POSETS ARISING AS 1-SKELETA OF SIMPLE POLYTOPES, THE NONREVISITING PATH CONJECTURE, AND POSET TOPOLOGY

PATRICIA HERSH

Abstract. Given any polytope P and any generic linear functional $\mathbf{c}$, one obtains a directed graph $G(P,\mathbf{c})$ from the 1-skeleton of P by orienting each edge $e(u,v)$ from u to v for $\mathbf{c}\cdot u < \mathbf{c}\cdot v$. For P a simple polytope and $G(P,\mathbf{c})$ the Hasse diagram of a lattice L, the join of any collection S of elements which all cover a common element u in L is proven to equal the sink of the smallest face of P containing u and all of the elements of S. The author conjectures for such $G(P,\mathbf{c})$ that no directed path in $G(P,\mathbf{c})$ ever revisits any facet of P. This would imply for such P and $\mathbf{c}$ that the simplex method for linear programming is efficient under all possible pivot rules. This conjecture is proven for 3-polytopes and for spindles.

For simple polytopes in which $G(P,\mathbf{c})$ is the Hasse diagram of a lattice L, the order complex of each open interval in L is proven homotopy equivalent to a ball or a sphere. Applications are given to the weak Bruhat order, the Tamari lattice, and the Cambrian lattices.

This paper concludes with an appendix by Dominik Preuß proving the monotone Hirsch conjecture for P a simple polytope and $G(P,\mathbf{c})$ the Hasse diagram of a lattice. This confirms one of the main consequences that the author's conjecture would have.

Keywords: nonrevisiting path conjecture, poset topology, weak order, simplex method, associahedron, permutahedron.

Data availability statement: Data sharing not applicable to this article as no datasets were generated or analysed during the current study.

1. Introduction

This paper undertakes a study of directed graphs $G(P,\mathbf{c})$ which are the 1-skeleta of simple polytopes P in $\mathbb{R}^d$ with each edge $e_{u,v}$ of the polytope oriented in the direction that "cost" increases, focusing on the case when $G(P,\mathbf{c})$ is also the Hasse diagram of a partially ordered set (poset). The **cost** of a vertex v in a polytope P is the dot product $\mathbf{c}\cdot v$ where $\mathbf{c}$ is a fixed vector in $\mathbb{R}^d$ known as the cost vector. A major reason there has been as much interest as there has been over the last several decades in the graphs $G(P,\mathbf{c})$ is the following connection to the simplex method for linear programming. The task of linear programming is to find the point $\mathbf{v}$ in a polytope (or polyhedron or more general convex body) where the cost $\mathbf{c}\cdot\mathbf{v}$ is maximized (or minimized); the simplex method accomplishes this by starting at a vertex of P and moving greedily along the edges of the 1-skeleton of P in a manner that increases (resp. decreases) cost at each step, or in other words along the directed edges of $G(P,\mathbf{c})$, until reaching the sink (resp. source) of $G(P,\mathbf{c})$. For this reason, there is particular interest in upper bounds on the diameter of $G(P,\mathbf{c})$ and on the length of the longest path in $G(P,\mathbf{c})$.

2010 *Mathematics Subject Classification.* 06A07, 52B05, 52B12, 05E45, 90C05, 05C12.

This work was supported by NSF grants DMS-1200730, DMS-1500987, DMS-1953931 and NSF conference grant DMS-1101740.

An important approach to proving such upper bounds was a conjecture known as the nonrevisiting path conjecture. This posited the existence of a path from any vertex u to any vertex v of P in the undirected version of the graph $G(P, \mathbf{c})$ with the property that the path would not revisit any facet (i.e. maximal face in the boundary of P) from which it had departed. While this conjecture was proven false by Santos in [39], there is still interest in finding classes of polytopes and cost vectors for which it does hold. We conjecture for P a simple polytope and $G(P, \mathbf{c})$ the Hasse diagram of a lattice that something far stronger than the nonrevisiting path conjecture should be true. See Conjecture 1.

We prove a number of results regarding the graphs $G(P, \mathbf{c})$ under the assumptions that P is a simple polytope and that $G(P, \mathbf{c})$ is the Hasse diagram of either a poset or more specifically of a lattice. We are hopeful that these results may be useful steps towards proving Conjecture 1. Our starting point for this work is a pair of observations, the former of which is well-known and the latter of which is proven later in the paper:

(1) Requiring "monotone paths" in the 1-skeleton of a polytope (namely paths where cost increases at each step) never to revisit any facets (a property we call the **nonrevisiting property**) implies that they also cannot revisit any faces of any dimension, by virtue of each face in a polytope being a finite intersection of facets.
(2) Requiring monotone paths in the 1-skeleton of a polytope P never to revisit any 1-dimensional faces is equivalent to requiring $G(P, \mathbf{c})$ to be the Hasse diagram of a finite partially ordered set (poset). This will be proven in Lemma 3.3.

Combining these observations, notice that if $G(P, \mathbf{c})$ has the nonrevisiting property, then $G(P, \mathbf{c})$ must be the Hasse diagram of a poset (in which case we say that $G(P, \mathbf{c})$ has the **Hasse diagram property**). Spurred on by these observations, this paper combines poset theoretic techniques with ideas from discrete geometry to prove structural results regarding $G(P, \mathbf{c})$ in the case when P is a simple polytope (see Section 2 for the definition) and $G(P, \mathbf{c})$ is the Hasse diagram of a poset or more specifically of a lattice (also reviewed in Section 2).

One motivation for studying polytopes with this nonrevisiting property is that it implies an upper bound of $n-d$ on the directed diameter of $G(P, \mathbf{c})$ where n is the number of facets in P and d is the dimension of P. In fact it forces the simplex method for linear programming to run on P with cost vector $\mathbf{c}$ in at most $n - d$ steps regardless of choice of pivot rule (another notion reviewed in Section 2). A second, quite different type of motivation for our work is that lattices whose Hasse diagrams may be realized as 1-skeleta of simple polytopes will turn out to have well-controlled topological structure, as will be proven in Theorem 1.2; this gives a new and unified explanation why several classes of posets all have all of their open intervals homotopy equivalent to balls or spheres. Turning this around, Theorem 1.2 may give some new insight into the question of which Hasse diagrams of posets may be realized as 1-skeleta of (simple) polytopes. See Remark 1.3 for more on this.

The famous Klee-Minty cubes (introduced in [26] and discussed more in Example 3.4) violate not only our nonrevisiting property, but also the Hasse diagram property. In fact, they violate this in a way that really seems to be at the heart of why the simplex method for linear programming can be so inefficient on Klee-Minty cubes.

We assume throughout this paper that $\mathbf{c}$ is a "generic" cost vector in $\mathbb{R}^d$, by which we mean that $\mathbf{c} \cdot u \neq \mathbf{c} \cdot v$ for each pair u, v of vertices of our polytope $P \subseteq \mathbb{R}^d$ that are the two endpoints of an edge in P. The resulting directed graph $G(P, \mathbf{c})$ is easily seen to be acyclic. While others have previously studied combinatorial questions related to the simplex method for linear programming and to diameter bounds on polytopes (see e.g. [39], [42], [26], [27], [10]), we are not aware of any other work in which $G(P, \mathbf{c})$ is assumed to be a Hasse diagram. This condition is not only necessary to have the nonrevisiting property, but is also a surprisingly useful input for many of our proofs.

Our focus on simple polytopes throughout this paper is in some sense not such a severe restriction, since Klee and Walkup proved in [27] that the Hirsch Conjecture (reviewed in Section 2) for simple polytopes would have implied the Hirsch Conjecture for all polytopes. Many of our results might also hold for non-simple polytopes, with obvious small modifications to the statements of the results. Our proofs do heavily rely upon our assumptions that our polytopes P are simple and that the graphs $G(P, \mathbf{c})$ are Hasse diagrams.

Now let us describe our main results, beginning with some terminology that this will require. Given a polytope P and a generic cost vector $\mathbf{c}$, each face F of P has a unique source and a unique sink in the restriction of $G(P, \mathbf{c})$ to F. Given a simple polytope P such that $G(P, \mathbf{c})$ is the Hasse diagram of a poset Q, then for any $u \in Q$ and any collection $a_1, \ldots, a_i$ of elements all covering u in Q, there is a unique i-dimensional face F in P containing u along with the outward edges from u to the elements $a_1, \ldots, a_i$. We then define the **pseudo-join** of $a_1, a_2, \ldots, a_i$ to be the unique sink of this face F. This definition is further explained and justified in Section 3. One might hope in the case that Q is a lattice that the pseudo-join of $a_1, \ldots, a_i$ would equal the join (namely the unique least upper bound) of $a_1, \ldots, a_i$.

In Theorem 4.7, we prove this equivalence of the join and the pseudo-join operations:

Theorem 1.1. *If P is a simple polytope and $\mathbf{c}$ is a generic cost vector such that $G(P, \mathbf{c})$ is the Hasse diagram of a lattice L, then the pseudo-join of any collection of atoms equals the join of this same collection of atoms. Moreover, this also holds for each interval $[u, v]$ in L.*

Theorem 1.1 combined with the Quillen Fiber Lemma also leads to the following topological property for lattices whose Hasse diagrams arise as 1-skeleta of simple polytopes:

Theorem 1.2. *If P is a simple polytope and $\mathbf{c}$ is a generic cost vector such that $G(P, \mathbf{c})$ is the Hasse diagram of a lattice L, then each open interval (u, v) in L has order complex which is homotopy equivalent to a ball or a sphere of some dimension. Therefore, the Möbius function $\mu_L(u, v)$ only takes values $0, 1$, and -1.*

This is proven as Theorem 4.11. Applications are given to the weak Bruhat order, the Tamari lattice, and to the c-Cambrian lattices in Theorems 5.6, 5.7 and 5.8, respectively. Posets meeting the hypotheses of Theorem 1.2 need not be shellable, as explained in Remark 4.12, so other methods are indeed needed to understand their topological structure.

Remark 1.3. People sometimes ask whether the Hasse diagrams for a given family of posets are realizable as 1-skeleta of polytopes (or more specifically of simple polytopes). Well-known examples having this property include the weak order, the Tamari lattice, and more recently the Cambrian lattices (arising from the theory of cluster algebras). Theorem 1.2 gives a necessary condition for a Hasse diagram of a lattice to be the 1-skeleton of a simple polytope.

In light of these and other results proven in this paper, we make the following conjecture:

Conjecture 1. *Given a simple polytope P and a generic cost vector $\mathbf{c}$ such that $G(P, \mathbf{c})$ is the Hasse diagram of a lattice, the directed paths in $G(P, \mathbf{c})$ cannot revisit any faces. That is, any directed path $v_1 \to v_2 \to \cdots \to v_k$ with v_1 and v_k in a face F has $v_i \in F$ for $1 \leq i \leq k$.*

Conjecture 1 would directly imply Theorem 1.1 above. We do prove part of Conjecture 1 (see Corollary 4.3), namely we prove the desired face nonrevisiting property for those faces F of P containing either the source or the sink of P. We use this together with other structural results in our paper to deduce Conjecture 1 in the case of 3-dimensional polytopes, carrying this out in Theorem 5.1. We also prove Conjecture 1 for spindles (see Definition 2.3) whose two distinguished vertices are the source and sink of the polytope, doing this in Theorem 5.2. In other words, we have shown that the established method that has led to counterexamples to the Hirsch Conjecture (see [39]) cannot yield counterexamples to Conjecture 1.

Given any simple polytope P and any generic cost vector $\mathbf{c}$ such that $G(P, \mathbf{c})$ is the Hasse diagram of a lattice, Conjecture 1 would imply that no directed path in $G(P, \mathbf{c})$ could have length more than $n - d$; that is, it would imply the Monotone Hirsch Conjecture (cf. Conjecture 4) for such P and $\mathbf{c}$. Dominik Preuß recently informed us that he has proven the Monotone Hirsch Conjecture for P a simple polytope with cost vector $\mathbf{c}$ such that $G(P, \mathbf{c})$ is the Hasse diagram of a lattice, motivated to do so by our work (private communication from Dominik Preuß). Preuß has provided his proof as Appendix A to our paper.

In Section 2, we review background. Section 3 introduces and develops seemingly new (or at least not widely known) notions to be used later. Section 4 gives the proofs of our main technical results. This includes Theorems 1.1 and 1.2 mentioned above as well as a number of consequences and related results. Section 5.1 gives applications to two important classes of polytopes, namely the 3-polytopes and the spindles. Section 5.2 shows how other well-known families of polytopes also fit into our framework. Section 5.3 turns to the case of zonotopes, where especially clean results are possible. Section 5.4 extends our results to more general acyclic orientations derived from shellings. Further questions and remarks appear in Section 6. The paper concludes with the aforementioned appendix by Dominik Preuß.

The author is indebted to Karola Mészáros for several invaluable conversations at early stages of this project. She also thanks Georgia Benkart, Stephanie van Willigenberg, Monica Vazirani, and the Banff International Research Station (BIRS) for conducting an inspiring workshop entitled Algebraic Combinatorixx for female researchers in algebraic combinatorics. In some sense, this project grew out of discussions that began at that workshop.

The author also thanks Louis Billera, Heather Dye, Gil Kalai, Nathan Reading, Victor Reiner, Francisco Santos, Bridget Tenner, and Günter Ziegler for extremely helpful discussions, references, and in some cases counterexamples to questions she raised while working on this project. She thanks the anonymous referees for lots of highly valuable feedback that made the paper much better. Finally, the author wishes to express her gratitude for the excellent mentoring of Victor Klee and the combinatorics group at the University of Washington during her term from 1999 to 2001 in a postdoctoral position there and for introducing her to the interesting topic of the Hirsch Conjecture back then.

2. Background

A **cover relation** $u \prec v$ in a finite partially ordered set (poset) Q is $u \leq v$ in Q with the requirement that $u \leq z \leq v$ implies either $z = u$ or $z = v$. The **Hasse diagram** of a finite poset Q is the directed graph with directed edges $u \to v$ if and only if $u \prec v$ in Q. If a poset has a unique minimal element, denote this element by $\hat{0}$. If a poset has a unique maximal element, denote this by $\hat{1}$. An **atom** in a poset Q with $\hat{0}$ is any $a \in Q$ satisfying $\hat{0} \prec a$. Likewise a **coatom** in a poset Q with $\hat{1}$ is any element c satisfying $c \prec \hat{1}$.

If $x, y \in Q$ have a unique least upper bound, this is called the **join** of x and y, denoted $x \vee y$. If $x, y \in Q$ have a unique greatest lower bound, this is the **meet** of x and y, denoted $x \wedge y$. A poset Q is a **lattice** if each pair of elements $x, y \in L$ have a meet and a join. Any partially ordered set Q gives rise to a **dual** partial order on the same set of elements, denoted Q^*, with $u \le v$ in Q^* if and only if $v \le u$ in Q.

Remark 2.1. A key example throughout this paper of poset duality will be as follows: whenever $G(P, \mathbf{c})$ is the Hasse diagram of a poset L, then the dual poset L^* will have as its Hasse diagram the directed graph $G(P, -\mathbf{c})$.

Denote by (u, v) the subposet of Q comprised of those $z \in Q$ satisfying $u < z < v$. This is known as the **open interval** from u to v. Likewise, we define the **closed interval** from u to v, denoted $[u, v]$, to be the suposet of elements $z \in Q$ satisfying $u \le z \le v$. Define the Möbius function of Q, denoted μ_Q, recursively by setting $\mu_Q(u, u) = 1$ for each $u \in Q$ and

$$\mu_Q(u, v) = -\sum_{u \le z < v} \mu_Q(u, z).$$

The **order complex** of a finite poset Q, denoted $\Delta(Q)$, is the simplicial complex whose i-faces are the chains $v_0 < \cdots < v_i$ of $i + 1$ comparable elements of Q. We let $\Delta(u, v)$ (or $\Delta_Q(u, v)$) denote the order complex of the open interval (u, v) in Q. By definition, a poset and its dual poset have the same order complex. It is well-known (by Hall's Theorem) that $\mu_Q(u, v) = \tilde{\chi}(\Delta(u, v))$ where $\tilde{\chi}$ is the **reduced Euler characteristic** of $\Delta(u, v)$, namely

$$\tilde{\chi}(\Delta(u, v)) = -1 + f_0(\Delta(u, v)) - f_1(\Delta(u, v)) + f_2(\Delta(u, v)) - \cdots$$

for $f_i(\Delta)$ the number of i-dimensional faces in Δ. Sometimes we will speak of the homotopy type of a poset or of a poset interval, by which we mean the homotopy type of the order complex of that poset or that poset interval. See e.g. [41] for further background on posets.

A **polytope** is any subset of $\mathbb{R}^d$ arising as the convex hull of a finite set of vertices in $\mathbb{R}^d$ for some d; equivalently, a polytope is any bounded set given by a system of non-strict linear inequalities, or in other words any bounded set expressible as $\{\mathbf{x} \in \mathbb{R}^d | A\mathbf{x} \le \mathbf{b}\}$ for some choice of constant $n \times d$ real matrix A and some choice of constant vector $\mathbf{b} \in \mathbb{R}^n$. We call a polytope a **d-polytope** if there is a d-dimensional affine space containing the polytope but there is not a $(d-1)$-dimensional affine space containing this same polytope. An excellent reference to learn about polytopes in more detail than we will give here is [43].

Any hyperplane H that intersects a polytope P nontrivially but has all points of P either contained in H or on one side of H is called a **bounding hyperplane** of P. The intersection of a bounding hyperplane with a polytope is called a **face** of the polytope. A maximal face in the boundary of a polytope is called a **facet**.

A polytope is **simplicial** if each face in its boundary is a simplex. A polytope P is **simple** if for each vertex $v \in P$ and each collection of i edges emanating outward from v, there is an i-dimensional face of P containing v and all these edges incident to v. Equivalently, a polytope P is simple if its dual polytope (briefly discussed next and defined more precisely e.g. in [43]) is a simplicial polytope.

The **face poset**, denoted $F(P)$, of a polytope P is the partial order on faces with $\sigma < \tau$ if and only if σ is in the boundary of τ. Each polytope P has a **dual polytope**, denoted P^* with face poset satisfying $F(P^*) = (F(P))^*$.

Remark 2.2. As a word of caution, note that in cases where $G(P, \mathbf{c})$ is the Hasse diagram of a poset Q, the face poset $F(P)$ of the polytope P is typically a completely different poset than Q. In this case negating the cost vector $\mathbf{c}$ while preserving the polytope P yields a directed graph $G(P, -\mathbf{c})$ on the 1-skeleton of P which is the Hasse diagram of Q^*.

Definition 2.3. A **spindle** is a polytope P with a distinguished pair of vertices u and v such that each facet of P includes either u or v. The dual polytope to a spindle is called a **prismatoid**, and it is characterized by the property that it has two distinguished facets such that every vertex belongs to one or the other of these two facets.

A **zonotope** is a polytope arising as a linear projection of a cube of some dimension. In other words, a zonotope is a Minkowski sum of line segments.

In linear programming, a **pivot rule** for the simplex method is a rule for choosing for each vertex v of $G(P, \mathbf{c})$ which outward (resp. inward) oriented edge from v to traverse in choosing a directed path to the sink (resp. source).

Now we recall the Hirsch Conjecture, the Nonrevisiting Path Conjecture, and the Monotone Hirsch Conjecture. We refer readers e.g. to [43] for a more in-depth discussion of all of these conjectures.

Conjecture 2 (Hirsch Conjecture)**.** *For $n > d \geq 2$, let $\Delta(d,n)$ denote the largest possible diameter of the graph of a d-polytope with n facets. Then $\Delta(d,n) \leq n - d$.*

Conjecture 3 (Nonrevisiting Path Conjecture)**.** *For any two vertices u, v of a d-dimensional polytope, there is a path from u to v which does not revisit any facet it has left before.*

The Nonrevisiting Path Conjecture, proposed by Klee and Wolfe, implies the Hirsch Conjecture. To see this implication, notice that any directed path from u to v of the type given by the Nonrevisiting Path Conjecture would involve at most $n - d$ edges, since each edge would depart a facet, with no facet departed more than once, and since the ending vertex v for the path would still belong to d facets; thus, each pair of vertices u, v would have a path of length at most $n - d$ between them. Counterexamples to the Hirsch Conjecture (and thereby also to the Nonrevisiting Path Conjecture) were first obtained by Francisco Santos in [39]:

Theorem 2.4 (Santos)**.** *The Hirsch Conjecture is false. Therefore, the Nonrevisiting Path Conjecture is also false.*

One may still ask for sufficient conditions on a polytope for these conjectures to hold and for our even stronger face nonrevisiting property to hold. This latter question is related to the Monotone Hirsch Conjecture for polytopes:

Conjecture 4 (Monotone Hirsch Conjecture)**.** *Let $H(d,n)$ be the smallest integer N such that for every d-polytope P in $\mathbb{R}^d$, every cost vector $\mathbf{c} \in \mathbb{R}^d$ in general position with respect to P, and every vertex $v \in P$, there exists a strict monotone path from v to the sink of P using at most N steps. Then $H(d,n) \leq n - d$.*

To see the connection, notice that any polytope P and cost vector $\mathbf{c}$ such that no directed path in $G(P,\mathbf{c})$ revisits any facet must have $G(P,\mathbf{c})$ of diameter at most $n-d$, by the same reasoning which showed that the Nonrevisiting Path Conjecture implies the Monotone Hirsch Conjecture for polytopes (see the discussion just after Conjecture 3 above).

Todd gave the first counterexample to the Monotone Hirsch Conjecture in [42]; this came from a 4-polytope with 8 facets and a vertex v requiring at least 5 steps for all monotone paths from v to the sink. However, this polytope has an edge from v to the source and then a directed edge from source to sink. From the standpoint of our own conjecture, it is perhaps worth noting that the existence of this directed edge prevents the directed graph in this case from being a Hasse diagram.

We conclude this section with some further background on poset topology. A map $f : P \to Q$ from a poset P to a poset Q is a **poset map** if $u \le v$ in P implies $f(u) \le f(v)$ in Q.

Theorem 2.5 (Quillen Fiber Lemma, [34])**.** *Let $f : P \to Q$ be a poset map such that for each $q \in Q$ the order complex $\Delta(f^{-1}_{\ge q})$ for $f^{-1}_{\ge q} = \{p \in P | f(p) \ge q\}$ is contractible. Then $\Delta(P) \simeq \Delta(Q)$.*

A **dual closure map** is a poset map $f : P \to P$ with $f(u) \le u$ and $f^2(u) = f(u)$ for all $u \in P$. Notice that any such f meets the contractibility requirement of the Quillen Fiber Lemma, by virtue of each $u \in im(f)$ being a cone point in $\Delta(f^{-1}_{\ge u})$. Thus, $\Delta(im(f)) \simeq \Delta(P)$ in this case.

Remark 2.6. The poset map f sending each element u in a finite lattice to the join of those atoms a satisfying $a \le u$ is a dual closure map with the further property that $f^{-1}(\hat{0}) = \{\hat{0}\}$. Thus, the Quillen Fiber Lemma yields $\Delta(P \setminus \{\hat{0}\}) \simeq \Delta(im(f) \setminus \{\hat{0}\})$ in this case.

3. The nonrevisiting property, the Hasse diagram property, pseudo-joins and pseudo-meets

In this section, we introduce or in some cases make more precise a few seemingly new notions which will play key roles throughout this paper.

Definition 3.1. A directed graph $G(P,\mathbf{c})$ on the 1-skeleton of a polytope P satisfies the **nonrevisiting property** if for each facet F and each directed path p_F that starts and ends at vertices in F, the path p_F must stay entirely within F. We say that $G(P,\mathbf{c})$ satisfies the nonrevisiting property for i-dimensional faces if for each i-face F and each directed path p_F that starts and ends at vertices of F, every vertex of p_F is in F.

Remark 3.2. This nonrevisiting property for facets is equivalent to the nonrevisiting property for all faces of every dimension since each face is an intersection of facets which means that to depart a face and revisit that face would require departing at least one facet containing the face and then revisiting this same facet.

The nonrevisiting property for 1-faces is equivalent to the Hasse diagram property:

Lemma 3.3. *The nonrevisiting property holds for 1-dimensional faces of a polytope P with respect to cost vector $\mathbf{c}$ if and only if $G(P,\mathbf{c})$ is the Hasse diagram of a poset.*

Proof. First we show how the nonrevisiting property for 1-faces implies $G(P,\mathbf{c})$ is a Hasse diagram. Acyclicity of $G(P,\mathbf{c})$ ensures that the directed paths indeed specify comparabilities in a partially ordered set. The nonrevisiting property for 1-dimensional faces guarantees that a directed path cannot visit the sink of a directed edge after departing from the source of that directed edge in a manner other than traversing that edge. This shows that each directed edge gives rise to a cover relation in the poset rather than simply an order relation.

The other direction follows directly from the definition of a cover relation in a poset. □

Example 3.4. An especially instructive and significant family of simple polytopes failing the Hasse diagram property are the Klee-Minty cubes (cf. [26]). Klee-Minty cubes are polytopes P that are realizations of d-dimensional hypercubes admitting a cost vector $\mathbf{c}$ with the property that a directed path exists in $G(P,\mathbf{c})$ that visits all 2^d vertices of P. These were historically the first examples of polytopes demonstrating that the simplex method for linear programming is not a polynomial time algorithm.

The d-dimensional Klee-Minty cube may be realized for any fixed ϵ satisfying $0 < \epsilon < 1/2$ as $\{(x_1,\dots,x_d) \in \mathbb{R}^d | 0 \le x_1 \le 1; \epsilon x_{i-1} \le x_i \le 1-\epsilon x_{i-1} \text{ for } i > 1\}$. For example, the 2-dimensional Klee-Minty cube in $\mathbb{R}^2$ with $\epsilon = 1/3$ is the convex hull of vertices $(0,0),(0,1),(1,1/3),(1,2/3)$; for cost vector $\mathbf{c} = (0.1,1)$, the resulting graph $G(p,\mathbf{c})$ has a directed path p from $(0,0)$ to $(1,1/3)$ to $(1,2/3)$ to $(0,1)$ that visits all 4 vertices. Notice that this 2-polytope also has a directed edge e from $(0,0)$ to $(0,1)$, which means that the aforementioned directed path p departs from the edge e at $(0,0)$ and later revisits e at $(0,1)$. This demonstrates that $G(P,\mathbf{c})$ is not a Hasse diagram in this case.

For further background and properties of Klee-Minty cubes, including a helpful illustration of a 3-dimensional Klee-Minty cube, we refer readers to [16].

Recall for a polytope P and generic cost vector $\mathbf{c}$ that the **source** of a face F is the vertex $v \in F$ minimizing $\mathbf{c} \cdot v$ while the **sink** of F is the vertex $w \in F$ maximizing $\mathbf{c} \cdot w$. The source of F is also the unique vertex in F only having outward oriented edges to other vertices of F while the sink is the unique vertex in F only having inward oriented edges to it from other vertices of F. The uniqueness assumption that is implicit in these notions above is justified as follows:

Remark 3.5. It is well known (cf. Theorem 3.7 in [43]) and straight-forward to see that there is a unique source and a unique sink in the directed graph on the 1-skeleton of any face F of a polytope P obtained by restricting $G(P,\mathbf{c})$ for a generic $\mathbf{c}$ to the face F.

Next is a simple observation that is surprisingly useful in various proofs later in the paper.

Remark 3.6. The directed graph $G(P,\mathbf{c})$ restricted to any 2-dimensional face F consists of two directed paths from the unique source of F to the unique sink of F. These paths are disjoint except at their two endpoints, i.e., the source and sink.

Next we introduce and justify existence of the following pair of notions which are dual to each other (the first of which was already introduced more informally in the introduction):

Definition 3.7. Consider any simple polytope P and any generic cost vector $\mathbf{c}$ such that $G(P,\mathbf{c})$ is the Hasse diagram of a poset Q. Define the **pseudo-join** of any collection S of atoms of an interval $[u,v]$ in Q, denoted $psj(S)$, to be the sink of the unique smallest face F_S of P that contains u and all of the elements of S. Define the **pseudo-meet** of any collection T of coatoms in an interval $[u,v]$ in Q, denoted $psm(T)$, to be the source vertex of the smallest face G_T containing v and all of the elements of T.

Next let us justify the existence (and uniqueness) of such pseudo-joins.

Lemma 3.8. *Consider a simple polytope P and a generic cost vector $\mathbf{c}$ such that $G(P, \mathbf{c})$ is the Hasse diagram of a poset Q. Any collection S of atoms in any interval $[u, v]$ in Q has a unique pseudo-join.*

Proof. For any vertex v in $G(P, \mathbf{c})$ and any collection S of neighboring vertices that cover v in Q, consider the edges $e_1, \ldots, e_i$ emanating outward from v whose other endpoints are the elements of S. A face of P will contain all of the vertices in $S \cup \{v\}$ if and only if it contains all of the edges $e_1, \ldots, e_i$. By the definition of simple polytope, there exists an i-face F containing v and all of these edges $e_1, \ldots, e_i$. This is necessarily the unique smallest face containing v and all of the edges $e_1, \ldots, e_i$, due to any such face needing to be at least i-dimensional combined with the fact that the intersection of any two faces containing all of these edges $e_1, \ldots, e_i$ will also contain all of these same edges. Since F has a unique sink, this makes the sink of F the desired pseudo-join. $\square$

4. Poset theoretic results regarding 1-skeleta of simple polytopes

In this section, we develop a series of general results about directed paths in the graphs $G(P, \mathbf{c})$ under the assumptions that P is a simple polytope, $\mathbf{c}$ is a generic cost vector, and $G(P, \mathbf{c})$ is the Hasse diagram of a poset Q. This leads to results later in this section on the topological structure of the order complexes of such posets.

Lemma 4.1. *Let P be a simple polytope with faces $F \subseteq G$ satisfying $dim(G) = dim(F) + 1$. Let $\mathbf{c}$ be a generic cost vector such that $G(P, \mathbf{c})$ is the Hasse diagram of a poset Q. Given vertices $v, w \in F$ with a directed path p_F from v to w fully contained in F, $G(P, \mathbf{c})$ cannot have an edge from v directed outward to some vertex $v' \in G \setminus F$ and an edge directed from some $w' \in G \setminus F$ to w.*

Proof. Suppose there is a directed path p_F of the type we aim to exclude. Since P is simple, each vertex u in p_F has exactly one edge e_u incident to it whose other endpoint is in $G \setminus F$. If e_u is oriented outward from u, denote this by $o(u) = +1$, whereas we say $o(u) = -1$ when e_u is oriented towards u. Our hypotheses give us that $v \in p_F$ has $o(v) = +1$ while $w \in p_F$ has $o(w) = -1$. This forces the existence of two consecutive vertices $v_1 \to v_2$ in the directed path p_F with $o(v_1) = +1$ and $o(v_2) = -1$.

We now show that the edges e_{v_1,x_1} and e_{v_2,x_2} from $v_1 \in F$ to some $x_1 \in G \setminus F$ and from some $x_2 \in G \setminus F$ to $v_2 \in F$ must both be contained in a single 2-dimensional face $F(e_{v_1}, e_{v_2})$ in G that also contains the edge e_{v_1,v_2}. The fact that P is simple implies that the pair of edges e_{v_1,v_2} and e_{v_1,x_1} are both contained in a 2-face $F(v_1, v_2, x_1)$. Moreover, $F(v_1, v_2, x_1) \not\subseteq F$ since $x_1 \notin F$. Likewise there exists a 2-face $F(v_1, v_2, x_2)$ containing e_{v_1,v_2} and e_{x_2,v_2} with $F(v_1, v_2, x_2) \not\subseteq F$ because $x_2 \notin F$. But the fact that P is simple implies that each edge e in F is contained in a unique 2-face σ in G such that $\sigma \not\subseteq F$; this follows from each upper interval in the face poset of a simple polytope being a Boolean lattice and in particular the interval $[e, G]$ being a Boolean lattice where all but one of its atoms is in $[e.F]$. Applying this observation to the edge e_{v_1,v_2} yields that $F(v_1, v_2, x_1)$ and $F(v_1, v_2, x_2)$ must both be this unique 2-face containing e_{v_1,v_2} and not contained in F. Thus, $F(v_1, v_2, x_1) = F(v_1, v_2, x_2)$ with these both being the unique 2-face which contains all three edges e_{v_1,v_2}, e_{v_1,x_1}, and e_{v_2,x_2}.

Observe then that v_1 must be the unique source for $F(v_1, v_2, x_1)$ and that v_2 must be the unique sink for $F(v_1, v_2, x_1)$. But there is a directed edge from v_1 to v_2 in $G(P, \mathbf{c})$, namely e_{v_1,v_2}. This edge must constitute one of the two directed paths from the source to the sink in the boundary of $F(v_1, v_2, x_1)$ that are guaranteed to exist by Remark 3.6. For $G(P, \mathbf{c})$ to be the Hasse diagram of a poset Q, e_{v_1,v_2} must give rise to a cover relation $v_1 \prec v_2$ in Q. However, the other directed path from v_1 to v_2 in the boundary of $F(v_1, v_2, x_1)$ gives a saturated chain from v_1 to v_2 in Q having x_1 as an intermediate element, a contradiction to e_{v_1,v_2} being a cover relation. □

With a little more work we also deduce the following:

Lemma 4.2. *Let P be a simple polytope and let $\mathbf{c}$ be a generic cost vector such that $G(P, \mathbf{c})$ is the Hasse diagram of a poset Q. Let F be any face of P containing the source vertex of P. Then each edge $e_{v,w}$ with $v \in F$ and $w \notin F$ must be oriented from v to w. Likewise for any face F' of P containing the sink vertex of P, any edge $e_{x,y}$ with $x \in F'$ and $y \notin F'$ must be oriented from y to x.*

Proof. Let F be a face of P containing the source vertex of P, so in other words containing $\hat{0} \in Q$. Suppose there is an edge $e_{w,v}$ oriented from $w \in P \setminus F$ to $v \in F$. Since $\hat{0}$ is the unique source in P (and hence in F), there must be a directed path $\hat{0} = v_0 \prec v_1 \prec v_2 \prec \cdots \prec v_k \prec v$ staying within F. Now we will apply Lemma 4.1 to the face F viewed as a codimension one face of the unique simple polytope G containing both $e_{w,v}$ and F.

The proof of the statement for any face F' containing the sink of P is similar. □

Corollary 4.3. *Let P be a simple polytope and let $\mathbf{c}$ be a generic cost vector such that $G(P, \mathbf{c})$ is the Hasse diagram of a poset. Then each face F of P which contains $\hat{0}$ has the property that there are no directed paths which revisit F. Likewise for each face G containing $\hat{1}$, there cannot be any directed paths that revisit G.*

Corollary 4.3 has the following important special case.

Corollary 4.4. *Let S be any set of atoms in a poset P whose Hasse diagram is $G(P, \mathbf{c})$ for P a simple polytope. Then each directed path of $G(P, \mathbf{c})$ from $\hat{0}$ to $psj(S)$ stays within the unique smallest face F_S of P containing $\hat{0}$ and all of the elements of S. Likewise for each set T of coatoms, each directed path from $psm(T)$ to $\hat{1}$ stays within the unique smallest face F_T containing $\hat{1}$ and all of the coatoms in T.*

This in turn implies the following:

Corollary 4.5. *Let P be a simple polytope and let $\mathbf{c}$ be a generic cost vector such that $G(P, \mathbf{c})$ is the Hasse diagram of a lattice L. Then the join of any collection S of atoms in L is contained in the unique smallest face F_S containing $\hat{0}$ and all of the atoms in S. Dually, the meet of any collection T of coatoms in L is contained in the unique smallest face G_T containing $\hat{1}$ and all of the coatoms in T.*

Proof. There is a directed path from any atom $a \in S$ to the join $J(S)$ of the set of elements of S since $J(S)$ is an upper bound for the elements of S. There is also a directed path from $J(S)$ to the pseudo-join $psj(S)$ by virtue of $psj(S)$ being an upper bound for the elements of S, hence being greater than or equal to the least upper bound $J(S)$. Concatenating these directed paths yields a directed path p from a to $psj(S)$. By Corollary 4.3, p must stay within F_S since $\hat{0} \in F_S$. In particular, this implies $J(S) \in F_S$. The proof for coatoms is similar. □

Next is a result that will serve as a key ingredient in the proof of Theorem 4.7. The proof of Theorem 4.6 is one of the trickiest proofs in the paper, so we give an especially high level of detail in the argument below.

Theorem 4.6. *Let P be a simple polytope and let $\mathbf{c}$ be a generic cost vector such that $G(P, \mathbf{c})$ is the Hasse diagram of a lattice L. Let F be a 2-face in P, let u be the source vertex in F, and let $x, y \in F$ be vertices both covering u in L. Then $psj(x, y) = x \vee y$. Dually, for any 2-face F' in P with sink v and elements $x', y' \in F'$ both covered by v in L, $psm(x', y') = x' \wedge y'$.*

Proof. If a 2-face σ in P with source u has $x, y \in \sigma$ satisfying $u \prec x$ and $u \prec y$ in L with $psj(x, y) \neq x \vee y$, let us say that σ has property (LT). Since $psj(x, y)$ is an upper bound for x and y, property (LT) implies $x \vee y <_L psj(x, y)$ where $<_L$ denotes the order relation for L. Given any $u \in L$, let $d(u)$ be the length of the longest saturated chain in L from $\hat{0}$ to u.

We will prove below that P cannot have any 2-faces with property (LT). The dual case will then follow by replacing the cost vector $\mathbf{c}$ with its negation $-\mathbf{c}$ which will reverse the direction of all the directed edges in $G(P, \mathbf{c})$, having the effect of replacing L by the dual poset L^* with the meet operation for L becoming the join operation for L^* and the pseudo-meet operation in $G(P, \mathbf{c})$ becoming the pseudo-join operation in $G(P, -\mathbf{c})$.

Supposing there were a 2-face with property (LT), choose such a 2-face F with source u with $d(u)$ as large as possible among all 2-faces with property (LT). Denote this largest value for $d(u)$ by $d^{\max}$. We will show that the existence of such a 2-face F forces the existence of such a 2-face F'' with source u'' also satisfying property (LT) with $d(u'') > d^{\max}$, thereby yielding a contradiction. Figure 1 may help in reading this upcoming argument.

Let x and y be the elements of F covering u in L. The Hasse diagram of $F \cap (u, psj(x, y))$ is obtained by removed the vertices u and $psj(x, y)$ from the boundary of F; this yields two disjoint nonempty paths, one containing x and the other containing y (by Remarks 3.5 and 3.6). Denote the vertices proceeding upward in L along one of these two paths from u to $psj(x, y)$ as $u, x_1, x_2, \ldots, x_r$ where $x_1 = x$ and $x_r = psj(x, y)$. Denote the vertices proceeding upward in L along the other path from u to $psj(x, y)$ as $u, y_1, y_2, \ldots, y_s$ where $y_1 = y$ and $y_s = psj(x, y)$.

The strict inequality $x \vee y <_L psj(x, y)$ implies either (a) $x \vee y \notin F$ or (b) $x \vee y \in \{x_1, \ldots, x_{r-1}, y_1, \ldots, y_{s-1}\}$. In the latter case, we may assume without loss of generality that $x \vee y \in \{x_1, \ldots, x_{r-1}\}$. Let us now show that in either case there must be some directed path that exits F at some y_i with $1 \leq i < r - 1$ and revisits F. This follows in case (a) from the existence of a directed path from $y = y_1$ to $x \vee y \notin F$ along with the existence of a directed path from $x \vee y$ to $psj(x, y)$, the latter of which exists because $psj(x, y)$ is an upper bound for x and y and hence must be greater than or equal to $x \vee y$. In case (b), we use the existence of a directed path from $y = y_1$ to $x \vee y \in \{x_1, \ldots, x_{r-1}\}$ together with the disconnectedness of the restriction of F to the vertex set $\{x_1, \ldots, x_{r-1}, y_1, \ldots, y_{s-1}\}$; this directed path from y_1 to $x \vee y$ must exit and reenter F to get from the connected component containing y_1 to the connected component containing $x \vee y$. In either case, we have now shown that there is a directed path that exits F at some y_i for $i \in \{1, 2, \ldots, s - 1\}$ and revisits F.

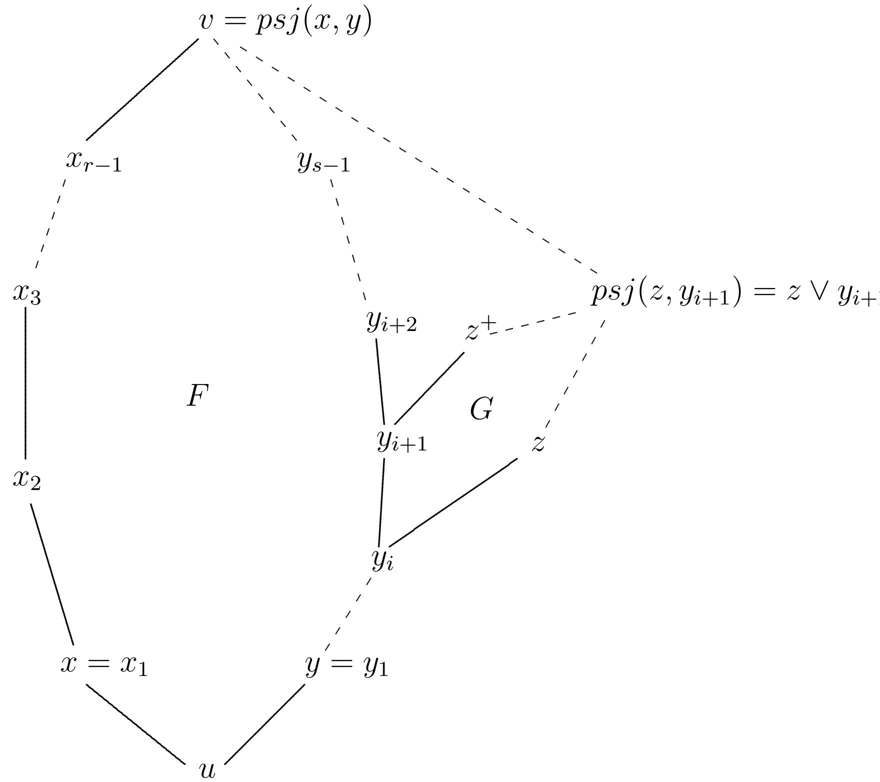

FIGURE 1. Diagram for proof of Theorem 4.6

Consider a directed path p_i that exits from F via an edge from some $y_i \in F$ to some $z \notin F$ and later revisits F. Choose such y_i and p_i with i as large as possible. See Figure 1. Our choice of F with source u such that $d(u) = d^{\max}$ allows us to deduce that $y_{i+1} \vee z = psj(y_{i+1}, z)$ as follows: otherwise y_i, y_{i+1} and z would all belong to a 2-face having source y_i and having property (LT) with $d(y_i) > d(u)$, contradicting our choice of F and u with $d(u)$ as large as possible. Thus, we may assume $y_{i+1} \vee z = psj(y_{i+1}, z)$. But the identity $y_{i+1} \vee z = psj(y_{i+1}, z)$ implies that the four vertices y_i, y_{i+1}, z and $y_{i+1} \vee z$ are all contained in a single 2-face G having $y_{i+1} \vee z$ as its sink. Since $z \in G$ and $z \notin F$, we must have $G \neq F$.

Since F and G are distinct 2-faces, they intersect in a proper face which is at most an edge. Thus, they share at most two vertices. F and G do share y_i and y_{i+1}, hence cannot share any further vertices. Since $y_{i+1} \vee z \in G$ for $y_{i+1} \vee z \notin \{y_i, y_{i+1}\}$, it follows that $y_{i+1} \vee z \notin F$. Similarly the elements of $G \cap (y_{i+1}, y_{i+1} \vee z]$ are not in F, so in particular the element z^+ covering y_{i+1} in G is not in F. But we have a directed path from z^+ to $y_{i+1} \vee z$ by virtue of z^+ being in the 2-face G with $y_{i+1} \vee z$ as its sink. Thus, $z^+ \leq_L y_{i+1} \vee z$. We also have $y_{i+1} \vee z \leq_L psj(x,y)$ because $psj(x,y)$ is an upper bound for y_{i+1} and z. Combining inequalities yields $z^+ \leq_L psj(x,y)$. Thus, we have a directed path from $z^+ \notin F$ to a vertex $psj(w,y) \in F$. Since we also have a directed edge from y_{i+1} to z^+, this directed edge and this directed path together exhibit the existence of a directed path exiting F at y_{i+1}, going through $z^+ \notin F$, and later revisiting F. This contradicts our choice of y_i with i as large as possible. $\square$

Theorem 4.7. *Let P be a simple polytope and let $\mathbf{c}$ be a generic cost vector such that $G(P, \mathbf{c})$ is the Hasse diagram of a lattice L. Given any $u \in L$ and any $a_1, \ldots, a_j \in L$ all covering u in L, $psj(a_1, \ldots, a_j) = a_1 \vee \cdots \vee a_j$. Dually, for any $v \in L$ and any $c_1, \ldots, c_j \in L$ all covered by v in L, $psm(c_1, \ldots, c_j) = c_1 \wedge \cdots \wedge c_j$.*

Proof. The $j = 1$ case holds tautologically. The $j = 2$ case is proven in Theorem 4.6. We will rely on these cases to prove the result for any fixed $j > 2$ (though the proof for $j > 2$ is not by induction on j). Given $u \in L$ and $\{a_1, \ldots, a_j\}$ a set of elements which all cover u in L, recall our notation $J(a_1, \ldots, a_j)$ for $a_1 \vee \cdots \vee a_j$. Let $F_{\{a_1,\ldots,a_j\}}$ be the unique smallest face containing all of the elements in $S \cup \{u\}$. Thus, u is the source of $F_{\{a_1,\ldots,a_j\}}$. Let $\leq_L$ denote the order relation on the lattice L.

First we will prove $psj(a_1, \ldots, a_j) \leq_L J(a_1, \ldots, a_j)$. To accomplish this, we will prove for each $x \in F_{\{a_1,\ldots,a_j\}}$ other than $psj(a_1, \ldots, a_j)$ that each $y \in F_{\{a_1,\ldots,a_j\}}$ that covers x satisfies $y \leq_L J(a_1, \ldots, a_j)$, assuming inductively that we have already proven this same statement for every $x' \in F_{\{a_1,\ldots,a_j\}}$ such that $x' <_L x$. To get started, we first check the base case of this inductive statement, namely the case of $x = u$. The definition of the join operation directly yields $a_i \leq_L J(a_1, \ldots, a_j)$ for $i = 1, \ldots, j$. This in turn implies $u \leq_L J(a_1, \ldots, a_j)$ since $u \leq_L a_i \leq_L J(a_1, \ldots, a_j)$ for $i = 1, 2, \ldots, j$. This completes the base case.

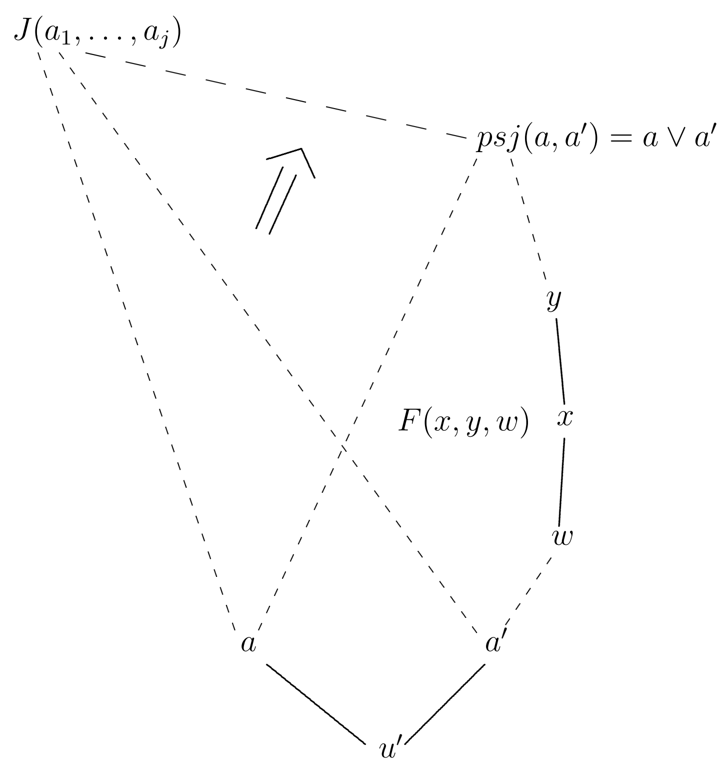

FIGURE 2. Inductive step for Theorem 4.7

Now to the inductive step where we assume $x > u$. See Figure 2. Consider any $y \in F_{\{a_1,\dots,a_j\}}$ which covers x. There must also exist some $w \in F_{\{a_1,\dots,a_j\}}$ covered by x, since we already have $x > u$. Consider the unique 2-face $F(x, y, w)$ containing x, y and w, a face which is guaranteed to exist because P is simple. Let u' be its source. Since $u' <_L x$, our inductive hypothesis allows us to assume $u' \le_L J(a_1, \dots, a_j)$ and that each element of $F_{\{a_1,\dots,a_j\}}$ covering u' is also bounded above by $J(a_1, \dots, a_j)$. In particular, we may assume that the two elements a, a' of $F(x, y, w)$ covering u' are bounded above by $J(a_1, \dots, a_j)$. But $psj(a, a') = a \vee a'$ by Theorem 4.6. Thus, $psj(a, a') \le_L J(a_1, \dots, a_j)$ due to $J(a_1, \dots, a_j)$ already being an upper bound for a and a' and hence for $a \vee a'$ as well. The inequality $psj(a, a') \le_L J(a_1, \dots, a_i)$ implies that every element of $F(x, y, w)$ is bounded above by $J(a_1, \dots, a_j)$, so in particular that $y \le_L J(a_1, \dots, a_j)$. This completes the inductive step, hence the proof that $psj(a_1, \dots, a_j) \le_L J(a_1, \dots, a_j)$.

We deduce the inequality $J(a_1, \dots, a_j) \le_L psj(a_1, \dots, a_j)$ from the fact that $psj(a_1, \dots, a_j)$ is an upper bound for all of the elements $a_1, \dots, a_j$. These weak inequalities

$$psj(a_1, \dots, a_j) \le J(a_1, \dots, a_j) \le psj(a_1, \dots, a_j)$$

combine to yield $psj(a_1, \dots, a_j) = J(a_1, \dots, a_j)$, as desired.

This also allows us to deduce the desired dual statement as follows. We negate the cost vector and observe that $G(P, -\mathbf{c})$ is the Hasse diagram of the dual poset to the poset we get from $G(P, \mathbf{c})$. Poset duality reverses the roles of the meet and join operations, reducing the desired statement about meets to the statement we have already proven about joins. □

A useful immediate consequence of Theorem 4.7 is as follows.

Corollary 4.8. *Let P be a simple polytope and let $\mathbf{c}$ be a generic cost vector such that $G(P, \mathbf{c})$ is the Hasse diagram of a lattice L. If $a_1, \dots, a_j$ all cover an element $u \in L$, then $a_1 \vee \cdots \vee a_j$ is in the unique smallest face of P containing $a_1, a_2, \dots, a_j, u$. Dually, for $v \in L$ and $c_1 \dots, c_j$ all covered by v, $c_1 \wedge \cdots \wedge c_j$ is in the unique smallest face of P containing $c_1, c_2, \dots, c_j, v$.*

Next we deduce a further property of pseudo-joins (and pseudo-meets) that will be helpful for understanding the topological structure of posets whose Hasse diagram is $G(P, \mathbf{c})$ for a simple polytope P and a generic cost vector $\mathbf{c}$.

Lemma 4.9. *Let P be a simple polytope and let $\mathbf{c}$ be a generic cost vector such that $G(P, \mathbf{c})$ is the Hasse diagram of a poset Q. If S and T are distinct sets of atoms in Q, then $psj(S) \ne psj(T)$. If S and T are distinct sets of coatoms in Q, then $psm(S) \ne psm(T)$. If Q is a lattice, then the same pair of statements holds for each interval $[u, v]$ in Q.*

Proof. Given a collection S of atoms in Q, once again let F_S be the smallest face of P containing the vertices in $S \cup \{\hat{0}\}$. Such a face F_S exists because P is simple. Note that $S \ne T$ implies $F_S \ne F_T$ since each face F_S is itself simple with $dim(F_S) = |S|$ and with the neighboring vertices to $\hat{0}$ in F_S being exactly the elements of S.

Now to our claim about pseudo-joins. First consider F_T that is a codimension one face of a face F_S, with both these faces including $\hat{0}$. Suppose also that both F_S and F_T have the same sink, denoted v. Then there is an edge directed from a vertex $v_S \in F_S \setminus F_T$ to $v \in F_T$, by virtue of v being the sink of F_S and having a neighbor in $F_S \setminus F_T$. There is also an edge directed from $\hat{0}$ to a vertex $v'_S \in F_S \setminus F_T$, by virtue of $\hat{0}$ being the source of F_S and having a neighbor in $F_S \setminus F_T$. There is also a directed path from $\hat{0}$ to v that stays within F_T because $\hat{0}$ and v are the source and sink vertices of F_T. Thus, Lemma 4.1 applies in this case, giving a contradiction to such faces F_S and F_T having the same sink.

Next we turn to the case of $F_T \subsetneq F_S$ with both faces containing $\hat{0}$ with F_T of codimension higher than one in F_S. Again suppose both F_S and F_T have the same sink. We use the existence of an intermediate face $F_{T'}$ with $F_T \subsetneq F_{T'} \subsetneq F_S$ to reduce as follows to the codimension one case above. Since F_T and F_S have the same source and sink, $F_{T'}$ must also have this same source and sink, enabling us to reduce to the lower codimension case of $F_T \subsetneq F_{T'}$. Doing this repeatedly yields the codimension one case above, allowing us to rule out this case as well.

Next we reduce the case of any two faces F_S and $F_{S'}$ both having source $\hat{0}$ and both having the same sink to the cases already ruled out above. Since the face $F_S \cap F_{S'}$ must also contain $\hat{0}$ as well as containing the common sink for F_S and $F_{S'}$, this implies $F_S \cap F_{S'}$ will also have this same sink. But $F_S \cap F_{S'} \subsetneq F_S$, allowing us to use the above cases to rule out this case. This completes the proof of the statement about distinct atoms in L.

Turning now to arbitrary intervals $[u, v]$ of Q, we now assume Q is a lattice. The fact that v is an upper bound for the set S of atoms of $[u, v]$ implies that the join of the elements of S is contained in $[u, v]$. Theorem 4.7 ensures that this join equals the pseudo-join of the elements of S, hence that the pseudo-join is also in the interval. These observations allow the argument above to be applied to the interval $[u, v]$.

The same proof applied to the dual poset yields the desired analogous statements for pseudo-meets of coatoms in all of Q as well as in any interval $[u, v]$. □

Corollary 4.10. *Let P be a simple polytope and let $\mathbf{c}$ be a generic cost vector such that $G(P, \mathbf{c})$ is the Hasse diagram of a poset Q. Then the subposet of Q consisting of all pseudo-joins of atoms (resp. pseudo-meets of coatoms) is a Boolean lattice $B_{|A|}$ for A the set of atoms (resp. coatoms) of Q. If Q is a lattice, then this same property holds for each interval $[u, v]$ in Q.*

Now to a topological consequence of the above results.

Theorem 4.11. *Let P be a simple polytope and let $\mathbf{c}$ be a generic cost vector such that the directed graph $G(P, \mathbf{c})$ is the Hasse diagram of a lattice Q. Then each nonempty open interval (u, v) has order complex that is homotopy equivalent to a ball or a sphere. Thus, $\mu_Q(u, v)$ equals $0, 1,$ or -1 for each $u \le v$.*

Proof. Given any nonempty open interval (u, v) in Q, we will define a surjective poset map $f : (u, v) \to B$ for $B = B_n \setminus \{\hat{0}, \hat{1}\}$ or $B = B_n \setminus \{\hat{0}\}$ where B_n denotes the Boolean lattice of subsets of $\{1, 2, \ldots, n\}$ ordered by containment. For $n \ge 2$, $B_n \setminus \{\hat{0}, \hat{1}\}$ has order complex homeomorphic to the sphere S^{n-2}, by virtue of being the barycentric subdivision of the boundary of the simplex Δ^{n-1}, while $B_n \setminus \{\hat{0}\}$ has contractible order complex due to having a cone point at $\hat{1}$.

Consider the map f sending each $z \in (u, v)$ to the pseudo-join of the set of atoms $a \in [u, z]$ Under our hypotheses, the join of a set of atoms of an interval $[u, v]$ equals the pseudo-join of this same set of atoms, by Theorem 4.7. This allows a reinterpretation of f as sending z to the join of the set of atoms of $[u, z]$, making it clear that f is a poset map. We proved that the pseudo-joins of distinct sets of atoms in an interval $[u, v]$ are themselves distinct in Lemma 4.9. This yields the desired Boolean lattice structure on the image of f as in Corollary 4.10. The fibers of f meet the contractibility requirement for the Quillen Fiber Lemma by virtue of each fiber having a unique smallest element, hence a cone point in the order complex of the fiber.

The claim that $\mu_Q(u,v) \in \{0,\pm 1\}$ for each $u \le v$ now follows directly from Hall's well-known interpretation for $\mu_Q(u,v)$ as the reduced Euler characteristic $\tilde{\chi}(\Delta(u,v))$ together with the fact that $\tilde{\chi}(K) = 0$ for K a ball and $\tilde{\chi}(K) = (-1)^d$ for K a d-sphere. □

Remark 4.12. The posets Q considered in Theorem 4.11 typically have order complex which is not shellable. To see this, note that a shelling would force every 2-dimensional face F in our polytope P to have one of the two directed paths from the source to the sink in the boundary of F to have exactly two edges in it. This follows from the fact that the boundary of F with its source and sink removed would be disconnected which would force one of its two connected components to have order complex that is 0-dimensional in order for a shelling of the order complex of Q to be possible. See e.g. [9] for background on shellability.

5. Applications to various classes of polytopes and regular CW balls

Next we apply our earlier results to several classes of polytopes (and regular CW balls). Rather than trying to give as comprehensive a list of applications as possible, we focus on some important families where our theory applies particularly naturally.

5.1. Applications to 3-polytopes and to spindles.

Theorem 5.1. *Let P be a simple polytope of dimension 3, and let $\mathbf{c}$ be a generic cost vector such that $G(P,\mathbf{c})$ is the Hasse diagram of a lattice L. Then $G(P,\mathbf{c})$ has the nonrevisiting property. That is, any directed path from u to v with u,v both contained in a face F must stay entirely in the face F.*

Proof. Acyclicity of $G(P,\mathbf{c})$ implies the nonrevisiting property for 0-dimensional faces. The fact that $G(P,\mathbf{c})$ is a Hasse diagram by definition implies the nonrevisiting property for 1-dimensional faces. Suppose there is a 2-dimensional face F in P and a directed path in $G(P,\mathbf{c})$ that departs F at $u \in F$ and re-enters F at $v \in F$. By Lemma 4.1, there cannot also be a directed path from u to v that stays entirely in F, since F is a codimension one face of P. Let a_1 and a_2 be the vertices of F that cover the source of F, chosen so that there is a directed path in F from a_1 to u. We also have a directed path in $G(P,\mathbf{c})$ from u to v; the fact that there is no directed path from u to v within F implies in this case that there is a directed path within F from a_2 to v. The existence of these three directed paths implies that v is an upper bound for a_1 and a_2. But we have just proven in Theorem 4.6 that the pseudo-join of a_1 and a_2 equals the join of a_1 and a_2. Thus, we deduce that the sink of F, namely the pseudo-join of a_1 and a_2, is less than or equal to v in L. This implies v is the sink of F, since $v \in F$. But this contradicts the fact that there is no directed path from u to v in F. Thus, we have a contradiction to the existence of a 2-face F and a path that starts and ends in F but does not stay entirely in F. □

Next we turn to the class of polytopes producing all known counterexamples to the Hirsch Conjecture, namely spindles. See Definition 2.3 for a review of the notion of a spindle.

Theorem 5.2. *Let P be a simple d-polytope with n facets. Suppose that P is a spindle with vertices u and v such that each facet of P contains either u or v. Let $\mathbf{c}$ be a generic cost vector such that $G(P, \mathbf{c})$ is the Hasse diagram of a poset having $\hat{0} = u$ and $\hat{1} = v$. Then $G(P, \mathbf{c})$ satisfies the face nonrevisiting property, implying that every directed path from u to v has at most $n - d$ edges. This gives an upper bound of $n - d$ on the distance from u to v.*

Proof. The definition of spindle ensures that each facet F of P includes either $\hat{0}$ or $\hat{1}$. But then Corollary 4.3 implies that there are no directed paths that depart from F and later revisit F for faces F that include $\hat{0}$ or $\hat{1}$. Since every facet in the spindle includes either u or v as a vertex, every facet has this nonrevisiting property. But every directed path from u to v departs a facet at each step. Since there are only n facets, and v is incident to d facets, there are at most $n-d$ facets that may be departed, hence at most $n-d$ steps in any directed path from u to v. This implies that the distance from u to v cannot be greater than $n - d$. □

This result has as a consequence that all of the known counterexamples to the Hirsch Conjecture (to date) fail to meet the hypotheses for Theorem 5.2:

Corollary 5.3. *Given any simple d-polytope with n facets that is a spindle with vertices u and v such that every facet includes either u or v, if the distance from u to v is greater than $n - d$, then there does not exist any generic cost vector $\mathbf{c}$ such that $G(P, \mathbf{c})$ is the Hasse diagram of a poset with u as source and v as sink.*

5.2. **Simple polytopes with Hasse diagrams of well-known lattices arising as their 1-skeleta.** We now turn to three well-known families of lattices, namely the weak order, the Tamari lattices, and the Cambrian lattices. Their Hasse diagrams will arise as 1-skeleta of the permutahedra, the associahedra, and the generalized associahedra, respectively.

Example 5.4. The permutahedron P_n is a simple polytope yielding weak order as follows. Let $(x_1, \ldots, x_n)$ be a point in $\mathbb{R}^n$ with distinct coordinates, most typically chosen with $x_i = i$ for $i = 1, \ldots, n$. Recall that

$$P_n(x_1, \ldots, x_n) = \mathrm{conv}\{(x_{\pi(1)}, \ldots, x_{\pi(n)}) | \pi \in S_n\}$$

is the canonical V-representation for P_n. Two of its vertices $x_u = (x_{u(1)}, \ldots, x_{u(n)})$ and $x_v = (x_{v(1)}, \ldots, x_{v(n)})$ for $u, v, \in S_n$ are connected by an edge if and only if $v = us_i$ for some adjacent transposition $s_i = (i, i+1)$ acting on values. If starting from $x_e = (x_1, \ldots, x_n)$, corresponding to the identity element e of S_n we orient the edges of $P_n(x_1, \ldots, x_n)$ from shorter towards longer permutations, then we obtain the weak order. Thus a cover relation $u \prec v$, for $v = us_i$ in weak order means that we introduced a descent involving values i and $i+1$ which were in positions k and l, $k < l$, in u, respectively. Then, taking the linear functional $\mathbf{c} = (c_1, \ldots, c_n) \in \mathbb{R}^n$ to be one with strictly descending coordinates $c_1 > c_2 > \cdots > c_n$ we obtain that $\mathbf{c} \cdot x_v - \mathbf{c} \cdot x_u = c_k(i+1) + c_l i - c_k i - c_l(i+1) = c_k - c_l > 0$. This verifies for each cover relation $u \prec v$ in weak order that $\mathbf{c} \cdot x_u < \mathbf{c} \cdot x_v$. See also Example 3.3 in [2].

Example 5.5. The associahedron A_n is another example of a simple polytope with a generic cost vector $\mathbf{c}$ yielding $G(A_n, \mathbf{c})$ as the Hasse diagram of a well known poset, namely the Tamari lattice. Consider the presentation for the associahedron introduced by Loday in [28]. The vertices of the associahedron are indexed by the unlabeled, rooted planar, binary trees with n leaves and $n-1$ internal nodes (i.e. non-leaf vertices). We associate to each such tree

t the polytope vertex $M(t) \in \mathbb{R}^{n-1}$ defined as follows. $M(t) = (a_1b_1, \dots, a_ib_i, \dots, a_{n-1}b_{n-1})$ where a_i is the number of leaves that are left descendants of the i-th internal node v_i of the tree t and b_i is the number of leaves that are right descendants of v_i within the tree t. One may check, for example, that the associahedron given by trees with 4 leaves has vertices $(3,2,1), (3,1,2), (1,4,1), (2,1,3)$, and $(1,2,3)$.

Turning now to the Tamari lattice, a cover relation $u \prec v$ in the Tamari lattice results from applying a single associativity relation in our rooted, binary, planar tree regarded as a parenthesization. Thus, v is obtained from u by replacing $((x,y),z))$ by $(x,(y,z))$ somewhere in the parenthesized expression, with the objects x, y, z either being individual letters or being larger bracketed expressions themselves. Notice that such an operation will have the impact within Loday's realization of the associahedron of replacing some pair (a_i, b_i) by $(a_i, b_i + b_{i+r})$ and replacing (a_{i+r}, b_{i+r}) by $(a_{i+r} - a_i, b_{i+r})$ while leaving all other a_j, b_j unchanged. Thus, $M(t)$ is unchanged except for having the coordinate a_ib_i replaced by $a_ib_i + a_ib_{i+r}$ and $a_{i+r}b_{i+r}$ replaced by $a_{i+r}b_{i+r} - a_ib_{i+r}$. We may again use any cost vector $\mathbf{c} = (c_1, \dots, c_n) \in \mathbb{R}^n$ with strictly descending coordinates $c_1 > c_2 > \cdots > c_n$ to deduce that $u \prec v$ implies $\mathbf{c} \cdot u < \mathbf{c} \cdot v$. See [9], [22] for further background on the Tamari lattice.

Theorem 5.6. *Each open interval (u,v) in the weak order has order complex which is homotopy equivalent to a ball or a sphere of some dimension.*

Proof. We obtain the Hasse diagram for weak order as the 1-skeleton of the permutahedron, which is a simple polytope, using a cost vector as in Example 5.4. A proof that the weak order is a lattice may be found in [6]. Thus, Theorem 4.11 applies. □

The homotopy type of the intervals in weak order was previously determined in [14], [15], and subsequently by a different method in [5].

Theorem 5.7. *Each open interval (u,v) in the Tamari lattice has order complex homotopy equivalent to a ball or a sphere of some dimension.*

Proof. The Tamari lattice has as its Hasse diagram the 1-skeleton of the associahedron with respect to any cost vector as in Example 5.5. A proof that the Tamari lattice is a lattice appears in [31]. Thus, the Tamari lattice meets all the conditions of Theorem 4.11. □

The homotopy type of the intervals in the Tamari lattice was previously determined by Björner and Wachs in [9], where they note that this result also essentially follows from work of Pallo in [32].

Next we combine several results from the literature in a manner suggested to us by Nathan Reading to deduce the following result which generalizes Theorem 5.7.

Theorem 5.8. *Each open interval (u,v) in any c-Cambrian lattice has order complex that is homotopy equivalent to a ball or a sphere of some dimension.*

Proof. See Proposition 3.1 in [20] for the fact that the Hasse diagram of the c-Cambrian lattice is obtained from the polytope $Asso_c^a(W)$ known as a generalized associahedron given by W (as defined e.g. in [20]) by choosing a suitable cost vector $\mathbf{c}$ and taking the directed graph that $\mathbf{c}$ induces on the 1-skeleton of the polytope. Just before Example 3.5 in [20], it is asserted that all of these polytopes are simple. This is proven as Theorem 3.4 in [21]. Thus, Theorem 4.11 applies in the case of all c-Cambrian lattices. □

Thus, we recover Reading's results on the homotopy type of intervals in the c-Cambrian lattice, thereby showing that all generalized associahedra can also be handled by our approach.

5.3. **The case of zonotopes.** Now we turn to another large class of examples of polytopes to which our results will apply, namely all simple polytopes which are zonotopes. Björner already determined the homotopy type of all open intervals for zonotopes in [5], but we nonetheless include this discussion so as to show how another large and important class of polytopes fits into our framework.

Proposition 5.9. *Any zonotope P and any generic linear functional $\mathbf{c}$ together will satisfy the nonrevisiting property, and hence the Hasse diagram property. Thus, $G(P,\mathbf{c})$ has directed diameter at most $n-d$ for n the number of facets in P and d the dimension of P.*

Proof. Any zonotope is a Minkowski sum of line segments. Departing a face while increasing the dot product with the cost vector $\mathbf{c}$ means traversing an edge in the direction of one of these line segments generating the zonotope. But we can never traverse an edge going in exactly the opposite direction to this while still increasing the dot product. By virtue of a zonotope being a Minkowski sum of line segments, it is not possible to return to the face without at some point traversing a parallel edge in the opposite direction. The proof of Proposition 5.10 will give another way of seeing why this nonrevisiting property holds.

For the last claim, simply observe that each edge in a directed path departs from a facet that may never be revisited and that the final vertex in a directed path will still belong to d facets. Thus, there can be at most $n-d$ steps since there are at most $n-d$ facets available to be departed at some stage in the directed path. □

Proposition 5.10. *If P is a simple polytope that is a zonotope and $\mathbf{c}$ is generic, then $G(P,\mathbf{c})$ is the Hasse diagram of a lattice.*

Proof. We deduce this fact by combining assorted known results, as explained next. It is well known that every zonotope may be obtained from a central hyperplane arrangement as follows. Any central hyperplane arrangement induces a subdivision of a unit sphere centered at the origin. If the arrangement is not essential, then restrict this sphere to a subspace through the origin of as high dimension as possible such that the arrangement restricted to that subspace is essential. Let the vertices of the resulting subdivision of the sphere be the vertices of a polytope. Taking the dual polytope to this, the result is a zonotope, and in fact every zonotope may be realized this way. The point is to make the hyperplanes perpendicular to the line segments comprising the Minkowski sum of line segments. See e.g. [43] or [19]. From this perspective, an edge of a zonotope departs a face by crossing one of these hyperplanes, namely one that is perpendicular to the direction of the edge being traversed. We can never revisit the face we just left without crossing the hyperplane in the opposite direction. But this would mean traversing an edge of the polytope the opposite direction to the edge we used to depart the face, contradicting $G(P,\mathbf{c})$ being induced by a cost vector $\mathbf{c}$. Thus, for P a simple zonotope and $\mathbf{c}$ generic, this implies that $G(P,\mathbf{c})$ must satisfy the nonrevisiting property (so in particular must be a Hasse diagram).

It is proven in [7] that the poset of regions given by a central, simplicial hyperplane arrangement is a lattice. Given a simple zonotope P and generic cost vector $\mathbf{c}$, the poset having $G(P,\mathbf{c})$ as its Hasse diagram is exactly the poset of regions of a central, simplicial hyperplane arrangement, hence is always a lattice. □

Theorem 5.11. *Whenever a zonotope is a simple polytope, then the poset given by $G(P, \mathbf{c})$ has each open interval homotopy equivalent to a ball or a sphere.*

Proof. The first thing to note is that the poset will always be a lattice in this case, by Proposition 5.10. The Hasse diagram property is proven for all zonotopes in Proposition 5.9. Thus, Theorem 4.11 applies to all simple zonotopes. □

5.4. **More general facial orientations of simple regular CW spheres.** Vic Reiner raised the question (personal communication) of whether we could use Proposition 5.3 from [1], a result that is recalled as Proposition 5.12 below, to generalize our results. Specifically, he suggested generalizing from our framework of acyclic orientations on 1-skeleta of simple polytopes given by cost vectors to more general acyclic orientations known as facial orientations; Reiner also suggested trying to prove results for a somewhat larger class of regular CW spheres than just the simple polytopes discussed so far.

Recall that a **facial orientation** of the 1-skeleton of a regular CW complex K is an orientation $\mathcal{O}$ of the 1-skeleton graph of K such that for each cell $\sigma \in K$, the restriction of $\mathcal{O}$ to the closure of σ has a unique source and a unique sink. It is well known that a shelling of a simplicial polytope is equivalent to a facial orientation of its dual polytope; the special case of line shellings is also well-known to yield precisely those facial orientations which are induced by cost vectors. One may easily construct examples demonstrating that not all facial orientations can be induced by cost vectors.

Proposition 5.12 (Proposition 5.3 of [1])**.** *Let X be a shellable regular CW sphere with P its face poset. There is a dual regular CW sphere, denoted X^*, with face poset P^*. Letting $G(P^*)$ denote the graph arising as the 1-skeleton of X^*, then the acyclic orientation $\mathcal{O}$ of $G(P^*)$ induced by any shelling order of X is a facial orientation on the graph of X^*.*

Our techniques do yield the following partial answer to Reiner's question. On the other hand, Example 5.15 in conjunction with Remark 5.14 constrains the extent to which a positive answer to Reiner's question is possible.

Theorem 5.13. *Let P be a simple polytope, and let $\mathcal{O}$ be a facial acyclic orientation on its 1-skeleton. Suppose that the directed graph on the 1-skeleton of P induced by $\mathcal{O}$ is the Hasse diagram of a lattice L. Then L has the following properties:*

1. *The pseudo-join of any collection $\{a_1, a_2, \dots, a_i\}$ of elements of L all covering a common element u will equal the join $a_1 \vee a_2 \vee \cdots \vee a_i$ of these same elements.*
2. *For S, T distinct collections of elements all covering a fixed element u, then the pseudo-join of the elements of S will not equal the pseudo-join of the elements of T.*
3. *Each open interval in L has order complex homotopy equivalent to a ball or a sphere.*

Proof. Statements (1), (2) and (3), respectively, were already proven for $\mathcal{O}$ induced by a cost vector in Theorem 4.7, Lemma 4.9 and Theorem 4.11, respectively. The same proofs still hold entirely unchanged for more general facial orientations. Checking this is left as a completely straightforward exercise for the interested reader. □

Remark 5.14. Our proof of Theorem 4.6 relies in an essential way on the property of polytopes that two distinct 2-dimensional faces cannot share both an edge and a vertex not in that edge. Our proofs of Theorem 4.7, Lemma 4.9 and Theorem 4.11 all rely upon

Theorem 4.6. This property of polytopes used in the proof of Theorem 4.6 does not hold for regular CW spheres in general, even with the further assumption that the regular CW sphere is simple. Example 5.15 exhibits this non-implication, showing that Theorem 5.13 cannot be extended from polytopes to regular CW spheres.

Example 5.15. Now we will construct a simple regular CW sphere with two 2-cells sharing an edge and also sharing a vertex that is disjoint from that edge. To this end, we give a regular CW decomposition of the boundary of a cylinder as follows. Begin by placing four vertices denoted v_1, v_2, v_3, v_4 clockwise about the boundary of the upper disk comprising the top of the cylinder. Now likewise put four vertices denoted w_1, w_2, w_3, w_4 clockwise about the boundary of the bottom disk comprising the bottom of the cylinder. Introduce edges e_{v_i,v_j} for each $i \neq j$ other than the pair $i = 1, j = 3$. Likewise introduce edges e_{w_i,w_j} for each $i \neq j$ other than the pair $i = 1, j = 3$. Also introduce edges e_{v_1,w_1} and e_{v_3,w_3}. The resulting subdivision of this 2-sphere, namely of the boundary of a cylinder, will also have the following six 2-cells. There are 2-cells with vertices $\{v_1, v_2, v_4\}$ and with $\{v_3, v_2, v_4\}$ covering the top disk, 2-cells with vertices $\{w_1, w_2, w_4\}$ and with $\{w_3, w_2, w_4\}$ covering the bottom disk, and 2-cells F and F' with vertex sets $\{v_1, v_2, v_3, w_1, w_2, w_3\}$ and with $\{v_1, v_4, v_3, w_1, w_4, w_3\}$ covering the remainder of the boundary of the cylinder. The 2-cells F and F' share the edge e_{v_1,w_1} and also the edge e_{v_3,w_3}. Thus, F and F' share four vertices with two of these four vertices comprising an edge.

6. Further questions and remarks

Remark 6.1. In [17], Curtis Greene raised the question of finding interesting classes of posets with each open interval having Möbius function $0, 1$, or -1. Theorems 4.11 and 5.13 speak to that question by giving large classes of such posets.

Remark 6.2. We refer readers to [25] for interesting, related work that takes a somewhat similar perspective to ours, work which provided some inspiration for parts of our work. In [25], Kalai proved that the combinatorial type of a simple polytope is determined by its 1-skeleton, also making use of a cost vector in this construction as well as utilizing the consequent sources and sinks of the various faces.

One might be tempted, in light of our results, to ask the following question:

Question 6.3. Let P be a simple polytope and let $\mathbf{c}$ be a generic cost vector such that $G(P, \mathbf{c})$ is the Hasse diagram for a poset. Does this imply that this poset is a lattice?

An affirmative answer would have allowed our hypotheses throughout much of this paper to be relaxed from lattice to poset. However, Francisco Santos has provided the following example, showing that the answer to Question 6.3 is negative in general.

Example 6.4 (Francisco Santos)**.** Start with an octahedron P with two antipodal vertices as source and sink, leaving four intermediate vertices v_1, v_2, v_3, v_4 connected with each other with the structure of a 4-cycle. Put two opposite vertices v_1, v_3 among these four vertices at a higher height than the other two, namely with $\mathbf{c} \cdot v_i > \mathbf{c} \cdot v_j$ for each $i \in \{1, 3\}$ and each $j \in \{2, 4\}$. Now truncate each of the six vertices by slicing by a generic hyperplane with slope chosen so as to make this a simple polytope with the Hasse diagram property (with each of the original vertices replaced by four new vertices). This will yield a simple polytope with

$G(P, \mathbf{c})$ a Hasse diagram for a poset that is not a lattice, since there will be a pair of vertices having two different least upper bounds; specifically, we may use one of the four vertices replacing v_2 together with one of the four vertices replacing v_4. We may choose such vertices so that we get one least upper bound in the quadrilateral replacing v_1 and another in the quadrilateral replacing v_3.

One may also construct a polytope P and a cost vector $\mathbf{c}$ such that $G(P, \mathbf{c})$ is the Hasse diagram of a poset with the pseudo-join of some collection of atoms which is not equal to the least upper bound of this same collection of atoms, as shown next.

Example 6.5. Start with a 3-dimensional cube and add a new vertex by coning over one of the facets of the cube that contains the vertex of the cube where the cost vector was maximized, positioning this new vertex so that it becomes the pseudo-join of all the atoms. This can be done by letting $\mathbf{c} = (100, 2, 1)$, letting the vertices of the cube be $(\pm 1, \pm 1, \pm 1)$ and taking as the cone point over a facet of this cube the vertex $(2, 0, 0)$. To make the 1-skeleton of the polytope obtained this way a Hasse diagram, we cut off the vertex $(2, 0, 0)$ of the cone with a hyperplane near this vertex with a slope for this slicing hyperplane chosen in such a way that one of the resulting four new vertices (replacing $(2, 0, 0)$) becomes the pseudo-join of all the atoms, while the vertex that was the least upper bound of the atoms in the original cube still remains the least upper bound of all the atoms.

This is not a simple polytope, but one may transform this into a simple polytope by shaving by a hyperplane at each node of degree higher than 3. However, that shaving operation will change which element is the join of the set of three atoms in such a way that indeed the join of the three atoms is the pseudo-join of the same set of three atoms, transforming this into a positive example of our result that joins equal pseudo-joins for simple polytopes.

To make our results more effective on naturally arising examples, it could also help to answer the following question:

Question 6.6. Is there a good way to recognize when $G(P, \mathbf{c})$ will be the Hasse diagram of a poset? Is there an effective way to determine when this poset will be a lattice?

Regarding the first part part of Question 6.6, Louis Billera has suggested considering the directed adjacency matrix A where he observed that the Hasse diagram property would imply that the trace of $A^T \cdot A^i$ would need to be 0 for each $i \geq 2$, letting A^T denote the transpose of A. From the standpoint of algorithmic efficiency, this requires an $n \times n$ matrix where n is the number of vertices of $G(P, \mathbf{c})$, which may be much larger than either the dimension or the number of facets in the polytope.

Remark 6.7. It may seem natural now to ask whether a simple polytope P together with a generic cost vector $\mathbf{c}$ such that $G(P, \mathbf{c})$ is the Hasse diagram of a poset will always satisfy the directed graph version of the nonrevisiting path conjecture. An obvious place to start is to ask whether any of the known counterexamples to the Hirsch Conjecture give rise to directed graphs $G(P, \mathbf{c})$ that meet the hypotheses of Conjecture 1 or at least are Hasse diagrams of posets, since these polytopes are all known to be counterexamples to the undirected version of the nonrevisiting path conjecture. Lemma 4.2 implied that there are no counterexamples of this type to Conjecture 1.

The original construction of Francisco Santos in [39] of a polytope violating the Hirsch Conjecture was a spindle (see Definition 2.3) but was not a simple polytope. His presentation for this polytope is essentially as an H-polytope, in that he gives the vertices of its dual polytope (from which the bounding hyperplanes of the original polytope may easily be deduced). Santos remarks on p. 389 in [39] that determining the vertices of this polytope seems computationally out of reach, which we note also makes determining the undirected graph of the 1-skeleton elusive. In particular, this makes the directed graph, $G(P, \mathbf{c})$ for any particular choice of $\mathbf{c}$, also computationally out of reach. A second type of computational challenge to thoroughly examining these examples would be the need to consider all possible generic cost vectors $\mathbf{c}$. There are exponentially many orientations on the 1-skeleton graph to consider (as a function of the number of graph edges), though in principle one would only need to consider the "good orientations" in the sense of Kalai from [25]. Thus, there are multiple substantial challenges to understanding this example in full, but in any case it will not yield a counterexample to Conjecture 1.

The later smaller counterexamples of Matschke, Santos, and Weibel to the Hirsch Conjecture appearing in [29] are simple polytopes that are spindles. All of these known examples of polytopes violating the Hirsch Conjecture result from d-polytopes which are spindles with n facets having the property that the known pair of vertices at distance greater than $n - d$ from each other are the two distinguished vertices in the spindle. Our Theorem 5.2 shows for these examples where P is a simple spindle that $G(P, \mathbf{c})$ is not a Hasse diagram. Thus, Theorem 5.2 shows that these constructions violating the Hirsch Conjecture do not also serve as counterexamples to our Conjecture 1.

Remark 6.8. In seeking more examples of polytopes fitting into our framework, one might be tempted to consider fiber polytopes (introduced in [3]); after all, the permutahedron and associahedron are both fiber polytopes and more specifically are monotone path polytopes, and both the permutahedron and associahedron do fit into our framework. However, every polytope P may be realized as a monotone path polytope as follows. Take the join of P with a point p. Project the resulting polytope P' to the real line by a linear map π in such a way that the fiber $\pi^{-1}(t)$ over each point t on the real line is either empty, the single point p, or has the combinatorial type of P. One may check that the fiber polytope resulting from the map $\pi : P' \to \mathbb{R}$ has the combinatorial type of P. Thus, monotone path polytopes are too general a class of polytopes to hope for our results to apply to all of them.

Focusing on generalized permutahedra (see [33]) still gives a class that will not always satisfy all of our hypotheses. To see this, note that generalized permutahedra sometimes have triangular faces, forcing the Hasse diagram property to fail for all choices of cost vector.

We conclude with one of the most natural and potentially important questions that stems from our work:

Question 6.9. Do any of the important classes of polytopes coming from real-world problems studied in operations research fit into our framework? In other words, are they simple polytopes whose 1-skeleta (with respect to a choice of cost vector) are Hasse diagrams of lattices? One can either ask this for a fixed cost vector, or preferably for all generic cost vectors for the given simple polytope.

Given an affirmative answer for some such class of polytopes, Conjecture 1 would assert that the simplex method for linear programming applied to this class of polytopes would require at most $n-d$ steps regardless of choice of pivot rule. It is plausible that not all of the hypotheses in Conjecture 1 are necessary, though they all seem like they would be helpful for proving Conjecture 1; if a version of Conjecture 1 with fewer hypotheses could be proven, this would of course increase the potential for an affirmative answer to Question 6.9.

We have ruled out some classes of real-world polytopes as potential candidates for Question 6.9. Specifically, we have found small examples of transportation polytopes, 0/1-polytopes, vertex decomposable polytopes, and traveling salesman polytopes that do not meet all of the hypotheses for Conjecture 1 in its present form by virtue of having triangular faces. However, many more classes arising in operations research remain.

An affirmative answer to Question 6.9 could contribute a new perspective for particular classes of polytopes to the understanding of the widely observed phenomenon that the simplex method typically is much more efficient in practice in real-world applications than is predicted by worst case analysis. The simplex method is already well known to have polynomial time average case complexity in many cases; Spielman and Teng proved the much stronger result that the simplex method has so-called polynomial smoothed complexity (see [40]), i.e. has polynomial complexity in cases that interpolate between average case and worst case.

One way to view much of our work in this paper is as an attempt to better understand conceptually what properties of a polytope P and cost vector $\mathbf{c}$ will suffice to ensure that directed paths in $G(P,\mathbf{c})$ may never revisit any faces that they have left, thereby also ensuring for such P and $\mathbf{c}$ that the simplex method is efficient even in the worst case, i.e. no matter what choices of pivots are made. Others such as in [10] also have recently examined (and made headway) on the closely related question of proving that particular classes of polytopes and cost vectors arising naturally in operations research will make linear programming efficient for all choices of pivot rule when certain parameters are held fixed. There is certainly much more to be done in this area, including the question of resolving our main conjecture which in spite of the result described in the upcoming appendix does still remain an open question.

7. Appendix by Dominik Preuss: Monotone Hirsch conjecture for simple polytopes whose 1-skeleta are Hasse diagrams of lattices

Here we prove that if $G(P,\mathbf{c})$ is the Hasse diagram of a lattice, any directed path in $G(P,\mathbf{c})$ has length at most $n-d$. In particular, P satisfies the monotone Hirsch conjecture.

To that end, recall from lattice theory that an element b of a lattice L is called *join-irreducible* iff $b=\bigvee_{a\in S} a$ for a set $S\subseteq L$ implies $b\in S$. It makes sense to define the join of the empty set as the unique minimal element of the lattice, hence that element is not join-irreducible by definition. We shall need the following property of join-irreducibles:

Lemma 7.1 ([41], Chapter 3.4)**.** *An element of a finite lattice is join-irreducible iff it covers precisely one element.*

As an immediate consequence of Lemma 1, we have:

Lemma 7.2. *Let P be a simple polytope and $\mathbf{c}$ a generic cost vector such that $G(P,\mathbf{c})$ is the Hasse diagram of a lattice. Then a vertex v is join-irreducible iff it is distinct from $\hat{0}$* and *is the unique source of a facet.*

Proof. If v is not equal to $\hat{0}$, then there must be a directed path from $\hat{0}$ to v, since $\hat{0}$ is a source of the whole graph. Hence there must be an edge which ends in v, which means that v covers at least one other lattice element. On the other hand, if v is the source of a facet F, then it must be the initial vertex of every incident edge that belongs to F. There are $d-1$ such edges, where d is the dimension of P (since P is simple). Since v is incident to a total of d edges (again, since P is simple), only one more edge remains that can end in v. Since there is at least one such edge, there is *precisely* one such edge. Hence v covers precisely one lattice element, and therefore, it is join-irreducible by Lemma 7.1.

Conversely, suppose v is join-irreducible. Then it covers precisely one lattice element/vertex u. Therefore, except for the edge $\{u, v\}$, all edges incident to v point away from v. Since P is simple, there is a facet F for every set X of $d-1$ edges incident to v such that X are precisely those edges incident to v that lie in F. Therefore, there is a facet in which all edges point away from v, implying that v is the source of that facet. □

For the proof of our main theorem, we need one further purely lattice-theoretic result:

Lemma 7.3. *If b, c are elements of a finite lattice L such that b covers c, then there is a join-irreducible $j \in L$ such that $c \vee j = b$ (where j is not necessarily distinct from b).*

Proof. Since every element of a finite lattice is a join of join-irreducibles, we may choose a minimal set S of join-irreducibles such that $b = c \vee \bigvee_{a \in S} a$. If S contains only one join-irreducible, we are done. Otherwise, choose $j \in S$ and consider the element $c \vee \bigvee_{a \in S \setminus \{j\}} a$. This must be $\geq c$ and $\leq b$, which means that it is either equal to b or to c, since b covers c. If it were equal to b, then this would contradict the minimality of S. If it were equal to c, then $c \vee j = c \vee (\bigvee_{a \in S \setminus \{j\}} a) \vee j = c \vee \bigvee_{a \in S} a = b$, and so $\{j\}$ would also be a proper subset of S whose join with c is b, again contradicting the minimality of S. □

Our main theorem now is:

Theorem 7.4. *Let P be a simple polytope and $\mathbf{c}$ a generic cost vector such that $G(P, \mathbf{c})$ is the Hasse diagram of a lattice. Then every directed path in $G(P, \mathbf{c})$ has length at most $n - d$, where n is the number of facets and d the dimension.*

Proof. Suppose we are given a directed path consisting of vertices $v_0, \ldots, v_m$. Then by Lemma 7.3, there is a join-irreducible j_i for every $i > 0$ such that $v_i = v_{i-1} \vee j_i$. In other words, $v_i = v_0 \vee \bigvee_{k=1}^{i} j_k$. It follows that all the j_k must be distinct, since $v_i = v_0 \vee \cdots \vee j_k \vee \cdots \vee j_{i-1} \vee j_i$ and $j_k = j_i$ imply together with the idempotence of the join operation that

$$\begin{aligned} v_i &= v_0 \vee \cdots \vee j_k \vee \cdots \vee j_{i-1} \vee j_i \\ &= v_0 \vee \cdots \vee j_k \vee \cdots \vee j_{i-1} \vee j_k \\ &= v_0 \vee \cdots \vee j_k \vee \cdots \vee j_{i-1} = v_{i-1} \end{aligned}$$

Therefore the length of a path is at most the number of join-irreducibles. But by Lemma 7.2, there is precisely one join-irreducible for every facet *except* the d facets whose source is $\hat{0}$, so the number of join-irreducibles is $n - d$. □

References

[1] C. Athanasiadis, P. Edelman and V. Reiner, Monotone paths on polytopes, *Math. Zeitschrift* **235** (2000), 315–334.

[2] L. Billera, M. Kapranov and B. Sturmfels, Cellular strings on polytopes, *Proc. Amer. Math. Soc.* **122** (1994), 549–555.
[3] L. Billera and B. Sturmfels, Fiber polytopes, *Annals of Mathematics* **135** (1992), 527–549.
[4] A. Björner, Orderings of Coxeter groups. Combinatorics and algebra (Boulder, Colo., 1983), 175–195, *Contemp. Math.*, **34**, Amer. Math. Soc., Providence, RI, 1984.
[5] A. Björner, Essential chains and homotopy, *Proc. Amer. Math. Soc.* **116** No. 4 (1992), 1179–1181.
[6] A. Björner and F. Brenti, Combinatorics of Coxeter groups, *Graduate Texts in Mathematics*, **231**, Springer, New York, 2005. xiv + 363 pp.
[7] A. Björner, P. Edelman, and G. Ziegler, Hyperplane arrangements with a lattice of regions, *Disc. and Comput. Geom.* **5** (1990), 263–288.
[8] A. Björner and M. Wachs, On lexicographically shellable posets, *Trans. Amer. Math. Soc.*, **277** (1983), no 1, 323–341.
[9] A. Björner and M. Wachs, Shellable and nonpure complexes and posets, II, *Trans. Amer. Math. Soc.* **349** (1997), 3945–3975.
[10] M. Blanchard, J.A. De Loera and Q. Louveaux, On the length of monotone paths in polyhedra, arXiv:2001.09575 preprint 2020.
[11] F. Chapoton, S. Fomin, and A. Zelevinsky, Polytopal realizations of the generalized associahedra, *Canadian Mathematical Bulletin* **45** (2002), 537–566.
[12] W. Cook, In pursuit of the Traveling Salesman: Mathematics at the Limits of Computation, Princeton University Press, 2011.
[13] J.A. De Loera and S. Klee, Transportation problems and simplicial polytopes that are not weakly vertex-decomposable, *Mathematics of Operations Research* **37** (2012), 670–674.
[14] P. Edelman, A partial order on the regions of $\mathbb{R}^n$ dissected by hyperplanes, *Trans. Amer. Math. Soc.* **283** (1984), 617–631.
[15] P. Edelman and J. Walker, The homotopy type of hyperplane posets, *Proc. Amer. Math. Soc.* **94** (1985), 221–225.
[16] B. Gärtner, M. Henk, and G. Ziegler, Randomized simplex algorithms on Klee-Minty cubes, *Combinatorica* **18** (3) (1998), 349-372.
[17] C. Greene, A class of lattices with Möbius function $\pm 1, 0$, *Europ. J. Combinatorics* **9** (1988), 225–240.
[18] P. Hersh and K. Mészáros, SB-labelings and posets with each interval homotopy equivalent to a sphere or a ball, *J. Combinatorial Theory, Ser. A* **152** (2017), 104–120.
[19] P. Hersh and E. Swartz, Coloring complexes and arrangements, *J. Algebraic Combinatorics* **27** (2008), no. 2, 205–214.
[20] C. Hohlweg, Permutahedra and associahedra: generalized associahedra from the geometry of finite reflection groups, 129–159, *Prog. Math. Phys.*, **299**, Birkhäuser/Springer, Basel, 2012.
[21] C. Hohlweg, C. Lange, H. Thomas, Permutahedra and generalized associahedra, *Advances in Mathematics* **226** (2011), no. 1, 608–640.
[22] S. Huang and D. Tamari, Problems of associativity: A simple proof for the lattice property of systems ordered by a semi-associative law, *J. Combin. Theory Ser. A* **13** (1972), 7–13.
[23] V. Kaibel and M. Pfetsch, Computing the face lattice of a polytope from its vertex-facet incidences, *Comput. Geom.* **23** (2002), 281–290.
[24] V. Kaibel and M. Pfetsch, Some algorithmic problems in polytope theory, *Algebra, geometry, and its software systems*, 23–47, Springer, Berlin, 2003.
[25] G. Kalai, A simple way to tell a simple polytope from its graph, *J. Combinatorial Theory, Ser. A* **49** (1988), 381–383.
[26] V. Klee and G. Minty, How good is the simplex algorithm?, in: *Inequalities III* (O. Sisha ed.), Academic Press, New York, 1972, 159–175.
[27] V. Klee and D. Walkup, The d-step conjecture for polyhedra of dimension $d < 6$, *Acta Math.* **133** (1967), 53–78.
[28] J.-L. Loday, Realization of the Stasheff polytope, *Arch. Math.* **83** (2004), 267–278.
[29] B. Matschke, F. Santos, and C. Weibel, The width of a 5-dimensional prismatoid, *Proc. London Math. Soc.* **110** (2015, 647–672.

[30] P. McMullen, The maximum numbers of faces of a convex polytope, *Mathematika* **17** (1970), 179–184.
[31] J. M. Pallo, Enumerating, ranking and unranking binary trees, *Computer J.* **29** (1986), 171–175.
[32] J. M. Pallo, An algorithm to compute the Möbius function of the rotation lattice of binary trees, *Informatique Theor. et. Applic./Theor. Informatics and Applic.* **27** (1993), 341–348.
[33] A. Postnikov, Permutahedra, associahedra, and beyond, *IMRN* **2009** (2009), 1026–1106.
[34] D. Quillen, Homotopy properties of the poset of non-trivial p-subgroups of a group, *Advances in Mathematics*, **28** (1978), 101–128.
[35] N. Reading, Lattices congruences of the weak order, *Order*, **21** (2004) no. 4, 315–344.
[36] N. Reading, Cambrian lattices, *Adv. Math.* **205** (2006), no. 2, 313–353.
[37] N. Reading, Lattice congruences, fans, and Hopf algebra, *J. Combin. Theory, Ser. A* **110** (2005), no. 2, 237–273.
[38] V. Reiner, The generalized Baues problem, chapter in: New Perspectives in Algebraic Combinatorics, *Mathematical Sciences Research Institute Publications* **38**, 293–336.
[39] F. Santos, A counterexample to the Hirsch conjecture, *Annals of Mathematics (2)* **176** (2012), 383–412.
[40] D. Spielman and Teng, Smoothed analysis of algorithms: why the simplex algorith usually takes polynomial time, *J. ACM* **51** (2004), 385–463.
[41] R. P. Stanley, Enumerative Combinatorics, Volume 1: Second Edition. Cambridge Studies in Advanced Mathematics, 49. Cambridge University Press, New York, 2012.
[42] M. Todd, The monotone bounded Hirsch conjecture is false for dimension at least 4, *Math. Operations Research* **5** (1980), 599–601.
[43] G. Ziegler, Lectures on Polytopes, Graduate Texts in Mathematics **152**, Berlin, New York, Springer, 1995.

Department of Mathematics, University of Oregon, Eugene, OR 97403
E-mail address: `plhersh@uoregon.edu`