COMBINATORIAL GENERATION VIA PERMUTATION LANGUAGES. VII. SUPERSOLVABLE HYPERPLANE ARRANGEMENTS

SOFIA BRENNER, JEAN CARDINAL, THOMAS MCCONVILLE, ARTURO MERINO, AND TORSTEN MÜTZE

ABSTRACT. For an arrangement \mathcal{H} of hyperplanes in \mathbb{R}^n through the origin, a region is a connected subset of $\mathbb{R}^n \setminus \mathcal{H}$. The graph of regions $G(\mathcal{H})$ has a vertex for every region, and an edge between any two vertices whose corresponding regions are separated by a single hyperplane from \mathcal{H} . We aim to compute a Hamiltonian path or cycle in the graph $G(\mathcal{H})$, i.e., a path or cycle that visits every vertex (=region) exactly once. Our first main result is that if \mathcal{H} is a supersolvable arrangement, then the graph of regions $G(\mathcal{H})$ has a Hamiltonian cycle. More generally, we consider quotients of lattice congruences of the poset of regions $P(\mathcal{H}, R_0)$, obtained by orienting the graph $G(\mathcal{H})$ away from a particular base region R_0 . Our second main result is that if \mathcal{H} is supersolvable and R_0 is a canonical base region, then for any lattice congruence \equiv on $P(\mathcal{H}, R_0) =: L$, the cover graph of the quotient lattice L/\equiv has a Hamiltonian path.

These paths and cycles are constructed by a generalization of the well-known Steinhaus-Johnson-Trotter algorithm for listing permutations. This algorithm is a classical instance of a combinatorial Gray code, i.e., an algorithm for generating a set of combinatorial objects by applying a small change in each step. When applying our two main results to well-known supersolvable arrangements, such as the coordinate arrangement, and the braid arrangement and its subarrangements, we recover a number of known Gray code algorithms for listing various combinatorial objects, such as binary strings, binary trees, triangulations, rectangulations, acyclic orientations of graphs, congruence classes of quotients of the weak order on permutations, and of acyclic orientation lattices. These were obtained earlier from the framework of zigzag languages of permutations proposed by Hartung, Hoang, Mütze, and Williams (Trans. Amer. Math. Soc., 2022). When applying our main results to the type B Coxeter arrangement and its subarrangements, we obtain a number of new Gray code algorithms for listing (pattern-avoiding) signed permutations, symmetric triangulations, acyclic orientations of certain signed graphs, and in general for combinatorial families of Coxeter type B, which generalizes the theory of zigzag languages to signed permutations. Our approach also yields new Hamiltonicity results for large classes of polytopes, in particular signed graphic zonotopes and the type B quotientopes of Padrol, Pilaud, and Ritter (Int. Math. Res. Not., 2023).

1. Introduction

1.1. Combinatorial generation. Given a family of combinatorial objects, such as permutations, binary strings, binary trees, partitions etc., the *combinatorial generation* problem asks to efficiently list all objects from the family, one after the other and each object exactly once. This algorithmic problem is closely related to the problems of counting, random sampling, and optimization. It is covered in depth in Knuth's seminal book 'The Art of Computer Programming

⁽Sofia Brenner) Institut für Mathematik, Universität Kassel, Germany

⁽Jean Cardinal) Computer Science Department, Université Libre de Bruxelles (ULB), Belgium

⁽Thomas McConville) DEPARTMENT OF MATHEMATICS, KENNESAW STATE UNIVERSITY, UNITED STATES

⁽Arturo Merino) Universidad de O'Higgins, Rancagua, Chile

⁽Torsten Mütze) Institut für Mathematik, Universität Kassel, Germany

E-mail addresses: sbrenner@mathematik.uni-kassel.de, jean.cardinal@ulb.be,

tmcconvi@kennesaw.edu, arturo.merino@uoh.cl, tmuetze@mathematik.uni-kassel.de.

An extended abstract of this paper has been accepted for presentation at SODA 2026.

Sofia Brenner and Torsten Mütze were supported by German Science Foundation grant 522790373. Jean Cardinal was supported by the Fonds de la Recherche Scientifique-FNRS under grant n° T003325F. Arturo Merino was supported by ANID FONDECYT Iniciación No. 11251528.

Vol. 4A' [Knu11]. In order for the generation algorithm to be efficient, it is often helpful to list the objects in a way such that any two consecutive objects differ only by a small, elementary change, sometimes called a flip. An example for this is the classical binary reflected Gray code for binary strings, in which all binary strings of a fixed length n are listed in such a way that any two consecutive strings differ only in a single bit [Gra53]. The term combinatorial Gray code is used more broadly for any listing of combinatorial objects subject to an appropriately defined flip operation between successive objects [Sav97, Müt23]. The flip operation also gives rise to a flip graph on the set of objects, in which two objects are adjacent whenever they differ by a single flip. The problem of finding a combinatorial Gray code amounts to finding a Hamiltonian path in the corresponding flip graph, i.e., a path that visits every vertex exactly once. A Gray code is cyclic if the first and last object also differ in a flip, and this corresponds to a Hamiltonian cycle in the flip graph.

1.2. The Steinhaus-Johnson-Trotter algorithm for permutations. An early, and now classical example of a cyclic Gray code is the so-called Steinhaus-Johnson-Trotter algorithm [Ste64, Joh63, Tro62 for listing all permutations of $[n] := \{1, \ldots, n\}$ by adjacent transpositions, i.e., any two consecutive permutations differ by an exchange of two entries at adjacent positions. The method is also known as 'plain changes' and has been used since a long time by bell ringers (see [Whi96]). An easy way to describe this listing of permutations is by the following greedy algorithm due to Williams [Wil13]: Start with the identity permutation, and then repeatedly apply an adjacent transposition to the last permutation in the list that involves the largest possible value so as to create a new permutation, which then gets added to the list. Equivalently, the construction can be described by induction on n. Given the listing L_{n-1} for permutations of length n-1, the listing L_n for permutations of length n is obtained by replacing every permutation π in L_{n-1} by n permutations of length n obtained by inserting the largest value n in all possible positions in π . Specifically, the value n moves alternatingly from right to left and from left to right in the successive permutations of L_{n-1} , in a zigzag pattern; see Figure 1 (a). This algorithm can be implemented *looplessly*, i.e., so that each permutation is generated in constant time, while using only linear memory to store the current permutation and a few auxiliary arrays.

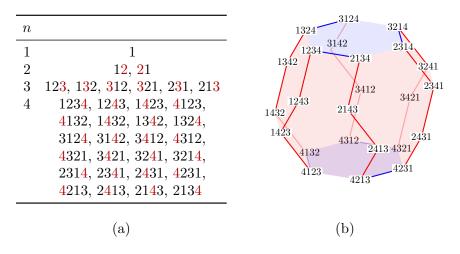


FIGURE 1. (a) Steinhaus-Johnson-Trotter listings of permutations for n = 1, ..., 4, with the largest element n highlighted; (b) visualization of the listing for n = 4 as a Hamiltonian cycle in the permutahedron.

The corresponding flip graph for this problem, namely the Cayley graph of the symmetric group generated by adjacent transpositions, is the well-known *permutahedron*, and the Steinhaus-Johnson-Trotter algorithm yields a Hamiltonian cycle in this graph, i.e., the first and last permutation also differ in an adjacent transposition; see Figure 1 (b).

1.3. Zigzag languages of permutations. In a recent series of papers, Mütze and coauthors showed that the simple greedy algorithm to generate the Steinhaus-Johnson-Trotter listing can be generalized to efficiently generate many other families of combinatorial objects that are in bijection to certain subsets of permutations. Specifically, a zigzag language is a subset of permutations of [n] that satisfies a certain closure property (see [HHMW22] for details), and via those bijections the flip operations on the permutations from the zigzag language translate to natural flip operations on the combinatorial objects. The algorithm to generate a zigzag language presented in [HHMW22] in each step greedily moves the largest possible value in the permutation over a number of neighboring smaller entries, so as to create a new permutation in the zigzag language, which then gets added to the list. The zigzagging framework applies to generate pattern-avoiding permutations [HHMW22], lattice quotients of the weak order on permutations [HM21], pattern-avoiding rectangulations [MM23], elimination trees of chordal graphs [CMM25], acyclic orientations of chordal graphs and hypergraphs and quotients of acyclic reorientation lattices [CHM+23], and pattern-avoiding binary trees [GMN24].

The main contribution of this paper is to treat many of the aforementioned results under a common umbrella, namely that of hyperplane arrangements. This yields an abstract unifying view, yet provides simpler and shorter proofs, and at the same time yields more general results, including new Hamiltonicity results for interesting flip graphs and polytopes, and concrete new Gray code algorithms for a variety of combinatorial objects. The following sections provide three complementary views on the families of objects we aim to generate. Much of the terminology and definitions are spelled out in more detail in Section 2.

1.4. Hyperplane arrangements and graph of regions. The geometric perspective that we will use is that of hyperplane arrangements, defined as a nonempty finite set of real hyperplanes through the origin. Given such an arrangement \mathcal{H} in \mathbb{R}^n , the set $\mathcal{R} = \mathcal{R}(\mathcal{H}) := \mathbb{R}^n \setminus \mathcal{H}$ is a collection of open convex sets called regions. The graph of regions $G(\mathcal{H})$ has $\mathcal{R}(\mathcal{H})$ as its vertex set, and an edge between any two vertices corresponding to regions that are separated by exactly one hyperplane from \mathcal{H} ; see Figure 2.

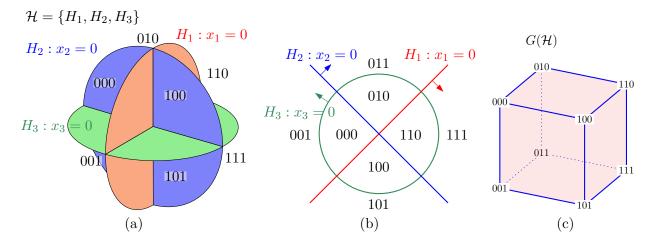


FIGURE 2. (a) The coordinate arrangement for n=3 with regions labeled by binary strings; (b) its stereographic projection from the south pole with the normal vectors for each hyperplane; (c) its graph of regions, realized as a polytope, namely the 3-dimensional hypercube.

Our goal is to construct a Hamiltonian path or cycle in $G(\mathcal{H})$, i.e., we aim to visit all regions, one after the other and each region exactly once. This problem generalizes a number of combinatorial generation problems and their corresponding flip graphs:

- The coordinate arrangement is the set of hyperplanes \mathcal{H} in \mathbb{R}^n whose normal vectors are the canonical base vectors $\{\overrightarrow{e_i} \mid i \in [n]\}$. The regions of this arrangement \mathcal{H} are in bijection with binary strings of length n, and in the graph of regions $G(\mathcal{H})$, any two regions differing in a single bit are joined by an edge, i.e., $G(\mathcal{H})$ is the graph of the n-dimensional hypercube; see Figure 2.
- The braid arrangement, also known as the type A Coxeter arrangement, is the set of hyperplanes \mathcal{H} in \mathbb{R}^n defined by the normal vectors $\{\vec{e_i} \vec{e_j} \mid 1 \leq i < j \leq n\}$. The regions of this arrangement \mathcal{H} are in bijection with permutations of [n], and in $G(\mathcal{H})$, any two regions differing in an adjacent transposition are joined by an edge, i.e., $G(\mathcal{H})$ is the graph of the (n-1)-dimensional permutahedron; see Figure 3.
- The type B Coxeter arrangement is the set of hyperplanes \mathcal{H} in \mathbb{R}^n defined by the normal vectors $\{\overrightarrow{e_i} \pm \overrightarrow{e_j} \mid 1 \leq i < j \leq n\} \cup \{\overrightarrow{e_i} \mid i \in [n]\}$. The regions of this arrangement \mathcal{H} are in bijection with signed permutations of [n]. A signed permutation is a permutation of [n] in which every entry has a positive or negative sign. In $G(\mathcal{H})$, any two regions differing either in an adjacent transposition or a sign change of the first entry are joined by an edge, i.e., $G(\mathcal{H})$ is the graph of the n-dimensional B-permutahedron; see Figure 4.
- Given a simple graph F = ([n], E), the graphic arrangement of F is the subarrangement of the braid arrangement consisting only of the hyperplanes with normal vectors $\{\vec{e_i} \vec{e_j} \mid \{i, j\} \in E\}$. The regions of this arrangement \mathcal{H} are in bijection with acyclic orientations of F, and in in $G(\mathcal{H})$, any two regions differing in reversing a single arc in the acyclic orientation of F are joined by an edge.

An interesting special case of this problem is when the arrangement is that of a reflection group. In that case, it was shown by Conway, Sloane, and Wilks [CSW89] that the graph of regions is always Hamiltonian, hence that a Gray code always exists.

We conclude this section with some easy observations about the graph $G(\mathcal{H})$. First of all, the number of regions $|\mathcal{R}(\mathcal{H})|$ is always even, i.e., $G(\mathcal{H})$ has an even number of vertices. This follows from the observation that the opposition map $x \mapsto -x$ on \mathbb{R}^n induces an involution without fixed points on the set of regions $\mathcal{R}(\mathcal{H})$. Furthermore, the graph $G(\mathcal{H})$ is bipartite, which can be shown by induction on the number of hyperplanes in \mathcal{H} . If the number of hyperplanes in

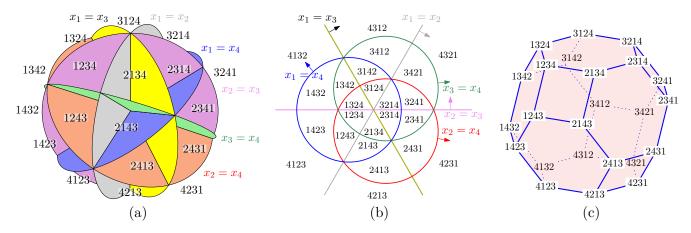


FIGURE 3. (a) The braid arrangement (=type A Coxeter arrangement) for n=4 projected to three dimensions with the regions labeled by permutations; (b) its stereographic projection; (c) its graph of regions, realized as a polytope, namely the 3-dimensional permutahedron of type A.

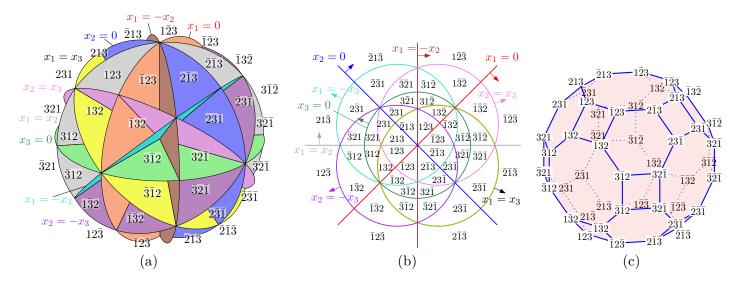


FIGURE 4. (a) The type B Coxeter arrangement for n=3 with the regions labeled by signed permutations (barred entries have a negative sign); (b) its stereographic projection; (c) its graph of regions, realized as a polytope, namely the 3-dimensional B-permutahedron.

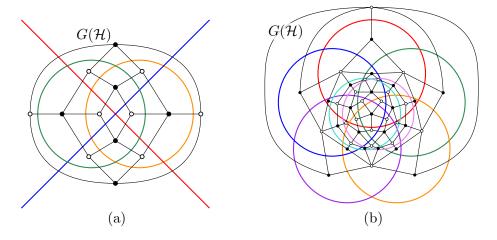


FIGURE 5. Stereographic projections of two hyperplane arrangements \mathcal{H} in \mathbb{R}^3 for which $G(\mathcal{H})$ does not admit a Hamiltonian path (nor cycle): (a) \mathcal{H} consists of 4 hyperplanes and the two partition classes of $G(\mathcal{H})$ have sizes 6 (black vertices) and 8 (white vertices); (b) \mathcal{H} consists of 7 hyperplanes and both partition classes of $G(\mathcal{H})$ have size 22.

the arrangement \mathcal{H} is odd, then the two partition classes of $G(\mathcal{H})$ have the same size (as the opposition map switches classes in this case). On the other hand, if the number of hyperplanes is even then the two partition classes of $G(\mathcal{H})$ can have different sizes, which rules out the existence of a Hamiltonian cycle, and if the size difference is more than 1 then it also rules out the existence of a Hamiltonian path; see Figure 5 (a). Even if the partition classes are balanced, a Hamiltonian path is not guaranteed; see Figure 5 (b).

1.5. Posets of regions and lattice congruences. By choosing one region $R_0 \in \mathcal{R}(\mathcal{H})$ as base region, one can orient the edges of the graph of regions $G(\mathcal{H})$ away from R_0 . Because $G(\mathcal{H})$ is bipartite, it becomes the cover graph of a graded poset $P(\mathcal{H}, R_0)$, called the poset of regions. For example, if \mathcal{H} is the aforementioned type A Coxeter arrangement, and R_0 is the

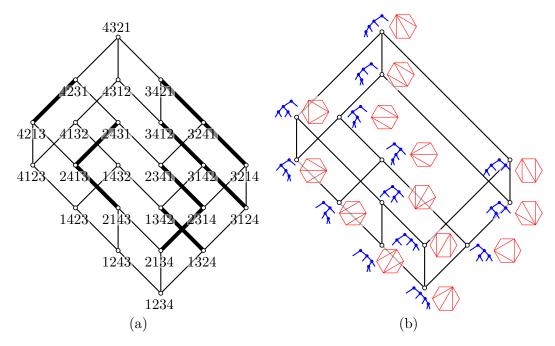


FIGURE 6. (a) The weak order of type A for n=4, obtained by orienting the permutahedron shown in Figure 3 (c) away from the base region $R_0=1234$. The bold lines connect pairs of permutations that are in the same equivalence class of the sylvester congruence. It is defined by the rewriting rule $_b_ca_ \equiv _b_ac_$ where a < b < c, which means that whenever we see a subsequence of three entries b, c, a with a < b < c in a permutation, where c and a are at consecutive positions, then the permutation obtained by transposing c and a is in the same equivalence class. This corresponds to destroying one occurrence of the pattern 231. Hence the 231-avoiding permutations are representatives of the equivalence classes and sit at the bottom of each class. (b) Quotient lattice for the congruence on the left, namely the Tamari lattice with n-vertex binary trees (blue) or triangulations of a convex (n+2)-gon (red) corresponding to the 231-avoiding permutations.

identity permutation, then $P(\mathcal{H}, R_0)$ is the weak order of type A, i.e., permutations ordered by their inversion sets; see Figure 6 (a). Similarly, for the type B Coxeter arrangement and R_0 the identity permutation, $P(\mathcal{H}, R_0)$ is the weak order of type B, i.e., signed permutations ordered by their inversion sets; see Figure 7 (a).

In these two cases, $P(\mathcal{H}, R_0)$ has the additional structure of a *lattice*, i.e., for any two elements $X, Y \in \mathcal{R}(\mathcal{H})$, there is a unique largest element below X and Y, called the *meet* $X \wedge Y$, and a unique smallest element above X and Y, called the *join* $X \vee Y$. If we have a lattice on a ground set P, a *lattice congruence* is an equivalence relation \equiv on P that is compatible with the meet and join operations, i.e., we require that if $X \equiv X'$ and $Y \equiv Y'$ then $X \wedge Y \equiv X' \wedge Y'$ and $X \vee Y \equiv X' \vee Y'$. The *quotient lattice* is the lattice on P/\equiv , i.e., on the equivalence classes on P formed by \equiv , with the order relation inherited from P. Consequently, the cover graph of P/\equiv is obtained from the cover graph of P by contracting the vertices in each equivalence class to a single vertex.

For the weak order of type A, there is a large number of distinct lattice congruences, and for several of them, the corresponding congruence classes encode interesting combinatorial objects (under suitable bijections), and the cover graphs of the resulting quotient lattices are interesting flip graphs. For example, the well-known $Tamari\ lattice$ is the quotient of the weak order of type A under the so-called sylvester congruence [HNT05, MHPS12]. Its elements can be

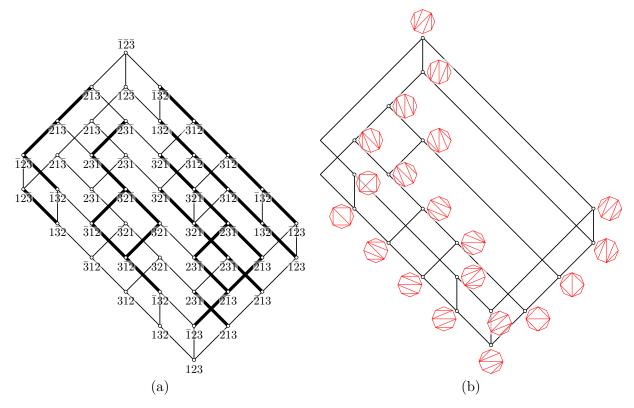


FIGURE 7. (a) The weak order of type B for n=3, obtained by orienting the B-permutahedron shown in Figure 4(c) away from the base region $R_0=123$. The bold lines connect pairs of signed permutations that are in the same equivalence class of the congruence defined by the rewriting rule $_b_ca_ \equiv _b_ac_$ where a < b < c and b > 0, which has to be applied to the full notation of each signed permutation $x = (x_1, \ldots, x_n)$, i.e., to the string $(\overline{x_n}, \overline{x_{n-1}}, \ldots, \overline{x_1}, x_1, x_2, \ldots, x_n)$. The pattern-avoiding signed permutations sit at the bottom of each equivalence class. (b) Quotient lattice, namely the type B Tamari lattice for the congruence on the left, with the point-symmetric triangulations of a convex (2n+2)-gon corresponding to the pattern-avoiding signed permutations (the correspondence is given by Theorem 29).

identified with 231-avoiding permutations, and hence they are in bijection with binary trees and triangulations of a convex polygon, all of which are Catalan families; see Figure 6. In the cover graph of the quotient lattice, edges are precisely between binary trees that differ in a tree rotation, i.e., we obtain the rotation graph of binary trees [STT88, Pou14]. In the corresponding triangulations, a tree rotation translates to a flip that removes one edge between two triangles and replaces it by the other diagonal in the resulting empty quadrilateral.

Different families of rectangulations can also be obtained as quotients of the weak order of type A [Gir12, LR12, Rea12, CSS18], and the flip operations are tree rotations on the dual twin binary trees. The *Boolean lattice*, with the hypercube as its cover graph, is also a lattice quotient of the weak order of type A. In fact, all lattice congruences of the weak order form a lattice themselves, with the order relation being refinement of the equivalence classes.

If \mathcal{H} is the graphic arrangement of a graph F and R_0 a fixed acyclic orientation of F, then $P(\mathcal{H}, R_0)$ is the reorientation poset of all acyclic orientations of F with respect to the reference orientation R_0 . Pilaud [Pil24] gave a necessary and sufficient condition on R_0 for this poset to be a lattice, and he also described the corresponding lattice congruences and quotients.

Reading [Rea06] defined a type B Tamari lattice via a lattice congruence of the weak order of type B. In Figure 7, we show the quotient lattice with the corresponding combinatorial

objects, namely point-symmetric triangulations of a convex polygon. Edges in the cover graph are between triangulations that differ in a flip of one or two symmetric edges.

1.6. **Polytopes.** Finally, the aforementioned flip graphs all arise as the *skeletons* of well-known combinatorial *polytopes*. The *zonotope* of an arrangement of hyperplanes \mathcal{H} is the polytope obtained as the Minkowski sum of line segments in the directions of the normal vectors of the hyperplanes in \mathcal{H} . Conversely, the *normal fan* of the zonotope is the fan defined by the arrangement \mathcal{H} .

The vertices of the zonotope are in bijection with the regions $\mathcal{R}(\mathcal{H})$ of the arrangement, and its skeleton is the graph of regions $G(\mathcal{H})$. For example, the hypercube is the zonotope of the coordinate arrangement; see Figure 2. Similarly, the *permutahedron* is the zonotope of the braid arrangement (see Figure 3), and the *B-permutahedron* is the zonotope of the type *B* Coxeter arrangement (see Figure 4). The *graphic zonotope* is the zonotope of the graphic arrangement of a graph.

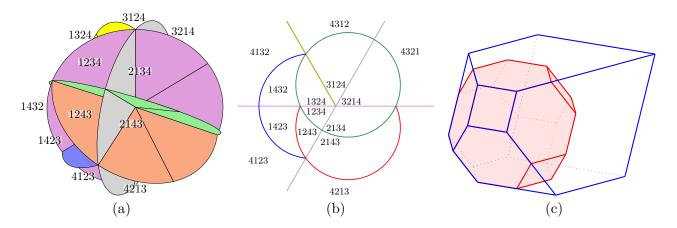


FIGURE 8. (a) Quotient fan obtained from the type A Coxeter fan in Figure 3 (a) under the sylvester congruence of Figure 6, and (b) its stereographic projection. It is the normal fan of the associahedron shown in (c), obtained by removing certain hyperplanes that bound the permutahedron inside.

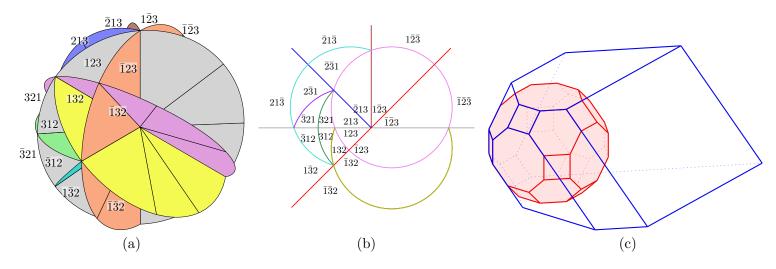


FIGURE 9. (a) Quotient fan obtained from the type B Coxeter fan in Figure 4 (a) under the congruence of Figure 7, and (b) its stereographic projection. It is the normal fan of the B-associahedron shown in (c), obtained by removing certain hyperplanes that bound the B-permutahedron inside.

A lattice congruence of type A acts on the normal fan of the permutahedron by gluing together regions that belong to the same equivalence class of the congruence. Pilaud and Santos [PS19] showed that the resulting quotient fan is the normal fan of a polytope they called quotientope; see Figure 8. Put differently, their result shows that the cover graph of any lattice quotient of the weak order on permutations can be realized geometrically as the skeleton of a polytope. Quotientopes thus generalize permutahedra, associahedra [Lod04], rectangulotopes [CP25], hypercubes etc. Padrol, Pilaud, and Ritter [PPR23] later gave a construction of quotientopes via Minkowski sums. Notably, their approach also works for all lattice congruences on the type B Coxeter arrangement, i.e., we obtain type B quotientopes, including the B-associahedron [Sim03]; see Figure 9. We refer to the recent expository paper of Pilaud, Santos, and Ziegler [PSZ23] for context and references about associahedra and their generalizations.

Table 1 gives an overview of the various structures discussed here and their connections.

Table 1. Correspondence between combinatorial objects encoded by different control of the contro	$_{ m fferent}$
hyperplane arrangements and their lattice congruences, posets and polytopes	s.

Arrangement	Combinatorial objects	pinatorial objects Poset	
coordinate	binary strings	Boolean lattice	hypercube
type A graphic	permutations binary trees, triangulations (diagonal/generic) rectangulations permutrees acyclic orientations of graphs	weak order of type A Tamari lattice lattice of rectangulations rotation lattice acyclic reorientation order	permutahedron of type A associahedron rectangulotopes [CP25] permutreehedra [PP18] graphic zonotope
type B	signed permutations symmetric triangulations acyclic orientations of signed graphs	weak order of type B Tamari lattice of type B acyclic reorientation order of type B	permutahedron of type B associahedron of type B signed graphic zonotope

1.7. Our results. The zigzagging procedure of the Steinhaus-Johnson-Trotter algorithm and its generalizations [HHMW22] exploit a well-studied property of hyperplane arrangements called supersolvability. This term was first used in the context of group theory, as a property of the lattice of subgroups of a group. Stanley [Sta72] later gave a lattice-theoretic definition, which was then used to qualify hyperplane arrangements. A convenient definition, in our context, is the following: A hyperplane arrangement \mathcal{H} is supersolvable if the set \mathcal{H} can be partitioned into two nonempty subsets \mathcal{H}_0 and \mathcal{H}_1 , where \mathcal{H}_0 is a supersolvable arrangement of lower rank, and for any pair of distinct hyperplanes in \mathcal{H}_1 , their intersection is contained in a hyperplane of \mathcal{H}_0 [BEZ90]. (We refer to Section 2.8 for complete definitions.) When a supersolvable arrangement \mathcal{H} is split into the two subarrangements \mathcal{H}_0 and \mathcal{H}_1 , then the regions of the arrangement \mathcal{H}_0 can be seen as equivalence classes of the regions of \mathcal{H} , and the regions in each such equivalence class are linearly ordered, i.e., they form a path in the graph of regions $G(\mathcal{H})$. The zigzagging method alternatingly moves back and forth along those paths, following a Hamiltonian cycle in the graph of regions $G(\mathcal{H}_0)$ of the lower-rank arrangement \mathcal{H}_0 that is built inductively; see Figure 10. As the number of regions $|\mathcal{R}(\mathcal{H}_0)|$ is always even, the last visited region is adjacent to the starting region, i.e., we obtain a Hamiltonian cycle in $G(\mathcal{H})$.

This leads to our first result.¹

Theorem 1. Let \mathcal{H} be a supersolvable hyperplane arrangement of rank $n \geq 2$. Then the graph of regions $G(\mathcal{H})$ has a Hamiltonian cycle of even length.

 $^{^{1}}$ After submitting this paper, we learned that Körber, Schnieders, Stricker, and Walizadeh [KSSW25] found an independent proof of Theorem 1.

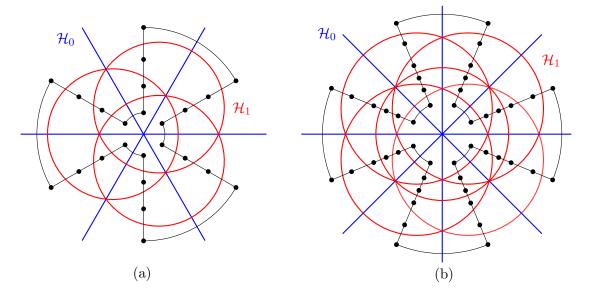


FIGURE 10. Illustration of the supersolvable partition $\mathcal{H} = \mathcal{H}_0 \cup \mathcal{H}_1$ of the Coxeter arrangements of (a) type A and (b) type B from Figures 3 and 4, respectively (as stereographic projections). The resulting Hamiltonian cycles in the graph of regions obtained from the zigzagging procedure are shown in black, and they correspond to the listings given in Figures 1 and 11, respectively.

Björner, Edelman and Ziegler [BEZ90] showed that the regions of a supersolvable arrangement \mathcal{H} can always be ordered into a lattice, for a suitable choice of base region $R_0 \in \mathcal{R}(\mathcal{H})$. Our second main result is that the cover graph of the quotient of any such lattice always has a Hamiltonian path, which can be found by the zigzagging procedure.

Theorem 2. Let \mathcal{H} be a supersolvable hyperplane arrangement, and let $L := P(\mathcal{H}, R_0)$ be its lattice of regions for a canonical base region R_0 . Then for any lattice congruence \equiv on L, the cover graph of L/\equiv has a Hamiltonian path.

(We refer to Section 2.10 for the definition of canonical base region.) The Hamiltonian paths and cycles in those results can be computed by the following simple greedy algorithm based on the supersolvable partition $\mathcal{H} = \mathcal{H}_0 \cup \mathcal{H}_1$: Start at a canonical base region, and then repeatedly apply the following rule where to move from the currently visited region: If possible, traverse a hyperplane from \mathcal{H}_1 that leads to a previously unvisited region, otherwise recursively select a hyperplane to traverse from \mathcal{H}_0 according to this rule.

- 1.8. **Applications of our results.** With the two results above, we recover a number of previously known Hamiltonicity results and Gray code algorithms, and we provide several new ones.
- 1.8.1. The coordinate arrangement and binary strings. The coordinate arrangement \mathcal{H} with normal vectors $\{\vec{e_i} \mid i \in [n]\}$ is clearly supersolvable. Indeed, for the required partition $\mathcal{H} = \mathcal{H}_0 \cup \mathcal{H}_1$ we can take \mathcal{H}_1 as the singleton hyperplane with normal vector $\vec{e_n}$. The resulting Hamiltonian cycle in the hypercube is the well-known binary reflected Gray code; see Figure 2.

It has the following greedy description ([Wil13]): Start with the all-0 string of length n, and then repeatedly flip the rightmost bit in the last string so as to create a new binary string.

1.8.2. The type A Coxeter arrangement and permutations. Recall that the braid arrangement, or type A Coxeter arrangement, is defined by the normal vectors $\{\vec{e_i} - \vec{e_j} \mid 1 \leq i < j \leq n\}$. The supersolvability of this arrangement \mathcal{H} is witnessed by letting \mathcal{H}_1 be the set of hyperplanes with normal vectors whose nth component is nonzero. Indeed, consider two of them, with normal

Table 2. The classical binary reflected Gray code for n = 1, ..., 4.

\overline{n}	
1	0, 1
2	00, 01, 11, 10
3	000, 001, 011, 010, 110, 111, 101, 100
4	0000, 0001, 0011, 0010, 0110, 0111, 0101, 0100,
	1100, 1101, 1111, 1110, 1010, 1011, 1001, 1000

vectors $\vec{e_i} - \vec{e_n}$ and $\vec{e_j} - \vec{e_n}$, respectively, with i < j < n. Then their intersection is a flat of codimension 2 lying in the hyperplane with normal vector $\vec{e_i} - \vec{e_j}$, which belongs to \mathcal{H}_0 . Theorem 1 yields the aforementioned Steinhaus-Johnson-Trotter listing of permutations (recall Section 1.2).

1.8.3. The type B Coxeter arrangement and signed permutations. Recall that the type B Coxeter arrangement is defined by the hyperplanes with normal vectors $\{\vec{e_i} \pm \vec{e_j} \mid 1 \le i < j \le n\} \cup \{\vec{e_i} \mid i \in [n]\}$. Similarly to before, the supersolvability of this arrangement \mathcal{H} is witnessed by letting \mathcal{H}_1 be the hyperplanes with normal vectors whose nth component is nonzero.

As mentioned before, the regions of \mathcal{H} are in bijection with signed permutations, i.e., permutations of [n] in which every entry has a positive or negative sign. Adjacencies in the graph $G(\mathcal{H})$ are adjacent transpositions or a sign change of the first entry. The Gray code for signed permutations obtained from Theorem 1 is shown in Figure 4, and it is cyclic for all n. Note that this Gray code is different from the Gray code given by Conway, Sloane, and Wilks [CSW89] and also from the one given by Korsh, LaFollette, and Lipschutz [KLL11].

For any positively signed element $i \in [n]$ we write $\overline{i} := -i$ for its negatively signed counterpart. Note that $\overline{i} = i$. We define $[\overline{n}] := \{\overline{1}, \overline{2}, \dots, \overline{n}\}$, and we denote a signed permutation as a string $\pi = (a_1, \dots, a_n)$ such that $a_i \in [\overline{n}] \cup [n]$ for $i = 1, \dots, n$ and $\{|a_1|, \dots, |a_n|\} = [n]$. It is useful in this context to introduce the so-called *full notation* of a signed permutation $\pi = (a_1, \dots, a_n)$, which is the string $\widehat{\pi} := (\overline{a_n}, \overline{a_{n-1}}, \dots, \overline{a_1}, a_1, a_2, \dots, a_n)$ of length 2n. Note that an adjacent transposition in π corresponds to a symmetric pair of adjacent transpositions in $\widehat{\pi}$. Furthermore, a sign change of the first entry of π corresponds to an adjacent transposition of the middle two

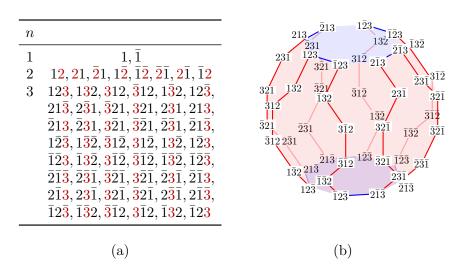


FIGURE 11. (a) Our new cyclic Gray code for signed permutations for n = 1, 2, 3; (b) visualization of the listing for n = 3 as a Hamiltonian cycle on the *B*-permutahedron.

entries of $\widehat{\pi}$. Therefore, our Gray code for signed permutations can be described greedily, using the full notation, as follows: Start with the identity permutation $(\overline{n}, \overline{n-1}, \dots, \overline{1}, 1, 2, \dots, n)$, and then repeatedly apply an adjacent transposition² to the last permutation in the list that involves the largest possible value so as to create a new permutation. This algorithm can be implemented looplessly, i.e., in time $\mathcal{O}(1)$ per generated signed permutation, while using only $\mathcal{O}(n)$ memory. We prepared an implementation in C++, available for download and experimentation on the Combinatorial Object Server website [cos].

1.8.4. Type A subarrangements and acyclic orientations of graphs. Recall from Section 1.4 that the graphic arrangement \mathcal{H} of a graph F = ([n], E) is a subarrangement of the braid arrangement, obtained by retaining only the hyperplanes corresponding to the edges of F, with normal vectors $\{\vec{e_i} - \vec{e_j} \mid \{i, j\} \in E\}$. It is well-known and not hard to see that the regions of \mathcal{H} are in bijection with acyclic orientations of the graph F, and that an adjacency in $G(\mathcal{H})$ corresponds to reversing a single arc in the acyclic orientation of F. Stanley [Sta72, Prop. 2.8] proved that the graphic arrangement of F is supersolvable if and only if F is a chordal graph, i.e., F has no induced cycle of length 4 or more. Applying Theorem 1, we thus obtain the Gray codes for acyclic orientations of chordal graphs studied in [SSW93] and [CHM⁺23].

Via Theorem 2 we also recover the Gray codes for quotients of acyclic reorientation lattices of chordal graphs, whose existence was proved by Cardinal, Hoang, Merino, Mička, and Mütze [CHM⁺23], addressing a problem raised by Pilaud [Pil24, Prob. 52].

1.8.5. Type B subarrangements and acyclic orientations of signed graphs. Zaslavsky [Zas82] showed that subarrangements of the type B Coxeter arrangement can be interpreted as the type B analogue of graphic arrangements but for signed graphs, in which every edge carries a sign. The conditions for supersolvability of this arrangement are known; see [Zas01, STT19]. In particular, these conditions identify a class of signed graphs that are the signed analogues of chordal graphs, those having a signed perfect elimination ordering. We refer to Sections 2.7 and 2.8 for the definitions. Theorem 1 directly yields new Gray codes for acyclic orientations of those signed graphs, and Hamiltonicity of the corresponding polytopes.

Corollary 3. For any signed graph that has a signed perfect elimination ordering, the corresponding signed graphic zonotope has a Hamiltonian cycle.

From Theorem 2 we also directly obtain Gray codes for quotients of acyclic reorientation lattices of signed graphs, generalizing the result of [CHM⁺23].

- 1.8.6. Hamiltonian paths on type A quotientopes. Recall from Section 1.6 that the quotientopes introduced by Pilaud and Santos are the polytopes whose skeletons are the cover graphs of quotients of lattice congruences of the weak order of type A. Hoang and Mütze [HM21] proved that all quotientopes admit a Hamiltonian path, a result that can be recovered as a special case from our Theorem 2. In particular, we recover a Hamiltonian path on the associahedron (via the sylvester congruence), which coincides with the Gray codes for binary trees by tree rotations due to Lucas, Roelants van Baronaigien, and Ruskey [LRvBR93]. Furthermore, we recover Hamiltonian cycles on rectangulotopes, which coincide with the cyclic Gray codes for diagonal and generic rectangulations due to Merino and Mütze [MM23].
- 1.8.7. Hamiltonian paths on type B quotientopes. As mentioned before, type B quotientopes were defined by Padrol, Pilaud, and Ritter [PPR23] as the polytopes whose skeletons are the cover graphs of quotients of lattice congruences of the weak order of type B. Applying Theorem 2 to the type B Coxeter arrangement, we thus obtain the following new result.

Corollary 4. The skeleton of any type B quotientope has a Hamiltonian path.

²possibly together with a forced symmetric transposition

In particular, we obtain a Hamiltonian cycle on the B-associahedron [Sim03]. Via a suitable bijection, this gives a new cyclic Gray code for point-symmetric triangulations of a convex (2n+2)-gon, in which any two consecutive triangulations differ in one flip of the edge through the center or a pair of symmetric flips not involving the center-edge; see Figures 12 and 13 (see Theorems 29 and 30 below).

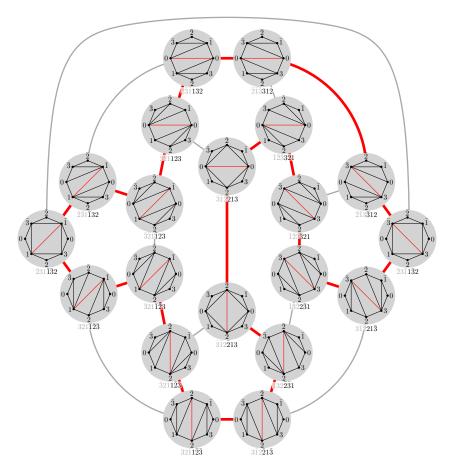


FIGURE 12. The skeleton of the B-associahedron, i.e., the cover graph of the lattice quotient shown in Figure 7, with the pattern-avoiding permutations in full notation. The Hamiltonian cycle corresponding to the Gray code in Figure 13 for n=3 is highlighted.

It is not known whether every cover graph of a quotient of the lattice of regions of a supersolvable hyperplane arrangement can be realized as the skeleton of a polytope. If so, then its skeleton admits a Hamiltonian path according to Theorem 2. In fact, Padrol, Pilaud, and Ritter [PPR23, Conj. 143] conjecture that a polytopal realization is possible if the lattice of regions is congruence uniform.

- $1.8.8.\ Zigzag\ languages\ of\ signed\ permutations.$ As another contribution, we develop a theory of zigzag languages of signed permutations, and obtain the type B analogue of the methods described in Section 1.3 (see Theorem 17 below). This yields Gray codes for generating several different pattern-avoiding signed permutations and the corresponding combinatorial objects in bijection to them.
- 1.9. **Outline of the paper.** In Section 2 we set up additional terminology and notation needed for the proofs of our two main theorems. The proofs of those theorems are presented in Section 3. In Section 4 we develop the theory of zigzag languages of signed permutations. In Section 5, we

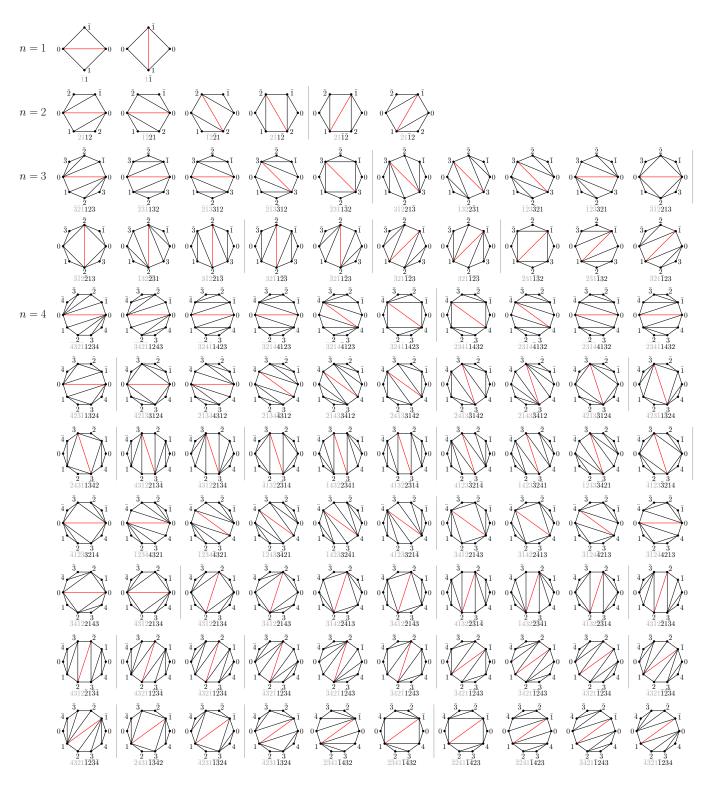


FIGURE 13. Our new cyclic Gray code for point-symmetric triangulations of a convex (2n+2)-gon for n=1,2,3,4, given by Theorems 29 and 30. Below each triangulation is the corresponding pattern-avoiding signed permutation. The vertical bars separate groups of descendants that are derived from the same parent permutation/triangulation in the previous listing.

describe three exceptional signed permutation patterns and establish bijections to other combinatorial objects. In particular, we describe our new Gray code for point-symmetric triangulations of a convex polygon, and the corresponding Hamiltonian cycle on the *B*-associahedron.

2. Background and terminology

2.1. Partial orders and lattices. A poset (P, \leq) is a set P together with a partial order \leq on P. We write < for the corresponding strict partial order. An element $x \in P$ is minimal if there is no $y \in P$ with y < x. Similarly, x is maximal if there is no $y \in P$ with x < y. A cover is a pair (x,y) of elements of P such that x < y and there does not exist z such that x < z < y. We write x < y if (x,y) is a cover. The cover relation \leq is the reflexive, transitive reduction of the order relation \leq . Conversely, the order relation \leq is the reflexive, transitive closure of the cover relation \leq . The cover graph of P is the (undirected) graph that has the elements of P as vertices, and as edges exactly the covers. A poset P is often visualized by a Hasse diagram, which is a drawing of the cover graph in the plane in which for every cover (x,y), the vertex y is embedded higher than x. A chain in P is a sequence of covers $x_0 \leq x_1 \leq \cdots \leq x_\ell$.

The downset and upset of an element $x \in P$ are defined as $x^{\downarrow} := \{y \mid y \leq x\}$ and $x^{\uparrow} := \{y \mid x \leq y\}$. The interval between two elements $x, y \in P$ is the subposet $[x, y] := x^{\uparrow} \cap y^{\downarrow} = \{z \mid x \leq z \leq y\}$. A poset P is graded if there is a rank function $\operatorname{rk}(x)$ mapping elements of P to integers such that $\operatorname{rk}(x) < \operatorname{rk}(y)$ if x < y and $\operatorname{rk}(y) = \operatorname{rk}(x) + 1$ if x < y.

Given two elements $x, y \in P$, their $meet \ x \wedge y$ is the maximum of $x^{\downarrow} \cap y^{\downarrow}$, if it exists and is unique. Dually, their $join \ x \vee y$ is the minimum of $x^{\uparrow} \cap y^{\uparrow}$, if it exists and is unique. The poset P is called a *lattice* if meets and joins exist for any pair of elements. It follows directly from this definition that any lattice has a unique minimum and maximum.

2.2. Lattice congruences. Let (L, \leq) be a lattice. For an equivalence relation \equiv on L and any $x \in L$, we write $[x] := \{y \in L \mid x \equiv y\}$ for the equivalence class of x, and we write $L/\equiv := \{[x] \mid x \in L\}$ for the set of all equivalence classes.

A lattice congruence is an equivalence relation \equiv on L such that for all $x, x', y, y' \in L$, if $x \equiv x'$ and $y \equiv y'$, then $x \wedge y \equiv x' \wedge y'$ and $x \vee y \equiv x' \vee y'$. The set of equivalence classes L/\equiv can be ordered by defining $X \leq Y$ for $X, Y \in L/\equiv$ if there are elements $x \in X$ and $y \in Y$ with $x \leq y$. Furthermore, this order on L/\equiv is a lattice, namely if [x] and [y] are equivalence classes, then their meet and join are $[x \wedge y]$ and $[x \vee y]$, respectively. We refer to L/\equiv with this lattice structure as the quotient lattice.

We recall some fundamental results about lattice congruences; see [Rea16, §9-5] for proofs.

Lemma 5. Let \equiv be an equivalence relation on a finite lattice L.

- (i) Every equivalence class is an interval.
- (ii) We have $X \triangleleft Y$ in L/\equiv if and only if there are elements $x \in X$ and $y \in Y$ such that $x \triangleleft y$ in L.
- 2.3. Hyperplane arrangements. We introduce the required terminology regarding hyperplane arrangements; for more details see [OT92] and [Sta07]. A hyperplane arrangement in \mathbb{R}^n is a nonempty finite set \mathcal{H} of hyperplanes through the origin. The rank of an arrangement \mathcal{H} is the dimension of the space spanned by the normal vectors of the hyperplanes, or equivalently, the codimension of the intersection of all the hyperplanes.

The regions $\mathcal{R} = \mathcal{R}(\mathcal{H})$ of an arrangement \mathcal{H} are the connected components of $\mathbb{R}^n \setminus \mathcal{H}$. The graph of regions $G(\mathcal{H})$ is the adjacency graph of the set of regions, i.e., it has $\mathcal{R}(\mathcal{H})$ as its vertex set, and an edge between any two regions that are separated by exactly one hyperplane from \mathcal{H} .

A flat of an arrangement \mathcal{H} in \mathbb{R}^n is a subspace of \mathbb{R}^n obtained by intersecting some subset of hyperplanes of \mathcal{H} . The intersection lattice $\mathcal{L} = \mathcal{L}(\mathcal{H})$ of the arrangement \mathcal{H} is the set of flats ordered in reverse inclusion order, completed by a minimal element $\hat{0} = \mathbb{R}^n$. The intersection lattice is a geometric lattice, namely the matroid lattice of the matroid defined by the arrangements of normal vectors of the hyperplanes of \mathcal{H} . Its atoms are the hyperplanes, and

the rank $\operatorname{rk}(X)$ of an element $X \in \mathcal{L}$ is equal to its $\operatorname{codimension} \operatorname{codim}(X) = n - \dim(X)$, so that $\operatorname{rk}(\widehat{0}) = 0$ and $\operatorname{rk}(H) = 1$ for all $H \in \mathcal{H}$. Its maximum is given by $\bigcap_{H \in \mathcal{H}} H$, and its rank in the lattice is precisely the rank of the arrangement \mathcal{H} . The meet of two elements $X, Y \in \mathcal{L}$ is equal to $X \wedge Y = \bigcap \{Z \in \mathcal{L} \mid Z \supseteq (X \cup Y)\}$, and the join is $X \vee Y = X \cap Y$.

- 2.4. **Zonotopes.** With every hyperplane arrangement \mathcal{H} , one can associate the zonotope $Z(\mathcal{H})$ defined as the Minkowski sum of line segments in the directions of the normal vectors of the hyperplanes in \mathcal{H} . The graph of regions $G(\mathcal{H})$ is isomorphic to the skeleton of $Z(\mathcal{H})$.
- 2.5. Coxeter arrangements. Interesting hyperplane arrangements are obtained by letting the normal vectors form a root system of a finite Coxeter group. The standard type A, B, and D root systems (the type C is essentially equivalent to the type B for our application) consist respectively of the following vectors; see [BB05] and [FR07]:

```
type A: \{\overrightarrow{e_i} - \overrightarrow{e_j} \mid i \neq j \in [n]\},
type B: \{\overrightarrow{e_i} \pm \overrightarrow{e_j} \mid i \neq j \in [n]\} \cup \{\pm \overrightarrow{e_i} \mid i \in [n]\},
type D: \{\overrightarrow{e_i} \pm \overrightarrow{e_j} \mid i \neq j \in [n]\}.
```

2.6. **Graphic arrangements.** Given a simple graph F = ([n], E), the graphic arrangement $\mathcal{H}(F)$ of F is the arrangement of hyperplanes with the normal vectors $\{\vec{e_i} - \vec{e_j} \mid \{i, j\} \in E\}$. The graphic arrangement is therefore a subarrangement of the type A Coxeter arrangement. These normal vectors form a realization of the graphic matroid of F, and the intersection lattice of $\mathcal{H}(F)$ is the lattice of flats of this matroid. An acyclic orientation of G is an assignment of a direction to each edge in E that does not create any directed cycles. The following is a well-known observation from Greene (see [GZ83]).

Lemma 6. The regions of the graphic arrangement $\mathcal{H}(F)$ of a graph F are in bijection with the acyclic orientations of F.

2.7. Signed graphic arrangements. A signed graph is a simple graph F = ([n], E) in which every edge in E is given a sign from $\{+, -\}$, i.e., we obtain a partition $E = E^+ \cup E^-$, where E^+ and E^- are the sets of edges with a positive and negative sign, respectively. The signed graphic arrangement $\mathcal{H}(F)$ of F is a subarrangement of the type D root system consisting of the hyperplanes with normal vectors

$$\{\overrightarrow{e_i} - \overrightarrow{e_j} \mid ij \in E^+\} \cup \{\overrightarrow{e_i} + \overrightarrow{e_j} \mid ij \in E^-\}.$$

Note that if $E=E^+$, then the signed graphic arrangement is a graphic arrangement. Also note that Zaslavsky [Zas82, Zas91] considered signed graphs with loops and half-edges, whose graphic arrangements are type B subarrangements. For simplicity, we restrict our definition to simple signed graphs.

An orientation of an edge e=ij of a signed graph consists of assigning a direction to the two half-edges composing e. If $e \in E^+$, then these two directions must be the same, either both towards i or both towards j. If $e \in E^-$, then the two directions must be opposite, exactly one towards i and one towards j. An orientation of the signed graph F is defined by orienting each of its edges.

A cycle in an oriented signed graph is an inclusion-minimal subgraph whose vertices and edges can be arranged in an alternating sequence $(v_1, e_1, v_2, e_2, \ldots, v_\ell, e_\ell)$ (possibly with repetitions) where for $i = 1, \ldots, \ell$, the edge e_i has the end vertices v_i and v_{i+1} , and from the two half-edges of e_{i-1} and e_i incident with v_i exactly one is oriented towards v_i (indices are considered cyclically modulo ℓ); see Figure 14. An orientation of a signed graph is acyclic if it does not contain any cycle.

A cycle in an oriented signed graph may consist of a simple cycle in the underlying unsigned graph, in which case it must be *balanced*, i.e., contain an even number of negative edges. However, unlike for unsigned graphs, a cycle in an oriented signed graph may reuse edges and vertices,

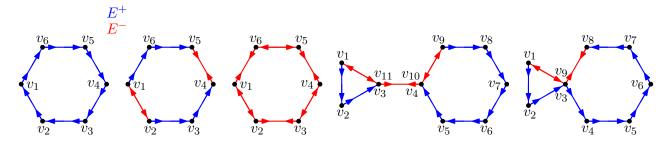


FIGURE 14. Different types of cycles in oriented signed graphs.

and may hence not be a cycle in the underlying unsigned graph. Specifically, it may consist of two unbalanced cycles connected by a path or sharing exactly one vertex [Zas91].

Zaslavsky [Zas91] proved the following analogue of Lemma 6.

Lemma 7. The regions of the signed graphic arrangement $\mathcal{H}(F)$ of a signed graph F are in bijection with the acyclic orientations of F.

2.8. Supersolvable arrangements. An arrangement of hyperplanes \mathcal{H} is supersolvable if its intersection lattice $\mathcal{L}(\mathcal{H})$ is supersolvable in the sense of Stanley [Sta72]. We will instead use the following characterization of supersolvable arrangements given by Björner, Edelman, and Ziegler [BEZ90, Thm. 4.3] as our definition. An arrangement \mathcal{H} of rank n is supersolvable if either $n \leq 2$, or $n \geq 3$ and \mathcal{H} is the disjoint union of two nonempty arrangements \mathcal{H}_0 and \mathcal{H}_1 , where \mathcal{H}_0 is a supersolvable arrangement of rank n-1, and for any pair of distinct hyperplanes $H, H' \in \mathcal{H}_1$, there exists a hyperplane $H'' \in \mathcal{H}_0$ such that $H \cap H' \subseteq H''$.

We proceed to discuss which of the aforementioned hyperplane arrangements are supersolvable. Among the Coxeter arrangements, those of types A and B are known to be the only supersolvable arrangements.

For graphic arrangements, supersolvability of the arrangement $\mathcal{H}(F)$ is equivalent to chordality of the graph F. Formally, a graph is *chordal* if it does not contain any induced cycle of length > 3. One can also characterize chordal graphs by the existence of a certain elimination ordering on their vertices. Specifically, a vertex v in a graph F is simplicial if the neighbors of v induce a complete subgraph of F. An ordering v_1, v_2, \ldots, v_n of the vertices of F is a *perfect elimination ordering* if v_i is simplicial in $F \setminus \{v_1, \ldots, v_{i-1}\}$ for all $i = 1, \ldots, n$. Note that a (reverse) perfect elimination ordering is obtained by repeatedly removing a simplicial vertex from F. It is well-known that a graph F is chordal if and only if it has a perfect elimination ordering. Stanley [Sta72, Prop. 2.8] proved the following; see also [Sta07, Cor. 4.10].

Theorem 8. The graphic arrangement $\mathcal{H}(F)$ of a graph F is supersolvable if and only if F is chordal.

A similar result is known for signed graphic arrangements. Specifically, a vertex v of a signed graph F with edge set $E = E^+ \cup E^-$ is signed simplicial if for any two distinct neighbors x, y of v, there is an edge between x and y that completes a balanced triangle, formally:

- if $vx, vy \in E^+$ or $vx, vy \in E^-$, then we have $xy \in E^+$;
- if $vx \in E^+$ and $vy \in E^-$, then we have $xy \in E^-$.

An ordering v_1, v_2, \ldots, v_n on the vertices of F is a signed perfect elimination ordering if v_i is signed simplicial in $F \setminus \{v_1, \ldots, v_{i-1}\}$. Note that a (reverse) signed perfect elimination ordering is obtained by repeatedly removing a signed simplicial vertex from F.

Zaslavsky [Zas01, Thm. 2.2] proved the following; see also [STT19, Thm. 4.19] for the explicit statement on signed graphs.

Theorem 9. If F is a signed graph graph with a signed perfect elimination ordering, then the signed graphic arrangement $\mathcal{H}(F)$ is supersolvable.

Note that Zaslavsky actually gave a complete characterization (if and only if) of signed graphs whose arrangements are supersolvable, which includes specific additional graphs. We omit the details here.

2.9. Structure of the graph of regions $G(\mathcal{H})$. For the following lemmas, we consider a supersolvable arrangement \mathcal{H} of rank $n \geq 3$ and the decomposition of \mathcal{H} into two subarrangements \mathcal{H}_0 and \mathcal{H}_1 given by the definition. We define the surjective map

$$\rho: \mathcal{R}(\mathcal{H}) \to \mathcal{R}(\mathcal{H}_0) \tag{1}$$

that maps every region of \mathcal{H} to the region of \mathcal{H}_0 that it is contained in.

We recall the following basic properties of supersolvable hyperplane arrangements that can be found, for instance, in [Ede84, BEZ90].

Lemma 10. Every hyperplane $H \in \mathcal{H}_1$ has the following properties:

- (i) We have $\operatorname{rk}(\mathcal{H}_0 \cup \{H\}) = n > \operatorname{rk}(\mathcal{H}_0) = n 1$.
- (ii) H splits every region in $\mathcal{R}(\mathcal{H}_0)$ into exactly two regions.

Proof. This proof is illustrated in Figure 15.

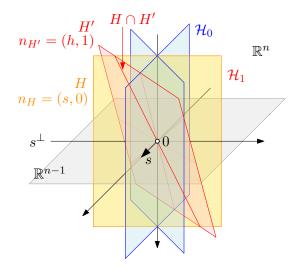


FIGURE 15. Illustration of the proof of Lemma 10.

We assume w.l.o.g. that the dimension of the ambient space equals $n = \text{rk}(\mathcal{H})$. Let $N(\mathcal{H}_0)$ denote the set of normal vectors of the hyperplanes in \mathcal{H}_0 . For vectors $x, y \in \mathbb{R}^n$, we write $x \cdot y$ for the standard scalar product of x and y. By the definition of supersolvability, we know that $\text{rk}(\mathcal{H}_0) = n - 1 < \text{rk}(\mathcal{H}) = n$. We may choose a basis of \mathbb{R}^n such that $\text{span}(N(\mathcal{H}_0)) \subseteq \mathbb{R}^{n-1} \times \{0\}$.

To prove (i), consider a hyperplane $H \in \mathcal{H}_1$, and suppose for the sake of contradiction that $\mathrm{rk}(\mathcal{H}_0 \cup \{H\}) = n-1$, i.e., H has a normal vector $n_H = (s,0) \in \mathbb{R}^n$ with $s \in \mathbb{R}^{n-1}$. Let $H' \in \mathcal{H}_1$ with normal vector $n_{H'} \notin \mathbb{R}^{n-1} \times \{0\}$. Such a hyperplane exists in \mathcal{H}_1 since $\mathrm{rk}(\mathcal{H}) = n > \mathrm{rk}(\mathcal{H}_0) = n-1$. We may assume w.l.o.g. that $n_{H'} = (h,1)$ for some $h \in \mathbb{R}^{n-1}$. Then we have

$$H \cap H' = \{(x, \lambda) \in \mathbb{R}^{n-1} \times \mathbb{R} \mid x \cdot s = 0 \land x \cdot h + \lambda = 0\}.$$

In particular, the map $H \cap H' \to s^{\perp}$ defined by $(x, \lambda) \mapsto x$ is surjective. Here, $s^{\perp} \subseteq \mathbb{R}^{n-1}$ denotes the set of vectors in \mathbb{R}^{n-1} orthogonal to s. Since \mathcal{H} is supersolvable, there exists a hyperplane $H'' \in \mathcal{H}_0$ with $H \cap H' \subseteq H''$. Let $n_{H''} = (t, 0) \in \mathbb{R}^n$ for some $t \in \mathbb{R}^{n-1}$ be one of the normal vectors of H''. Let $x \in s^{\perp}$. By the above, there is $\lambda \in \mathbb{R}$ such that $(x, \lambda) \in H \cap H' \subseteq H''$. Hence $(x, \lambda) \cdot (t, 0) = x \cdot t = 0$. This shows $s^{\perp} \subseteq t^{\perp}$, so $s^{\perp} = t^{\perp}$. But this implies $H = H'' \in \mathcal{H}_0$, a contradiction, so (i) is proved.

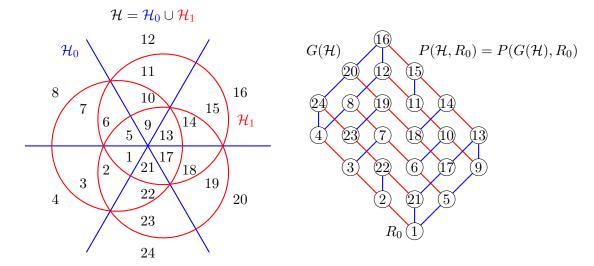


FIGURE 16. The graph of regions $G(\mathcal{H})$ of the arrangement of Figure 3. The lattice of regions $P(\mathcal{H}, R_0)$ is obtained by picking the region labeled 1 as canonical base region R_0 . The resulting lattice is isomorphic to the one shown in Figure 6 (a).

To prove (ii), note that every region $R \in \mathcal{R}(\mathcal{H}_0)$ is of the form $R = B \times \mathbb{R}$ for some $B \subseteq \mathbb{R}^{n-1}$ as $N(\mathcal{H}_0) \subseteq \mathbb{R}^{n-1} \times \{0\}$. Let $H \in \mathcal{H}_1$ with normal vector $n_H \in \mathbb{R}^n$. By (i), we may assume that $n_H = (h, 1)$ for some $h \in \mathbb{R}^{m-1}$. Then $\{x \in \mathbb{R}^n \mid x \cdot n_H > 0\} \cap R = \{(x, \lambda) \in B \times \mathbb{R} \mid x \cdot h + \lambda > 0\}$ is a proper subset of R. Hence H splits R into two regions.

The following key lemma captures the recursive structure of the graph of regions $G(\mathcal{H})$; see Figure 16.

Lemma 11. The graph of regions $G(\mathcal{H})$ has the following properties:

- (i) For any $R \in \mathcal{R}(\mathcal{H}_0)$, the subgraph of $G(\mathcal{H})$ induced by the regions in $\rho^{-1}(R)$ is a path of length $|\mathcal{H}_1|$.
- (ii) If two distinct regions $R, R' \in \mathcal{R}(\mathcal{H}_0)$ are adjacent in $G(\mathcal{H}_0)$, then in $G(\mathcal{H})$ the only edges between the two paths $\rho^{-1}(R)$ and $\rho^{-1}(R')$ are between the first two and the last two vertices, and possibly between other pairs of vertices at the same distance from the end vertices.
- (iii) If two distinct regions $R, R' \in \mathcal{R}(\mathcal{H}_0)$ are not adjacent in $G(\mathcal{H}_0)$, then in $G(\mathcal{H})$ there are no edges between the two paths $\rho^{-1}(R)$ and $\rho^{-1}(R')$.

Proof. We assume w.l.o.g. that the dimension of the ambient space equals $n = \text{rk}(\mathcal{H})$.

We first prove (i), and this part of the proof is illustrated in Figure 17 (a). For every region $R \in \mathcal{R}(\mathcal{H}_0)$ we can choose normal vectors $n_H \in \mathbb{R}^n$ for all $H \in \mathcal{H}_0$ such that R is the convex cone $R = \{x \in \mathbb{R}^n \mid x \cdot n_H > 0 \text{ for all } H \in \mathcal{H}_0\}$. By Lemma 10 (ii), every hyperplane $H \in \mathcal{H}_1$ splits R into two regions. For any $x \in R$, consider the ray $r_x(\lambda) := \lambda x \in R$, where λ is a real-valued parameter in the range $\lambda \in (0, \infty)$, which intersects each of the hyperplanes $H \in \mathcal{H}_1$. Specifically, for very small $\lambda > 0$ the point $r_x(\lambda)$ lies on one side of H, and for very large λ it lies on the other side of H. Consequently, the ray $r_x(\lambda)$ defines a total ordering of the hyperplanes in \mathcal{H}_1 , namely the order in which they are intersected by $r_x(\lambda)$ as the parameter λ increases. We claim that this total ordering is independent of the direction vector $x \in R$. Indeed, suppose for the sake of contradiction that there are points $x, y \in R$, values $0 < \lambda_{x,1} < \lambda_{x,2}, 0 < \lambda_{y,1} < \lambda_{y,2}$ and two hyperplanes $H, H' \in \mathcal{H}_1$ such that $r_x(\lambda_{x,1}) \in H$, $r_x(\lambda_{x,2}) \in H'$, $r_y(\lambda_{y,1}) \in H'$, and $r_y(\lambda_{y,2}) \in H$. Then for the direction vector $z = \nu x + (1 - \nu)y \in R$ for some suitable $\nu \in (0,1)$ and $\lambda > 0$ we would have $r_z(\lambda) \in H \cap H'$, i.e., some point in the intersection $H \cap H'$ that is not contained in any hyperplane from H_0 , a contradiction. It follows that for every $x \in R$, the ray $r_x(\lambda)$ intersects the hyperplanes of \mathcal{H}_1 in the same order, independent of x. We distinguish

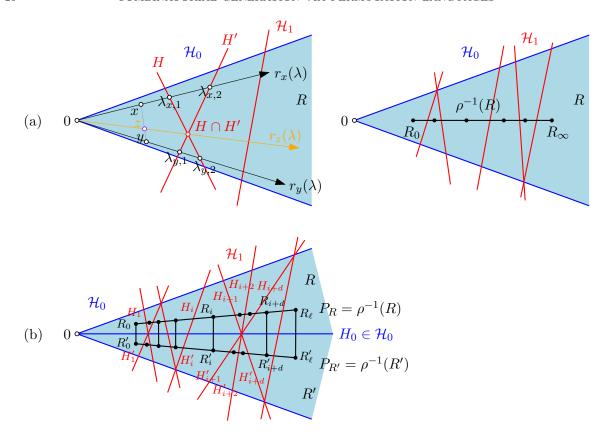


FIGURE 17. Illustration of the proof of Lemma 11.

two special regions $R_0, R_\infty \in \rho^{-1}(R)$, namely the one intersected by $r_x(\lambda)$ for very small $\lambda > 0$ and the one intersected for very large λ , respectively. By what we said before, each region in $\rho^{-1}(R) \setminus \{R_0, R_\infty\}$ is bounded by exactly two hyperplanes from \mathcal{H}_1 , and R_0, R_∞ are bounded by exactly one hyperplane from \mathcal{H}_1 . This completes the proof of (i).

The proof of (ii) is illustrated in Figure 17 (b). Let $R, R' \in \mathcal{R}(\mathcal{H}_0)$ be adjacent, i.e., these two regions are separated by exactly one hyperplane from \mathcal{H}_0 , which we denote by H_0 . By (i), $P_R := \rho^{-1}(R)$ and $P_{R'} := \rho^{-1}(R')$ are paths of length $\ell := |\mathcal{H}_1|$. Let $P_R = (R_0, \ldots, R_\ell)$ and $P_{R'} = (R'_0, \ldots, R'_\ell)$ be the sequences of regions along each of the paths. Each edge of the paths corresponds to a hyperplane from \mathcal{H}_1 , so for $i = 1, \ldots, \ell$ we let $H_i \in \mathcal{H}_1$ denote the hyperplane separating R_{i-1} and R_i , and $H'_i \in \mathcal{H}_1$ the hyperplane separating R'_{i-1} and R'_i . As any region R_i , $i \in \{0, \ldots, \ell\}$, is separated from any region R'_j , $j \in \{0, \ldots, \ell\}$, by $H_0 \in \mathcal{H}_0$, we have that R_i and R'_j are adjacent in $G(\mathcal{H})$ if and only if they are not separated by any hyperplanes from \mathcal{H}_1 .

We observe the following for regions R_i and R'_i that are adjacent in the graph $G(\mathcal{H})$. First of all we must have $\{H_1, \ldots, H_i\} = \{H'_1, \ldots, H'_i\}$ and $\{H_{i+1}, \ldots, H_\ell\} = \{H'_{i+1}, \ldots, H'_\ell\}$. Now let d > 0 be the smallest integer such that $H_{i+d} = H'_{i+1}$. If d > 1, then as H_{i+1} and H_{i+d} appear in reversed order along P_R and $P_{R'}$, then these two hyperplanes must intersect, and their intersection lies in H_0 . Furthermore, by a similar argument and the linearity of the hyperplanes we must have $(H_{i+1}, H_{i+2}, \ldots, H_{i+d}) = (H'_{i+d}, \ldots, H'_{i+2}, H'_{i+1})$, i.e., the next d hyperplanes appear in the opposite order along the two paths P_R and $P_{R'}$. Note that this conclusion is trivially true if d = 1. For any d > 0 we obtain that R_{i+d} and R'_{i+d} are separated by no hyperplane from \mathcal{H}_1 (only by $H_0 \in \mathcal{H}_0$), and therefore they are adjacent in the graph $G(\mathcal{H})$. Furthermore, one can check that R_j for $j = i+1, \ldots, i+d$, is separated from R'_k , $k = 0, \ldots, i+d-1$, by at least one hyperplane from \mathcal{H}_1 , in addition to $H_0 \in \mathcal{H}_0$, and therefore R_j and R'_k are not adjacent in the

graph $G(\mathcal{H})$. Observing that R_0 and R'_0 are adjacent in $G(\mathcal{H})$, and applying the aforementioned observation inductively for $i = 0, \ldots, \ell$ proves (ii).

To prove (iii), note that as R, R' are not adjacent in $G(\mathcal{H}_0)$, they are separated by at least two hyperplanes $H, H' \in \mathcal{H}_0$. It follows that any two regions $P \in \rho^{-1}(R)$ and $P' \in \rho^{-1}(R')$ are also separated by both H and H', and therefore P and P' are not adjacent in $G(\mathcal{H})$.

2.10. Lattice of regions. For any graph G and any two vertices x, y in G, we write $d_G(x, y)$ for the distance between x and y in G, i.e., for the length of the shortest path connecting x and y. Let G = (V, E) be a connected bipartite graph, and let $x_0 \in V$ be one of its vertices. We let $P(G, x_0)$ be the graded poset on the ground set V whose cover graph is obtained from G by orienting all edges away from x_0 . Formally, we define the rank of any $y \in V$ in $P(G, x_0)$ as $\operatorname{rk}(y) := d_G(x_0, y)$. As G is assumed to be bipartite, there are no odd cycles and hence no edges between vertices of the same rank, so this is well-defined.

Using this definition, by choosing one region $R_0 \in \mathcal{R}$ as base region, one can orient the edges of the graph of regions $G(\mathcal{H})$ away from R_0 . As $G(\mathcal{H})$ is always connected and bipartite, this defines a graded poset $P(\mathcal{H}, R_0) := P(G(\mathcal{H}), R_0)$, called the poset of regions; see Figure 16. It is easy to see that in this poset the rank of any $R \in \mathcal{R}$ is given by the number of hyperplanes from \mathcal{H} that separate R_0 and R, i.e.,

$$\operatorname{rk}(R) = |\{H \in \mathcal{H} \mid R_0 \text{ and } R \text{ lie on opposite sides of } H\}.$$

Björner, Edelman, and Ziegler [BEZ90] proved that for a supersolvable arrangement \mathcal{H} , there always exists a region $R_0 \in \mathcal{R}(\mathcal{H})$ such that the poset $P(\mathcal{H}, R_0)$ is a lattice.

To state the next lemma, we define the notion of a *canonical* region for a supersolvable arrangement \mathcal{H} of rank $n \geq 1$. If $n \leq 2$, then every region of $\mathcal{R}(\mathcal{H})$ is canonical. If $n \geq 3$, we consider the map ρ defined in (1), and for any $R \in \mathcal{R}(\mathcal{H})$ we define the fiber $\varphi(R) := \rho^{-1}(\rho(R)) \subseteq \mathcal{R}(\mathcal{H})$. Then a region $R_0 \in \mathcal{R}(\mathcal{H})$ is canonical if $\rho(R_0)$ is canonical for \mathcal{H}_0 and if R_0 is the first or last vertex of the path induced by the fiber $\varphi(R_0)$ in $G(\mathcal{H})$ (recall Lemma 11 (i)).

Lemma 12 ([BEZ90, Thm. 4.6]). If R_0 is a canonical region, then the poset $P(\mathcal{H}, R_0)$ is a lattice.

3. Proofs of Theorems 1 and 2

3.1. **Graphic suspension.** For a graph G = (V, E) and subset of vertices $U \subseteq V$ we write G[U] for the subgraph of G induced by U. Given two graphs H = (V, E) and F = (V', E'), the Cartesian product $H \times F$ is the graph on the vertex set $V \times V'$ with edge set

$$\{\{(x, x'), (y, y')\} \mid (\{x, y\} \in E \land x' = y') \lor (x = y \land \{x', y'\} \in E')\}.$$

We write P_{ℓ} for the path of length ℓ on the vertices $0, 1, \dots, \ell$, and we write $G \simeq H$ to indicate that two graphs G and H are isomorphic.

Let G and H=(V,E) be graphs, and let $\ell \geq 1$ be an integer. We say that G is ℓ -suspended over H if G is a spanning subgraph of the Cartesian product $H \times P_{\ell}$ such that $G[\{v\} \times P_{\ell}] \simeq P_{\ell}$ for all $v \in V$, and furthermore $G[V \times \{0\}] \simeq H$ and $G[V \times \{\ell\}] \simeq H$. In words, G is a spanning subgraph of the Cartesian product of H with a path of length ℓ in which all copies of the path are present, and the first and last copy of H along the paths are present, whereas edges of the other intermediate copies of H along the paths may or may not be present. The notion of suspension is introduced to provide a graphic translation and clarification of the structure described in [BEZ90]. Specifically, using this definition, Lemma 11 can be restated more concisely as follows.

Lemma 13. The graph $G(\mathcal{H})$ is (isomorphic to a graph that is) ℓ -suspended over $G(\mathcal{H}_0)$ for $\ell := |\mathcal{H}_1|$.

Lemma 14. Suppose that G is ℓ -suspended over H. Then for any two vertices x, y of H and any $i \in \{0, ..., \ell\}$ we have $d_G((x, 0), (y, i)) = d_G((x, \ell), (y, \ell - i)) = d_H(x, y) + i$.

Proof. This is immediate from the definition of suspension via the Cartesian product. \Box

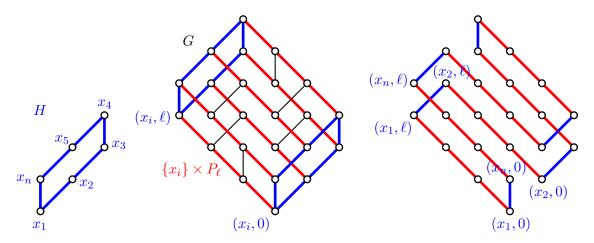


FIGURE 18. Illustration of Lemma 15, continuing the example from Figure 16 (also see Figure 10).

We also need a second (non-equivalent) definition of suspension, which does not use Cartesian products. Given a sequence $(\lambda(x))_{x\in V}$ of integers $\lambda(x)\geq 1$, we say that G is λ -suspended over H=(V,E) if G is the union of the paths $((x,0),(x,1),\ldots,(x,\lambda(x)))$ and the edges $\{(x,0),(y,0)\}$ and $\{(x,\lambda(x)),(y,\lambda(y))\}$ for all $\{x,y\}\in E$. In words, G is λ -suspended over H if it consists of two disjoint copies of H, which are connected via disjoint paths between corresponding pairs of vertices, where the lengths of the paths are specified by the function λ . Note that if G is λ -suspended for the constant function $\lambda:=\ell$ for all $x\in V$, then it is ℓ -suspended, and furthermore it is the ℓ -suspended graph with the fewest possible edges.

The following easy lemma asserts that suspended graphs behave particularly nicely with regards to Hamiltonian paths/cycles.

Lemma 15. Suppose that G is λ -suspended over H. If H has a Hamiltonian path, then G has a Hamiltonian path.

Furthermore, if G is ℓ -suspended over H and H has a Hamiltonian cycle of even length, then G has a Hamiltonian cycle of even length.

Lemma 15 and its proof are illustrated in Figure 18.

Proof. Let V be the vertex set of H and let (x_1, \ldots, x_n) be a Hamiltonian path in H. Then for each odd index k, replace x_k by the path $((x_k, 0), (x_k, 1), \ldots, (x_k, \lambda(x_k)))$ in G. Similarly, for each even index k, replace x_k by the path $((x_k, \lambda(x_k)), (x_k, \lambda(x_k) - 1), \ldots, (x_k, 0))$ in G. Consecutive pairs of end vertices from these subpaths are of the form $\{(x_k, 0), (x_{k+1}, 0)\}$ if k is even and $\{(x_k, \lambda(x_k)), (x_{k+1}, \lambda(x_{k+1}))\}$ if k is odd, respectively, i.e., these are edges in G. We conclude that the resulting sequence of vertices in G is a Hamiltonian path.

To prove the second part of the lemma assume that n is even and that (x_1, \ldots, x_n) is Hamiltonian cycle in H, i.e., a Hamiltonian path such that x_n and x_1 are adjacent vertices. Then the above argument yields a Hamiltonian path from $(x_1, 0)$ to $(x_n, 0)$ in G, so the pair $\{(x_1, 0), (x_n, 0)\}$ is an edge in $G[V \times \{0\}] \simeq H$. We thus also obtain a Hamiltonian cycle in G of length $(\ell + 1) \cdot n$, which is even.

3.2. Congruences on lattices with suspended cover graphs. The following crucial proposition describes how a congruence acts on a lattice whose cover graph is suspended.

Proposition 16. Suppose that G is ℓ -suspended over H = (V, E), and let x_0 be a vertex of H such that $L_H := P(H, x_0)$ and $L_G := P(G, (x_0, 0))$ are both lattices. Then $L_G[V \times \{0\}]$ and $L_G[V \times \{\ell\}]$ are two sublattices isomorphic to L_H , and $(x, 0) < (x, 1) < \cdots < (x, \ell)$ is a chain for each $x \in V$. Furthermore, for any lattice congruence \equiv on L_G we have the following:

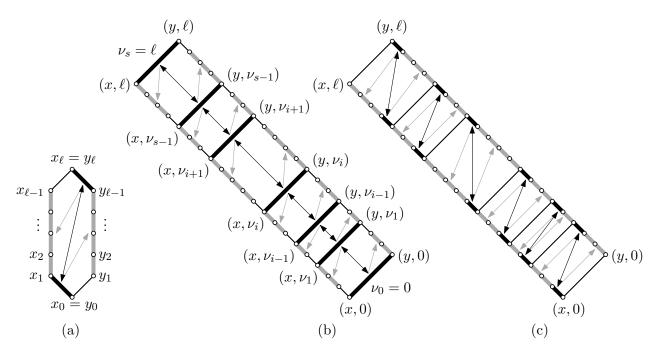


FIGURE 19. Illustration of the proof of Proposition 16. The bold lines indicate pairs of elements that are in the same equivalence class. The arrows indicate (directed) forcing constraints between pairs of equivalent pairs.

- (i) The restriction $\equiv^* := \{(x,y) \mid (x,0) \equiv (y,0)\}$ is a lattice congruence on L_H .
- (ii) For any equivalence class X of \equiv , the projection $p(X) := \{x \mid (x, i) \in X\}$ is an equivalence class of \equiv^* .
- (iii) The quantity $\lambda(x) := |\{[(x,i)] \mid i = 0, \dots, \ell\}| 1$, defined for any $x \in V$, is either equal to 0 for all $x \in V$ or strictly positive for all $x \in V$.
- (iv) If the cover graph of L_H/\equiv^* has a Hamiltonian path, then the cover graph of L_G/\equiv has a Hamiltonian path.

Of course, by symmetry the poset $P(G,(x_0,\ell))$ is also a lattice.

Proof. For any $x \in V$ and $r \in \{0, ..., \ell\}$ we define $C_{\lambda}(x) := \{(x, i) \mid 0 \le i \le \lambda\}$ and $C(x) := C_{\ell}(x)$. The fact that $L_G[V \times \{0\}]$ and $L_G[V \times \{\ell\}]$ are sublattices isomorphic to L_H and that the elements of C(x) form the chain $(x, 0) \le (x, 1) \le \cdots \le (x, \ell)$ for each $x \in V$ follows directly from Lemma 14 and our definition of rank in $L_G = P(G, (x_0, 0))$.

We now prove (i). For any two $x, y \in V$ we have

$$(x,0) \wedge (y,0) = (x \wedge y,0).$$
 (2)

Now consider four elements $x, x', y, y' \in V$ satisfying $x \equiv^* x'$ and $y \equiv^* y'$. From the definition of restriction, we have $(x,0) \equiv (x',0)$ and $(y,0) \equiv (y',0)$. Applying the definition of lattice congruence to \equiv , we obtain that $(x,0) \wedge (y,0) \equiv (x',0) \wedge (y',0)$. Applying (2) to this relation yields $(x \wedge y,0) \equiv (x' \wedge y',0)$, from which we obtain $x \wedge y \equiv^* x' \wedge y'$ with the definition of restriction. The argument for joins is dual, proving that \equiv^* is indeed a lattice congruence of L_H .

The proof of (ii) and (iii) needs the following auxiliary claim (a), which is immediate from the definition of lattice congruence; see Figure 19 (a): Consider the lattice L formed by two chains $x_0 \leqslant x_1 \leqslant \cdots \leqslant x_\ell$ and $y_0 \leqslant y_1 \leqslant \cdots \leqslant y_\ell$ with $x_0 = y_0$, $x_\ell = y_\ell$, and $x_i \neq y_j$ for any $i, j \in \{1, \ldots, \ell-1\}$, and let \equiv be a lattice congruence on L. If $x_0 \equiv x_1$, then we have $y_\ell \equiv y_{\ell-1} \equiv \cdots \equiv y_1$. If $y_\ell \equiv y_{\ell-1}$, then we have $x_0 \equiv x_1 \equiv \cdots \equiv x_{\ell-1}$.

Given $\{x,y\} \in E$ and a sequence $\nu = (\nu_0, \dots, \nu_s)$ where $s \ge 1$ and $0 = \nu_0 < \nu_1 < \dots < \nu_s = \ell$, where $s \ge 1$, the (ℓ, ν) -ladder is the ℓ -suspended graph over $\{x,y\}$ in which precisely the edges $\{(x,\nu_i),(y,\nu_i)\}$ are present for all $i=0,\dots,s$; see Figure 19 (b).

To prove (ii) and (iii), consider a cover x < y in L_H . Then the interval $[(x,0),(y,\ell)]$ in L_G contains exactly the elements $C(x) \cup C(y)$ and its cover graph is an (ℓ,ν) -ladder L for some sequence $\nu = (\nu_0, \dots, \nu_s)$ determined by G. From our auxiliary claim (a), we obtain the following about the ladder L; see Figure 19 (b): If $(x,\nu_i) \equiv (y,\nu_i)$, then we have $(x,\nu_{i-1}) \equiv (y,\nu_{i-1})$ and $(x,\nu_{i+1}) \equiv (y,\nu_{i+1})$, and furthermore $(x,\nu_i) \equiv (x,\nu_i-1) \equiv \cdots \equiv (x,\nu_{i-1}+1)$ and $(y,\nu_i) \equiv (y,\nu_i+1) \equiv \cdots \equiv (y,\nu_{i+1}-1)$. Applying this observation exhaustively, we note that either $(x,\nu_i) \not\equiv (y,\nu_i)$ for all $i=0,\dots,s$, or $(x,\nu_i) \equiv (y,\nu_i)$ for all $i=0,\dots,s$, and in the latter case for every $i \in \{0,\dots,\ell\}$ there is some $j \in \{0,\dots,\ell\}$ with $(x,i) \equiv (y,j)$ and for every $j \in \{0,\dots,\ell\}$ there is some $i \in \{0,\dots,\ell\}$ with $(x,i) \equiv (y,j)$. This proves property (ii).

To prove (iii), note that the auxiliary claim (a) applied to the ladder L yields that $(x,0) \equiv (x,1) \equiv \cdots \equiv (x,\ell)$ if and only if $(y,0) \equiv (y,1) \equiv \cdots \equiv (y,\ell)$; see Figure 19 (c). Consequently, the number of equivalence classes of \equiv along every chain C(x), $x \in V$, is either equal to 1 for all $x \in V$ or strictly larger than 1 for all $x \in V$, which proves (iii).

It remains to prove (iv). If $\lambda(x) = 0$ for all $x \in V$, then we have $L_G \simeq L_H$ and therefore $L_G/\equiv \simeq L_H/\equiv^*$, so the claim is trivial. On the other hand, if $\lambda(x) \geq 1$ for all $x \in V$, which is the only other possibility by (iii), then the cover graph of L_G/\equiv has an induced subgraph that is λ -suspended over L_H/\equiv^* by (ii) and Lemma 5. Using the assumption that the cover graph of L_H/\equiv^* has a Hamiltonian path, applying the first part of Lemma 15 proves that the cover graph of L_G/\equiv also has a Hamiltonian path.

This completes the proof.

3.3. **Proofs of Theorems 1 and 2.** We are now ready to present the proof of our two main theorems.

Proof of Theorem 1. The proof is by induction on the rank $n \geq 2$ of \mathcal{H} . If n = 2, then the graph of regions $G(\mathcal{H})$ is a cycle of length $2|\mathcal{H}|$ and thus clearly has a Hamiltonian cycle of even length. If $n \geq 3$, then consider the decomposition of \mathcal{H} into two subarrangements \mathcal{H}_0 and \mathcal{H}_1 . As \mathcal{H}_0 is a supersolvable arrangement of rank n-1, we know by induction that $G(\mathcal{H}_0)$ has a Hamiltonian cycle of even length. Applying Lemma 13 and the second part of Lemma 15 yields that $G(\mathcal{H})$ has a Hamiltonian cycle of even length.

Proof of Theorem 2. The proof is by induction on the rank $n \ge 1$ of \mathcal{H} . If n=1, then the cover graph of L is a single edge. If n=2, then the cover graph of L is a cycle of length $2|\mathcal{H}|$. By Lemma 5 (i), the equivalence classes of \equiv are intervals, so in the two aforementioned cases the cover graph of L/\equiv is either a cycle, a single edge, or a single vertex, hence it has a Hamiltonian path.

If $n \geq 3$, then consider the decomposition of \mathcal{H} into two subarrangements \mathcal{H}_0 and \mathcal{H}_1 . Let $R_0 \in \mathcal{R}(\mathcal{H})$ be a canonical base region, then $\rho(R_0)$ is canonical for \mathcal{H}_0 . We know that $L' := P(\mathcal{H}_0, \rho(R_0)) = P(G(\mathcal{H}_0), \rho(R_0))$ and $L = P(\mathcal{H}, R_0) = P(G(\mathcal{H}), R_0)$ are lattices by Lemma 12. Furthermore, we know that $G(\mathcal{H})$ is (isomorphic to a graph that is) ℓ -suspended over $G(\mathcal{H}_0)$ for $\ell := |\mathcal{H}_1|$ by Lemma 13.

For any $R \in \mathcal{R}(\mathcal{H})$, we let $p(R) \in \mathcal{R}(\mathcal{H})$ be the minimum of the chain on the fiber $\varphi(R)$ in L. Let \equiv be lattice congruence on L. By Proposition 16 (i) the restriction $\equiv^* := \{(\rho(X), \rho(Y)) \mid X = p(X) \text{ and } Y = p(Y) \text{ and } X \equiv Y\}$ is a lattice congruence on L'. As \mathcal{H}_0 is a supersolvable arrangement of rank n-1, we obtain by induction that the cover graph of L'/\equiv^* has a Hamiltonian path. Therefore, applying Proposition 16 (iv) yields that the cover graph of L also has a Hamiltonian path. This completes the proof.

4. Zigzag languages of signed permutations

In this section we develop a theory of zigzag languages of signed permutations, analogous to the theory for (unsigned) permutations presented in [HHMW22] and discussed in Section 1.3. We also show applications to generating pattern-avoiding signed permutations and the corresponding combinatorial objects. With this paper we hope to encourage further study of pattern-avoiding signed permutations by the pattern-avoidance community.

4.1. Patterns in permutations and signed permutations. Recall from Section 1.8.3 that we use bars to denote negative signs in a signed permutation. We write B_n for the set of signed permutations of [n] in full notation (length 2n), and we will use this representation throughout. We index the entries in a signed permutation $\pi \in B_n$ by integers from $[\overline{n}] \cup [n]$, i.e., we consider π as a string $\pi = (a_{\overline{n}}, \ldots, a_{\overline{1}}, a_1, \ldots, a_n)$, where $a_{\overline{i}} = \overline{a_i}$ holds for all $i \in [\overline{n}] \cup [n]$. Notably, the index 0 does not exist. Furthermore, we write S_n for the set of (unsigned) permutations of length n. We write ε for the empty string, i.e., we have $B_0 = \{\varepsilon\}$ and $S_0 = \{\varepsilon\}$. The identity signed permutation is $\mathrm{id}_n := \overline{n} \cdots 2\overline{1}12 \cdots n \in B_n$.

We say that a string π of integers contains a pattern $\tau \in S_k$, if there is a subsequence of k (not necessarily consecutive) entries in π in the same relative order as in τ . Otherwise we say that π avoids τ . We apply this definition in both cases when $\pi \in S_n$ is a permutation of length n, or when $\pi \in B_n$ is a signed permutation in full notation, where in the second case the signs of entries are taken into account when comparing relative positions. For example, $\pi = \overline{42311324} \in B_4$ contains the pattern $\tau = 231$, as witnessed by the subsequence $\overline{2}3\overline{3}$, or alternatively by the subsequences $\overline{2}1\overline{3}$ and $\overline{2}1\overline{3}$. We write $S_n(\tau)$ for the set of all permutations from S_n that avoid τ , and similarly $B_n(\tau)$ for the set of all signed permutations from B_n that avoid τ .

Generalizing the notion of pattern-containment for signed permutations, we also allow underlining or overlining some of entries in a pattern $\tau \in S_n$, with the interpretation that an underlined entry of the pattern has to match a strictly positive entry in the signed permutation, whereas an overlined entry of the pattern has to match a strictly negative entry in the signed permutation. We refer to such a pattern as a marked (signed) permutation pattern. For example, in the signed permutation π from before, the first two occurrences of the pattern 231 are also occurrences of the pattern 231, whereas the third occurrence is also an occurrence of the pattern 2\overline{3}1. Note that there are some trivial equivalences of patterns, in the sense that the set of avoiders are identical. Specifically, if an entry in a marked pattern is underlined, then all larger entries may also be underlined. Similarly, if an entry is overlined, then all smaller entries may also be overlined. For example, the patterns 231 and 231 are equivalent in this sense, and so are the patterns $2\overline{3}1$, $\overline{2}\overline{3}1$, $\overline{2}\overline{3}1$ and $\overline{2}\overline{3}\overline{1}$.

4.2. **Jumps in signed permutations.** Given a signed permutation $\pi \in B_n$, a left hop of the value $a_i > 0$, $i \neq \overline{n}$, is a transposition of the value a_i with the value b directly to its left, i.e., either with $b := a_{i-1}$ if $i \neq 1$ or $b := a_{\overline{1}}$ if i = 1, subject to the constraint that $a_i > b$. Similarly, a right hop of the value $a_i > 0$, $i \neq n$, is a transposition of the value a_i with the value b directly to its right, i.e., either with $b := a_{i+1}$ if $i \neq \overline{1}$ or $b := a_1$ if $i = \overline{1}$, subject to the constraint that $a_i > b$. A left jump of $a_i > 0$ by d steps is sequence of d consecutive left hops of the value a_i . Similarly, a right jump of $a_i > 0$ by d steps is a sequence of d consecutive right hops of the value a_i . In all these operations the modifications are applied to the permutation in full notation in a symmetric way, so that the property $a_{\overline{i}} = \overline{a_i}$ for all $i \in [\overline{n}] \cup [n]$ is maintained. For a subset $L_n \subseteq B_n$ and a permutation $\pi \in L_n$, a jump in π is minimal w.r.t. L_n if no jump of the same value by fewer steps yields a signed permutation from L_n . For example, if $L_4 := B_4(231)$, then in the permutation $\pi = \overline{21}43\overline{3412} \in L_4$ a right jump of the value 4 by 5 steps is minimal, yielding the permutation $\overline{421}3\overline{3}124 \in L_4$ as each of the intermediate permutations $\overline{21}34\overline{4312}$, $\overline{21433412}$, $\overline{241}3\overline{3412}$ contains the marked pattern $\underline{231}$.

We propose the following algorithm to generate a set of signed permutations by minimal jumps. An analogous algorithm for (unsigned) permutations was first presented in [HHMW22].

Algorithm J (Greedy minimal jumps for signed permutations). This algorithm attempts to greedily generate a set of signed permutations $L_n \subseteq B_n$ using minimal jumps starting from an initial signed permutation $\pi_0 \in L_n$.

- **S1.** [Initialize] Visit the initial signed permutation π_0 .
- **S2.** [Jump] Generate an unvisited signed permutation from L_n by performing a minimal jump of the largest possible positive value in the most recently visited signed permutation. If no such jump exists, or the jump direction is ambiguous, then terminate. Otherwise visit this signed permutation and repeat S2.
- 4.3. Zigzag languages of signed permutations. For any $\pi \in B_{n-1}$ and integer $i \in [\overline{n}] \cup [n]$ we write $c_i(\pi) \in B_n$ for the signed permutation obtained by inserting the largest value n and smallest value \overline{n} at positions i and \overline{i} , respectively. In particular, $c_{\overline{n}}(\pi) = n \pi \overline{n}$ and $c_n(\pi) = \overline{n} \pi n$. Conversely, for any $\pi \in B_n$ we write $p(\pi)$ for the permutation in B_{n-1} obtained by deleting the extremal values n and \overline{n} . Clearly, we have $p(c_i(\pi)) = \pi$. For example, for $\pi \in 1\overline{2}2\overline{1}$ we have

$$(c_{\bar{3}}(\pi), c_{\bar{2}}(\pi), c_{\bar{1}}(\pi), c_{1}(\pi), c_{2}(\pi), c_{3}(\pi)) = (31\bar{2}2\bar{1}\bar{3}, 13\bar{2}2\bar{3}\bar{1}, 1\bar{2}3\bar{3}2\bar{1}, 1\bar{2}\bar{3}32\bar{1}, 1\bar{3}\bar{2}23\bar{1}, \bar{3}1\bar{2}2\bar{1}3).$$

A set $L_n \subseteq B_n$ is a zigzag language, if either n=0 and $L_0=\{\varepsilon\}$, or if $n\geq 1$ and $L_{n-1}:=\{p(\pi)\mid \pi\in L_n\}$ is a zigzag language such that for every $\pi\in L_{n-1}$ we have that $c_{\overline{n}}(\pi)=n\,\pi\,\overline{n}\in L_n$ and $c_n(\pi)=\overline{n}\,\pi\,n\in L_n$.

Theorem 17. Given any zigzag language of signed permutations L_n and initial permutation $\pi_0 = \mathrm{id}_n$, Algorithm J visits every signed permutation from L_n exactly once.

The proof is analogous to the proof of Theorem 1 in [HHMW22].

Proof. Given a zigzag language L_n , we define a sequence $J(L_n)$ of all signed permutations from L_n , and we prove that Algorithm J generates the permutations of L_n exactly in this order. For any $\pi \in L_{n-1}$ we let $\overrightarrow{c}(\pi)$ be the sequence of all $c_i(\pi) \in L_n$ for $i = \overline{n}, \overline{n-1}, \ldots, \overline{1}, 1, \ldots, n$, starting with $c_{\overline{n}}(\pi)$ and ending with $c_n(\pi)$, and we let $\overleftarrow{c}(\pi)$ denote the reverse sequence, i.e., it starts with $c_n(\pi)$ and ends with $c_{\overline{n}}(\pi)$. In words, those sequences are obtained by inserting the new largest value n from left to right, or from right to left, respectively, into π , in all possible positions that yield a signed permutation from L_n , skipping the positions that yield a permutation that is not in L_n . The sequence $J(L_n)$ is defined recursively as follows: If n = 0 then we define $J(L_0) := \varepsilon$, and if $n \ge 1$ then we consider the finite sequence $J(L_{n-1}) := (\pi_1, \pi_2, \ldots)$ and define

$$J(L_n) := \overleftarrow{c}(\pi_1), \overrightarrow{c}(\pi_2), \overleftarrow{c}(\pi_3), \overrightarrow{c}(\pi_4), \ldots,$$

i.e., this sequence is obtained from the previous sequence by inserting the new largest value n in all possible positions alternatingly from right to left, or from left to right.

A straightforward induction shows that Algorithm J generates the signed permutations in L_n precisely in the order $J(L_n)$. We omit the details. The crucial observation is that a minimal jump in a signed permutation $\pi \in L_{n-1}$ is also a minimal jump in the two permutations $c_{\overline{n}}(\pi) = n \pi \overline{n}$ and $c_n(\pi) = \overline{n} \pi n$, because the values n and \overline{n} are inserted at the boundaries.

Remark 18. It is easy to see that in the ordering of signed permutations $J(L_n)$, the first and last permutation differ in a minimal jump if and only if $|L_i|$ is even for all i = 1, ..., n-1. Consequently, if this condition holds, then the Gray code produced by Algorithm J is actually cyclic.

The following lemma provides easily checkable sufficient conditions on a signed permutation pattern τ so that $B_n(\tau)$ is a zigzag language.

We say that an infinite sequence of sets L_0, L_1, \ldots with $L_i \subseteq B_i$ for all $i \ge 0$ is hereditary if $L_{i-1} = p(L_i)$ holds for all $i \ge 1$. We say that a signed permutation pattern τ is tame, if $B_n(\tau)$, $n \ge 0$, is a hereditary sequence of zigzag languages.

Lemma 19. If in a marked pattern $\tau \in S_k$, the smallest value 1 is underlined or not at the leftmost or rightmost position, and the largest value k is overlined or not at the leftmost or rightmost position, then τ is tame. In particular, if $\tau \in S_k$ has neither the smallest value 1 nor the largest value k at the leftmost or rightmost position, then it is tame.

Four of the signed permutation patterns shown in Table 3 classify as tame by Lemma 19.

Proof. Consider a marked pattern $\tau \in S_k$ satisfying the conditions of the lemma. We prove that $B_n(\tau), n \geq 0$, is a hereditary sequence of zigzag languages by induction on n. The induction basis n=0 is trivial, as $B_0(\tau)=\{\varepsilon\}$. For the induction step, we assume that $B_{n-1}(\tau)$ is a zigzag language. If $\pi \in B_n$ avoids τ , then clearly $p(\pi) \in B_{n-1}$ also avoids τ , as removing the values n and \overline{n} cannot create an occurrence of the pattern τ . If follows that $B_{n-1}(\tau) \supseteq p(B_n(\tau))$. We next show that if $\pi \in B_{n-1}$ avoids τ , then both $c_{\overline{n}}(\pi) = n \pi \overline{n}$ and $c_n(\pi) = \overline{n} \pi n$ avoid τ . As $p(c_{\overline{n}}(\pi)) = \pi$ and $p(c_n(\pi)) = \pi$, this proves that $B_{n-1}(\tau) = p(B_n(\tau))$ and it also establishes the zigzag property. Indeed, any occurrence of τ in one of the signed permutations $c_{\overline{n}}(\pi)$ and $c_n(\pi)$ must contain \overline{n} or n, and those entries must match the smallest and largest values in τ with the correct sign, respectively, which must be at the leftmost or rightmost position in τ . This, however, is ruled out by the assumptions on τ stated in the lemma.

TABLE 3. Various signed permutation patterns τ and corresponding counting sequences and combinatorial objects. The third column indicates whether Algorithm J succeeds to generate the set of all pattern-avoiding signed permutations. A pattern being tame is sufficient for this property by Theorem 17.

au	tame	Alg J	$ B_n(\tau) , n = 0, 1, \dots, 6$	OEIS	combinatorial objects
$\begin{array}{c} 2\underline{1} \\ \underline{2}\overline{1} \end{array}$	no		1, 2, 4, 8, 16, 32, 64	A000079	binary strings
$2\overline{1}$	no	yes	1, 1, 2, 6, 24, 120, 720	A000142	permutations \rightarrow Sec. 5.1
$23\underline{1}$	yes	yes	1, 2, 8, 40, 224, 1344, 8448	A151374, A052701	two-colored binary trees \rightarrow Sec. 4.3.1
213	no	yes	1, 2, 4, 8, 16, 32, 64	A000079	binary strings \rightarrow Sec. 5.2
$\underline{2}31$	no	yes	1, 2, 6, 20, 70, 252, 924	A000984	symmetric triangulations \rightarrow Sec. 5.3
$\bar{2}31$	no		—	—II—	
$2\bar{3}\bar{1}$	no		1, 2, 4, 10, 34, 154, 874	A003422	
$\bar{2}\underline{3}1$	no		1, 2, 6, 22, 94, 462, 2606	A193763	
$23\overline{1}$	no		1, 2, 4, 9, 23, 65, 197	A014137	
$2\underline{3}1$	no		<u>—п—</u>	——II——	
3142	yes	yes	1, 2, 7, 32, 169, 974, 5947	A115197	
3412	yes	yes	1, 2, 7, 33, 183, 1118, 7281	A086618	
$\underline{3}41\bar{2}$	yes	yes	1, 2, 7, 34, 209, 1546, 13327	A002720	

Remark 20. So far we have only considered avoidance of a single pattern. However, the theory generalizes straightforwardly to avoiding multiple patterns. This is based on the observation that the intersection of two zigzag languages is again a zigzag language (also the hereditary property is preserved). Consequently, if we have tame patterns $\tau_1, \ldots, \tau_\ell$, then the set of signed permutations avoiding each of $\tau_1, \ldots, \tau_\ell$ is also a zigzag language that can be generated by Algorithm J.

Note that Theorem 17 provides a sufficient condition for Algorithm J to generate a set of signed permutations $L_n \subseteq B_n$ exhaustively, namely to be a zigzag language. To close this section, we consider one of the tame patterns mentioned in Table 3 and present the resulting listing of combinatorial objects obtained from applying Algorithm J.

4.3.1. The tame pattern 231. Let $\tau := 231$. We present a bijection between signed permutations avoiding the pattern τ and binary trees whose vertices are colored in one of two colors, red or blue, say. We write \mathcal{T}_n for the sets of binary trees with n vertices, and \mathcal{T}'_n for the set of two-colored binary trees with n vertices.

Theorem 21. Let $\tau := 23\underline{1}$. There is a bijection $g : \mathcal{T}'_n \to B_n(\tau)$, and therefore $|B_n(\tau)| = 2^n C_n$, where $C_n := \frac{1}{n+1} \binom{2n}{n}$ are the Catalan numbers.

Proof. The bijection g uses as a building block the following well-known bijection $f: \mathcal{T}_n \to S_n(231)$ between binary trees and 231-avoiding permutations; see Figure 22. We consider each tree $T \in \mathcal{T}_n$ as a binary search tree, with the vertices labeled bijectively with integers $1, \ldots, n$ such that the label of every vertex is larger than all labels in the left subtree and smaller than all labels in the right subtree. Given a tree T with root r and left and right subtrees T^L and T^R , we define $f(T) := (r, f(T^L), f(T^R))$, i.e., we first write the label of the root, then recursively all labels in the left subtree, then recursively all labels in the right subtree. One can easily check that this mapping f is indeed a bijection between \mathcal{T}_n and $S_n(231)$, and that it induces a one-to-one correspondence between tree rotations and minimal jumps in 231-avoiding permutations (for details, see [HHMW22, Sec. 3.3]).

To construct the desired bijection $g: \mathcal{T}'_n \to B_n(\tau)$, note that $\tau = 23\underline{1} = \underline{231}$, i.e., when searching for occurrences of this pattern in a signed permutation in full notation, one only needs to consider the positive entries, whereas all negative entries can be completely ignored. Hence, the mapping g can be defined as follows: Given a two-colored tree $T' \in \mathcal{T}'_n$, we first consider the underlying uncolored tree $T \in \mathcal{T}_n$, and construct the permutation $\pi := f(T) \in S_n(231)$. In the second step, we define which n out of the 2n entries of g(T') receive a negative sign. Specifically, the entry at position $i \in [n]$ of g(T') receives a negative sign if vertex i is colored red in T', and a positive sign if vertex i is colored blue. Note that in the first case this forces the entry at position i to have a positive sign, and in the second case this forces entry at position i to have a negative sign. In the third step, we fill the string π into the n positions of g(T') with a positive sign from left to right, which defines the signed permutation g(T') completely (the values of the negatively signed entries are forced). It can be checked directly that this mapping g is indeed a bijection between \mathcal{T}'_n and $B_n(\tau)$.

It is known that $|\mathcal{T}_n| = C_n$, and since each tree with n vertices can be colored in 2^n ways, this proves the claimed counting formula.

As τ is tame by Lemma 19, we may use Algorithm J to generate $B_n(\tau)$; see Figure 20. The resulting Gray code is cyclic by Remark 18. The minimal jumps performed by Algorithm J translate to the following three types of operations on the two-colored trees \mathcal{T}'_n : Vertex 1 switches its color; vertices i and i+1 exchange their colors; a tree rotation combined with a cyclic rotation of colors among an interval of vertices.

5. Three exceptional patterns

Interestingly, we found three patterns τ that are not tame and for which $B_n(\tau)$ is not a zigzag language (i.e., Theorem 17 does not apply), and yet Algorithm J succeeds in generating the set $B_n(\tau)$ exhaustively. We were not able to treat those exceptional patterns in a unified way, and we were also not able to find more such exceptional patterns for which Algorithm J would work. Consequently, in this section we describe each of these patterns separately. The last one is particularly important, as it corresponds to the type B associahedron.

5.1. The pattern $2\overline{1}$ and (unsigned) permutations. Let $\tau := 2\overline{1}$. Signed permutations avoiding this pattern have all negative entries on the set of positions $[\overline{n}]$, and all positive entries on the set of positions [n], and no other constraints, i.e., they are in bijection with (unsigned) permutations. The pattern τ is not tame and $B_n(\tau)$ is not a zigzag language, and yet Algorithm J

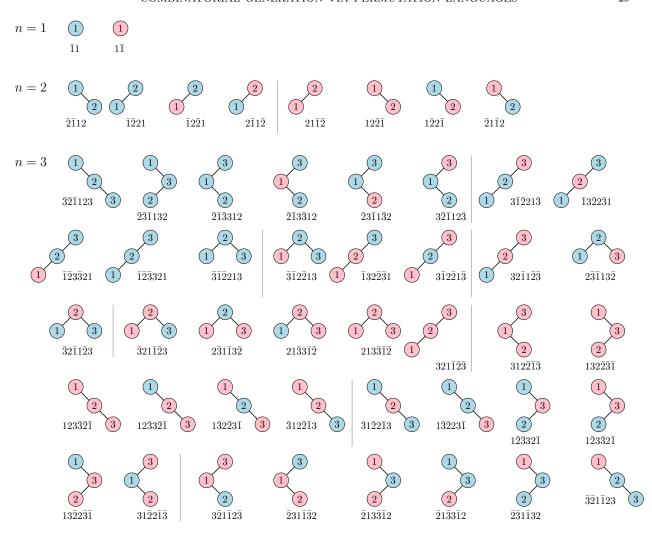


FIGURE 20. Listings of signed permutations $B_n(\tau)$ for n = 1, 2, 3 for the tame pattern $\tau = 23\underline{1}$ generated by Algorithm J and corresponding two-colored binary trees.

succeeds in generating the set $B_n(\tau)$ when initialized with $\pi_0 := \mathrm{id}_n$, and the corresponding ordering of permutations is the Steinhaus-Johnson-Trotter ordering.

5.2. The pattern 213 and binary strings. Let $\tau := 213$. We present a bijection between signed permutations avoiding the pattern τ and binary strings of fixed length.

Theorem 22. Let $\tau := 213$. There is a bijection $h : B_n(\tau) \to \{0,1\}^n$, and therefore $|B_n(\tau)| = 2^n$.

Proof. The following properties can be easily verified by induction. The structure of a signed permutation $\pi \in B_{n-1}(\tau)$ is either

$$\pi = \mathrm{id}_{n-1}$$
 (case 1),

or for some $k \in \{0, 1, \dots, n-2\}$ we have

$$\pi = (n - k - 1, n - k, \dots, n - 2, n - 1, \pi', \overline{n - 1}, \overline{n - 2}, \dots, \overline{n - k - 1})$$
 (case 2)

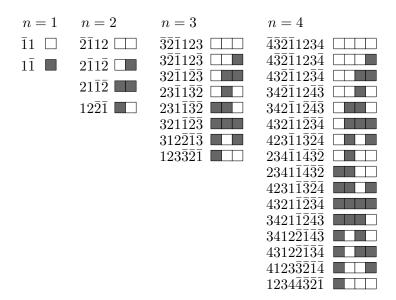


FIGURE 21. Listings of signed permutations $B_n(\tau)$ for n = 1, 2, 3, 4 for the non-tame pattern $\tau = 213$ generated by Algorithm J and corresponding binary strings. Each 0-bit is visualized as a white square, and each 1-bit as a black square. The resulting ordering is the classical binary reflected Gray code.

with $\pi' \in B_{n-k-2}(\tau)$. For every such permutation π there are exactly two indices $i \in [\overline{n}] \cup [n]$ such that $c_i(\pi) \in B_n(\tau)$, namely, in case 1

$$c_{\overline{n}}(\pi) = (n \, \pi, \overline{n})$$
 (\rightarrow case 2),
 $c_n(\pi) = (\overline{n}, \pi, n) = (\overline{n}, \mathrm{id}_{n-1}, n) = \mathrm{id}_n$ (\rightarrow case 1);

and in case 2

$$c_{\overline{n}}(\pi) = (n, \pi, \overline{n}) \qquad (\to \text{case 2}),$$

$$c_{\overline{n-k-1}}(\pi) = (n-k-1, \dots, n-1, n, \pi', \overline{n}, \overline{n-1}, \dots, \overline{n-k-1}) \quad (\to \text{case 2}).$$

In other words, in both cases the new largest value n can be inserted at the leftmost position, and directly after the maximum increasing prefix.

The bijection $h: B_n(\tau) \to \{0,1\}^n$ can thus be defined as follows: Given $\pi \in B_n(\tau)$ and $i \in [n]$, consider $\pi' := p^{n-i}(\pi) \in B_i(\tau)$, and if the largest value i in π' is at the leftmost position, then the ith bit of $h(\pi)$ is defined to be 1, and 0 otherwise.

The pattern τ is not tame and $B_n(\tau)$ is not a zigzag language, and yet Algorithm J succeeds in generating the set $B_n(\tau)$ when initialized with $\pi_0 := \mathrm{id}_n$; see Figure 21. The resulting ordering of binary strings is the classical binary reflected Gray code, i.e., minimal jumps performed by the algorithm translate to single bitflips.

However, there is no one-to-one correspondence between minimal jumps in τ -avoiding permutations and bitflips in the hypercube. In particular, the initial signed permutation id_n admits only one possible minimal jump (the value n jumping all the way to the other side). Consequently, the listing of τ -avoiding signed permutations is not cyclic under minimal jumps, even though the binary reflected Gray code is cyclic, i.e., the last binary string differs only in a single bit from the first string.

5.3. The pattern 231 and point-symmetric triangulations. For the rest of this section, we consider the permutation pattern $\tau := 231$. This pattern is not tame and $B_n(\tau)$ is not a zigzag language, and yet Algorithm J can be used to generate the set $B_n(\tau)$ exhaustively. In

order to prove this (Theorem 30 below), we precisely describe the structure of τ -avoiding signed permutations. This also helps us to establish a bijection between signed permutations in $B_n(\tau)$ and point-symmetric triangulations of a convex (2n+2)-gon, with the property that minimal jumps in the permutations are in one-to-one correspondence to flips in the triangulations (see Section 5.3.4 and Theorem 29 therein). As a consequence, Algorithm J not only lists all signed permutation of $B_n(\tau)$, but the listing actually corresponds to a Hamiltonian path on the B-associahedron [Sim03, Rea06]; see Figures 7, 12 and 13.

The bijection is essentially the same as the one described in [Rea06, Sec. 7]. However, we describe it purely combinatorially, without reference to lattice-theoretic concepts. Our goal is to give a simple, explicit description that can easily be exploited for the purpose of Gray codes.

5.3.1. Triangulations. We consider a set of n+2 points in the plane, placed equidistantly on the unit circle and labeled $0,1,\ldots,n,\bar{0}$ in counterclockwise (ccw) order. We write Δ_n for the set of triangulations on this point set. A flip in a triangulation removes one of its inner edges and replaces it by the other diagonal of the resulting empty quadrilateral. It is well-known that there is a bijection between Δ_n and $S_n(231)$, and that both families of objects are counted by the Catalan numbers, i.e., we have $|\Delta_n| = |S_n(231)| = C_n = \frac{1}{n+1} {2n \choose n}$.

The bijection $g: \Delta_n \to S_n(231)$ will be used as a building block later, so we define it in the following; see Figure 22: Given a triangulation $T \in \Delta_n$, we consider the triangle in T seen through the edge $0\bar{0}$. To obtain g(T), we process each of the n the triangles one by one, recording one vertex of each triangle. Specifically, for every (ccw) triangle abc seen through the edge ac, we record the label of the third vertex b, and we then recursively compute $g(T^L)$ and $g(T^R)$, where T^L and T^R are the subtriangulations of T seen through the left edge ab and the right edge bc, respectively. Therefore, if b is the third vertex of the triangle in T containing the edge $0\bar{0}$, then $g(T) := (b, g(T^L), g(T^R))$.

Lemma 23 ([HHMW22]). The mapping $g: \Delta_n \to S_n(231)$ is a bijection that induces a one-to-one correspondence between flips in triangulations and minimal jumps in 231-avoiding permutations.

Note that the composition $f^{-1} \circ g : \Delta_n \to \mathcal{T}_n$, where $f : \mathcal{T}_n \to S_n(231)$ is the bijection described in the proof of Theorem 21, is the geometric duality mapping that assigns to every triangulation its corresponding dual tree (seen through the edge $0\bar{0}$). Furthermore, note that flips in triangulations are in one-to-one correspondence with tree rotations under this bijection.

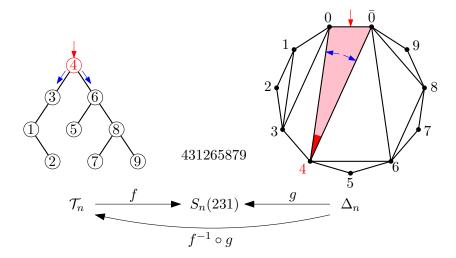


FIGURE 22. Bijections between binary trees, 231-avoiding permutations and triangulations of a convex polygon.

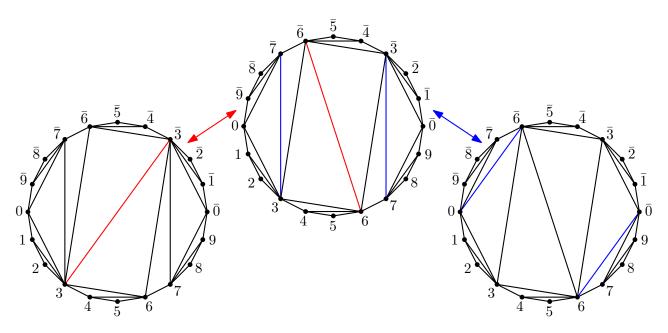


FIGURE 23. A symmetric triangulation, and the two possible flip types.

5.3.2. Point-symmetric triangulations. We consider a set of 2n+2 points in the plane, placed equidistantly on the unit circle and labeled $0, 1, \ldots, n, \overline{0}, \overline{1}, \ldots, \overline{n}$ in ccw order; see Figure 23. Note that the points i and \bar{i} are opposite to each other on the circle. A symmetric triangulation is a triangulation of this point set that is point-symmetric w.r.t. the origin, i.e., whenever an edge uvis present, then the edge $\overline{u}\,\overline{v}$ must also be present. We write Δ_n^* for the set of all symmetric triangulations on this point set. Note that every such triangulation has a unique edge $u\overline{u}$ through the origin, which we call the central edge. A flip in a symmetric triangulation either consists of a single flip of the central edge, or a pair of symmetric flips not involving the central edge.

Observe that a symmetric triangulation is uniquely determined by the subtriangulation on one side of the central edge, and that there n+1 different central edges $i\bar{i}, i=0,\ldots,n$. It follows that $|\Delta_n^*| = (n+1)|\Delta_n| = (n+1)C_n = \binom{2n}{n}$, i.e., symmetric triangulations are counted by the central binomial coefficients.

5.3.3. Structure of τ -avoiding signed permutations. Our first lemma, illustrated in Figure 24, describes the structure of 231-avoiding (unsigned) permutations. In the following, for a string x = (x_1,\ldots,x_s) , we write |x|:=s for its length.

Lemma 24. For a permutation $\pi = (a_1, \ldots, a_n) \in S_n$, let $s \in [n]$ be the index of the value 1 in π , and define $\pi^L = (x_s, \ldots, x_1) := (a_1, \ldots, a_s)$ and $\pi^R := (a_{s+1}, \ldots, a_n)$. Then $\pi = (\pi^L, \pi^R)$ avoids 231 if and only if the following conditions hold:

- (i) We have $x_s > x_{s-1} > \cdots > x_1 = 1$, i.e., the values to the left of 1 appear in decreasing order. (ii) There is a decomposition $\pi^R = (\pi_1^R, \dots, \pi_s^R)$ where $x_1 < \pi_1^R < x_2 < \pi_2^R < \cdots < x_{s-1} < \pi_{s-1}^R < x_s < \pi_s^R$, such that each of the subpermutations π_j^R , $j = 1, \dots, s$, avoids 231.

The inequality $x_j < \pi_j^{\mathrm{R}} < x_{j+1}$ means that all values of the subpermutation π_j^{R} are larger than x_j and smaller than x_{j+1} .

Proof. We first assume that $\pi \in S_n$ avoids 231. To prove (i), note that if there are indices i < j < s with $a_i < a_j$, then $a_i a_j a_s = a_i a_j 1$ is an occurrence of the pattern 231, a contradiction. To prove (ii), we first establish the following auxiliary claim: For any $j \in [s]$, we have that in π^{R} all entries smaller than x_j are to the left of all entries larger than x_j . Indeed, if there is a $j \in [s]$ and indices $s+1 \le k < \ell \le n$ such that $a_k > x_j > a_\ell$, then $x_j a_k a_\ell$ is an occurrence of 231, a

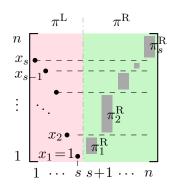


FIGURE 24. Illustration of the structure of a 231-avoiding permutation as captured by Lemma 24. Specifically, the figure shows the corresponding permutation matrix, i.e., a point in column i and row j indicates that the permutation has the value j at position i.

contradiction. This proves the decomposition of π into subpermutations π_i^R , $j \in [s]$, claimed in (ii). Clearly, each of these subpermutations must also avoid 231.

Now assume that $\pi \in S_n$ satisfies (i) and (ii). If $a_i a_j a_k$, i < j < k, is an occurrence of 231 in π , then we have $a_k < a_i < a_j$, and we distinguish three possible cases: (1) The values a_i and a_j are both in π^{L} : this is impossible because of (i). (2) The value a_i is in π^{L} and a_j , a_k are in π^{R} : this is impossible because of the decomposition into subpermutations stated in (ii). (3) The values a_i, a_j, a_k are all in π^R : because of the inequalities $\pi_1^R < \pi_2^R < \cdots < \pi_{s-1}^R < \pi_s^R$ and $a_i, a_j > a_k$ the values a_i and a_j must be belong to the same subpermutation as a_k , i.e., a_i, a_j, a_k are all in $\pi_{\widehat{j}}^{R}$ for some $\widehat{j} \in [s]$, which is impossible because of the 231-avoidance condition stated in (ii).

The next lemma, illustrated in Figure 25 (a), describes the structure of τ -avoiding signed permutations. For any string $x=(a_1,\ldots,a_k)$ of integers, we define $x':=(\overline{a_k},\ldots,\overline{a_1})$, i.e., this operation reverses the string and complements all signs.

Let $\pi = (a_{\overline{n}}, \dots, a_{\overline{1}}, a_1, \dots, a_n) \in B_n$. We define $N(\pi) := \{i \in [n] \mid a_i < 0\}$ and

$$s_{\pi} := \begin{cases} \max N(\pi) & \text{if } N(\pi) \neq \emptyset, \\ 0 & \text{else.} \end{cases}$$
 (3a)

In words, s_{π} is the largest positive index for which the corresponding entry of π is negative, or 0 if no such index exists. Furthermore, we let $p_{\pi} \in [\overline{n}] \cup [n]$ be the position of the value 1 in π . Lastly, we define

$$t_{\pi} := \max\{s_{\pi}, p_{\pi}\}. \tag{3b}$$

Lemma 25. For a signed permutation $\pi = (a_{\overline{n}}, \ldots, a_{\overline{1}}, a_1, \ldots, a_n) \in B_n$, let $s := s_{\pi}$ be as defined in (3a). Define $\pi^{\mathrm{M}} := (a_{\overline{s}}, \ldots, a_{\overline{1}}, a_1, \ldots, a_s)$, $\pi^{\mathrm{R}} := (a_{s+1}, \ldots, a_n)$ and $\pi^{\mathrm{L}} := (a_{\overline{n}}, \ldots, a_{\overline{s+1}}) = a_{s+1}$ $(\pi^R)'$, and let $x_s, x_{s-1}, \ldots, x_1$ denote the sequence of positive entries of π^M from left to right. Then $\pi = (\pi^L, \pi^M, \pi^R)$ avoids τ if and only if the following conditions hold:

- (i) We have $x_s > x_{s-1} > \cdots > x_1$, i.e., the positive entries of π^{M} appear in decreasing order. (ii) There is a decomposition $\pi^{\mathrm{R}} = (\pi_0^{\mathrm{R}}, \pi_1^{\mathrm{R}}, \dots, \pi_s^{\mathrm{R}})$ where $\pi_0^{\mathrm{R}} < x_1 < \pi_1^{\mathrm{R}} < x_2 < \pi_2^{\mathrm{R}} < \cdots < x_{s-1} < \pi_{s-1}^{\mathrm{R}} < x_s < \pi_s^{\mathrm{R}}$, such that each of the subpermutations π_j^{R} , $j = 0, \dots, s$, avoids 231.

We emphasize that the pattern 231 mentioned in (ii) has no underscore. Note that $|\pi^{\rm M}| = 2s_{\pi}$ and $|\pi^{R}| = |\pi^{L}| = n - s_{\pi}$.

Proof. We first assume that $\pi \in B_n$ avoids τ . To prove (i), note that if there are indices i < jwith $x_i < x_j$, then $x_i x_j a_s = x_i x_j \overline{x_s}$ is an occurrence of the pattern τ , a contradiction. To

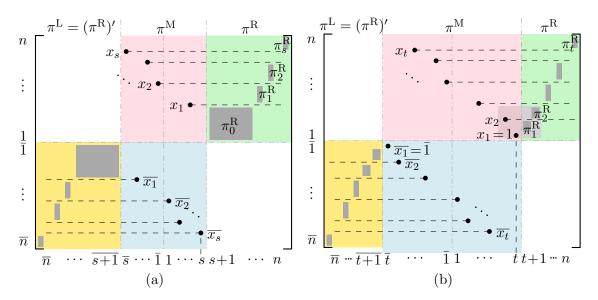


FIGURE 25. Illustration of the structure of a τ -avoiding signed permutation as described by (a) Lemma 25 and (b) Lemma 26.

prove (ii), we first establish the following auxiliary claim: For any $j \in [s]$, we have that in π^R all entries smaller than x_j are to the left of all entries larger than x_j . Indeed, if there is a $j \in [s]$ and indices $s+1 \le k < \ell \le n$ such that $a_k > x_j > a_\ell$, then $x_j a_k a_\ell$ is an occurrence of τ , a contradiction. This proves the decomposition of π into subpermutations π_j^R , $j=0,\ldots,s$, claimed in (ii). As all values in π^R are positive by definition, any occurrence of 231 would also be an occurrence of τ , so each of the subpermutations π_j^R , $j=0,\ldots,s$, must avoid 231.

We now assume that $\pi \in B_n$ satisfies (i) and (ii). If $a_ia_ja_k$, i < j < k, is an occurrence of τ in π , then we have $a_k < a_i < a_j$ and $a_i > 0$, so a_i, a_j, a_k are all in π^M or π^R . We only need to consider three possible cases: (1) The values a_i and a_j are both in π^M : this is impossible because of (i). (2) The value a_i is in π^M and a_j, a_k are in π^R : this is impossible because of the decomposition into subpermutations stated in (ii). (3) The values a_i, a_j, a_k are all in π^R : because of the inequalities $\pi_0^R < \pi_1^R < \pi_2^R < \cdots < \pi_{s-1}^R < \pi_s^R$ and $a_i, a_j > a_k$ the values a_i and a_j must belong to the same subpermutation as a_k , i.e., a_i, a_j, a_k are all in π_j^R for some $\hat{j} \in \{0, \dots, s\}$, which is impossible because of the 231-avoidance condition stated in (ii).

We also need the following alternative structural description of τ -avoiding signed permutations; see Figure 25 (b).

Lemma 26. For a signed permutation $\pi = (a_{\overline{n}}, \dots, a_{\overline{1}}, a_1, \dots, a_n) \in B_n$, let $s := s_{\pi}$, $p := p_{\pi}$ and $t := t_{\pi}$ be as defined in (3). Define $\pi^{\mathrm{M}} := (a_{\overline{t}}, \dots, a_{\overline{1}}, a_1, \dots, a_t)$, $\pi^{\mathrm{R}} := (a_{t+1}, \dots, a_n)$ and $\pi^{\mathrm{L}} := (a_{\overline{n}}, \dots, a_{\overline{t+1}}) = (\pi^{\mathrm{R}})'$, and let x_t, x_{t-1}, \dots, x_1 denote the sequence of positive entries of π^{M} from left to right. Then $\pi = (\pi^{\mathrm{L}}, \pi^{\mathrm{M}}, \pi^{\mathrm{R}})$ avoids τ if and only if the following conditions hold:

- (i) We have $x_t > x_{t-1} > \cdots > x_1 = 1$, i.e., the positive entries of π^M appear in decreasing order, ending with 1.
- (ii) In the decomposition $\pi^{R} = (\pi_1^{R}, \dots, \pi_s^{R})$ where $x_1 < \pi_1^{R} < x_2 < \pi_2^{R} < \dots < x_{t-1} < \pi_{t-1}^{R} < x_t < \pi_t^{R}$, each of the subpermutations π_j^{R} , $j \in [t]$, avoids 231.

By Lemma 26, a τ -avoiding signed permutation π can be viewed as a 231-avoiding permutation σ (recall Lemma 24), combined with the reversed and complemented permutation σ' such that the positive values to the left and including 1 of σ and their negated counterparts in σ' are interleaved arbitrarily (in a symmetric way).

Proof. Note that the value 1 is always in π^{M} , because the positive values of π are all in either π^{M} or π^{R} , and if the value 1 were in π^{R} , then we would have s > t, a contradiction to the definition (3b).

If $s \geq p$, then the claims of the lemma are identical to Lemma 25. On the other hand, if s < p, then we apply Lemma 25 to obtain a decomposition $\pi = (\pi^{L^*}, \pi^{M^*}, \pi^{R^*})$ with $\pi^{R^*} = (\pi_0^{R^*}, \pi_1^{R^*}, \dots, \pi_s^{R^*})$, and we then apply Lemma 24 to further decompose $\pi_0^{R^*}$, yielding a decomposition that has the desired properties.

We note that $|\pi^{\mathcal{M}}| = t_{\pi}$ and $|\pi^{\mathcal{R}}| = |\pi^{\mathcal{L}}| = n - t_{\pi}$. Furthermore, as the value 1 belongs to $\pi^{\mathcal{M}}$ we have $\pi^{\mathcal{M}} \neq \varepsilon$.

For $\pi \in B_n(\tau)$, we refer to (π^L, π^M, π^R) and $\pi^R = (\pi_1^R, \dots, \pi_t^R)$ as given by Lemma 26 as staircase decomposition of π , and to the sequence $x_t > \dots > x_1$ as staircase values. Note that the length of π^M is 2t and the length of π^R and π^L is n-t.

From Lemma 26, we can directly deduce the following structure of minimal jumps inside a τ -avoiding signed permutation; see Figure 27.

Lemma 27. Let $\pi = (a_{\overline{n}}, \dots, a_n) \in B_n(\tau)$, and consider its staircase decomposition $\pi = (\pi^L, \pi^M, \pi^R)$ and $\pi^R = (\pi_1^R, \dots, \pi_t^R)$ with staircase values $x_t > \dots > x_1$, as given by Lemma 26. A minimal jump of a value $a_i > 0$ in π is exactly one of the following:

- (i) If a_i is in π^M and its neighboring entry b in π^M is negative (either $b = a_{i-1}$ or $b = a_{i+1}$), it is an adjacent transposition $a_i \leftrightarrow b$ within π^M .
- (ii) If a_i is in π^M and $a_{i+1} > 0$, i.e., $a_i = x_j$ for some $j \in [t]$, then it is a right jump ending directly to the left of π_j^R .
- (iii) If a_i is the leftmost entry of π_j^R for some $j \in [t]$, then it is a left jump ending directly to the left of x_j .
- (iv) If a_i is in π_j^R for some $j \in [s]$, then it is a jump within the subpermutation π_j^R that is minimal w.r.t. avoidance of the pattern 231.

Note that left and right jumps as in (i) are inverse to each other, and the same is true for left and right jumps as in (iv). On the other hand, right jumps as in (ii) are the inverse of the left jumps in (iii). We refer to a minimal jump as in (i) or (iv) as a *close* minimal jump, and to a jump as in (ii) or (iii) as a *far* minimal jump. Note that close minimal jumps do not change the lengths of the three parts of the staircase decomposition, whereas a far jump either decreases or increases (in case (ii) and (iii), respectively) the length of the middle part $\pi^{\rm M}$ by 2.

For a signed permutation $\pi \in B_{n-1}(\tau)$, it is easy to see (using Lemma 26) that $c_n(\pi) \in B_n(\tau)$, i.e., inserting n at the rightmost position and \overline{n} at the leftmost position yields again a permutation that avoids τ . A similar statement is false for $c_{\overline{n}}(\pi)$, which prevents $B_n(\tau)$ from being a zigzag language. For example, for $\pi = \overline{2}\overline{1}12 \in B_2(\tau)$ we have $c_{\overline{3}}(\pi) = 3\overline{2}\overline{1}12\overline{3}$, and the substring $12\overline{3}$ is an occurrence of τ .

The following lemma describes the leftmost possible position to insert the value n into a τ -avoiding signed permutation $\pi \in B_{n-1}(\tau)$, so that the resulting signed permutation is again τ -avoiding. Furthermore, importantly for Algorithm J, minimal jumps are preserved under such insertions.

Lemma 28. Let $\pi \in B_{n-1}(\tau)$, let $t := t_{\pi}$ be as defined in (3), and consider the staircase decomposition $\pi = (\pi^{L}, \pi^{M}, \pi^{R})$ as given by Lemma 26. Then the smallest integer $i \in [\overline{n}] \cup [n]$ such that $c_{i}(\pi)$ avoids τ equals $i = \overline{t+1}$, i.e., we have $\sigma := c_{\overline{t+1}}(\pi) = (\pi^{L}, n, \pi^{M}, \overline{n}, \pi^{R}) \in B_{n}(\tau)$. Furthermore, the staircase decomposition of σ is $(\sigma^{L}, \sigma^{M}, \sigma^{R}) = (\pi^{L}, (n, \pi^{M}, \overline{n}), \pi^{R})$, and any minimal jump in π is also a minimal jump in σ .

Proof. Combine Lemmas 26 and 27.

We emphasize that a close minimal jump in π with d steps is still a close minimal jump by d steps in σ , whereas a far minimal jump in π turns into a far minimal jump by d+1 steps in σ ,

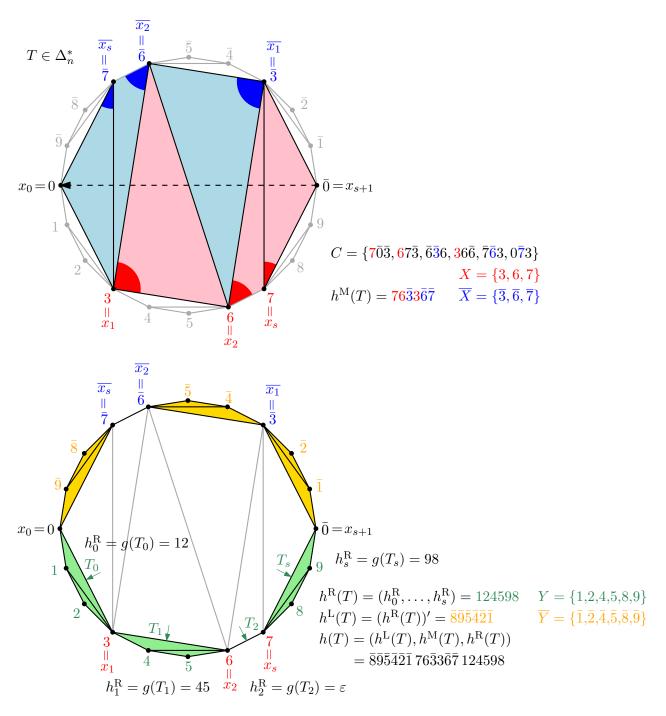


FIGURE 26. Definition of the bijection h between symmetric triangulations and τ -avoiding signed permutations. The triangulation is the same as in Figure 23.

because the value \overline{n} is inserted at position t+1. For example, $\pi:=\bar{3}\bar{1}2\bar{2}13\to\bar{1}32\bar{2}\bar{3}1$ is a far minimal left jump of 3 by 4 steps in π , and we have $t_{\pi}=2$, so $\sigma:=c_{\overline{t+1}}(\pi)=c_{\bar{3}}(\pi)=\bar{3}4\bar{1}2\bar{2}1\bar{4}3$, and $\sigma=\bar{3}4\bar{1}2\bar{2}1\bar{4}3\to 4\bar{1}32\bar{2}\bar{3}1\bar{4}$ is a far minimal left jump of 3 by 5 steps in σ .

5.3.4. A bijection between Δ_n^* and $B_n(\tau)$. The following definitions are illustrated in Figure 26. Consider a symmetric triangulation $T \in \Delta_n^*$. Let C be the set of triangles in T that have points on both sides of the line $\bar{0}0$. We refer to them as crossing triangles of T. Let

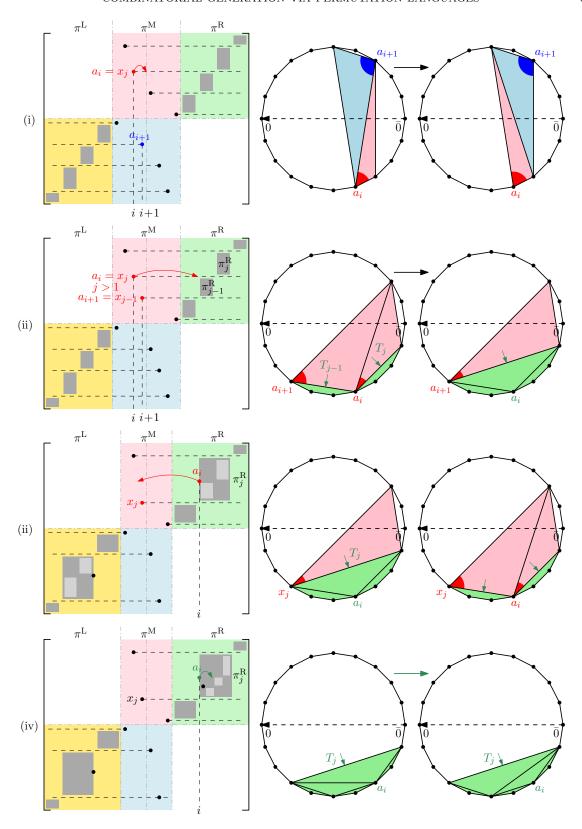


FIGURE 27. Correspondence between minimal jumps in τ -avoiding signed permutations as described by Lemma 27 (left) and flips in symmetric triangulations (right) under the bijection h.

 $X = \{x_1 < x_2 < \dots < x_s\}$ be the set of positive vertices of triangles in C, excluding the vertices $x_0 := 0$ and $x_{s+1} := \bar{0}$. Note that we have |C| = 2s.

We first define a permutation $h^{\mathrm{M}}(T)$ of the integers in $X \cup \overline{X}$ as follows. We sweep through the triangles in C in the order in which they intersect the ray from $\overline{0}$ to 0, starting with the triangle $x_s\overline{x_1}\overline{0}$ and ending with the triangle $0x_1\overline{x_s}$. For every triangle in C, we add one of its vertices to the string $h^{\mathrm{M}}(T)$, according to the following rule: If the triangle is abc with $a,b \in [n] \cup \{\overline{0}\}$, $b=\overline{0}$ or a < b, and $c \in [\overline{n}]$, we add the vertex a. If the triangle is abc with $a,b \in [\overline{n}] \cup \{0\}$, a=0 or a < b, and $c \in [n]$, we add the vertex b. One can check that $(h^{\mathrm{M}}(T))' = h^{\mathrm{M}}(T)$. Furthermore, if central edge is $u\overline{u}$, then u is the smallest/rightmost positive value in the left half of $h^{\mathrm{M}}(T)$, and \overline{u} is the largest/leftmost negative value in the right half of $h^{\mathrm{M}}(T)$. In the special case were the central edge is $0\overline{0}$, we have $h^{\mathrm{M}}(T) = \varepsilon$.

We now define a permutation $h^{\mathrm{R}}(T)$ of the integers in $Y := [n] \setminus X$ as follows. For $j = 0, \ldots, s$, we let T_j denote the subtriangulation of T of non-crossing triangles seen through the edge $x_j x_{j+1}$. We define $h^{\mathrm{R}}(T) := (h_0^{\mathrm{R}}, \cdots, h_s^{\mathrm{R}})$ with $h_j^{\mathrm{R}} := g(T_j)$, where g is the mapping defined in Section 5.3.1, and in evaluating $g(T_j)$ we consider the (sub)triangulation T_j through the edge $x_j x_{j+1}$. Furthermore, we define $h^{\mathrm{L}}(T) := (h^{\mathrm{R}}(T))'$, which is a permutation of the integers in $\overline{Y} = [\overline{n}] \setminus \overline{X}$.

Combining the results from the previous steps, we define $h(T) := (h^{L}(T), h^{M}(T), h^{R}(T))$.

Theorem 29. The mapping $h: \Delta_n^* \to B_n(\underline{2}31)$ is a bijection that induces a one-to-one correspondence between flips in symmetric triangulations and minimal jumps in $\underline{2}31$ -avoiding signed permutations.

The theorem shows that the two flip graphs on symmetric triangulations and τ -avoiding signed permutations are isomorphic. Figure 12 shows the B-associahedron for n=3, labeled with the symmetric triangulations and τ -avoiding signed permutations corresponding to each other under the bijection h.

Proof. The fact that h is a bijection can be verified directly, using Lemmas 23 and 25. The correspondence between flips and minimal jumps follows from Lemma 27. Specifically, the flips corresponding to the four different types of jumps (i)—(iv) described in the lemma are illustrated in Figure 27 (i)—(iv), respectively.

Theorem 30. Given the set $L_n = B_n(231)$ and initial permutation $\pi_0 = \mathrm{id}_n$, Algorithm J visits every signed permutation from L_n exactly once. Furthermore, the last permutation differs from the first one in a minimal jump.

The resulting orderings of τ -avoiding signed permutations and symmetric triangulations are shown in Figures 12 and 13 for n = 1, 2, 3, 4.

Proof. The proof is a straightforward adaption of the proof of Theorem 17. We use again the property that if $\pi \in B_{n-1}(\tau)$, then $c_n(\pi) \in B_n(\tau)$ and any minimal jump in π is also a minimal jump in $c_n(\tau)$. As a similar statement does not hold for $c_{\overline{n}}(\pi)$, we apply Lemma 28 instead. Using that $|B_n(\tau)| = {2n \choose n}$ is even, an easy induction shows that the last permutation in $J(B_n(\tau))$ is $\overline{n} \cdots \overline{32}1\overline{1}23 \cdots n$ (recall Remark 18), i.e., it differs from the first permutation id_n in a minimal jump of 1.

ACKNOWLEDGEMENTS

We thank the reviewers of the extended abstract of this paper for numerous helpful comments and suggestions.

References

[BB05] A. Björner and F. Brenti. Combinatorics of Coxeter groups, volume 231 of Graduate Texts in Mathematics. Springer, New York, 2005.

- [BEZ90] A. Björner, P. H. Edelman, and G. M. Ziegler. Hyperplane arrangements with a lattice of regions. Discrete Comput. Geom., 5(3):263–288, 1990.
- [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.
- [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.
- [cos] The Combinatorial Object Server: Generate colored permutations. http://www.combos.org/cperm.
- [CP25] J. Cardinal and V. Pilaud. Rectangulotopes. European J. Combin., 125:Paper No. 104090, 25, 2025.
- [CSS18] J. Cardinal, V. Sacristán, and R. I. Silveira. A note on flips in diagonal rectangulations. Discrete Math. Theor. Comput. Sci., 20(2):Paper No. 14, 22 pp., 2018.
- [CSW89] J. H. Conway, N. J. A. Sloane, and A. R. Wilks. Gray codes for reflection groups. Graphs Combin., 5(4):315–325, 1989.
- [Ede84] P. H. Edelman. A partial order on the regions of \mathbb{R}^n dissected by hyperplanes. Trans. Amer. Math. Soc., 283(2):617–631, 1984.
- [FR07] S. Fomin and N. Reading. Root systems and generalized associahedra. In *Geometric combinatorics*, volume 13 of *IAS/Park City Math. Ser.*, pages 63–131. Amer. Math. Soc., Providence, RI, 2007.
- [Gir12] S. Giraudo. Algebraic and combinatorial structures on pairs of twin binary trees. J. Algebra, 360:115–157, 2012.
- [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.
- [Gra53] F. Gray. Pulse code communication, 1953. March 17, 1953 (filed Nov. 1947). U.S. Patent 2,632,058.
- [GZ83] C. Greene and T. Zaslavsky. On the interpretation of Whitney numbers through arrangements of hyperplanes, zonotopes, non-Radon partitions, and orientations of graphs. Trans. Amer. Math. Soc., 280(1):97–126, 1983.
- [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.
- [HM21] H. P. Hoang and T. Mütze. Combinatorial generation via permutation languages. II. Lattice congruences. Israel J. Math., 244(1):359–417, 2021.
- [HNT05] F. Hivert, J.-C. Novelli, and J.-Y. Thibon. The algebra of binary search trees. *Theoret. Comput. Sci.*, 339(1):129–165, 2005.
- [Joh63] S. M. Johnson. Generation of permutations by adjacent transposition. *Math. Comp.*, 17:282–285, 1963.
- [KLL11] J. Korsh, P. LaFollette, and S. Lipschutz. A loopless implementation of a Gray code for signed permutations. *Publ. Inst. Math. (Beograd) (N.S.)*, 89(103):37–47, 2011.
- [Knu11] D. E. Knuth. The Art of Computer Programming. Vol. 4A. Combinatorial Algorithms. Part 1. Addison-Wesley, Upper Saddle River, NJ, 2011.
- [KSSW25] V. Körber, T. Schnieders, J. Stricker, and J. Walizadeh. Hamiltonian cycles in simplicial and supersolvable hyperplane arrangements. https://arxiv.org/abs/2508.14538, 2025.
- [Lod04] J.-L. Loday. Realization of the Stasheff polytope. Arch. Math. (Basel), 83(3):267–278, 2004.
- [LR12] S. Law and N. Reading. The Hopf algebra of diagonal rectangulations. *J. Combin. Theory Ser. A*, 119(3):788–824, 2012.
- [LRvBR93] 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.
- [MHPS12] F. Müller-Hoissen, J. M. Pallo, and J. Stasheff, editors. Associahedra, Tamari lattices and related structures, volume 299 of Progress in Mathematics. Birkhäuser/Springer, Basel, 2012. Tamari Memorial Festschrift.
- [MM23] A. Merino and T. Mütze. Combinatorial generation via permutation languages. III. Rectangulations. Discrete Comput. Geom., 70(1):51–122, 2023.
- [Müt23] T. Mütze. Combinatorial Gray codes—an updated survey. *Electron. J. Combin.*, DS26(Dynamic Surveys):99 pp., 2023.
- [OT92] P. Orlik and H. Terao. Arrangements of hyperplanes, volume 300 of Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Springer-Verlag, Berlin, 1992.
- [Pil24] V. Pilaud. Acyclic reorientation lattices and their lattice quotients. Ann. Comb., 28(4):1035–1092, 2024.
- [Poul4] L. Pournin. The diameter of associahedra. Adv. Math., 259:13–42, 2014.
- [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.
- [PSZ23] V. Pilaud, F. Santos, and G. M. Ziegler. Celebrating Loday's associahedron. Arch. Math. (Basel), 121(5-6):559-601, 2023.
- $[Rea06] \qquad \text{N. Reading. Cambrian lattices. } \textit{Adv. Math.}, \, 205(2):313-353, \, 2006.$

[Rea12] N. Reading. Generic rectangulations. European J. Combin., 33(4):610–623, 2012.

[Rea16] N. Reading. Lattice theory of the poset of regions. In *Lattice theory: special topics and applications*. Vol. 2, pages 399–487. Birkhäuser/Springer, Cham, 2016.

[Sav97] C. Savage. A survey of combinatorial Gray codes. SIAM Rev., 39(4):605–629, 1997.

[Sim03] R. Simion. A type-B associahedron. *Adv. in Appl. Math.*, 30(1-2):2–25, 2003. Formal power series and algebraic combinatorics (Scottsdale, AZ, 2001).

[SSW93] C. D. Savage, M. B. Squire, and D. B. West. Gray code results for acyclic orientations. In Proceedings of the Twenty-fourth Southeastern International Conference on Combinatorics, Graph Theory, and Computing (Boca Raton, FL, 1993), volume 96, pages 185–204, 1993.

[Sta72] R. P. Stanley. Supersolvable lattices. Algebra Universalis, 2:197–217, 1972.

[Sta07] R. P. Stanley. An introduction to hyperplane arrangements. In *Geometric combinatorics*, volume 13 of *IAS/Park City Math. Ser.*, pages 389–496. Amer. Math. Soc., Providence, RI, 2007.

[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.

[STT19] D. Suyama, M. Torielli, and S. Tsujie. Signed graphs and the freeness of the Weyl subarrangements of type B_{ℓ} . Discrete Math., 342(1):233–249, 2019.

[Tro62] H. F. Trotter. Algorithm 115: Perm. Commun. ACM, 5(8):434–435, 1962.

[Whi96] A. T. White. Fabian Stedman: the first group theorist? Amer. Math. Monthly, 103(9):771–778, 1996.

[Wil13] A. Williams. The greedy Gray code algorithm. In Algorithms and data structures, volume 8037 of Lecture Notes in Comput. Sci., pages 525–536. Springer, Heidelberg, 2013.

[Zas82] T. Zaslavsky. Signed graphs. Discrete Appl. Math., 4(1):47–74, 1982.

[Zas91] T. Zaslavsky. Orientation of signed graphs. European J. Combin., 12(4):361–375, 1991.

[Zas01] T. Zaslavsky. Supersolvable frame-matroid and graphic-lift lattices. European J. Combin., 22(1):119–133, 2001.