (**Theorem 11**).

1.7. Rhombic strips. One of the key concepts introduced in [ACF⁺25] is that of a *rhombic strip*, which is a plane spanning subgraph of the cover graph of the face lattice of some polytope such that every face in the embedding is a *diamond*, i.e., a 4-cycle that spans three consecutive levels of the face lattice; see Figure 2 (f). Furthermore, the rhombic strip wraps around at the left and right boundary like on a sphere with $\emptyset$ and the entire polytope P at the south and north pole, respectively. Note that if $G(L(P))$ admits a rhombic strip, then for any fixed level/dimension $k \in \{1, \dots, d-1\}$, $d := \dim(P)$, we may consider all diamonds whose two middle vertices lie on level k . In Figure 2 (f), these cyclic orderings of diamonds for $k = 1, 2, 3$ are drawn blue, yellow and green, respectively. The plane embedding of the rhombic strip forces a cyclic ordering of these diamonds, which forces a cyclic ordering of all k -dimensional faces of P , with the property that any two consecutive faces have a common subspace of dimension $k+1$ (this is the top vertex of the corresponding diamond) and a common subspace of dimension $k-1$ (this is the bottom vertex of the corresponding diamond), i.e., we obtain a Gray code ordering of the faces of dimension k . As the rhombic strip contains the faces of all dimensions, we obtain such orderings simultaneously for all $k = 1, \dots, d-1$. In particular, for $k = 1$ we obtain a cyclic ordering of the vertices of P , where any two consecutive vertices are joined by an edge, which is precisely a Hamiltonian cycle in the skeleton $G(P)$, i.e., a rhombic strip in $G(L(P))$ implies a Hamiltonian cycle in $G(P)$. Furthermore, by the earlier remarks about the polar polytope P^* , a rhombic strip in $G(L(P))$ implies a Hamiltonian cycle in $G(P^*)$. Combining these observations, we see that a necessary condition on P for $G(L(P))$ to admit a rhombic strip is that $G(P)$ and $G(P^*)$ both admit a Hamiltonian cycle. Note that for 3-dimensional polytopes P , the skeleton $G(P^*)$ is simply the dual graph of $G(P)$.

Akitaya, Cardinal, Felsner, Kleist, and Lauff [ACF⁺25] proved that in order to resolve Conjecture 2, it is sufficient to find a rhombic strip in the cover graph of the face lattice of the hypercube. Our proof of Conjecture 2 follows this route, by constructing a rhombic strip in the face lattice of the hypercube of any dimension (**Theorem 7**). Furthermore, we characterize 3-dimensional polytopes P for which $G(L(P))$ admits a rhombic strip (**Theorem 19**), and this characterization is a strengthening of the aforementioned necessary condition that $G(P)$ and its dual graph must admit a Hamiltonian cycle. Using this characterization, we provide infinitely many examples of 3-dimensional polytopes P for which the necessary condition holds but $G(L(P))$ does not admit a rhombic strip (**Theorem 20** and **Proposition 21**, illustrated in Figures 30 and 32, respectively).

1.8. Efficient Gray code algorithms. Several of our aforementioned results yield new and efficient Gray code algorithms for listing various combinatorial objects. Specifically, the constructions for establishing Hamiltonian cycles in the cover graphs of face lattices of hypercubes, permutahedra, B -permutahedra, associahedra and cyclic polytopes (Theorems 6, 8, 9, 12 and 14 respectively) can be turned straightforwardly into loopless algorithms to compute those cycles, i.e., they work in optimal time $\mathcal{O}(1)$ per visited vertex. Furthermore, the algorithms only store the currently visited vertex in memory, plus a few additional data structures. In particular, from the permutahedron we obtain a loopless Gray code algorithm for listing all ordered set partitions of $[n]$, where consecutive partitions differ in either merging or splitting two adjacent sets/blocks in the ordered partition; see Figure 3 (b2)–(b4). Furthermore, from the associahedron we obtain a loopless Gray code algorithm for listing all dissections of a convex n -gon, where consecutive dissections differ in either adding or removing a single diagonal; see Figure 4. Implementations of these Gray codes in C++ are available for download and experimentation on the Combinatorial Object Server website [cosa], [cosb] and [cosc], respectively.

The construction of a facet-Hamiltonian cycle in the B -permutahedron described in the proof of Conjecture 2 (Theorem 11) also translates to a polynomial-time and -space algorithm for computing the cycle. An implementation in C++ can be found at [cosd].

1.9. Key ideas and methods. The constructions of Hamiltonian cycles in the face lattices of most of the polytopes listed after Conjecture 1 (Theorems 6, 8, 9, 12, 22 and 24) can be seen as applications of the zigzag framework for combinatorial generation pioneered in [HHMW22] and extensively used in [HM21, MM23, CMM25, CHM⁺23, GMN24], thus further extending the reach of this framework. The idea is to construct a Hamiltonian cycle by induction on the dimension, where in the induction step $d \rightarrow d + 1$, each face in the d -dimensional polytope gives rise to a sequence of ‘children’, namely a sequence of faces in the $(d + 1)$ -dimensional polytope which form a path in the cover graph of the face lattice. Consequently, those children faces can be visited consecutively, alternatingly in forward (=zig) or backward (=zag) order, and the transitions between the children sequences are inherited inductively from the previous Hamiltonian cycle.

On the other hand, our result for cyclic polytopes (Theorem 14) is obtained by restricting the binary reflected Gray code to the subset of bitstrings that describe faces of the polytope. Furthermore, our results for 3-dimensional polytopes (Theorems 18, 19, 20 and Proposition 21), are proved by direct combinatorial arguments on 3-connected planar graphs, using Steinitz’ theorem.

As outlined before, Conjecture 2 is proved following the approach suggested in [ACF⁺25], by constructing a rhombic strip in the face lattice of the hypercube of any dimension. This is based on the observation that the B -permutahedron is obtained by truncating every (non-trivial) face of the hypercube; see Figure 18. Therefore, a vertex of the B -permutahedron corresponds

to a maximal chain in the face lattice of the hypercube, and a facet of the B -permutahedron corresponds to the union of all chains through a certain face of the hypercube, i.e., the union of the upset and downset of this face, forming an ‘hourglass’ centered at this face. Consequently, sweeping a maximal chain from left to right through the rhombic strip enters and leaves each ‘hourglass’, i.e., each facet of the B -permutahedron, exactly once.

1.10. Outline of this paper. In Section 2, we provide some terminology and notation that will be used throughout this paper. We prove our results for the different types of polytopes starting with the elementary ones, namely simplices, hypercubes, permutahedra, associahedra, and cyclic polytopes in Sections 3, 4, 5, 7, 8 respectively, followed by the 3-dimensional polytopes in Section 9, before proceeding to more advanced ones, namely graph associahedra and quotientopes in Sections 10 and 11, respectively. An interlude is Section 6, where we present the proof of Conjecture 2, i.e., we establish the existence of facet-Hamiltonian cycles in B -permutahedra.

In Section 12, we conclude with some open questions and ideas on how to tackle Conjecture 1 in general.

2. PRELIMINARIES

Let $(P, <)$ be a poset. We say that two distinct elements $x, y \in P$ with $x < y$ are in a *cover relation*, denoted $x \lessdot y$, if there is no $z \in P$ with $x < z < y$. The *cover graph* of P , denoted $G(P)$, has as vertices all elements of P , and an edge (x, y) for every cover relation $x \lessdot y$. The *downset* of some $x \in P$ is the set of all $y \in P$ for which $y \leq x$. Similarly, the *upset* of some $x \in P$ is the set of all $y \in P$ for which $x \leq y$. An *interval* $[x, y]$ is the set of all $z \in P$ with $x \leq z \leq y$, i.e., it is the intersection of the upset of x and the downset of y . A *chain* in P is a sequence of elements $(x_1, \dots, x_k)$ from P such that $x_1 \lessdot x_2 \lessdot \dots \lessdot x_k$. It corresponds to a monotonically increasing path in the cover graph. For any $x, y \in P$, the *join* $x \vee y$ of x and y denotes the unique smallest element $z \in P$, such that $x \leq z$ and $y \leq z$, and the *meet* $x \wedge y$ of x and y denotes the unique largest element $z \in P$ such that $z \leq x$ and $z \leq y$ (if such elements exist). If for all $x, y \in P$, both the join $x \vee y$ and meet $x \wedge y$ exist and are unique, then $(P, <)$ is called a *lattice*.

A poset $(P, <)$ is *graded* if there is a function $\rho : P \rightarrow \mathbb{Z}$ such that $\rho(y) = \rho(x) + 1$ if (x, y) is a cover relation in P . The function ρ is called *rank function*, and $\rho(x)$ is called the *rank* of $x \in P$. A poset is *M_3 -free* if there are no five distinct elements a, b_1, b_2, b_3, c with $a \lessdot b_1 \lessdot c$, $a \lessdot b_2 \lessdot c$ and $a \lessdot b_3 \lessdot c$.

A *polytope* P is the convex hull of a finite set of points in $\mathbb{R}^d$. Equivalently, it is a bounded intersection of half-spaces. Its *dimension* is the dimension of the smallest affine space containing it. If P is d -dimensional, we sometimes refer to it as a *d -polytope*. We write $G(P)$ for the graph of the (1-)skeleton of P . A *face* of P is the intersection of P with a hyperplane such that all of P lies on the same side or on the hyperplane. Each face is itself a polytope and thus has a dimension. The combinatorial structure of P is captured by its *face lattice*, denoted $L(P)$, which is the inclusion order of all faces of P . Figure 2 (a) shows a 3-dimensional polytope, and part (c) of the figure shows its face lattice. The face lattice has the empty set $\emptyset$ as its unique minimum, and the full polytope P as its unique maximum. We refer to these two special faces of P as *trivial* faces. The face lattice is graded, where we can take as the rank function the dimensions of the faces. The unique minimum $\emptyset$ of $L(P)$ is assigned the rank (dimension) -1 . Furthermore, the 0-dimensional faces are called *vertices* and the 1-dimensional faces are called *edges*. If P is d -dimensional, then the faces of dimension $d - 1$ are called *facets*. We write $F_k(P)$ for the set of k -dimensional faces of P . Given two distinct faces $F, G \in L(P)$, we write $F \subset G$ if F is a subface of G , i.e., a subset of G that is a face of P . Furthermore, we write $F \subsetneq G$ if F

and G form a cover relation in $L(P)$, which means that $F \subset G$ and the dimension/rank of G is exactly one higher than that of F . A *flag* of P is a chain in $L(P)$ that starts at $\emptyset$ and ends at P , i.e., this chain contains exactly one face from each rank.

For a 3-dimensional polytope P , we write $V(P) := F_0(P)$, $E(P) := F_1(P)$, $F(P) := F_2(P)$ for the sets of vertices, edges, and 2-dimensional faces of P , respectively. Similarly, for a plane graph H , we write $V(H)$, $E(H)$, and $F(H)$ for the set of vertices, edges, and faces of H , respectively, and we refer to the elements of these sets as *cells* (so in the setting of plane graphs the word face is strictly reserved for the 2-dimensional cells). The projection of a 3-dimensional polytope onto the plane through a point just outside one of its faces yields a plane graph, called its *Schlegel diagram*. By Steinitz' theorem, the Schlegel diagrams of 3-dimensional polytopes are precisely the 3-connected plane graphs. For a plane graph H , we write $L(H)$ for the inclusion order of its cells $\{\emptyset\} \cup V(H) \cup E(H) \cup F(H) \cup \{H\}$, which includes the two *trivial* cells $\emptyset$ and H . While $L(H)$ is a graded poset, it is in general not a lattice. For example, if H is a cycle, then any two edges e, e' on the cycle have the inner face f and the outer face $\bar{f}$ covering both of them, i.e., $e, e' \subset f$ and $e, e' \subset \bar{f}$. For a polytope P , the face lattice $L(P)$ is known to be M_3 -free [Zie95, Thm. 2.7 (iii)]. Similarly, if H is 2-connected, then $L(H)$ is M_3 -free. The unbounded face of a plane graph H is referred to as the *outer face*, and the remaining faces are called *inner* faces.

A *Hamiltonian cycle* in a graph is a cycle that visits every vertex exactly once. A Hamiltonian cycle of $G(L(P))$ corresponds to a cyclic listing of all faces of P such that any two consecutive faces form a cover relation in $L(P)$; see Figure 2 (d1)+(d2).

A *facet-Hamiltonian cycle* of a polytope P is a cycle C in the skeleton graph $G(P)$ with the property that every facet of P has a nonempty and connected intersection with C ; see Figure 2 (e1)+(e2). In terms of the face lattice $L(P)$, such a cycle enters and leaves the downset of every facet exactly once.

Given a graded poset P with a unique minimum and maximum, a *rhombic strip* is a spanning subgraph of the cover graph $G(P)$ embedded on the sphere subject to the following conditions; see Figure 2 (f):

- the height of vertices in the embedding is given by their rank, so that the unique minimum and maximum are embedded at the south and north pole, respectively;
- edges are drawn as straight lines and there are no crossings between them;
- every face is a *diamond*, i.e., a 4-cycle (a, b, c, d) , where a and c have the same rank k , and b and d have ranks $k - 1$ and $k + 1$, respectively.

The following lemma formalizes the observations made in Section 1.7.

Lemma 3. *Let $d \geq 2$ and let P be a graded poset with ranks $-1, \dots, d$ that has a unique minimum and maximum and that admits a rhombic strip R . Then we have the following:*

- For every rank $k = 0, \dots, d - 1$ and any two elements x, x' that are consecutive in the cyclic ordering of rank k elements given by R , there is an element y (of rank $k + 1$) such that both $x \lessdot y$ and $x' \lessdot y$ are edges in R and an element z (of rank $k - 1$) such that both $z \lessdot x$ and $z \lessdot x'$ are edges in R .*
- For any element x of rank $k \in \{-1, \dots, d - 1\}$ of P , there is at least one edge $x \lessdot y$ in R , and for any element x of rank $k \in \{0, \dots, d\}$ of P , there is at least one edge $z \lessdot x$ in R .*
- If P is M_3 -free, then for any element x of rank $k \in \{0, \dots, d - 1\}$ of P , there are at least two edges $x \lessdot y$ and $x \lessdot y'$ or at least two edges $z \lessdot x$ and $z' \lessdot x$ in R .*

As a consequence of (i), we have the following:

- If P is the face lattice of a polytope Q , i.e., $P = L(Q)$, then $G(Q)$ and $G(Q^*)$ both have a Hamiltonian cycle.*

(v) If P is the inclusion order of the cells of a 2-connected plane graph H , i.e., $P = L(H)$, then H and its dual graph both have a Hamiltonian cycle.

Proof. Part (i) follows by considering the plane embedding of all diamonds that have two vertices on rank k , as described in Section 1.7. Part (ii) is a direct consequence of (i). Part (iii) follows directly from (ii), using that if there was exactly one edge $x < y$ and exactly one edge $z < x$, then there would be two diamonds (x', z, x, y) and (x, z, x'', y) that create a copy of M_3 .

Parts (iv) and (v) follow by applying (i) to the ranks $k = 1$ and $k = d - 1$. $\square$

Given strings x and y , we write xy for their concatenation. This operation extends to sets of strings Y in the natural way, specifically $xY := \{xy \mid y \in Y\}$. Similarly, for a sequence of strings $Y = (y_1, \dots, y_n)$, we define $xY := (xy_1, \dots, xy_n)$. For a string x and an integer n , we write x^n for the n -fold concatenation of x with itself. Given a sequence $x = (x_1, \dots, x_n)$, we write $\text{rev}(x) = (x_n, x_{n-1}, \dots, x_1)$ for the reversed sequence.

Given a positive integer n , we define $[n] := \{1, \dots, n\}$, and given two positive integers $a \leq b$, we define $[a, b] := \{n \in \mathbb{N} \mid a \leq n \leq b\}$ and $]a, b[:= [a, b] \setminus \{a, b\}$.

3. SIMPLICES

3.1. The polytope and its face lattice. The *Boolean lattice* Q_n is the inclusion order of all 2^n subsets of $[n]$. We identify subsets of $[n]$ by their characteristic vectors of length n , i.e., we think of $G(Q_n)$ as the graph that has $\{0, 1\}^n$ as its vertex set, and an edge between any two bitstrings that differ in a single bit. The *n -simplex* Δ_n is the convex hull of the n unit vectors in $\mathbb{R}^n$; see Figure 7. It has n vertices, and any k -subset of vertices forms a $(k - 1)$ -face, for $k = 1 \dots, n$, which is itself a simplex. Specifically, the 1-simplex is a point, the 2-simplex is a line segment, the 3-simplex is a triangle, the 4-simplex is a tetrahedron, etc. Furthermore, the face lattice of the n -simplex Δ_n is the Boolean lattice Q_n .




FIGURE 7. The 3- and 4-simplex and their face lattices, the Boolean lattices of dimension 3 and 4, respectively.

3.2. Hamiltonian cycles and rhombic strips in the face lattice. The face lattice of the simplex is known to admit a Hamiltonian cycle and a rhombic strip.

Theorem 4 (Folklore). *For any $n \geq 2$, the graph $G(L(\Delta_n)) = G(Q_n)$ has a Hamiltonian cycle.*

Theorem 5 (Folklore). *For any $n \geq 2$, the graph $G(L(\Delta_n)) = G(Q_n)$ has a rhombic strip.*

Even though the notion of rhombic strips in general posets was introduced only recently in [ACF⁺25], the concept was implicitly used long before that in special cases. In particular, we

refer to Theorem 5 as a folklore result because it follows directly from classical and well-known constructions of Venn diagrams, a connection that we will briefly explain at the end of this section.

We revisit two proofs for Theorem 5, and later combine the techniques presented here for constructing rhombic strips for the face lattice of the hypercube. The following proofs are based on two classical constructions of the binary reflected Gray code, and in fact yield the same rhombic strip, under an isomorphism that renames elements of the ground set.

Proof 1 of Theorem 5. A rhombic strip is a spanning subgraph of $G(Q_n)$ with certain additional properties, and hence it has $\{0,1\}^n$ as its vertex set. We specify such a subgraph by the set of edges of $G(Q_n)$ that it contains. Specifically, we construct a rhombic strip $R_n \subseteq G(Q_n)$ by induction on n . The notations used in this construction are illustrated in Figures 8 and 9. The construction also maintains two chains $C_n = (x_0, \dots, x_n) \subseteq R_n$ (red in the figures) and $D_n = (y_0, \dots, y_n) \subseteq R_n$ (blue in the figures) with $x_0 = y_0 = 0^n$ and $x_n = y_n = 1^n$, where C_n and D_n contains the leftmost and rightmost vertices across all ranks, respectively, and both chains are connected by edges $Z_n := \{(x_i, y_{i+1}) \mid i = 1, \dots, n-2\} \subseteq R_n$ (in the gray-shaded area in the figures). We refer to the edges in Z_n as *zipper edges*, as they connect the left and right side of the plane drawing to make a rhombic strip on the sphere, like a zipper that connects two pieces of fabric.

For the base case $n = 2$ of the construction we take $R_2 := G(Q_2)$, $C_2 := (00, 01, 11)$ and $D_2 := (00, 10, 11)$, and in this case we have $Z_2 = \emptyset$.

For the induction step, let $R_n, C_n = (x_0, \dots, x_n), D_n = (y_0, \dots, y_n)$ be given for some $n \geq 2$. We construct the rhombic strip $R_{n+1} \subseteq G(Q_{n+1})$ as follows: We remove the zipper edges from R_n , and consider the plane graph $R' := R_n \setminus Z_n$. We take a copy of $0R'$, i.e., the graph R' in which all vertices are prefixed with a 0-bit, maintaining the same embedding. In addition, we take a mirrored copy of $1R'$, i.e., the graph R' in which all vertices are prefixed with a 1-bit and the embedding is mirrored along a vertical line. The second (mirrored) graph is embedded to the right of the first one and one unit higher than the first one (note that prefixing with a 1-bit increases the ranks). We connect the two graphs by the edges

$$E := \{(0y_i, 1y_i) \mid i = 0, \dots, n\}$$

and the zipper edges

$$Z_{n+1} := \{(0x_i, 1x_i) \mid i = 1, \dots, n-1\},$$

which yields the rhombic strip

$$R_{n+1} := 0R' \cup 1R' \cup E \cup Z_{n+1}.$$





FIGURE 8. Schematic illustration of proof 1 of Theorem 5.

FIGURE 9. The rhombic strips obtained from proof 1 of Theorem 5 for $n = 2, 3, 4, 5$.

Furthermore, the new leftmost and rightmost chains are

$$C_{n+1} := (0C_n, 1^{n+1}) \text{ and } D_{n+1} := (0^{n+1}, 1C_n),$$

respectively.

One can check that the vertices on the chains C_n and D_n have the explicit form $x_i = 0^{n-i}1^i$ for $i = 0, \dots, n$ and $y_i = 10^{n-i}1^{i-1}$ for $i = 1, \dots, n$, respectively. $\square$

Proof 2 of Theorem 5. An *x-monotone path*¹ in a drawing of a graph is a path with the property that its vertices are embedded with strictly increasing abscissa values. We construct a rhombic strip $R_n \subseteq G(Q_n)$ by induction on n . The notations used in this construction are illustrated in Figures 10 and 11. The construction maintains a partition $R_n = P_n \cup A_n \cup B_n \cup Z_n$ into four sets with the following properties:

- P_n is an *x-monotone Hamiltonian path*¹ of $G(Q_n)$ that starts at the vertex 0^n and ends at a vertex of Hamming distance 1 (red in the figures).
- The sequence $C_n = (x_0, \dots, x_n)$ of leftmost vertices across all ranks is a chain $C_n \subseteq R_n$ (cyan in the figures), and the sequence $(0^n, D_n) = (0^n, y_1, \dots, y_n)$ of rightmost vertices across all ranks is a chain $(0^n, D_n) \subseteq R_n$ (pink in the figures), and the two chains are connected by the zipper edges $Z_n := \{(x_i, y_{i+1}) \mid i = 1, \dots, n-2\} \subseteq R_n$ (in the gray-shaded area in the figures).
- The edges A_n and B_n are the edges of R_n strictly above and below P_n , respectively (in the blue-shaded and yellow-shaded areas in the figures).

For the base case $n = 2$ of the construction we take $R_2 := G(Q_2)$ with $P_2 := (00, 10, 11, 01)$, which gives $C_2 = (00, 10, 11)$, $D_2 = (01, 11)$, $Z_2 = \emptyset$, $A_2 = \emptyset$ and $B_2 = \{(00, 01)\}$.

¹It can be shown that this path follows the binary reflected Gray code, though with a reflected labeling of bits compared to the sequence defined in Section 8.





FIGURE 10. Schematic illustration of proof 2 of Theorem 5.

FIGURE 11. The rhombic strips obtained from proof 2 of Theorem 5 for $n = 2, 3, 4, 5$.

For the induction step, let P_n, C_n, D_n, A_n, B_n be given for some $n \geq 2$. We construct the rhombic strip R_{n+1} as follows: We consider the plane graphs $R^\times := R_n \setminus (A_n \cup Z_n) = P_n \cup B_n$ and $R_\times := R_n \setminus (B_n \cup Z_n) = P_n \cup A_n$. We take a copy of $0R^\times$ and a copy of $1R_\times$, maintaining the same embedding. Both graphs are embedded at the same positions horizontally, but the second graph is placed one unit higher than the first one, so that the corresponding vertices on

the two copies of P_n line up vertically (note again that prefixing with 1 increases the ranks). Let $P_n = (p_1, \dots, p_N)$, $N = 2^n$, be the sequence of vertices on the Hamiltonian path P_n . Then we connect the two graphs by the edges

$$E := \{(0p_i, 1p_i) \mid i = 1, \dots, N\}$$

and the zipper edges

$$Z_{n+1} := 1Z_n \cup \{(1p_N, 10^n)\},$$

which yields the rhombic strip

$$R_{n+1} := 0R^\times \cup 1R_\times \cup E \cup Z_{n+1}.$$

The new Hamiltonian path is

$$P_{n+1} := (0p_1, 1p_1, 1p_2, 0p_2, 0p_3, 1p_3, \dots, 0p_{N-1}, 1p_{N-1}, 1p_N, 0p_N), \quad (2)$$

which becomes x -monotone by slight perturbations of the vertical edges. Furthermore, the new leftmost and rightmost chains are

$$C_{n+1} := (0^{n+1}, 1C_n) \text{ and } D_{n+1} := (0p_N, 1D_n),$$

respectively (with the exception of 0^{n+1} , which is not part of D_{n+1}), and the sets of edges above and below P_{n+1} are

$$\begin{aligned} A_{n+1} &:= 1A_n \cup \{(1p_i, 1p_{i+1}) \mid i = 2, 4, 6, \dots, N-2\} \text{ and} \\ B_{n+1} &:= 0B_n \cup \{(0p_i, 0p_{i+1}) \mid i = 1, 3, 5, \dots, N-1\}, \end{aligned} \quad (3)$$

respectively. Note that in (2) and (3) we have used that $N = 2^n$ is even.

One can check that the vertices x_i and y_i on the chains C_n and D_n have the explicit form $x_i = 1^i 0^{n-i}$ for $i = 0, \dots, n$ and $y_i = 1^{i-1} 0^{n-i} 1$ for $i = 1, \dots, n$, respectively. $\square$

Maybe surprisingly, both aforementioned proofs actually yield the same rhombic strip, up to reversal of the vertex labels, which can be checked for $n = 2, 3, 4, 5$ by carefully comparing Figures 9 and 11, and which can easily be shown in general using induction.

There is an interesting connection between rhombic strips in $G(Q_n)$ and Venn diagrams. Specifically, the dual graph of a rhombic strip of $G(Q_n)$ is an *n -Venn diagram* [RW97], i.e., a diagram of n simple closed curves in the plane with only finitely many intersection points, such that each of the 2^n possible intersection patterns is represented by exactly one of the regions in the diagram; see Figure 12. Furthermore, those diagrams are *monotone*, which means that for all $k = 1, \dots, n-1$ and every region inside of exactly k of the curves, there is a neighboring




FIGURE 12. The dual graph of a rhombic strip of $G(Q_n)$ is a monotone n -Venn diagram. Vertices and regions are encoded by their characteristic vectors of length n .

region inside of exactly $k - 1$ curves, and another neighboring region inside of exactly $k + 1$ curves (this is a consequence of Lemma 3 (ii)).

The two aforementioned proofs can be dualized into the language of Venn diagrams, and amount to constructing a monotone $(n + 1)$ -Venn diagram inductively from a monotone n -Venn diagram. Such constructions have already been provided in Venn's original paper [Ven80] from 1880.

4. HYPERCUBES

4.1. The polytope and its face lattice. The n -(hyper)cube is the convex hull of the point set $\{0, 1\}^n$; see Figure 13. Equivalently, it is the n -fold Cartesian product of the unit interval $[0, 1]$. The skeleton of this polytope is isomorphic to the cover graph of the Boolean lattice Q_n , and so we use Q_n to also denote the polytope. Each face F of Q_n can be encoded by a ternary string $x_F \in \{0, 1, -\}^n$, where the set of vertices contained in F is given by replacing in x_F all occurrences of $-$ by either 0 or 1; see Figure 13. The dimension of the face F is given by the number of $-$ s in the string x_F . For example, $x_F = 01-- = \{0100, 0101, 0110, 0111\}$ represents a 2-dimensional face F (a quadrilateral), and $-^n = Q_n$ is the entire polytope.




FIGURE 13. The 2- and 3-cube and their face lattices, with a Hamiltonian cycle in the cover graph highlighted. The cycles are the same as in Figure 3 (a2) and (a3).

The cover graph of the face lattice $G(L(Q_n))$ has as vertex set $\{0, 1, -\}^n \cup \{\emptyset\}$, and as edges all pairs of ternary strings that differ in replacing a single 0 or 1 by $-$, or vice versa, plus the edges $(\emptyset, x)$ for all $x \in \{0, 1\}^n$.

4.2. Hamiltonian cycles and rhombic strips in the face lattice. The following two theorems are the main results of this section.

Theorem 6. *For any $n \geq 1$, the graph $G(L(Q_n))$ has a Hamiltonian cycle.*

The construction described in the following proof is illustrated in Figure 13 for $n = 2, 3$ and in Figure 3 (a1)–(a4) for $n = 1, 2, 3, 4$.

Proof. We can list all 3^n strings from the set $\{0, 1, -\}^n$ using the ternary reflected Gray code in such a way that any two consecutive strings differ in a flip $0 \leftrightarrow -$ or $- \leftrightarrow 1$. This listing is defined inductively as $\Gamma_1 := (0, -, 1)$, and for $n \geq 2$ as

$$\Gamma_n := (0\Gamma_{n-1}, -\text{rev}(\Gamma_{n-1}), 1\Gamma_{n-1}).$$

In words, the listing Γ_n consists of three copies of the previous listing Γ_{n-1} , where all strings in the first copy are prefixed with 0, all strings in the second copy are prefixed with -, and all strings in the third copy are prefixed with 1, and the order of strings in the second copy is reversed. It is easy to check that the sequence Γ_n starts with the string 0^n and ends with the string 1^n , i.e., with two faces of rank 0 in the face lattice. Consequently, $(\Gamma_n, \emptyset)$ is the desired Hamiltonian cycle in $G(L(Q_n))$. $\square$

Theorem 7. *For any $n \geq 1$, the graph $G(L(Q_n))$ has a rhombic strip.*

For proving this result, we combine the techniques developed in the two proofs of Theorem 5 presented in Section 3. We use induction and embed three copies of the previous structures, one prefixed with 0, one prefixed with -, and one prefixed with 1. The two structures prefixed with 0 and - are connected as in the second proof via ‘stacking’, and the two structures prefixed with - and 1 are connected as in the first proof via ‘mirroring’. An efficient Gray code algorithm implementing this construction in C++ can be found at [cosd].

Proof. We construct a rhombic strip $R_n \subseteq G(L(Q_n))$ by induction on n . The notations used in this construction are illustrated in Figures 14 and 15. We define the abbreviation $B'_n := \{(\emptyset, x) \mid x \in \{0, 1\}^n\}$ for the set of edges between the minimum $\emptyset$ of $L(Q_n)$ and all rank 0 faces. The construction maintains a partition $R_n = P_n \cup A_n \cup B_n \cup B'_n \cup Z_n$ into five sets with the following properties:

- P_n is an x -monotone Hamiltonian path of $G(L(Q_n)) \setminus \{\emptyset\}$ that starts and ends at rank 0 vertices (red in the figures).
- The sequence $(\emptyset, C_n) = (\emptyset, x_0, \dots, x_n)$ of leftmost vertices across all ranks is a chain $(\emptyset, C_n) \subseteq R_n$ (cyan in the figures), and the sequence $(\emptyset, D_n) = (\emptyset, y_0, \dots, y_n)$ of rightmost vertices across all ranks is a chain $(\emptyset, D_n) \subseteq R_n$ (pink in the figures), and the two chains are connected by the zipper edges $Z_n := \{(x_i, y_{i+1}) \mid i = 0, \dots, n-2\} \subseteq R_n$ (in the gray-shaded area in the figures).
- The edges A_n and $B_n \cup B'_n$ are the edges of R_n strictly above and below P_n , respectively (in the blue-shaded and yellow-shaded areas in the figures).

For the base case $n = 1$ of the construction we take $R_1 := G(L(Q_1))$ with $P_1 := (0, -, 1)$, which gives $C_1 = (0, -)$, $D_1 = (1, -)$, $Z_1 = \emptyset$, $A_1 = \emptyset$ and $B_1 = \emptyset$.

For the induction step, let $P_n, C_n = (x_0, \dots, x_n), D_n = (y_0, \dots, y_n), A_n, B_n$ be given for some $n \geq 1$. We construct the rhombic strip R_{n+1} as follows: We consider the plane graphs:

$$\begin{aligned} R^\times &:= R_n \setminus (A_n \cup B'_n \cup Z_n) = P_n \cup B_n \\ R_\times &:= R_n \setminus (B_n \cup B'_n \cup Z_n) = P_n \cup A_n \\ R' &:= R_n \setminus (B'_n \cup Z_n) = P_n \cup A_n \cup B_n. \end{aligned}$$

We take a copy of $0R^\times$ and a copy of $-R_\times$, maintaining the same embedding, and a vertically mirrored copy of $1R'$. The first two graphs are embedded at the same positions horizontally, but the second graph is placed one unit higher than the first one, so that the corresponding vertices on the two copies of P_n line up vertically (note that prefixing with - increases the ranks). The third (mirrored) graph is embedded to the right of the first two and at the same height as




FIGURE 14. Schematic illustration of the proof of Theorem 7.

the first one. Let $P_n := (p_1, \dots, p_N)$, $N = 3^n$, be the sequence of vertices on the Hamiltonian path P_n . Then we connect the first and second graph by the edges

$$E := \{(0p_i, -p_i) \mid i = 1, \dots, N\},$$

the second and third graph by the edges

$$F := \{(1y_i, -y_i) \mid i = 0, \dots, n\},$$

and the third and first graph by the zipper edges

$$Z_{n+1} := \{(1x_i, -x_i) \mid i = 0, \dots, n-1\},$$

which yields the rhombic strip


$$R_{n+1} := 0R^\times \cup -R_\times \cup 1R' \cup E \cup F \cup Z_{n+1} \cup B'_{n+1}.$$

The new Hamiltonian path (for $G(L(Q_{n+1})) \setminus \{\emptyset\}$) is

$$P_{n+1} := (0p_1, -p_1, -p_2, 0p_2, 0p_3, -p_3, \dots, -p_{N-1}, 0p_{N-1}, 0p_N, -p_N, 1 \text{ rev}(P_n)), \quad (4)$$

which becomes x -monotone, again, by slight perturbations of the vertical edges in the first part. Furthermore, the new leftmost and rightmost chains are

$$C_{n+1} := (0p_1, -C_n) \text{ and } D_{n+1} := (1C_n, -^{n+1}),$$

FIGURE 15. The rhombic strips obtained from the proof of Theorem 7 for $n = 1, 2, 3, 4$.

respectively (with the exception of $\emptyset$, which is neither part of C_{n+1} nor D_{n+1}), and the sets of edges above and below P_{n+1} (in addition to B'_{n+1}) are

$$\begin{aligned} A_{n+1} &:= -A_n \cup 1A_n \cup \{(-p_i, -p_{i+1}) \mid i = 2, 4, 6, \dots, N-1\} \cup (F \setminus \{(1p_N, -p_N)\}) \text{ and} \\ B_{n+1} &:= 0B_n \cup 1B_n \cup \{(0p_i, 0p_{i+1}) \mid i = 1, 3, 5, \dots, N-2\}, \end{aligned} \quad (5)$$

respectively. Note that in (4) and (5) we have used that $N = 3^n$ is odd.

One can check that the vertices x_i and y_i on the chains C_n and D_n , respectively, have the explicit form $x_i = -^i 0^{n-i}$ for $i = 0, \dots, n$ and $y_i = 1-^i 0^{n-1-i}$ for $i = 0, \dots, n-1$. $\square$

5. PERMUTAHEDRA

5.1. The polytope and its face lattice. The *permutahedron* Π_n is the convex hull of $\{(\pi(1), \dots, \pi(n)) \mid \pi \in S_n\}$, where S_n denotes the set of all permutations on $[n]$; see Figure 16. The permutahedron is $(n-1)$ -dimensional, and its edges connect pairs of permutations that differ in an adjacent transposition.² More generally, faces of the permutahedron can be described as follows: an (ordered) *set partition* of $[n]$ is a sequence $A_1, \dots, A_k$ of nonempty disjoint subsets $A_i \subseteq [n]$ whose union is $[n]$, where $k \in [n]$. We write such a partition as a string $A_1|A_2|\dots|A_k$, and we refer to the sets A_i , $i \in [k]$, as *blocks*, and to the characters $|$ as *bars*. Each k -dimensional face F of Π_n corresponds to a set partition $x_F = A_1|\dots|A_{n-k}$ with $n-k$ blocks, where the set of permutations contained in F is given by permuting the elements in each block of the partition x_F arbitrarily. For example $x_F = 25|1|34$ encodes the 2-dimensional face F of Π_5 that contains the permutations $\{25134, 52134, 25143, 52143\}$ (a quadrilateral). The facets of Π_n are set partitions with two blocks $A_1|A_2$, i.e., these can be identified with nonempty proper

²By this we mean a transposition of adjacent positions, i.e., $\pi(i) \leftrightarrow \pi(i+1)$, which is true if vertices are labeled by the inverse of the permutation that defines the vertex coordinates in our definition via the convex hull.

subsets $A_1 \subseteq [n]$, and there are $2^n - 2$ of them. The total number of faces of the permutahedron (excluding $\emptyset$) is given by the Fubini numbers or ordered Bell numbers (OEIS A000670).




FIGURE 16. The 2-dimensional permutahedron Π_3 and its face lattice, with a Hamiltonian cycle in the cover graph highlighted. The cycle is the same as in Figure 3 (b3).

The cover graph $G(L(\Pi_n))$ has as vertex set all set partitions of $[n]$, and edges between pairs of partitions that differ in removing a bar or adding a bar, which results in joining two blocks to one or splitting a block into two, respectively, plus the edges $(\emptyset, \pi)$ for all $\pi \in S_n$.

5.2. Hamiltonian cycles in the face lattice. The following theorem and the analogous result for B -permutahedra (stated in Theorem 9 below) are the two main results of this section.

Theorem 8. *For any $n \geq 2$, the graph $G(L(\Pi_n))$ has a Hamiltonian cycle.*

The Hamiltonian cycles constructed in this proof are shown in Figure 3 (b2)–(b4) for $n = 2, 3, 4$. An efficient Gray code algorithm implementing this construction in C++ can be found at [cosa].

Proof. For any set partition $x = A_1|A_2|\cdots|A_k$ of $[n-1]$ and for $i = 0, \dots, k$, we let $\check{c}_i(x)$ be the set partition of $[n]$ obtained by adding the singleton set $\{n\}$ as a new block after A_i and before A_{i+1} (at the beginning if $i = 0$ and the end if $i = k$), i.e.,

$$\check{c}_i(x) := A_1|\cdots|A_i|\{n\}|A_{i+1}|\cdots|A_k.$$

Furthermore, for $i = 1, \dots, k$, we let $\hat{c}_i(x)$ be the set partition of $[n]$ obtained by adding the element n to the set A_i , i.e.,

$$\hat{c}_i(x) := A_1|\cdots|A_{i-1}|A_i \cup \{n\}|A_{i+1}|\cdots|A_k.$$

We then define a sequence $\vec{c}(x)$ of set partitions of $[n]$ by

$$\vec{c}(x) := (\check{c}_0(x), \hat{c}_1(x), \check{c}_1(x), \hat{c}_2(x), \check{c}_2(x), \dots, \hat{c}_{k-1}(x), \check{c}_{k-1}(x), \hat{c}_k(x), \check{c}_k(x)).$$

In words, $\vec{c}(x)$ is obtained from the set partition x by either adding $\{n\}$ as a new singleton block between two existing blocks, or by joining the element n with one of the existing blocks in x , alternatingly from left to right. Consequently, the sequence $\vec{c}(x)$ has length $2k + 1$, and consists alternatingly of k -faces and $(k + 1)$ -faces of Π_n . In fact, any two consecutive faces in the sequence $\vec{c}(x)$ form an edge in $G(L(\Pi_{n+1}))$, i.e., the sequence describes a path in $G(L(\Pi_n))$ that alternates between rank r and $r + 1$, where $r = n - 1 - k$ is the rank of x in $L(\Pi_{n-1})$. We write $\overleftarrow{c}(x) := \text{rev}(\vec{c}(x))$ for the reverse sequence/path. For example, for $x = 25|1|34 \in L(\Pi_5)$ we have the paths

$$\vec{c}(x) = (6|25|1|34, 25\bar{6}|1|34, 25|6|1|34, 25|\bar{1}6|34, 25|1|\bar{6}|34, 25|1|34\bar{6}, 25|1|34|6) \text{ and}$$

$$\overleftarrow{c}(x) = (25|1|34|\bar{6}, 25|1|34\bar{6}, 25|1|\bar{6}|34, 25|\bar{1}6|34, 25|6|1|34, 25\bar{6}|1|34, 6|25|1|34)$$

in $G(L(\Pi_6))$.

To prove the theorem, we construct a path P_n in $G(L(\Pi_n))$ that visits all faces of Π_n except $\emptyset$ and that starts and ends at rank 0 faces (i.e., two permutations), and therefore $(P_n, \emptyset)$ is the desired Hamiltonian cycle in $G(L(\Pi_n))$.

The path P_n is constructed inductively as follows: For the base case $n = 2$ of the construction we take $P_2 := 1|2, 12, 2|1$.

For the induction step, let $P_{n-1} =: (x_1, \dots, x_N)$ be the path in $G(L(\Pi_{n-1}))$. Recall that $L(\Pi_{n-1})$ is bipartite, with P_{n-1} starting and ending in the same part, so N is always odd. We define

$$P_n := (\overleftarrow{c}(x_1), \overrightarrow{c}(x_2), \overleftarrow{c}(x_3), \overrightarrow{c}(x_4), \dots, \overleftarrow{c}(x_N)). \quad (6)$$

Note that both the first and last entry in each of the subsequences $\overleftarrow{c}(x_1)$ and $\overleftarrow{c}(x_N)$ have rank 0, and so P_n starts and ends with rank 0 faces. We observe that if x and y are adjacent in $G(L(\Pi_{n-1}))$, then the last entry of $\overleftarrow{c}(x)$ and the first entry of $\overrightarrow{c}(y)$ (which are $\{n\}|x$ and $\{n\}|y$, respectively), are adjacent in $G(L(\Pi_n))$. Similarly the last entry of $\overrightarrow{c}(x)$ and the first entry of $\overleftarrow{c}(y)$ (which are $x|\{n\}$ and $y|\{n\}$, respectively) are adjacent. From this it follows inductively that P_n as defined in (6) is indeed a path in $G(L(\Pi_n))$ that visits all faces of Π_n except $\emptyset$. This completes the proof. $\square$

One can easily prove by induction the following two noteworthy properties of the path P_n in $G(L(\Pi_n))$ constructed in the proof of Theorem 8. Firstly, for $P_n =: (x_1, \dots, x_N)$ and all $i = 1, \dots, N$, if $x_i = A_1|A_2|\dots|A_k$, then we have $x_{N+1-i} = A_k|\dots|A_2|A_1$, i.e., the set partitions x_i and x_{N+1-i} differ only in reversing the order of their blocks. Secondly, let us consider the subsequence P'_n of P_n given by all permutations, i.e., set partitions with n blocks. For a permutation $x \in S_{n-1}$, let $\overleftarrow{c}'(x)$ denote the subsequence of $\overleftarrow{c}(x)$ of permutations from S_n , i.e., $\overleftarrow{c}'(x) = (\check{c}_{n-1}(x), \check{c}_{n-2}(x), \dots, \check{c}_1(x), \check{c}_0(x))$. Then for $P'_{n-1} =: (x_1, \dots, x_N)$ we have $P'_n = (\overleftarrow{c}'(x_1), \overleftarrow{c}'(x_2), \overleftarrow{c}'(x_3), \dots, \overleftarrow{c}'(x_N))$, i.e., the next sequence of permutations is obtained by repeatedly inserting the new largest value n from right to left in the previous list of permutations. This of course does not yield a Hamiltonian path in $G(\Pi_n)$.

5.3. B -permutahedra. The B -permutahedron $\overline{\Pi}_n$ is the convex hull of all signed permutations, i.e., the convex hull of $\{(\pm\pi(1), \dots, \pm\pi(n)) \mid \pi \in S_n\}$; see Figure 6 and Figure 17.

The B -permutahedron is n -dimensional. Its edges connect pairs of permutations that either differ in an adjacent transposition, preserving all signs, or in a complementation of the sign of the first entry. Each k -dimensional face F of $\overline{\Pi}_n$, $k \geq 1$, corresponds to a signed set partition of one of the following two types:




FIGURE 17. The 2-dimensional B -permutahedron $\overline{\Pi}_2$ and its face lattice, with a Hamiltonian cycle in the cover graph highlighted. Negative signs are indicated by overlining, and boxes indicate type 2 faces, i.e., both signs for the corresponding entries. The cycle is the same as the one constructed in the proof of Theorem 9.

- (type 1) a set partition $x_F = A_1|A_2|\cdots|A_{n-k}$ of $[n]$ with $n-k$ blocks where each $i \in [n]$ carries a positive or negative sign;
 (type 2) a set partition $x_F = \overline{A_1}|A_2|\cdots|A_{n-k+1}$ of $[n]$ with $n-k+1$ blocks where each $i \in [n] \setminus A_1$ carries a positive or negative sign, and all $i \in A_1$ carry both signs (positive and negative).

For type 1, the set of signed permutations contained in F is given by permuting all elements with their corresponding signs in each block of x_F arbitrarily. For type 2, one also iterates over all ways of selecting one of the two possible signs for all elements in the first block A_1 .

A type 1 face $A_1|A_2|\cdots|A_k$ has the cover relations

$$\begin{aligned} A_1|A_2|\cdots|A_k &\subset \overline{A_1}|A_2|\cdots|A_k, \quad \text{and} \\ A_1|A_2|\cdots|A_k &\subset A_1|\cdots|A_i \cup A_{i+1}|\cdots|A_k \quad \text{for } i = 1, \dots, k-1, \end{aligned}$$

whereas a type 2 face $\overline{A_1}|A_2|\cdots|A_k$ has the cover relations

$$\begin{aligned} \overline{A_1}|A_2|\cdots|A_k &\subset \overline{A_1 \cup A_2}|\cdots|A_k, \quad \text{and} \\ \overline{A_1}|A_2|\cdots|A_k &\subset \overline{A_1}|A_2|\cdots|A_i \cup A_{i+1}|\cdots|A_k \quad \text{for } i = 2, \dots, k-1. \end{aligned}$$

Theorem 9. *For any $n \geq 1$, the graph $G(L(\overline{\Pi}_n))$ has a Hamiltonian cycle.*

The Hamiltonian cycle constructed in this proof is shown in Figure 17 for $n = 2$.

Proof. For a type 1 signed set partition $x = A_1|\cdots|A_k$ of $[n-1]$ we define

$$\begin{aligned} \check{c}_i^+(x) &:= A_1|\cdots|A_i|\{n\}|A_{i+1}|\cdots|A_k, \\ \check{c}_i^-(x) &:= A_1|\cdots|A_i|\{\overline{n}\}|A_{i+1}|\cdots|A_k, \quad i = 0, \dots, k, \\ \hat{c}_i^+(x) &:= A_1|\cdots|A_{i-1}|A_i \cup \{n\}|A_{i+1}|\cdots|A_k, \\ \hat{c}_i^-(x) &:= A_1|\cdots|A_{i-1}|A_i \cup \{\overline{n}\}|A_{i+1}|\cdots|A_k, \quad i = 1, \dots, k, \end{aligned}$$

where the overline indicates a negative sign, and we further define

$$\hat{c}_0^\pm(x) := \overline{\{n\}}|A_1|\cdots|A_k.$$

For a type 2 signed set partition $y = \overline{A_1}|A_2|\cdots|A_k$ of $[n-1]$ we define $\check{c}_i^+(y)$, $\check{c}_i^-(y)$ for $i = 1, \dots, k$ as before, and $\hat{c}_i^+(y)$, $\hat{c}_i^-(y)$ for $i = 2, \dots, k$ as before, plus the additional

$$\hat{c}_1^\pm(y) := \overline{A_1 \cup \{n\}}|A_2|\cdots|A_k.$$

In each of these two cases, we define a sequence $c(x)$ and $c(y)$ of signed set partitions of $[n]$ by

$$\begin{aligned} c(x) &:= (\check{c}_k^+(x), \hat{c}_k^+(x), \dots, \hat{c}_1^+(x), \check{c}_0^+(x), \hat{c}_0^\pm(x), \check{c}_0^-(x), \hat{c}_1^-(x), \dots, \hat{c}_k^-(x), \check{c}_k^-(x)) \text{ and} \\ c(y) &:= (\check{c}_k^+(y), \hat{c}_k^+(y), \dots, \hat{c}_2^+(y), \check{c}_1^+(y), \hat{c}_1^\pm(y), \check{c}_1^-(y), \hat{c}_2^-(y), \dots, \hat{c}_k^-(y), \check{c}_k^-(y)), \end{aligned}$$

respectively. Note that both sequences define paths in $G(L(\overline{\Pi}_n))$.

To prove the theorem, we construct a path P_n in $G(L(\overline{\Pi}_n))$ that visits all faces of $\overline{\Pi}_n$ except $\emptyset$ and that starts and ends at rank 0 faces (i.e., two signed permutations), and therefore $(P_n, \emptyset)$ is the desired Hamiltonian cycle in $G(L(\overline{\Pi}_n))$. The path P_n is constructed inductively as follows: For the base case $n = 1$ of the construction we take $P_1 := 1, \overline{1}, \overline{1}$. For the induction step, let $P_{n-1} := (x_1, \dots, x_N)$ be the path in $G(L(\overline{\Pi}_{n-1}))$. Then we define

$$P_n := (c(x_1), \text{rev}(c(x_2)), c(x_3), \dots, c(x_N)).$$

It can be checked straightforwardly that the path P_n has the required properties. $\square$




FIGURE 18. (a) Truncating the simplex yields the permutahedron. (b) Truncating the hypercube yields the B -permutahedron. The facet-Hamiltonian cycles in the permutahedron and B -permutahedron are obtained via Lemma 10 from the rhombic strips shown in Figures 9 and 15, respectively. The highlighted vertex corresponds to the leftmost chain in the rhombic strip. Each of the highlighted four-cycles yields two possibilities for the facet-Hamiltonian cycle to walk around this face, corresponding to two options of sweeping the chain through the rhombic strip.

6. TRUNCATION

Given a polytope P , the *omnitruncation* of P , denoted $\text{tr}(P)$, is the simple polytope of the same dimension as P that has one vertex for each flag of P . Geometrically, it is obtained by truncating all the (non-trivial) faces of P ; see Figure 18.

It is well-known that truncating the simplex yields the permutahedron, i.e., we have $\text{tr}(\Delta_n) = \Pi_n$, and truncating the hypercube yields the B -permutahedron, i.e., we have $\text{tr}(Q_n) = \overline{\Pi}_n$.

In their paper [ACF⁺25], the authors established the following connection between rhombic strips in the face lattice $L(P)$ of a polytope P and facet-Hamiltonian cycles in the omnitruncated polytope $\text{tr}(P)$.

Lemma 10 ([ACF⁺25]). *The graph $G(\text{tr}(P))$ has a facet-Hamiltonian cycle if and only if $G(L(P))$ has a rhombic strip.*

The facet-Hamiltonian cycle in $G(\text{tr}(P))$ is obtained by sweeping a chain from left-to-right through the rhombic strip of $G(L(P))$, changing one element of the chain at a time, moving it across a diamond. If several elements in the chain can be changed, then we have freedom in constructing several different facet-Hamiltonian cycles; see Figure 18.

Applying Lemma 10 for $P := \Delta_n$ and using Theorem 5 thus proves that the permutahedron Π_n has a facet-Hamiltonian cycle; see Figure 18 (a). In fact, we obtain many different such cycles, corresponding to the different constructions discussed in Section 3 that establish Theorem 5.

Similarly, applying Lemma 10 for $P := Q_n$ and using Theorem 7 yields the following result, which affirmatively resolves Conjecture 1 raised by Akitaya, Cardinal, Felsner, Kleist and Lauff [ACF⁺25]; see Figure 18 (b).

Theorem 11. *For any $n \geq 2$, the B -permutahedron $G(\overline{\Pi}_n)$ has a facet-Hamiltonian cycle.*

7. ASSOCIAHEDRA

7.1. The polytope and its face lattice. The combinatorial model of the associahedron that we will use in this paper does not use binary trees, as mentioned in the introduction, but instead the corresponding dual graphs, namely triangulations of a convex n -gon; see Figure 1 (c). Specifically, the *associahedron* A_n is the polytope whose face lattice $L(A_n)$ is the refinement order of dissections of a convex n -gon; see Figure 19. This polytope is $(n - 3)$ -dimensional. The




FIGURE 19. The 2-dimensional associahedron A_5 and its face lattice, with a Hamiltonian cycle in the cover graph highlighted. The cycle is the same as in Figure 4 for $n = 5$.

vertices of A_n are in bijection with triangulations of a convex n -gon, and edges correspond to flips in a triangulation, where a flip removes an inner edge between two triangles, and replaces it by the other diagonal of the resulting quadrilateral. More generally, the k -dimensional faces are given by all dissections of the n -gon with exactly $(n - k - 3)$ inner edges. In the face lattice $L(A_n)$, the cover relations are between dissections that differ in adding or removing a

single (inner) edge. The total number of faces of the associahedron (excluding $\emptyset$) is given by the little Schröder numbers (OEIS A001003).

7.2. Hamiltonian cycles in the face lattice. The following is the main result of this section.

Theorem 12. *For any $n \geq 4$, the graph $G(L(A_n))$ has a Hamiltonian cycle.*

The Hamiltonian cycles constructed in this proof are shown in Figure 4 for $n = 3, 4, 5, 6$. An efficient Gray code algorithm implementing this construction in C++ can be found at [cosb].

Proof. We label the points $1, \dots, n$ in counterclockwise order. For a given dissection X , we consider the set of edges incident with the point n , and we let $v_1, \dots, v_k$ be the neighbors of n in increasing order. Clearly, we have $k \geq 2$ and $v_1 = 1$ and $v_k = n - 1$. For $i = 2, \dots, k$, we let $\hat{c}_i(X)$ be the dissection of the $(n + 1)$ -gon obtained from X by expanding the vertex n into two vertices n and $n + 1$ connected by an edge, where $n + 1$ retains $v_1, \dots, v_{i-1}$ as neighbors, and n retains $v_i, \dots, v_k$ as neighbors; see Figure 20 (a). Furthermore, for $i = 2, \dots, k$, we let $\check{c}_i(X)$ be the dissection obtained from $\hat{c}_i(X)$ by adding the edge $(v_i, n + 1)$. Note that $\check{c}_i(X)$ for $i = 2, \dots, k - 1$ is also obtained from $\hat{c}_{i+1}(X)$ by adding the edge (v_i, n) . Lastly, we let $\check{c}_1(X)$ be the dissection obtained from $\hat{c}_2(X)$ by adding the edge $(v_1, n) = (1, n)$.




FIGURE 20. Illustration of the proof of Theorem 12. Only edges incident with n or $n + 1$ are shown in the figure, whereas all others are omitted for clarity.

We define the sequence $\vec{c}(X)$ of dissections of an $(n+1)$ -gon by

$$\vec{c}(X) := (\check{c}_1(X), \hat{c}_2(X), \check{c}_2(X), \hat{c}_3(X), \check{c}_3(X), \dots, \hat{c}_{k-1}(X), \check{c}_{k-1}(X), \hat{c}_k(X), \check{c}_k(X)),$$

and we note that it has length $2k-1$ and describes a path in $G(L(A_{n+1}))$ that alternates between ranks r and $r+1$, where r is the rank of X in $L(A_n)$; see Figure 20 (b). We write $\overleftarrow{c}(X) := \text{rev}(\vec{c}(X))$ for the reverse sequence/path.

To prove the theorem, we construct a path P_n in $G(L(A_n))$ that visits all faces of A_n except $\emptyset$ and that starts and ends at rank 0 faces (i.e., two triangulations), and therefore $(P_n, \emptyset)$ is the desired Hamiltonian cycle in $G(L(A_n))$.

The path P_n is constructed inductively as follows: For the base case $n=4$ of the construction we take P_4 consisting of the three dissections of a 4-gon shown in Figure 4.

For the induction step, let $P_n = (X_1, \dots, X_N)$ be the path in $G(L(A_n))$, i.e., each X_i is a dissection of an n -gon. Then we define

$$P_{n+1} := (\overleftarrow{c}(X_1), \vec{c}(X_2), \overleftarrow{c}(X_3), \vec{c}(X_4), \dots, \overleftarrow{c}(X_N)). \quad (7)$$

Note that both the first and last entry in each of the subsequences $\overleftarrow{c}(X_1)$ and $\overleftarrow{c}(X_N)$ have rank 0, and so P_{n+1} starts and ends with rank 0 faces. We observe that if X and Y are adjacent in $G(L(A_n))$, then the last entry of $\overleftarrow{c}(X)$ and the first entry of $\vec{c}(Y)$ are adjacent in $G(L(A_{n+1}))$, and similarly the last entry of $\vec{c}(X)$ and the first entry of $\overleftarrow{c}(Y)$ are adjacent. From this it follows inductively that P_{n+1} as defined in (7) is indeed a path in $G(L(A_{n+1}))$ that visits all faces of A_{n+1} except $\emptyset$. This completes the proof. $\square$

8. CYCLIC POLYTOPES

Cyclic polytopes play an important role in polyhedral combinatorics. According to the upper bound theorem, they maximize the number of faces for a given dimension and number of vertices. Let us recall their definition.

8.1. The polytope and its face lattice. The *moment curve* in $\mathbb{R}^d$ is defined as $f : \mathbb{R} \rightarrow \mathbb{R}^d$, $t \mapsto f(t) = (t, t^2, \dots, t^d)$. The *cyclic polytope* in $\mathbb{R}^d$ on n vertices, denoted $C(n, d)$ for $n > d \geq 2$, is defined as the convex hull of n consecutive points $f(1), \dots, f(n)$ on the moment curve. The polytope $C_{n,d}$ is d -dimensional, and for $n = d+1$ it is isomorphic to the simplex Δ_n . Each face of $C_{n,d}$ corresponds to a subset $F \subseteq [n]$, namely $i \in F$ if and only if the corresponding point $f(i)$ belongs to the face. It is known that $C_{n,d}$ is a *simplicial* polytope, i.e., all of its facets are $(d-1)$ -dimensional simplices (on d points).

To describe the face lattice of cyclic polytopes, it is essential to describe the facets. As the polytopes are simplicial, all other faces are then obtained simply by taking all possible subsets of the facets. The facets of cyclic polytopes are given by *Gale's evenness conditions* [Zie95, Thm. 0.7]. Specifically, a subset $F = \{i_1 < i_2 < \dots < i_d\} \subseteq [n]$ is a facet of $C_{n,d}$ if and only if any two consecutive elements in $[n] \setminus F$ are separated by an even number of elements in F . In other words, any maximal subsequence of consecutive elements in F that excludes 1 and n has even length.

8.2. Description as a subposet of the Boolean lattice. In the following, we describe each face $F \subseteq [n]$ of $C_{n,d}$ by its corresponding characteristic vector, i.e., a bitstring x_F of length n with a 1-bit at position i if $i \in F$ and a 0-bit at position i if $i \notin F$. For any $x \in \{0, 1\}^n$ we let $w(x)$ be the Hamming weight of x . We refer to a sequence of 1s in x that is preceded and followed by a 0-bit as a *run*, and to the number of 1s in the sequence as the *length* of the run. Note that a sequence of 1s starting at the first position or ending at the last position is *not* a run. We say that a run is *even* or *odd*, depending on the parity of its length, and we write $o(x)$

for the number of odd runs in x . For example, $x = 110001111001000$ contains two runs, one even (of length 4) and one odd (of length 1), and therefore $o(x) = 1$. By Gale's evenness condition,

$$F_{n,d} := \{x \in \{0,1\}^n \mid w(x) = d \wedge o(x) = 0\}$$

is the set of facets of $C_{n,d}$. As $C_{n,d}$ is simplicial, the face lattice $L(C_{n,d})$ (without the trivial face at the top) is isomorphic to the downset of $F_{n,d}$ in the Boolean lattice Q_n . We can describe this downset combinatorially, as follows. For any $x \in \{0,1\}^n$ we introduce the parameter $\beta(x) := w(x) + o(x)$, which is the sum of the Hamming weight of x and the number of odd runs in x . For example, for $x = 110001111001000$ from before we have $w(x) = 7$ and $o(x) = 1$, and therefore $\beta(x) = 7 + 1 = 8$. The following lemma is illustrated in Figure 21.




FIGURE 21. Illustration of Lemma 13. The Hamiltonian cycle highlighted in the face lattice $L(C_{5,3})$ is obtained from Lemma 16.

Lemma 13. *The face lattice $L(C_{n,d})$ without the trivial face at the top is isomorphic to the subposet of Q_n induced by the set*

$$X_{n,d} := \{x \in \{0,1\}^n \mid \beta(x) \leq d\}. \quad (8)$$

In words, characteristic vectors of faces of $C_{n,d}$ are precisely the bitstrings whose β -value is bounded from above by d . Note that for all $x \in X_{n,d}$ with $w(x) = d$ we have $o(x) = 0$, and more generally if $w(x) = d - s$ then we have $o(x) \leq s$.

Proof. Firstly, we need to show that if $x \in F_{n,d}$, then any y in the downset of x satisfies $\beta(y) \leq d$. For this it is enough to prove the following claim: Let $x' \in \{0,1\}^n$ and suppose that y' is obtained from x' by flipping a single 1-bit to a 0-bit, then we have $\beta(y') \leq \beta(x')$. Clearly, we have $w(y') = w(x') - 1$. If the flipped 1-bit is contained in an even run in x' , then in y' this will create one odd run and at most one even run. On the other hand, if the 1-bit is contained in an odd run in x' , then in y' this will create either at most two odd runs or at most two even runs. Lastly, if the

1-bit is part of a maximal leading or terminal sequence of 1s in x' , then this will create at most one even or odd run. In each case, y' has at most one odd run more than x' , implying $o(y') \leq o(x') + 1$. Overall, we have $\beta(y') = w(y') + o(y') \leq (w(x') - 1) + (o(x') + 1) = w(x') + o(x') = \beta(x')$, as desired.

Secondly, we need to argue that if $y \in \{0, 1\}^n$ satisfies $\beta(y) \leq d$, then y is contained in the downset of some $x \in F_{n,d}$. If $w(y) = d$ then $\beta(y) \leq d$ implies $o(y) = 0$ and therefore $y \in F_{n,d}$. Otherwise, we have $w(y) < d$. To proceed, it is enough to prove the following claim: For any $y' \in \{0, 1\}^n \setminus \{1^n\}$, one can flip a 0-bit in y' to a 1-bit to obtain a bitstring x' such that $\beta(x') = \beta(y') + 1$ if $o(y') = 0$ and $\beta(x') \leq \beta(y')$ if $o(y') > 0$. Repeatedly applying such flips to some $y' \in \{0, 1\}^n$ with $\beta(y') \leq d$ and $w(y') < d$ increases the weight and thus leads to a y in the upset of y' with $\beta(y) \leq d$ and $w(y) = d$, as covered in the first case. The proof of the claim goes as follows. If $o(y') = 0$, then flipping the first 0-bit in y' gives a bitstring x' without any odd runs, so $o(x') = o(y') = 0$ and $w(x') = w(y') + 1$, yielding $\beta(x') = w(x') + o(x') = (w(y') + 1) + o(y') = \beta(y') + 1$. On the other hand, if $o(y') > 0$, then flipping the 0-bit preceding one of the odd runs in y' gives a bitstring x' that has at least one odd run less than y' , so $o(x') \leq o(y') - 1$ and $w(x') = w(y') + 1$, yielding $\beta(x') = w(x') + o(x') \leq (w(y') + 1) + (o(y') - 1) = \beta(y')$, as claimed.

This completes the proof of the lemma. $\square$

8.3. Hamiltonian cycles in the face lattice. The following is the main result of this section.

Theorem 14. *For any $n > d \geq 2$, the graph $G(L(C_{n,d}))$ has a Hamiltonian cycle.*

By Lemma 13, our goal is to list the bitstrings in the set $X_{n,d}$ defined in (8) by flipping a single bit in each step (also, the trivial face at the top of the face lattice still needs to be inserted). We will do this by restricting the binary reflected Gray code in Q_n to the subset $X_{n,d}$. This sublist technique is very powerful and has been exploited in several previous papers (see [Müt23, Sec. 3.16–3.18+4.1]). However, the fact that it works for the set $X_{n,d}$ is new and somewhat surprising, as in all previous applications it only gives less restrictive Gray codes (allowing up to 2 or 3 flipped bits between consecutive bitstrings, whereas we insist on single flips).

The following definitions and properties are illustrated in Figure 22. We define

$$E_n := \{x = (x_1, \dots, x_n) \in \{0, 1\}^n \mid \exists k \geq 0 \text{ even s.t. } (x_1, \dots, x_{k+1}) = 1^k 0\}, \text{ and}$$

$$O_n := \{x = (x_1, \dots, x_n) \in \{0, 1\}^n \mid \exists k \geq 1 \text{ odd s.t. } (x_1, \dots, x_{k+1}) = 1^k 0\},$$

i.e., these are the sets of bitstrings of length n with an even or odd number of leading 1s, respectively. For example, we have $E_3 = \{000, 001, 011, 010, 110\}$ and $O_3 = \{100, 101\}$. Clearly, $E_n \cup O_n \cup \{1^n\}$ is a partition of $\{0, 1\}^n$.

For any $x \in \{0, 1\}^{n-1}$ and $a \in \{0, 1\}$ we have

$$\beta(ax) = \begin{cases} \beta(x) & \text{if } a = 0 \text{ and } x \in E_{n-1} \cup \{1^{n-1}\}, \\ \beta(x) + 1 & \text{if } a = 1 \text{ or } x \in O_{n-1}. \end{cases} \quad (9)$$

Note that the condition stated in the second case of (9) is precisely the negation of the condition stated in the first case.

The binary reflected Gray code is defined as $\Gamma_1 := (0, 1)$ and

$$\Gamma_n := 0\Gamma_{n-1}, 1 \text{ rev}(\Gamma_{n-1}) \quad (10)$$

for $n \geq 1$. In words, the listing Γ_n consists of two copies of the previous listing Γ_{n-1} , where all bitstrings in the first copy are prefixed with 0 and all bitstrings in the second copy are prefixed with 1, and the order of bitstrings in the second copy is reversed. We let p_n be the position of the string 1^n in Γ_n . For $b = 0, \dots, n-1$ we write $\ell_{n,b}$ for the last bitstring x in Γ_n before 1^n with

$\beta(x) = b$. Similarly, for $b = 1, \dots, n-1$ we let $f_{n,b}$ denote the first bitstring x in Γ_n after 1^n with $\beta(x) = b$.




FIGURE 22. Illustration of the binary reflected Gray code, with the vertices aligned vertically according to their β -values. Note that the β -value changes by at most 1 in every step along the Gray code, as described by Lemma 15 (iii). The gray boxes highlight all bitstrings in O_n , located after position p_n in the listing Γ_n by Lemma 15 (i). Their β -values increase by +1 according to Lemma 15 (ii) in the induction step, i.e., the content of the gray boxes is shifted up by one unit when prefixing with 0. The vertices $f_{n,b}$ and $\ell_{n,b}$, whose explicit form is given by Lemma 15 (iv), are highlighted. Dashed lines connect certain pairs of vertices differing in a single bit.

For any bitstring $x \in \{0,1\}^n$ we write $x^- \in \{0,1\}^{n-1}$ for the string obtained from x by removing the first bit. The following lemma captures four crucial properties of the binary reflected Gray code that we will exploit for our construction later; see Figure 22.

Lemma 15. *The binary reflected Gray code $\Gamma_n = (x_1, \dots, x_{2^n})$ has the following properties:*

- (i) *For any $i < p_n$ we have $x_i \in E_n$ and for any $i > p_n$ we have $x_i \in O_n$. In words, all bitstrings before 1^n have an even number of leading 1s and all bitstrings after 1^n have an odd number of leading 1s.*
- (ii) *For any bitstring x_i we have*

$$\beta(x_i) = \begin{cases} \beta(x_i^-) & \text{if } i \leq p_{n-1}, \\ \beta(x_i^-) + 1 & \text{if } i > p_{n-1}. \end{cases}$$

- (iii) *For $i = 1, \dots, 2^n - 1$ we have $\beta(x_{i+1}) - \beta(x_i) \in \{-1, 0, +1\}$. In words, the β -values of any two consecutive bitstrings in Γ_n differ by at most 1.*

(iv) For $b = 0, \dots, n-1$ we have

$$\ell_{n,b} = \begin{cases} 1^b 0^{n-b} & \text{if } b \text{ is even,} \\ 1^{b-1} 0^{n-b} 1 & \text{if } b \text{ is odd,} \end{cases} \quad (11a)$$

and for $b = 1, \dots, n-1$ we have

$$f_{n,b} = \begin{cases} 1^b 0^{n-b} & \text{if } b \text{ is odd,} \\ 1^{b-1} 0^{n-b} 1 & \text{if } b \text{ is even.} \end{cases} \quad (11b)$$

We emphasize that in part (ii) of the lemma, the bitstring x_i is part of the Gray code Γ_n , but the index p_{n-1} is (by definition) determined with respect to the Gray code Γ_{n-1} . Consequently, by (10) the bitstring at position p_{n-1} in Γ_n is the bitstring 01^{n-1} .

Proof. To prove part (i), we argue by induction on n . The induction basis $n = 1$ is trivial. For the induction step, consider the definition (10). Clearly, all bitstrings in $0\Gamma_{n-1}$ belong to the set E_n , as they have zero leading 1s, which is an even number. By induction, we have $\Gamma_{n-1} = E'_{n-1}, 1^{n-1}, O'_{n-1}$, where E'_{n-1} and O'_{n-1} are certain orderings of the elements of E_{n-1} and O_{n-1} , respectively. Consequently, we have $1 \operatorname{rev}(\Gamma_{n-1}) = 1 \operatorname{rev}(O'_{n-1}), 1^n, 1 \operatorname{rev}(E'_{n-1})$, i.e., all bitstrings before 1^n belong to the set E_n , and all bitstrings after 1^n belong to the set O_n . This proves part (i).

Part (ii) follows from (10), part (i) of the lemma (applied to Γ_{n-1}) and (9).

To prove part (iii), we argue by induction on n . Using (10) and part (ii) settles the claim for all $i = 1, \dots, 2^n - 1$ except for $i = p_{n-1}$ and $i = 2^{n-1}$. For the transitions at these two indices we continue the computation as follows. For $i = p_{n-1}$ we obtain $\beta(x_i) = \beta(01^{n-1}) = n-1$ and $\beta(x_{i+1}) = \beta(0f_{n-1,n-2}) = n-1$ with the help of (11b), so $\beta(x_{i+1}) - \beta(x_i) = 0$. For $i = 2^{n-1}$ we obtain $\beta(x_i) = \beta(010^{n-2}) = 2$ and $\beta(x_{i+1}) = \beta(110^{n-2}) = 2$, so $\beta(x_{i+1}) - \beta(x_i) = 0$. This proves part (iii).

To prove (iv), first observe that the definition (10) implies $\ell_{n,0} = 0^n$, $\ell_{n,1} = 0^{n-1}1$ and

$$\begin{aligned} \ell_{n,b} &= 1f_{n-1,b-1} \text{ for } b = 2, \dots, n-1; \\ f_{n,b} &= 1\ell_{n-1,b-1} \text{ for } b = 1, \dots, n-1. \end{aligned}$$

From those recursive relations the explicit formulas follow easily by induction. $\square$

Lemma 16. For any $d \geq 1$ and $n \geq d+1$, let $\Gamma_{n,d}$ be the sublist of Γ_n induced by the set $X_{n,d}$ defined in (8), i.e., the list obtained from Γ_n by deleting all bitstrings x with $\beta(x) > d$. Then in $\Gamma_{n,d}$ any two cyclically consecutive bitstrings differ in a single bit, except the consecutive strings $\ell_{n,d}$ and $f_{n,d}$, which differ in a transposition $0 \leftrightarrow 1$.

In Figure 22, the bitstrings are arranged vertically according to their β -values. Consequently, the sublist $\Gamma_{n,d}$ of Γ_n is obtained by taking all bitstrings from left to right whose height is at most d , or equivalently, by cutting out subsequences of bitstrings at height more than d , taking shortcuts along the dashed edges. The resulting sublist $\Gamma_{5,3}$ is highlighted in Figure 21.

As mentioned in Section 1.8, the sublist $\Gamma_{n,d}$ of Γ_n can be computed in constant time per vertex, despite the fact that sometimes many intermediate vertices are skipped when taking a shortcut. This is achieved by establishing and exploiting explicit successor rules between shortcut vertices, analogous to what has been done for similar generation algorithms based on sublists of the binary reflected Gray code; see for example [TL73]. For details, see our C++ implementation [cosc].

Proof. Let $\Gamma_n = (x_1, \dots, x_{2^n})$, and consider a maximal subsequence of consecutive indices $k \in \{1, \dots, 2^n\}$ with $\beta(x_k) > d$. By Lemma 15 (iii), the preceding index i and the following index j

of this subsequence satisfy $\beta(x_i) = \beta(x_j) = d$, i.e., we have $\beta(x_k) > d$ for all $i < k < j$. If $x_i = \ell_{n,d}$ and $x_j = f_{n,d}$, then x_i and x_j differ in a transposition $0 \leftrightarrow 1$ by Lemma 15 (iv). If not, then we have $d \geq 2$ and we consider two cases. If $i < j < p_{n-1}$ or $p_{n-1} < i < j$, then the claim follows by induction on n , using Lemma 15 (ii). It remains to consider the case $i < p_{n-1} < j$. The assumption that $\beta(x_k) > d$ for all $i < k < j$, Lemma 15 (ii), and the fact that $0\Gamma_{n-1}$ ends with $y := 010^{n-2}$ which satisfies $\beta(y) = 2$, imply that $x_i = 0\ell_{n-1,d}$ and $x_j = 0f_{n-1,d-1}$. Therefore, if d is even we obtain from Lemma 15 (iv) that $x_i = 01^d 0^{n-d-1}$ and $x_j = 01^{d-1} 0^{n-d}$, i.e., x_i and x_j differ only in the bit at position $d+1$. On the other hand, if d is odd, we obtain $x_i = 01^{d-1} 0^{n-d-1} 1$ and $x_j = 01^{d-2} 0^{n-d} 1$, i.e., x_i and x_j differ only in the bit at position d . $\square$

We are now in position to prove Theorem 14.

Proof of Theorem 14. Combine Lemmas 13 and 16. We claim that the two strings $\ell_{n,d}$ and $f_{n,d}$ that are consecutive in the listing $\Gamma_{n,d}$ both correspond to facets of $C_{n,d}$, so we can insert the trivial face $C_{n,d}$ at the top of the face lattice $L(C_{n,d})$ between these two in order to obtain a Hamiltonian cycle in $G(L(C_{n,d}))$; see the bottom part of Figure 21. Indeed, $x \in \{\ell_{n,d}, f_{n,d}\}$ satisfies $\beta(x) = d$ by definition, and $o(x) = 0$ by (11), implying that $w(x) = d$. $\square$

9. 3-DIMENSIONAL POLYTOPES

9.1. Hamiltonian cycles in the face lattice. We consider 3-dimensional polytopes via their Schlegel diagrams, which are the 3-connected plane graphs. We establish the following auxiliary statement, which actually holds even for 2-connected plane graphs.

Lemma 17. *Let $H = (V, E)$ be a 2-connected plane graph, let $(e_1, v_1, e_2, v_2, \dots, v_{k-1}, e_k, v_k)$ be the sequence of edges and vertices in counterclockwise order along the outer face $\bar{f}$, and let f be the inner face incident with e_k . Then $G(L(H)) \setminus \{\bar{f}, \emptyset, H\}$ has a Hamiltonian path that starts at f , visits each of the pairs e_i, v_i for $i = 1, \dots, k-1$ consecutively, and ends with the triple v_{k-1}, e_k, v_k .*

Proof. This proof is illustrated schematically in Figure 23, and for a concrete example in Figure 24. The proof uses the notion of the *weak dual graph*, which is obtained from the dual graph by removing the vertex that corresponds to the outer face $\bar{f}$. Note that if a plane graph is 2-connected, then its weak dual graph is connected.

We prove the statement by induction on the number of faces of H . If H has only one face, then H is a cycle $(e_1, v_1, e_2, v_2, \dots, v_{k-1}, e_k, v_k)$, bounding the only inner face f , and then $P := (f, e_1, v_1, e_2, v_2, \dots, v_{k-1}, e_k, v_k)$ is the desired path; see Figure 23 (a).

For the induction step, suppose that H has more than one face, and let $(e_1, v_1, e_2, v_2, \dots, e_k, v_k)$ and f be as in the lemma. In the weak dual graph D of H , we pick a vertex f' different from f that is also adjacent to $\bar{f}$ in the dual graph; see Figure 23 (b). Note that f' is a face in the primal graph H , and this choice ensures that removing the face f' from H yields again a 2-connected graph. Let $(v_{s-1}, e_s, v_s, e_{s+1}, v_{s+1}, \dots, v_{t-1}, e_t, v_t)$ for $1 \leq s \leq t < k$ be the sequence of vertices and edges between f' and the outer face. Furthermore, let $(d_1, u_1, d_2, u_2, \dots, d_{r-1}, u_{r-1}, d_r)$, $r \geq 1$, be the sequence of edges and vertices between v_{s-1} and v_t in clockwise order around f' . Consider the 2-connected graph

$$H' := H \setminus \{e_s, v_s, e_{s+1}, v_{s+1}, \dots, v_{t-1}, e_t\},$$

which has the same faces as H except f' , and which has the sequence

$$(e_1, v_1, e_2, v_2, \dots, v_{s-2}, e_{s-1}, v_{s-1}, d_1, u_1, d_2, u_2, \dots, d_{r-1}, u_{r-1}, d_r, v_t, e_{t+1}, v_{t+1}, \dots, v_{k-1}, e_k, v_k)$$




FIGURE 23. Illustration of the proofs of Lemma 17 and Theorem 18. Vertices, edges and faces are drawn as bullets, circles and squares, respectively.

of edges and vertices in counterclockwise order along the outer face. Moreover, we denote by $\bar{d}$ the outer face of H' . By induction, $G(L(H')) \setminus \{\bar{d}, \emptyset, H'\}$ has a Hamiltonian path P' that starts at f , visits the pairs e_i, v_i for $i = 1, \dots, s-1$ and $i = t+1, \dots, k$ consecutively, as well as the pairs d_i, u_i for $i = 1, \dots, r-1$ and the pair d_r, v_t , and that ends with the triple v_{k-1}, e_k, v_k . To obtain the desired Hamiltonian path P for $G(L(H)) \setminus \{\bar{f} \cup f', \emptyset, H\}$, we insert the sequence

$$(f', e_s, v_s, e_{s+1}, v_{s+1}, \dots, v_{t-1}, e_t)$$

between d_r and v_t in P' . It can be checked directly that P has the required properties.

This completes the proof of the lemma. $\square$

Theorem 18. *For any 2-connected plane graph H , the graph $G(L(H))$ has a Hamiltonian cycle. Consequently, for any 3-dimensional polytope P , the graph $G(L(P))$ has a Hamiltonian cycle.*

Proof. We take the path $P = (f, \dots, v_{k-1}, e_k, v_k)$ in $G(L(H))$ guaranteed by Lemma 17. We turn the path into a Hamiltonian cycle of $G(L(H))$ by adding the outer face $\bar{f}$ and the trivial




FIGURE 24. Illustration of the path constructed inductively as described in the proof of Lemma 17 for the polytope P from Figure 2. The resulting Hamiltonian cycle in $G(L(P))$ is shown in parts (d1) and (d2) of that figure.

cells $\emptyset$ and H , by replacing the last two entries e_k, v_k of P by the sequence $(\emptyset, v_k, e_k, \bar{f}, H)$, which makes the resulting sequence cyclic, as H and f are adjacent in $G(L(H))$; see Figure 23 (c). $\square$

9.2. Rhombic strips in the face lattice. The next result, illustrated in Figure 25, gives a characterization for when the inclusion order $L(H)$ of the cells of a plane graph H admits a rhombic strip. In particular, this characterizes the 3-dimensional polytopes whose face lattice admits a rhombic strip. Recall from Lemma 3 (v) that a necessary condition for $L(H)$ to have a rhombic strip is that H and its dual graph both admit a Hamiltonian cycle. However, it turns out that these necessary conditions are not sufficient.

Given a graph H and Hamiltonian cycle C in H , we refer to the edges of $H \setminus C$ as *C-chords*.

Theorem 19. *Let H be a 2-connected plane graph. Then $G(L(H))$ has a rhombic strip if and only if H has a Hamiltonian cycle $C = (v_1, \dots, v_n)$ that satisfies one of the following two equivalent conditions:*

- (i) *The cycle C can be (vertex-)partitioned into two paths A and B such that every C -chord has one endpoint on A and the other on B .*
- (ii) *There are no three C -chords $(v_{i_1}, v_{i_2}), (v_{i_3}, v_{i_4}), (v_{i_5}, v_{i_6})$ with $1 \leq i_1 < i_2 \leq i_3 < i_4 \leq i_5 < i_6 \leq n+1$, where $v_{n+1} = v_1$.*

It is easy to check that the existence of a Hamiltonian cycle C as in Theorem 19 (i) implies that the dual graph of H also has a Hamiltonian cycle.

Proof. In the first part of the proof, we show that $G(L(H))$ has a rhombic strip if and only if H has a Hamiltonian cycle C that satisfies (i).

‘ $\Rightarrow$ ’: We assume that $L(H)$ has a rhombic strip R . Let C be the Hamiltonian cycle in H defined by the cyclic ordering of all vertices of H in R (recall Lemma 3 (i)+(v)). By Lemma 3 (i),




FIGURE 25. Two illustrations of the same graph and Hamiltonian cycle that satisfies condition (i) in Theorem 19. In (a), the C -chords outside of the cycle C wrap around the top and bottom boundary (dashed). In (b), the cycle C wraps around the left and right boundary, and the C -chords inside and outside of the cycle are shown below and above it, respectively.

for any edge $e = (u, v) \in C$, both cover relations $u \subseteq e$ and $v \subseteq e$ are present as edges in R . The sequence of vertices and edges along C and the cover relations between them thus forms a cycle in the rhombic strip R that alternates between ranks 0 and 1. By planarity and Lemma 3 (ii), for any C -chord $e = (u, v)$, exactly one of the cover relations $u \subseteq e$ or $v \subseteq e$ is present as an edge in R . Let f and f' be the two faces incident with e in H . By Lemma 3 (iii), both cover relations $e \subseteq f$ and $e \subseteq f'$ must be present as edges in R , and therefore f and f' appear consecutively in the cyclic ordering of all faces of H in R (recall Lemma 3 (i)), i.e., there is a diamond (f, e, f', H) in R . As this observation applies to every C -chord e , we obtain that no face of H is bounded by more than two C -chords, and thus the cyclic ordering of all faces of H given by R can be partitioned into two contiguous parts A' and B' , where A' consists of the faces inside of C and B' of the faces outside of C , such that any two consecutive faces in A' and B' cover the one C -chord that separates them, and the two pairs of boundary entries of A' and B' each cover one edge of C that separates them. Removing these two separating edges from C (keeping their end vertices) yields the partition into paths A and B with the desired properties.

‘ $\Leftarrow$ ’: The notations used in this part of the proof are illustrated in Figure 26. We assume that H has a Hamiltonian cycle C satisfying condition (i). Let e and e' be the two edges connecting the paths A and B to the cycle C .

We specify the rhombic strip of $L(H)$ by defining a cyclic ordering of the vertices, edges and faces of H , which determines the order in which they are embedded on the sphere, at three different heights according to their rank. Furthermore, we specify which cover relations of $L(H)$ are added as (straight) edges to the rhombic strip, ensuring that there are no edge crossings and that all faces are diamonds.

Specifically, vertices of H and edges on C are ordered according to the cycle C , and all cover relations between them are added to the rhombic strip, yielding a cycle that alternates between ranks 0 and 1 (the red cycle in Figure 26). Faces of H and edges of $H \setminus C$ are ordered according to the cycle D in the dual graph given by the duals of the edges in $H \setminus C \cup \{e, e'\}$. All cover




FIGURE 26. Illustration of the proof of Theorem 19. Every C -chord c has an arrow pointing to the vertex v such (v, c) is an edge in the rhombic strip. The top part shows the same graph as Figure 25, and the bottom part shows the resulting rhombic strip. Vertices, edges and faces are drawn as bullets, circles and squares, respectively. For clarity, the trivial cells at the bottom and top are not shown in the rhombic strip.

relations between them are added to the rhombic strip, yielding a cycle that alternates between ranks 1 and 2 (the blue cycle in Figure 26).

The orderings of the edges of $C \setminus \{e, e'\}$ and $H \setminus C$ are interleaved as follows: Let X and Y denote the sets of C -chords inside or outside of C , respectively.

Any edge $\hat{e} \in A$ is inserted between the two edges of $X \cup \{e, e'\}$ that bound the same face f inside of C , whereas any edge $\hat{e} \in B$ is inserted between the two edges of $Y \cup \{e, e'\}$ that bound the same face f outside of C . In both cases, the cover relation $\hat{e} < f$ is added as an edge to the rhombic strip. In the example of Figure 26, the edge $\hat{e} = (4, 5) \in A$ is inserted between the two edges $(4, 9), (7, 9) \in X \cup \{e, e'\}$, which all bound the face $f = (4, 5, 6, 7, 9)$, so the cover relation $(4, 5) = \hat{e} < f = (4, 5, 6, 7, 9)$ is added. The edge $\hat{e} = (10, 11) \in B$ is inserted between the two edges $(5, 10), (5, 11) \in Y \cup \{e, e'\}$, which all bound the face $f = (5, 10, 11)$, so the cover relation $(10, 11) = \hat{e} < f = (5, 10, 11)$ is added.

Any edge $\hat{e} \in X$ is inserted between the two edges of $A \cup \{e, e'\}$ that have the same endpoint v as $\hat{e}$, whereas any edge $\hat{e} \in Y$ is inserted between the two edges of $B \cup \{e, e'\}$ that have the same endpoint v as $\hat{e}$. In both cases, the cover relation $v < \hat{e}$ is added as an edge to the rhombic strip. In the example of Figure 26, the edge $\hat{e} = (7, 9) \in X$ is inserted between the two edges $(6, 7), (7, 8) \in A \cup \{e, e'\}$, which all share the endpoint $v = 7$, so the cover relation $7 = v < \hat{e} = (7, 9)$ is added. The edge $\hat{e} = (1, 11) \in Y$ is inserted between the two




FIGURE 27. A second rhombic strip constructed for the graph H from Figure 26 by swapping the roles of A and B .

edges $(10, 11), (11, 12) \in B \cup \{e, e'\}$, which all share the endpoint $v = 11$, so the cover relation $11 = v \triangleleft \hat{e} = (1, 11)$ is added.

In the second part of the proof, we show the equivalence between conditions (i) and (ii).

(i) $\Rightarrow$ (ii): Suppose that the cycle C satisfies (i), and let A and B be the paths that partition C as stated in (i). Every C -chord requires that one of its endpoints belongs to A and the other to B . Consequently, if there were three C -chords $(v_{i_1}, v_{i_2}), (v_{i_3}, v_{i_4}), (v_{i_5}, v_{i_6})$ with $1 \leq i_1 < i_2 \leq i_3 < i_4 \leq i_5 < i_6 \leq n + 1$, then the membership of vertices with respect to the paths A and B along the cycle C would alternate at least three times, which is impossible.

(i) $\Leftarrow$ (ii): Suppose that the cycle C satisfies (ii). Let X and Y denote the sets of all C -chords inside or outside of C , respectively. For $Z \in \{X, Y\}$ there are two disjoint subpaths A_Z and B_Z of C such that each C -chord from Z has one endpoint on A_Z and the other on B_Z , and the pair of first vertices $(f(A_Z), f(B_Z))$ and the pair of last vertices $(\ell(A_Z), \ell(B_Z))$ of A_Z and B_Z are each connected by such a C -chord from Z . We distinguish two cases:

(1) One of the two paths A_X, B_X has a nonempty intersection with one of the paths A_Y, B_Y . Condition (ii) rules out that one of these four paths has a nonempty intersection with both paths from the other pair. Consequently, the desired partition of C is given by extending the union of the two paths that have a nonempty intersection with the union of the remaining two paths into maximal disjoint paths A and B .

(2) All four paths A_X, B_X, A_Y, B_Y are disjoint. In this case the desired partition of C is given by extending $A_X \cup A_Y$ and $B_X \cup B_Y$ into maximal disjoint paths A and B .

This completes the proof of the theorem. $\square$

By swapping the roles of the paths A and B in condition (i) of Theorem 19, we see that any such Hamiltonian cycle C in H actually gives rise to two different rhombic strips for $G(L(H))$; see Figure 27.

The 3-dimensional permutahedron and associahedron admit Hamiltonian cycles satisfying the conditions of Theorem 19, so we immediately obtain a rhombic strip for each of them; see Figures 28 and 29, respectively.




FIGURE 28. Rhombic strip for the 3-dimensional permutahedron. For clarity, the trivial cells at the bottom and top of the face lattice are not shown. Also, the bars between any two entries of the number partitions at the bottom are omitted.

Using Theorem 19 (ii), we now construct an infinite family of planar graphs H , such that H and its dual graph both admit a Hamiltonian cycle, but $G(L(H))$ does not admit a rhombic strip. The *truncated tetrahedron* is the polytope obtained from the tetrahedron by truncating every vertex to a triangle. Its skeleton is the graph H on 12 vertices depicted in Figure 30 (a).

We turn H into an infinite family of plane graphs H_s , $s \geq 0$, by repeatedly truncating a vertex. Specifically, for $s \geq 0$, the graph H_s is obtained from H by subdividing the edge (v_4, v_5) by s additional vertices $a_1, a_2, \dots, a_s$ and the edge (v_6, v_5) by s vertices $b_1, b_2, \dots, b_s$, and by adding





FIGURE 29. Rhombic strip for the 3-dimensional associahedron.

FIGURE 30. (a) Skeleton H_0 of the truncated tetrahedron with one of its Hamiltonian cycles, and (b) the graph H_s of iterated truncations at the vertex v_5 , with a Hamiltonian cycle in its dual graph.




FIGURE 31. The two Hamiltonian cycles C_1 and C_2 in the graph H_s . The C_i -chords violating condition (ii) in Theorem 19 are dashed.

the edges (a_i, b_i) for $i = 1, \dots, s$; see Figure 30 (b). Equivalently, H_s is obtained from H_{s-1} by truncating the vertex v_5 . Clearly, H_s is a 2-connected cubic plane graph.

Theorem 20. *For every $s \geq 0$, the 2-connected plane graph H_s and its dual graph both have a Hamiltonian cycle, but $G(L(H_s))$ does not admit a rhombic strip.*

Proof. One can check directly that H_s and its dual graph both have a Hamiltonian cycle; see Figure 31 and 30 (b), respectively. In fact, up to symmetry, H_0 has a unique Hamiltonian cycle C , depicted in Figure 30 (a). Specifically, such a cycle contains exactly two consecutive edges from each of the triangles t_1, t_2, t_3, t_4 . It follows that, up to symmetry, H_s has the two distinct Hamiltonian cycles C_1 and C_2 shown in Figure 31.

We proceed to show that C_1 and C_2 both violate condition (ii) of Theorem 19, and consequently $G(L(H_s))$ does not admit a rhombic strip. For the cycle C_1 shown in Figure 31 (a), the C_1 -chords (v_1, v_3) , (v_4, v_6) and (v_7, v_9) violate this condition. For the cycle C_2 shown in Figure 31 (b), we can instead take the C_2 -chords (v_{11}, v_{10}) , (v_7, v_6) and (b_1, x) with $x := v_5$ if $s = 1$ and $x := b_2$ if $s \geq 2$. $\square$

With the help of a computer, we determined the smallest 3-connected plane graph H such that both H and its dual graph have a Hamiltonian cycle, but $G(L(H))$ does not admit a rhombic strip; see Figure 32 (a). The graph H has 7 vertices and is the incidence graph of the Fano plane. Also, there are no 2-connected graphs with 6 vertices with those properties.

Proposition 21. *The graph H and its dual graph both have a Hamiltonian cycle, but $G(L(H))$ does not admit a rhombic strip.*

Proof. Up to symmetry, H has two distinct Hamiltonian cycles, $C_1 = (v_1, v_2, v_3, v_4, v_5, v_6, v_7)$ and $C_2 = (v_1, v_6, v_2, v_3, v_4, v_5, v_7)$; see Figure 32 (b)+(c). A Hamiltonian cycle in the dual graph of H is shown in Figure 32 (b). For the cycle C_1 , the C_1 -chords (v_1, v_6) , (v_6, v_3) , (v_3, v_1) violate condition (ii) in Theorem 19. Similarly, for the cycle C_2 , the C_2 -chords (v_3, v_5) , (v_5, v_6) , (v_6, v_3) violate this condition. Consequently, $G(L(H))$ does not admit a rhombic strip. $\square$




FIGURE 32. The smallest plane graph H such that both H and its dual graph are Hamiltonian, but $G(L(H))$ does not admit a rhombic strip. The C_i -chords violating condition (ii) in Theorem 19 are dashed.

10. GRAPH ASSOCIAHEDRA OF CHORDAL GRAPHS

10.1. The polytope and its face lattice. Let $H = (V, E)$ be a graph. A *tube* T of H is a nonempty subset of vertices $T \subset V$, such that the induced subgraph $H[T]$ is connected. Two tubes T and T' are *compatible* if they are either nested, i.e., $T \subset T'$ or $T \supset T'$, or non-adjacent, i.e., $H[T \cup T']$ is not connected. A *tubing* of H is a family of pairwise compatible tubes that includes the vertex sets of each connected component of H . The *graph associahedron* $A(H)$ of H is the polytope whose face lattice is isomorphic to the reverse inclusion order of all tubings of H . In particular, the vertices of $A(H)$ are the inclusion-maximal tubings of H (where the number of tubes equals the number of vertices of H), and they are in one-to-one correspondence with elimination forests of H .

10.2. Hamiltonian cycles in the face lattice. The following is the main result of this section. A graph is *chordal* if it has no induced cycles of length at least 4.

Theorem 22. *For any chordal graph H with at least one edge, the graph $G(L(A(H)))$ has a Hamiltonian cycle.*

If H is a perfect matching with n edges, then $A(H)$ is the hypercube Q_n . If H is a complete graph on n vertices, then $A(H)$ is the permutahedron Π_n . If H is a path on n vertices, then $A(H)$ is the (standard) associahedron A_{n+2} . All these graphs are chordal, and therefore Theorem 22 generalizes Theorems 6, 8 and 12 presented before.

For a vertex v of a graph H we write $H - v$ for the graph obtained by deleting v (and all incident edges) from H . A *clique* in H is an induced subgraph of H that is complete. A well-known characterization of chordal graphs that we will be using is that they admit a perfect elimination order (PEO). This is a total ordering of the vertices of H such that every vertex induces a clique with the neighboring vertices that come before it in the ordering. Formally, a *PEO graph* $H = ([n], E)$ is one that satisfies one of the following two recursive conditions: $n = 1$, i.e., H is a single vertex graph; or $n > 1$, the graph $H - n$ is a PEO graph and the induced subgraph of the vertex n and its neighbors is a clique in H .

The following lemma is an immediate consequence of the definition of compatible tubes.

Lemma 23. *Let H be a graph, let C be a clique in H , let $\mathcal{T}$ be a tubing of H , and let $T, T' \in \mathcal{T}$ be two distinct tubes with $C \cap T \neq \emptyset$ and $C \cap T' \neq \emptyset$. Then T and T' are nested, i.e., we have $T \subset T'$ or $T \supset T'$.*

Proof of Theorem 22. We consider H as a PEO graph $H = ([n], E)$, and we assume w.l.o.g. that if it has exactly one edge and $n > 2$, then the vertex n is not incident to this edge.

The following definitions are illustrated in Figure 33. Let $\mathcal{T}$ be a tubing of $H - n$, and let $T_1, T_2, \dots, T_k$ be the tubes of $\mathcal{T}$ containing at least one neighbor of n in H . By Lemma 23 they are nested, so we can assume that $T_1 \subset T_2 \subset \dots \subset T_k$. Furthermore, let $\mathcal{R} := \mathcal{T} \setminus \{T_1, \dots, T_k\}$ be the remaining set of tubes that do not contain any neighbor of n .

If n is an isolated vertex of H , then we have $\mathcal{R} = \mathcal{T}$ and $k = 0$, and we define $\check{c}_0(\mathcal{T}) := \mathcal{R} \cup \{\{n\}\} = \mathcal{T} \cup \{\{n\}\}$. Otherwise, if n is not an isolated vertex, we define

$$\check{c}_i(\mathcal{T}) := \mathcal{R} \cup \bigcup_{j=1}^i \{T_j\} \cup \bigcup_{j=i}^k \{T_j \cup \{n\}\}, \quad i = 0, \dots, k,$$

$$\hat{c}_i(\mathcal{T}) := \mathcal{R} \cup \bigcup_{j=1}^{i-1} \{T_j\} \cup \bigcup_{j=i}^k \{T_j \cup \{n\}\}, \quad i = 1, \dots, k,$$

where $T_0 \cup \{n\}$ in the equation for $\check{c}_0(\mathcal{T})$ should be interpreted as $\{n\}$.




FIGURE 33. Illustration of the proof of Theorem 22. The vertex n and all tubes containing it are highlighted. The bold zigzag path shows edges in $G(L(A(H)))$, labeled by the tube that is added/removed along this edge.

Note that both $\check{c}_i(\mathcal{T})$ and $\hat{c}_i(\mathcal{T})$ are tubings of H , where the latter has the same number of tubes as $\mathcal{T}$, and the former has one more tube than $\mathcal{T}$. We define the sequence $\vec{c}(\mathcal{T})$ of tubings of H by

$$\vec{c}(\mathcal{T}) := (\check{c}_0(\mathcal{T}), \hat{c}_1(\mathcal{T}), \check{c}_1(\mathcal{T}), \hat{c}_2(\mathcal{T}), \check{c}_2(\mathcal{T}), \dots, \hat{c}_{k-1}(\mathcal{T}), \check{c}_{k-1}(\mathcal{T}), \hat{c}_k(\mathcal{T}), \check{c}_k(\mathcal{T})),$$

we note that it has length $2k + 1$ and describes a path in $G(L(A(H)))$ that alternates between ranks r and $r + 1$, where r is the rank of $\mathcal{T}$ in $L(A(H - n))$; see Figure 33. We write $\overleftarrow{c}(X) := \text{rev}(\overrightarrow{c}(X))$ for the reverse sequence/path.

To prove the theorem, we construct a path P_n in $G(L(A(H)))$ that visits all faces of $A(H)$ except $\emptyset$ and that starts and ends at rank 0 faces (i.e., two inclusion-maximal tubings of H with n tubes each), and therefore $(P_n, \emptyset)$ is the desired Hamiltonian cycle in $G(L(A(H)))$.

The path P_n is constructed by induction on n as follows: For the base case of the construction where H is a single edge connecting vertices 1 and 2 we take $P_2 := \{\{1\}, \{1, 2\}\}, \{\{1, 2\}\}, \{\{1, 2\}, \{2\}\}$.

For the induction step, let $P_n = (\mathcal{T}_1, \dots, \mathcal{T}_N)$ be the path in $G(L(A(H - n)))$. Then we define

$$P_{n+1} := (\overleftarrow{c}(\mathcal{T}_1), \overrightarrow{c}(\mathcal{T}_2), \overleftarrow{c}(\mathcal{T}_3), \overrightarrow{c}(\mathcal{T}_4), \dots, \overleftarrow{c}(\mathcal{T}_N)).$$

It can be checked straightforwardly that the path P_{n+1} has the required properties. The main observation is that for every tubing $\mathcal{T}$ of H , there is a uniquely defined tubing $\mathcal{T}'$ of $H - n$ and a unique integer i such that $\mathcal{T} = \check{c}_i(\mathcal{T}')$ or $\mathcal{T} = \hat{c}_i(\mathcal{T}')$. $\square$

11. QUOTIENTOPES

11.1. The polytopes and Hamiltonian cycles in the face lattice. We start by extending some of the terminology and notation introduced in Section 2.

Given a lattice $(P, <)$, a *lattice congruence* is an equivalence relation $\equiv$ on P that preserves the join and meet operations, i.e., if $x \equiv x'$ and $y \equiv y'$, then $x \vee y \equiv x' \vee y'$ and $x \wedge y \equiv x' \wedge y'$. For any $x \in P$ we write $[x] := \{y \in P \mid x \equiv y\}$ for the equivalence class of x , and we define the set of all equivalence classes as $P/\equiv := \{[x] \mid x \in P\}$. It is well-known and easy to show that every equivalence class is an interval. This set inherits a lattice structure, where for any $X, Y \in P/\equiv$, we define $X < Y$ if there exist elements $x \in X$ and $y \in Y$ such that $x < y$. We refer to $P/\equiv$ with this lattice structure as *quotient lattice*.

In this paper we focus on lattice congruences of the so-called weak order on permutations. Formally, the *inversion set* of a permutation $\pi = a_1 \cdots a_n \in S_n$ is defined as

$$\text{inv}(\pi) := \{(a_i, a_j) \mid 1 \leq i < j \leq n \mid a_i > a_j\},$$

and the *weak order* on S_n is defined by $\pi < \rho$ if and only if $\text{inv}(\pi) \subset \text{inv}(\rho)$; see the top part of Figure 34. The weak order is a lattice, namely, given two permutations π and ρ , their meet $\sigma := \pi \wedge \rho$ is given by $\text{inv}(\sigma) = \text{inv}(\pi) \cap \text{inv}(\rho)$, and their join $\sigma' := \pi \vee \rho$ has as its inversion set $\text{inv}(\sigma')$ the transitive closure of $\text{inv}(\pi) \cup \text{inv}(\rho)$. Cover relations in the weak order are adjacent transpositions, so the cover graph of the weak order on S_n is exactly the skeleton of the permutahedron Π_n . More generally, Pilaud and Santos [PS19] proved that for any lattice congruence $\equiv$ of the weak order, the cover graph of $S_n/\equiv$ can be realized geometrically as the skeleton of a polytope, called *quotientope* (see also [PPR23]), which we denote by $\Pi_n/\equiv$. Via suitable congruences, quotientopes generalize permutahedra, associahedra and hypercubes, and many other combinatorial polytopes, including permutreehedra [PP18] and rectangulotopes [CP25].

The following is the main result of this section.

Theorem 24. *For any $n \geq 1$ and any lattice congruence $\equiv$ of the weak order on S_n , the graph $G(L(\Pi_n/\equiv))$ has a Hamiltonian cycle, unless it is a single edge.*

The case that $G(L(\Pi_n/\equiv))$ is a single edge occurs if and only if $\equiv$ has only one equivalence class (in particular, if $n = 1$).




FIGURE 34. Top: The weak order on S_n for $n = 2, 3, 4$, together with a lattice congruence, namely the sylvester congruence. The bold edges indicate pairs of permutations that are equivalent. Bottom: The facial weak order on P_n for $n = 2, 3, 4$, with the extensions of the lattice congruences from the top. Set partitions corresponding to vertices, edges, 2- and 3-dimensional faces are indicated by white, red, blue and green bullets, respectively.

11.2. The face lattice of the permutahedron. Recall from Section 5 that the face lattice $L(\Pi_n)$ of the permutahedron Π_n can be modelled as the refinement order on set partitions of $[n]$. Specifically, we denote all set partitions of $[n]$ by P_n , and we write $\subset$ for the refinement order on P_n , which corresponds to the inclusion order on the faces of Π_n ; see Figure 35. I.e., $x \subset y$ for $x, y \in P_n$ if and only if x is obtained from y by splitting a block into two/adding one bar, or equivalently, y is obtained from x by taking the unions of two blocks/removing one bar. For example, we have $2|14|3|5 \subset 124|3|5 \subset 124|35$ and $2|14|3|5 \not\subset 14|2|35$. Set partitions with n blocks can be identified naturally with permutations in S_n .

11.3. The facial weak order. Finding an ‘explicit’ combinatorial model of the faces of quotientopes is an open problem in general. For the interesting special case of simple quotientopes, whose cover graphs are regular, the faces can be modeled by Schröder separating trees [BNP26].




FIGURE 35. The face lattices of the 2- and 3-dimensional permutahedron, namely the refinement orders on set partitions of $[n]$, for $n = 3$ and $n = 4$, respectively, plus the trivial face $\emptyset$ at the bottom. The gray bubbles are the equivalence classes of the lattice congruences shown in Figure 34.

In our paper we use a representation of the faces as equivalence classes of a lattice congruence of an order on the set P_n that is known as *facial weak order* [KLN⁺01, DHP18, DHMP22], which is different from the refinement order discussed in the previous section.

To define the facial weak order, we need some definitions; see the bottom part of Figure 34. We say that two blocks $A, B \subseteq [n]$ of a set partition are *separated* if $\max A < \min B$ or $\min A > \max B$. In the first case we write $A < B$ and in the second case $A > B$. If $A < B$, then we refer to A as the *smaller* block and to B as the *larger* block. To define the facial weak order, we start with the weak order on S_n , i.e., permutations ordered by containment of their inversion sets. In terms of set partitions, the cover relations are $\pi < \rho$ for partitions $\pi = a_1 | \dots | a_i | a_{i+1} | \dots | a_n \in S_n \subseteq P_n$ and $\rho = a_1 | \dots | a_{i-1} | a_{i+1} | a_i | a_{i+2} | \dots | a_n \in S_n \subseteq P_n$ where $a_i < a_{i+1}$. We extend this order to the entire set P_n by repeatedly adding, for every pair of set partitions $x = A_1 | \dots | A_i | A_{i+1} | \dots | A_k \in P_n$ and $y = A_1 | \dots | A_{i+1} | A_i | \dots | A_k \in P_n$ that have a separated pair of consecutive blocks A_i and A_{i+1} such that $A_i < A_{i+1}$, the comparabilities $x < z < y$ for $z := A_1 | \dots | A_i \cup A_{i+1} | \dots | A_k$. In the first round, all set partitions with $n - 1$ blocks are added (they encode edges of the permutahedron), and more generally, in the r th round all set partitions with $n - r$ blocks are added (they encode r -dimensional faces of the permutahedron). Note that the cover relations in the facial weak order also correspond to

adding/removing a bar in a set partition, but the order is different from the refinement order discussed in the previous section. Furthermore, while in the face lattice of the permutahedron, every addition/removal of a bar is a cover relation, in the facial weak order, a bar can only be added to a block so as to create two separated smaller blocks, and we can only remove a bar between two separated blocks. In the first case, we go down in the order when the smaller block is placed left of the added bar and the larger block is placed right of the bar, and we go up when placing them in the opposite order. In the second case, we go up in the order when the block left of the removed bar is smaller than the block right of the bar, and we go down when they are in the opposite order. For example, we have $12|34 \prec 1234 \prec 34|12$ in the facial weak order, and $12|34, 34|12 \subset 1234$ in the refinement order. On the other hand, while $x := 13|24 \subset 1234 =: y$ is a cover relation in the refinement order, x and y are incomparable in the facial weak order. In other words, the cover graph of the facial weak order is a (proper) spanning subgraph of the cover graph of the refinement order. The facial weak order is known to be a lattice [KLN⁺01].

The key observation is that every face of the permutahedron Π_n corresponds to an interval in the weak order on S_n , more precisely, that the set of vertices (permutations) belonging to this face forms an interval in the weak order. Namely, for a face of Π_n represented by the set partition $x = A_1 | \cdots | A_k \in P_n$, we let $x_0 \in S_n$ and $x_1 \in S_n$ be the permutations obtained from x by sorting the entries of each block A_i , $i \in [k]$, increasingly or decreasingly, respectively. Then the face x of Π_n contains precisely the permutations of the interval $[x_0, x_1]$ in the weak order on S_n .

Using this observation, every lattice congruence of the weak order on S_n can be extended in a natural way to a lattice congruence of the facial weak order on P_n . Specifically, given an equivalence relation $\equiv$ on the set S_n , we define an equivalence relation $\equiv^+$ on the set P_n as follows. For two set partitions $x, y \in P_n$ (representing faces of Π_n), we consider the intervals $[x_0, x_1]$ and $[y_0, y_1]$ in the weak order on S_n , and we define $x \equiv^+ y$ if and only if $x_0 \equiv y_0$ and $x_1 \equiv y_1$.

The following lemma provides an alternative, but equivalent, definition of this notion of extension.

Lemma 25. *For any two set partitions $x, y \in P_n$ we have $x \equiv^+ y$ if and only if for every $x' \in [x_0, x_1]$ there exists a $y' \in [y_0, y_1]$ such that $x' \equiv y'$, and for every $y' \in [y_0, y_1]$ there exists an $x' \in [x_0, x_1]$ such that $x' \equiv y'$.*

Proof. Let $x, y \in P_n$ be such that $x \equiv^+ y$, i.e., we have $x_0 \equiv y_0$ and $x_1 \equiv y_1$. Furthermore, consider an arbitrary element $x' \in [x_0, x_1]$. Since $x_0 \leq x' \leq x_1$, we clearly have $x' = x_0 \vee x'$ and $x' = x' \wedge x_1$. By the definition of lattice congruence, $\equiv$ preserves join and meet operations, and thus we have

$$x' = x_0 \vee x' = x_0 \vee (x' \wedge x_1) \equiv y_0 \vee (x' \wedge y_1) =: y'.$$

Note that $y_0 \leq y_1$ and $x' \wedge y_1 \leq y_1$, implying that $y' \leq y_1$. Furthermore, we have $y_0 \leq y'$ since y' is the join of y_0 and some other element. Combining these observations proves that $y' \in [y_0, y_1]$, as desired. The same argument also works with reversing the roles for x' and y' , so one direction of the equivalence is proved.

To prove the other direction, we assume that for every $x' \in [x_0, x_1]$ there exists a $y' \in [y_0, y_1]$ such that $x' \equiv y'$, and vice versa. Let $x' \in [x_0, x_1]$ and $y' \in [y_0, y_1]$ be such that $x' \equiv y_0$ $y' \equiv x_0$. Using that $\equiv$ preserves the meet operation, we get $x_0 = x_0 \wedge x' \equiv y' \wedge y_0 = y_0$, yielding $x_0 \equiv y_0$. An analogous argument with meets replaced by joins shows that $x_1 \equiv y_1$, proving the other direction of the equivalence. $\square$

Theorem 26 ([DHP18, Cor. 4.13+Thm. 4.22]). *For any $n \geq 1$ and any lattice congruence $\equiv$ of the weak order on S_n , its extension $\equiv^+$ is a lattice congruence of the facial weak order on P_n .*

The set $P_n/\equiv^+$ is in bijection with the set of faces of the quotientope $\Pi_n/\equiv$ (excluding the trivial face $\emptyset$).

Note, however, that not every lattice congruence of the facial weak order is an extension of a lattice congruence of the weak order.

In view of Theorem 26, our goal is to list equivalence classes of lattice congruences of the facial weak order on P_n . In fact, the following Hamiltonicity result is valid for *any* lattice congruence $\equiv$ of the facial weak order on P_n (in particular for those obtained by extending a congruence of the weak order on S_n). We refer to a path in the cover graph of $P_n/\equiv$ as *grounded*, if it starts and ends with an equivalence class that contains a permutation.

Theorem 27. *For any $n \geq 1$ and any lattice congruence $\equiv$ of the facial weak order on P_n , the cover graph of the lattice quotient $P_n/\equiv$ has a Hamiltonian path. If the congruence extends a lattice congruence of the weak order on S_n , then the path is grounded.*

11.4. Proofs of Theorems 27 and 24. The facial weak order on P_n has a simple recursive structure. Specifically, for a set partition $x = A_1 | \cdots | A_k \in P_{n-1}$, we define

$$\begin{aligned} \check{c}_i(x) &:= A_1 | \cdots | A_i | \{n\} | A_{i+1} | \cdots | A_k & \text{for } i = 0, \dots, k, \\ \hat{c}_i(x) &:= A_1 | \cdots | A_{i-1} | A_i \cup \{n\} | A_{i+1} | \cdots | A_k & \text{for } i = 1, \dots, k, \end{aligned}$$

and

$$\vec{c}(x) := (\check{c}_0(x), \hat{c}_1(x), \check{c}_1(x), \hat{c}_2(x), \check{c}_2(x), \dots, \hat{c}_{k-1}(x), \check{c}_{k-1}(x), \hat{c}_k(x), \check{c}_k(x))$$

and $\overleftarrow{c}(x) := \text{rev}(\vec{c}(x))$. The sequence $\overleftarrow{c}(x)$ is a chain in the facial weak order on P_n . Furthermore, consider set partitions $x, y \in P_{n-1}$ with $x = A_1 | \cdots | A_k$ and $y = A_1 | \cdots | A_i \cup A_{i+1} | \cdots | A_k$. Then




FIGURE 36. Recursive structure of the facial weak order.

the two chains $\overleftarrow{c}(x)$ and $\overleftarrow{c}(y)$ have exactly the following ‘ladder-like’ cover relations; see Figure 36. If $x < y$, i.e., if $A_i < A_{i+1}$, then we have

$$\begin{aligned} \check{c}_j(x) &\leq \check{c}_j(y) \text{ for } j = 0, \dots, i-1, \text{ and } \check{c}_j(x) < \check{c}_{j-1}(y) \text{ for } j = i+1, \dots, k, \\ \hat{c}_j(x) &\leq \hat{c}_j(y) \text{ for } j = 1, \dots, i-1, \text{ and } \hat{c}_j(x) < \hat{c}_{j-1}(y) \text{ for } j = i+1, \dots, k. \end{aligned} \quad (12a)$$

On the other hand, if $y < x$, i.e., if $A_i > A_{i+1}$, then we have

$$\begin{aligned} \check{c}_j(y) &< \check{c}_j(x) \text{ for } j = 0, \dots, i-1, \text{ and } \check{c}_{j-1}(y) < \check{c}_j(x) \text{ for } j = i+1, \dots, k, \\ \hat{c}_j(y) &< \hat{c}_j(x) \text{ for } j = 1, \dots, i, \quad \text{and } \hat{c}_{j-1}(y) < \hat{c}_j(x) \text{ for } j = i+2, \dots, k. \end{aligned} \quad (12b)$$

For any set partition $x \in P_n$, we write $p(x)$ for the partition obtained from x by removing the element n from the block that contains it. If x has $\{n\}$ as a singleton block, then the block is removed entirely. Clearly, we have $p(\check{c}_i(x)) = x$ and $p(\hat{c}_i(x)) = x$.

Proposition 28. *For any $n \geq 1$ and any lattice congruence $\equiv$ of the facial weak order on P_n , we have the following:*

- (i) *The restrictions $\equiv^* := \{(x, y) \mid x, y \in P_{n-1} \text{ and } x|_{\{n\}} \equiv y|_{\{n\}}\}$ and $\equiv' := \{(\pi, \rho) \mid \pi, \rho \in S_{n-1} \text{ and } \pi|_{\{n\}} \equiv \rho|_{\{n\}}\}$ are lattice congruences of the facial weak order on P_{n-1} and the weak order on S_{n-1} , respectively, and the former extends the latter.*
- (ii) *For any equivalence class $X \in P_n/\equiv$, the projection $p(X) := \{p(x) \mid x \in X\}$ is an equivalence class of $\equiv^*$.*
- (iii) *For any $x = A_1 | \dots | A_k \in P_{n-1}$ we define $\overleftarrow{c}(x) := (c_1(x), \dots, c_{2k+1}(x))$ and $\lambda(x) := |\{[c_i(x)] \mid i = 1, \dots, 2k+1\}|$, which is the number of equivalence classes intersected by the chain $\overleftarrow{c}(x)$. The quantity $\lambda(x)$ is either equal to 1 for all $x \in P_{n-1}$ or strictly larger than 1 for all $x \in P_{n-1}$.*
- (iv) *If the cover graph of $P_{n-1}/\equiv^*$ has a grounded Hamiltonian path H^* , then the cover graph of $P_n/\equiv$ has grounded a Hamiltonian path H .*

Proof. The proofs of (i)–(iii) are straightforward calculations using the basic definitions of lattice congruences, exploiting the ladder structure of the facial weak order captured by (12) and shown in Figure 36; for details, see [BCM⁺26, Prop. 16 (i)–(iii)].

It remains to prove (iv). If $\lambda(x) = 1$ for all $x \in P_{n-1}$, then $P_n/\equiv$ and $P_{n-1}/\equiv^*$ are isomorphic, so the claim is trivial. Otherwise, by (iii) we have $\lambda(x) \geq 2$ for all $x \in P_{n-1}$, i.e., each of the chains $\overleftarrow{c}(x)$ intersects at least two equivalence classes in $P_n/\equiv$. Given a Hamiltonian path $H^* = ([x_1], \dots, [x_N])$ of the cover graph of $P_{n-1}/\equiv^*$, where $x_i \in P_{n-1}$ and $N = |P_{n-1}/\equiv^*|$, we construct a Hamiltonian path H of the cover graph of $P_n/\equiv$ as follows; see Figure 37: For each x_i , $i \in [N]$, the chain $\overleftarrow{c}(x_i)$ intersects $\lambda(x_i) \geq 2$ many equivalence classes of the congruence $\equiv$, and by (ii) those equivalence classes are independent of the representative x_i of the class $[x_i]$ of the congruence $\equiv^*$. Thus we let $\overleftarrow{c}'(x_i)$ be a subsequence of $\overleftarrow{c}(x_i)$, chosen so that it contains exactly one representative from each of the intersected equivalence classes of $\equiv$. Since $\lambda(x_i) \geq 2$ we may choose the subsequence $\overleftarrow{c}'(x_i)$ so that it contains the first and last entry of $\overleftarrow{c}(x_i)$, namely the set partitions $x_i|_{\{n\}}$ and $\{n\}|x_i$. We also define $\overrightarrow{c}'(x_i) := \text{rev}(\overleftarrow{c}'(x_i))$. We then define

$$H' := (\overleftarrow{c}'(x_1), \overrightarrow{c}'(x_2), \overleftarrow{c}'(x_3), \overrightarrow{c}'(x_4), \dots)$$

as a sequence of representatives of all equivalence classes in $P_n/\equiv$, and we let H be the corresponding sequence of equivalence classes (apply $[\]$ to each entry of H'). By construction, H is a Hamiltonian path in the cover graph of $P_n/\equiv$.

It remains to show that H is grounded, using the assumption that H^* is grounded. We can assume that x_1 and x_N are permutations, i.e., $x_1, x_N \in S_{n-1}$. The first entry of H' is the first




FIGURE 37. (Grounded) Hamiltonian paths in the quotient lattices of the facial weak order for the lattice congruences from Figure 34, constructed as described in the proof of Proposition 28 (iv).

entry of $\overleftarrow{c}'(x_1)$, which is $x_1|n \in S_n$, so the first entry of H indeed contains a permutation. Furthermore, the last entry of H' is the first or last entry of $\overleftarrow{c}'(x_N)$, which is $x_N|n \in S_n$ or $n|x_N \in S_n$, respectively (the first case occurs if N is even and the second case occurs if N is odd), so the last entry of H indeed contains a permutation.

This completes the proof. $\square$

Proof of Theorem 27. The statement follows by induction on $n \geq 1$, using Proposition 28 (i)+(iv). $\square$

Proof of Theorem 24. Let $\equiv$ be any lattice congruence of the weak order on S_n , and let $\equiv^+$ be its extension to a lattice congruence of the facial weak order on P_n . By Theorem 26, the equivalence classes of $P_n/\equiv^+$ are in bijection with the faces of the quotientopes $\Pi_n/\equiv$, so the Hamiltonian path H in the cover graph of $P_n/\equiv^+$ obtained from Theorem 27 visits every face of the quotientope $\Pi_n/\equiv$ exactly once. Furthermore, since cover relations in the facial weak order on P_n are also cover relations in the refinement order on P_n , which is the face lattice of Π_n minus the trivial face $\emptyset$, we obtain a listing of the faces of $\Pi_n/\equiv$ along cover relations of $G(L(\Pi_n/\equiv))$. Furthermore, as the path H is grounded, the listing starts and ends at a rank 0 face, i.e., a permutation. Adding the trivial face $\emptyset$ thus yields the desired Hamiltonian cycle in $G(L(\Pi_n/\equiv))$. $\square$

12. OUTLOOK AND OPEN QUESTIONS

- Our proof of Conjecture 1 for 3-dimensional polytopes (Theorem 18) essentially uses *shellings* (see Lemma 17), i.e., a particular order to build the polytope by adding the facets one after the other, while maintaining some invariant on a Hamiltonian cycle in the face lattice of the polytopal complex constructed so far. This suggests a general strategy for tackling Conjecture 1, namely to combine a particularly nice shelling order with a particular invariant for the Hamiltonian cycle. Unfortunately, so far we were not able to push this idea beyond

the 3-dimensional case, mainly due to our limited knowledge about shellings. In view of our results regarding special classes of polytopes, worthwhile next targets might be 4-dimensional polytopes, 0/1-polytopes (for example, the uniform matroid polytope), or graph associahedra of non-chordal graphs.

- Using the techniques developed in [BCM⁺26], the methods presented in Section 11 can be applied more generally to quotientopes arising from any supersolvable hyperplane arrangement (such as the braid arrangement that gives rise to the permutahedron). This proves in particular that the face lattice of any of the type B quotientopes presented in [PPR23] admits a Hamiltonian cycle. This includes type B associahedra, which are also known as cyclohedra, i.e., graph associahedra of cycles.
- A necessary condition for the existence of a Hamiltonian cycle in a bipartite graph is the existence of a perfect matching. Is it true that the cover graph of the face lattice of every polytope admits a perfect matching? If yes, this can be seen as further evidence for Conjecture 1. If no, then this would be a counterexample to the conjecture.
- The face lattice of the 3-dimensional permutahedron Π_4 has a rhombic strip; see Figure 28. Does the cover graph of the face lattice of the permutahedron $G(L(\Pi_n))$ admit a rhombic strip for $n \geq 5$?
- The face lattice of the 3-dimensional associahedron A_6 has a rhombic strip; see Figure 29. Does the cover graph of the face lattice of the associahedron $G(L(A_n))$ admit a rhombic strip for $n \geq 7$?

Huemer, Hurtado, Noy and Omaña-Pulido [HHNOP09] constructed a Hamiltonian cycle in a supergraph G_n^+ of $G(L(A_n))$, obtained by adding edges between pairs of dissections at rank k that differ in removing an inner edge and replacing it by another edge inside the subpolygon that is the union of the two polygons on both sides of the removed edge. We note that in $G(L(A_n))$ such pairs of dissections at rank k always have a common neighbor at rank $k+1$, but not necessarily a common neighbor at rank $k-1$. Furthermore, not all pairs of dissections at rank k in $G(L(A_n))$ that have a common neighbor on rank $k+1$ are connected by an edge in G_n^+ . The authors also considered the subgraphs of G_n^+ obtained by restricting to fixed rank k , and provided Hamiltonian cycles for them, but this does not solve our problem.

- As a partial step towards the previous two questions about the permutahedron and associahedron: Are there listings of the faces of fixed rank k , such that any two consecutive faces have a common superface at rank $k+1$ and a common subface at rank $k-1$?

ACKNOWLEDGEMENTS

We thank the reviewers of this paper for numerous helpful comments and suggestions that greatly improved the presentation. In particular, we thank Vincent Pilaud for suggesting the facial weak order as a convenient tool for studying face lattices of quotientopes, and for his enthusiasm to share his knowledge and ideas on how to attack Conjecture 1 with us. Lastly, we thank Christoph Hertrich for suggesting cyclic polytopes as an interesting family of polytopes to investigate.

REFERENCES

- [ACF⁺25] H. Akitaya, J. Cardinal, S. Felsner, L. Kleist, and R. Lauff. Facet-Hamiltonicity. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2025)*, pages 5051–5064. SIAM, Philadelphia, PA, 2025.
- [BBM⁺26] N. Behrooznia, S. Brenner, A. Merino, T. Mütze, C. Rieck, and F. Verciani. Listing faces of polytopes. In *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 6212–6222. SIAM, Philadelphia, PA, 2026.

- [BCM⁺26] S. Brenner, J. Cardinal, T. McConville, A. Merino, and T. Mütze. Combinatorial generation via permutation languages. VII. Supersolvable hyperplane arrangements. *European J. Combin.*, 135:Paper No. 104367, 39 pp., 2026.
- [Ber22] B. A. Berendsohn. The diameter of caterpillar associahedra. In A. Czumaj and Q. Xin, editors, *18th Scandinavian Symposium and Workshops on Algorithm Theory, SWAT 2022, June 27-29, 2022, Tórshavn, Faroe Islands*, volume 227 of *LIPIcs*, pages 14:1–14:12. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
- [BNP26] E. Barnard, J.-C. Novelli, and V. Pilaud. Separating trees and simple congruences of the weak order. *European J. Combin.*, 136:Paper No. 104375, 25 pp., 2026.
- [CD06] M. P. Carr and S. L. Devadoss. Coxeter complexes and graph-associahedra. *Topology Appl.*, 153(12):2155–2168, 2006.
- [CHM⁺23] J. Cardinal, H. P. Hoang, A. Merino, O. Mička, and T. Mütze. Combinatorial generation via permutation languages. V. Acyclic orientations. *SIAM J. Discrete Math.*, 37(3):1509–1547, 2023.
- [CLPL18] J. Cardinal, S. Langerman, and P. Pérez-Lantero. On the diameter of tree associahedra. *Electron. J. Combin.*, 25(4):Paper No. 4.18, 13 pp., 2018.
- [CMM25] J. Cardinal, A. Merino, and T. Mütze. Combinatorial generation via permutation languages. IV. elimination trees. *ACM Trans. Algorithms*, 21(1):13:1–13:41, 2025.
- [cosa] The Combinatorial Object Server: Generate integer and set partitions. <http://www.combos.org/part>.
- [cosb] The Combinatorial Object Server: Generate k-ary trees and dissections. <http://www.combos.org/kary>.
- [cosc] The Combinatorial Object Server: Generate faces of cyclic polytopes. <http://www.combos.org/cyclic>.
- [cosd] The Combinatorial Object Server: Generate signed permutations. <http://www.combos.org/sperm>.
- [CP25] J. Cardinal and V. Pilaud. Rectangulotopes. *European J. Combin.*, 125:Paper No. 104090, 25 pp., 2025.
- [CPVP21] J. Cardinal, L. Pournin, and M. Valencia-Pabon. Bounds on the diameter of graph associahedra. *Procedia Comput. Sci.*, 195:239–247, 2021.
- [CPVP24] J. Cardinal, L. Pournin, and M. Valencia-Pabon. The rotation distance of brooms. *European J. Combin.*, 118:Paper No. 103877, 16 pp., 2024.
- [CS25a] J. Cardinal and R. Steiner. Inapproximability of shortest paths on perfect matching polytopes. *Math. Program.*, 210(1-2, Ser. B):147–163, 2025.
- [CS25b] J. Cardinal and R. Steiner. Shortest paths on polymatroids and hypergraphic polytopes. *Comb. Theory*, 5(3):Paper No. 3, 29 pp., 2025.
- [CSJ09] S. Cleary and K. St. John. Rotation distance is fixed-parameter tractable. *Inform. Process. Lett.*, 109(16):918–922, 2009.
- [CSJ10] S. Cleary and K. St. John. A linear-time approximation for rotation distance. *J. Graph Algorithms Appl.*, 14(2):385–390, 2010.
- [CSSV25] L. F. I. Cunha, I. Sau, U. S. Souza, and M. Valencia-Pabon. Computing distances on graph associahedra is fixed-parameter tractable. In K. Censor-Hillel, F. Grandoni, J. Ouaknine, and G. Puppis, editors, *52nd International Colloquium on Automata, Languages, and Programming, ICALP 2025, July 8-11, 2025, Aarhus, Denmark*, volume 334 of *LIPIcs*, pages 63:1–63:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025.
- [CSSV26] L. F. I. Cunha, I. Sau, U. S. Souza, and M. Valencia-Pabon. Computing distances is FPT on graph associahedra and $W[2]$ -hard on hypergraphic polytopes, 2026. <https://arxiv.org/abs/2504.18338>.
- [Dev09] S. L. Devadoss. A realization of graph associahedra. *Discrete Math.*, 309(1):271–276, 2009.
- [DHMP22] A. Dermenjian, C. Hohlweg, T. McConville, and V. Pilaud. The facial weak order on hyperplane arrangements. *Discrete Comput. Geom.*, 67(1):166–202, 2022.
- [DHP18] A. Dermenjian, C. Hohlweg, and V. Pilaud. The facial weak order and its lattice quotients. *Trans. Amer. Math. Soc.*, 370(2):1469–1507, 2018.
- [Dor26] J. Dorfer. Flip distance of triangulations of convex polygons / rotation distance of binary trees is NP-complete, 2026. <https://arxiv.org/abs/2602.22874>.
- [FR94] K. Fukuda and V. Rosta. Combinatorial face enumeration in convex polytopes. *Comput. Geom.*, 4(4):191–198, 1994.

- [GMN24] P. Gregor, T. Mütze, and Namrata. Combinatorial generation via permutation languages. VI. Binary trees. *European J. Combin.*, 122:Paper No. 104020, 40 pp., 2024.
- [HHMW22] E. Hartung, H. P. Hoang, T. Mütze, and A. Williams. Combinatorial generation via permutation languages. I. Fundamentals. *Trans. Amer. Math. Soc.*, 375(4):2255–2291, 2022.
- [HHNOP09] C. Huemer, F. Hurtado, M. Noy, and E. Omafia-Pulido. Gray codes for non-crossing partitions and dissections of a convex polygon. *Discrete Appl. Math.*, 157(7):1509–1520, 2009.
- [HM21] H. P. Hoang and T. Mütze. Combinatorial generation via permutation languages. II. Lattice congruences. *Israel J. Math.*, 244(1):359–417, 2021.
- [HN99] F. Hurtado and M. Noy. Graph of triangulations of a convex polygon and tree of triangulations. *Comput. Geom.*, 13(3):179–188, 1999.
- [IKK⁺22] T. Ito, N. Kakimura, N. Kamiyama, Y. Kobayashi, and Y. Okamoto. Shortest reconfiguration of perfect matchings via alternating cycles. *SIAM J. Discrete Math.*, 36(2):1102–1123, 2022.
- [Joh63] S. Johnson. Generation of permutations by adjacent transposition. *Math. Comp.*, 17:282–285, 1963.
- [KLN⁺01] D. Krob, M. Latapy, J.-C. Novelli, H.-D. Phan, and S. Schwer. Pseudo-permutations I: First combinatorial and lattice properties. In *Proceedings of the 13th International Conference on Formal Power Series and Algebraic Combinatorics (FPSAC 2001)*, 2001.
- [KP02] V. Kaibel and M. E. Pfetsch. Computing the face lattice of a polytope from its vertex-facet incidences. *Comput. Geom.*, 23(3):281–290, 2002.
- [LRR93] J. M. Lucas, D. Roelants van Baronaigien, and F. Ruskey. On rotations and the generation of binary trees. *J. Algorithms*, 15(3):343–366, 1993.
- [LX25] H. Li and G. Xia. An $\mathcal{O}(3.82^k)$ time FPT algorithm for convex flip distance. *Discrete Comput. Geom.*, 73(3):702–718, 2025.
- [MM23] A. Merino and T. Mütze. Combinatorial generation via permutation languages. III. Rectangulations. *Discrete Comput. Geom.*, 70(1):51–122, 2023.
- [MP15] T. Manneville and V. Pilaud. Graph properties of graph associahedra. *Sém. Lothar. Combin.*, 73:Art. B73d, 31 pp., 2015.
- [Müt23] T. Mütze. Combinatorial Gray codes—an updated survey. *Electron. J. Combin.*, DS26(Dynamic Surveys):Paper No. 1. 99 pp., 2023.
- [Pos09] A. Postnikov. Permutohedra, associahedra, and beyond. *Int. Math. Res. Not. IMRN*, 2009(6):1026–1106, 2009.
- [Pou14] L. Pournin. The diameter of associahedra. *Adv. Math.*, 259:13–42, 2014.
- [Pou17] L. Pournin. The asymptotic diameter of cyclohedra. *Israel J. Math.*, 219(2):609–635, 2017.
- [PP18] V. Pilaud and V. Pons. Permutrees. *Algebr. Comb.*, 1(2):173–224, 2018.
- [PPR23] A. Padrol, V. Pilaud, and J. Ritter. Shard polytopes. *Int. Math. Res. Not. IMRN*, (9):7686–7796, 2023.
- [PS19] V. Pilaud and F. Santos. Quotientopes. *Bull. Lond. Math. Soc.*, 51(3):406–420, 2019.
- [Rea05] N. Reading. Lattice congruences, fans and Hopf algebras. *J. Combin. Theory Ser. A*, 110(2):237–273, 2005.
- [RG96] J. Richter-Gebert. *Realization spaces of polytopes*, volume 1643 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1996.
- [RW97] F. Ruskey and M. Weston. A survey of Venn diagrams. *Electron. J. Combin.*, 4(1):Dynamic Survey 5, 10 HTML documents, 1997.
- [Sav97] C. Savage. A survey of combinatorial Gray codes. *SIAM Rev.*, 39(4):605–629, 1997.
- [Ste64] H. Steinhaus. *One hundred problems in elementary mathematics*. Basic Books, Inc., Publishers, New York, 1964. With a foreword by Martin Gardner.
- [STT88] D. D. Sleator, R. E. Tarjan, and W. P. Thurston. Rotation distance, triangulations, and hyperbolic geometry. *J. Amer. Math. Soc.*, 1(3):647–681, 1988.
- [TL73] D. T. Tang and C. N. Liu. Distance-2 cyclic chaining of constant-weight codes. *IEEE Trans. Computers*, C-22:176–180, 1973.
- [Tro62] H. F. Trotter. Algorithm 115: Perm. *Commun. ACM*, 5(8):434–435, 1962.
- [Ven80] J. Venn. On the diagrammatic and mechanical representation of propositions and reasonings. *Phil. Mag. S. 5.*, 9(59):1–18, 1880.
- [Zie95] G. M. Ziegler. *Lectures on polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.