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 $G(P,\mathbf{c})$ by taking the 1-skeleton of P and orienting each edge $e(u,v)$ from u to v for $\mathbf{c}\cdot u < \mathbf{c}\cdot v$. This paper raises the question of finding sufficient conditions on a polytope P and generic cost vector $\mathbf{c}$ so that the graph $G(P,\mathbf{c})$ will not have any directed paths which revisit any face of P after departing from that face. This is in a sense equivalent to the question of finding conditions on P and $\mathbf{c}$ under which the simplex method for linear programming will be efficient under all choices of pivot rules. Conditions on P and $\mathbf{c}$ are given which provably yield a corollary of the desired face nonrevisiting property and which are conjectured to give the desired property itself. This conjecture is proven for 3-polytopes and for spindles having the two distinguished vertices as source and sink; this shows that known counterexamples to the Hirsch Conjecture will not provide counterexamples to this conjecture.

A part of the proposed set of conditions is that $G(P,\mathbf{c})$ be the Hasse diagram of a partially ordered set, which is equivalent to requiring nonrevisiting of 1-dimensional faces. This opens the door to the usage of poset-theoretic techniques. This work also leads to a result for simple polytopes in which $G(P,\mathbf{c})$ is the Hasse diagram of a lattice L that the order complex of each open interval in L is homotopy equivalent to a ball or a sphere of some dimension. Applications are given to the weak Bruhat order, the Tamari lattice, and more generally to the Cambrian lattices, using realizations of the Hasse diagrams of these posets as 1-skeleta of permutahedra, associahedra, and generalized associahedra.

Keywords: nonrevisiting path conjecture, strict monotone Hirsch conjecture, polytope, poset topology, Tamari lattice, weak order, associahedron, permutahedron

1. Introduction

The focus of this paper is to show how poset-theoretic ideas may be applied to give new insights into elusive questions regarding polytopes that are motivated in a rather direct way by major open questions from operations research regarding linear programming. In particular, we consider questions motivated by linear programming whose solution could yield strong upper bounds on the diameters of interesting classes of polytopes, potentially giving new insight into why linear programming is so efficient in cases of interest. Throughout this paper, we assume that we have a simple polytope $P \subseteq \mathbb{R}^d$ and a "generic" cost vector $\mathbf{c} \in \mathbb{R}^d$, by which we mean that $\mathbf{c}\cdot u \neq \mathbf{c}\cdot v$ for u, v distinct vertices of P. Given such a vector $\mathbf{c}$, we obtain a directed graph $G(P,c)$ on the 1-skeleton of P by orienting each edge $e_{u,v}$ from u to v for $\mathbf{c}\cdot u < \mathbf{c}\cdot v$. By construction this graph will be acyclic, by which we mean that it will not have any directed cycles.

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

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

The starting point for our work is the following observation. Requiring $G(P, \mathbf{c})$ to be the Hasse diagram (defined in Section 2) of a partially ordered set is equivalent to requiring that the directed paths in $G(P, \mathbf{c})$ may never revisit any 1-dimensional face after leaving it. To make this more precise, we say that a directed path visits a face when it visits any vertex of that face. This sort of face nonrevisiting property for faces of all dimensions (or equivalently for facets) in a polytope forces linear programming to run very efficiently on that polytope. This observation led us to believe that it would be useful from the standpoint of linear programming to better understand polytopes P and cost vectors $\mathbf{c}$ such that $G(P, \mathbf{c})$ is a Hasse diagram of a partially ordered set (poset).

Indeed, one of the main overarching ideas of this paper is to raise the following question and begin to answer it: are there useful sufficient conditions on a polytope P and generic cost vector $\mathbf{c}$ under which directed paths in $G(P, \mathbf{c})$ may never revisit any face of P after departing from that face? This question appears as Question 2.5. This is motivated by consequences discussed shortly that this sort of face nonrevisiting property will have towards giving upper bounds on the diameter of the directed graph $G(P, \mathbf{c})$. These upper bounds in turn may contribute towards a better understanding of when to expect the simplex method for linear programming to be efficient, as we soon explain. Whenever a polytope P and generic cost vector $\mathbf{c}$ together satisfy the above face nonrevisiting property, this would give an upper bound of $n - d$ on the length of the longest directed path in $G(P, \mathbf{c})$ where d is the dimension of P and n is the number of facets (maximal boundary faces) in P.

In this paper, we conjecture that the following conditions suffice for the desired sort of face nonrevisiting:

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, then the directed paths in $G(P, \mathbf{c})$ can never revisit any faces they have left. That is, any directed path $v_1 \to v_2 \to \cdots \to v_k$ with v_1 and v_k both in a face F must have $v_i \in F$ for $1 \le i \le k$.*

We will prove several results which we suspect constitute significant steps towards positively resolving this conjecture. These upcoming results will allow us to prove Conjecture 1 in the case of 3-dimensional polytopes in Theorem 5.1. In Theorem 5.2, we prove Conjecture 1 for spindles whose two distinguished vertices are the source and sink of the polytope. Theorem 5.2 implies that the known counterexamples to the Hirsch Conjecture cannot be simple polytopes with $G(P, \mathbf{c})$ the Hasse diagram of a poset with the two distinguished vertices of the spindle as source and sink. In particular, this implies that the known examples of d-dimensional spindles with n facets having the property that the distance between the two distinguished vertices v_1 and v_2 is greater than $n - d$ cannot also meet the hypotheses for Conjecture 1 with respect to cost vectors having v_1 and v_2 as source and sink. In other words, we prove in Theorem 5.2 that none of the known counterexamples to the Hirsch Conjecture yield counterexamples to Conjecture 1 for cost vectors having the distinguished vertices at distance more than $n - d$ apart as source and sink. Examples of various families of polytopes meeting the hypotheses of Conjecture 1 are given in Section 5.

Given a polytope P and generic cost vector $\mathbf{c}$, each face F of P will have a unique "source" in the restriction of $G(P, \mathbf{c})$ to F, namely a vertex v whose directed edges to other vertices of F all point outward from v. Each face F also will have a unique "sink", namely a vertex w whose edges to other vertices of F all point inward to w. Given an i-dimensional face F

(usually called an i-face) of a simple polytope P such that $G(P, \mathbf{c})$ is a Hasse diagram of a poset L, then for u the source vertex of F, and elements $a_1, \ldots, a_i \in F$ all covering u in L, define the "pseudo-join" of $a_1, a_2, \ldots, a_i$ to be the unique sink of F. This notion of pseudo-join is a geometric construct that one might hope would equal the join $a_1 \vee a_2 \vee \cdots \vee a_i$ (namely the unique least upper bound) of this same collection $\{a_1, \ldots, a_i\}$ of atoms in the event that L is a lattice. This equality would indeed hold provided that directed paths do not revisit faces after departing from them. We refer readers to Section 2 for further background review e.g. regarding posets as well as polytopes.

Our first main result, Theorem 1.1, confirms this expectation that the join of a set of atoms equals the pseudo-join of the same set of atoms in the setting of polytopes P and generic cost vectors $\mathbf{c}$ such that $G(P, \mathbf{c})$ fits the general framework focused upon in much of this paper:

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 is proven as Theorem 4.7. The fact that this result holds is more subtle than it might appear at first glance. Santos' result in [37] that the nonrevisiting path conjecture (see Conjecture 3) is false demonstrates how statements that would be very useful for proving Theorem 1.1 that might sound like they should obviously hold are in fact false.

Recall that the aim of linear programming is to maximize (or to minimize) $\mathbf{c} \cdot \mathbf{x}$ for a fixed cost vector $\mathbf{c}$ over all possible choices of $\mathbf{x} \in P$ for P a polytope, i.e., for P a bounded subset of $\mathbb{R}^d$ given by a system of (weak) linear inequalities. One particularly famous example of a linear programming problem is the traveling salesman problem (see e.g. [11]). The simplex method for linear programming finds the vertex of P where this maximum (resp. minimum) is achieved by greedily following directed edges of $G(P, \mathbf{c})$ until reaching the unique sink (resp. source) of the directed graph $G(P, \mathbf{c})$. 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). We will henceforth focus on maximizing $\mathbf{c} \cdot \mathbf{v}$ for $\mathbf{v} \in \mathbf{P}$, since the minimization problem is completely equivalent.

From the viewpoint of linear programming, it seems natural and useful to ask for conditions on P and $\mathbf{c}$ that would ensure that no directed path can revisit any face it has left. After all, having this nonrevisiting property on a d-polytope P with n facets would guarantee that all possible pivot rules would be efficient in the sense that every directed path in the graph $G(P, \mathbf{c})$ would reach the vertex $\mathbf{v}$ where $\mathbf{c} \cdot \mathbf{v}$ is maximized in at most $n - d$ steps.

Theorem 1.1 gives some evidence for Conjecture 1. Examining the example of Klee-Minty cubes may give further intuition as well as evidence for the pertinence of the Hasse diagram requirement to helping ensure that the desired face nonrevisiting property will hold. Specifically, we note that the famous Klee-Minty cubes (introduced in [24]) violate not only our face nonrevisiting property but more specifically our requirement that $G(P, \mathbf{c})$ be the Hasse diagram of a poset in a way that really seems to be at the heart of why the simplex method may be so inefficient on Klee-Minty cubes – recall that the Klee-Minty cubes are polytopes P and cost vectors $\mathbf{c}$ with P a realization of a d-dimensional cube such that a directed path exists in $G(P, \mathbf{c})$ that visits all 2^d vertices of P. These were historically the first examples demonstrating that the simplex method is not always efficient. Careful examination of this important family of examples led us to to regard these examples as quite suggestive that the

Hasse diagram requirement on $G(P, \mathbf{c})$ may indeed help preclude long directed paths while being a reasonable checkable condition and one that many important families of polytopes may satisfy. For further background and properties of Klee-Minty cubes, including a helpful illustration of a 3-dimensional Klee-Minty cube, we refer readers to [14].

In some sense, our focus on simple polytopes is not such a severe restriction in seeking a better understanding of which polytopes will satisfy strong upper bounds on their diameters. After all, Klee and Walkup did prove in [25] that the Hirsch Conjecture for simple polytopes would have implied it for all polytopes. It seems plausible that some of our results could hold also for polytopes that need not be simple (perhaps with minor modifications to their statements). However, the proofs that we give throughout this paper do heavily utilize the hypothesis that our polytopes are simple. The combination of requiring P to be a simple polytope and $G(P, \mathbf{c})$ to be a Hasse diagram seems to be quite powerful when taken together.

Remark 1.2. Our upcoming results hint at the distinct possibility for simple polytopes with $G(P, \mathbf{c})$ the Hasse diagram of a lattice that requiring none of the directed paths ever to depart and revisit any low dimensional face of P may force this same face nonrevisiting property for the higher dimensional faces as well. Specifically, our result regarding faces $F \subseteq G$ for F a codimension one face in G appearing as Lemma 4.1 is designed to facilitate upward propagation in dimension within proofs. Indeed this does yield our proof of Conjecture 1 for 3-polytopes appearing as Theorem 5.1. We are hopeful that with more effort and additional insights Lemma 4.1 should have further applications to higher dimensions as well.

Our next main result, proven as Theorem 4.11, hints that perhaps situations where not all pivot rules are efficient could in some cases be detectable using poset topology.

Theorem 1.3. *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, 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.*

The posets meeting the hypotheses of Theorem 1.3 often are not shellable, necessitating other methods besides shellability to determine topological structure.

Remark 1.4. As a different sort of motivation for our work, we note that people sometimes ask whether the Hasse diagram of a given poset can be realized as the 1-skeleton of a polytope; for instance this was asked and later answered for the Tamari lattice. Our work might also give a new perspective on potential necessary conditions for this to be possible, particularly if Theorem 1.3 could be generalized beyond simple polytopes.

In Section 2, we review background, including the Hirsch Conjecture, the Nonrevisiting Path Conjecture, and the Strict Monotone Hirsch Conjecture. Section 3 introduces seemingly new notions (or at least not widely known notions) to be used later. Section 4 gives the proofs of our main technical results. This includes the two results mentioned above as well as a number of corollaries and related results. Section 5.1 gives applications to two important classes of polytopes, namely the 3-polytopes and the spindles. Section 5.2 gives well-known families of polytopes that will fit into our framework, namely permutahedra, associahedra and generalized associahedra. The posets on their 1-skeleta are weak order, the Tamari lattice and Cambrian lattices, respectively. Section 5.3 turns to the case of zonotopes, where especially clean results are possible. Section 5.4 generalizes our results from orientations on the 1-skeleton of a simple polytope induced by a generic cost vector to more general acyclic orientations derived from shellings of the dual simplicial polytope. Further questions and remarks appear in Section 6.

2. Background

A **cover relation** $u \prec v$ in a finite partially ordered set (poset) Q is $u \le v$ in Q with the requirement that $u \le z \le v$ 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 poset Q has a **dual poset**, denoted Q^* with $u \le v$ in Q^* if and only if $v \le u$ in Q.

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 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. [38] for further background on posets.

A **polytope** is any set 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 weak 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.

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.

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. For K a polyhedral complex, the order complex of the face poset of K is the first barycentric subdivision of K, and in particular is homeomorphic to K. Each polytope P has a **dual polytope**, denoted P^*, with $F(P^*) = (F(P))^*$. Another way to define what it means for a polytope P to be a simple polytope is that its dual polytope is a simplicial polytope.

Definition 2.1. 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. See e.g. [39] for further background on polytopes.

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 2.2 (Quillen Fiber Lemma). *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 $\Delta(P) \simeq \Delta(Q)$.*

Theorem 2.2 was proven in [32]. Recall that a **dual closure map** is a poset map $f : P \to P$ with $f(u) \le u$ such that $f^2(u) = f(u)$. 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 the order complex of $f^{-1}_{\ge u}$. Thus, $\Delta(im(f)) \simeq \Delta(P)$.

Remark 2.3. 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 which has the 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.

Now let us recall the Hirsch Conjecture, the Nonrevisiting Path Conjecture, and the Strict Monotone Hirsch Conjecture. We refer readers e.g. to [39] for a more in-depth discussion of all of these conjectures.

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

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 [37]:

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. We believe the requirement that $G(P, \mathbf{c})$ be a Hasse diagram would be a useful possibility as one such condition to consider, in conjunction with other properties such as $G(P, \mathbf{c})$ being the Hasse diagram of a lattice. Some evidence is provided by our upcoming results to suggest that this Hasse diagram property for $G(P, \mathbf{c})$ together with this lattice requirement might very well suffice for a simple polytope P to guarantee our face nonrevisiting property for $G(P, \mathbf{c})$, namely for our Conjecture 1.

Question 2.5. What are sufficient conditions on a polytope P and cost vector $\mathbf{c}$ so that directed paths in $G(P, \mathbf{c})$ may never revisit any face they have left? Notice that this is equivalent to asking for conditions under which directed paths may never revisit any facet they have left, by virtue of each face being an intersection of facets.

Question 2.5 is closely related to the Strict Monotone Hirsch Conjecture:

Conjecture 4 (Strict Monotone Hirsch Conjecture). *Let P be a d-dimensional polytope with n facets, and let $\mathbf{c}$ be a generic linear functional. Then there is a directed path in $G(P, \mathbf{c})$ from the source of P to the sink of P of length at most $n - d$.*

To see the connection, notice that any set of conditions on a polytope P and cost vector $\mathbf{c}$ that would guarantee that no directed path in $G(P, \mathbf{c})$ revisits any face after leaving it would imply that $G(P, \mathbf{c})$ has diameter at most $n - d$, by the same reasoning (recalled above) showing that the Nonrevisiting Path Conjecture implies the Hirsch Conjecture.

3. Pseudo-joins and pseudo-meets, the non-revisiting property, and the Hasse diagram property

In this section, we introduce some potentially new notions which will be quite useful later in the paper. First we 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 $\mathbf{c} \cdot w$. Equivalently, the source of F will be the unique vertex of $G(P, \mathbf{c})$ in F only having outward oriented edges to other vertices of F while the sink will the unique vertex of $G(P, \mathbf{c})$ in F only having inward oriented edges to it from other vertices of F. The implicit uniqueness assumption in these notions is justified as follows:

Remark 3.1. It is well known (cf. Theorem 3.7 in [39]) and straight-forward to see that the directed graph on the 1-skeleton of any face F of a polytope P obtained by restricting $G(P, \mathbf{c})$ for $\mathbf{c}$ generic to the face F has a unique source and unique sink.

Remark 3.2. The uniqueness of source and sink for each face implies in particular that 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. We specifically point out this property for 2-faces because we will use it repeatedly later.

For simple polytopes, the following related notions will figure prominently in this paper:

Definition 3.3. 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 L. Define the **pseudo-join** of any collection S of atoms of an interval $[u, v]$ in L to be the sink of the unique smallest face F_S of P that contains u and all of the elements of S. Denote this sink by $psj(S)$. Define the **pseudo-meet** of any collection T of coatoms in an interval $[u,v]$ in L to be the source vertex of the smallest face G_T containing v and all of the coatoms in T. Denote this source by $psm(T)$.

The existence of pseudo-joins (and likewise pseudo-meets) is justified by the following facts:

(1) For any vertex v and any collection of edges $e_1,\dots,e_i$ emanating outward from v in a simple polytope, there exists an i-face containing v and all of the edges $e_1,\dots,e_i$. This face will be the unique smallest face containing v and all these edges, due to having the minimal possible number of edges emanating out from v for an i-face.

(2) Whenever faces F and G both contain a collection of atoms, then $F\cap G$ will be another face also containing all these atoms.

Definition 3.4. A directed graph $G(P,\mathbf{c})$ on the 1-skeleton of a polytope P satisfies the **non-revisiting property** if for each facet F and each directed path p_F that starts and ends at vertices in F, this implies that p_F must stay entirely within F. We say that $G(P,\mathbf{c})$ satisfies the non-revisiting 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, this implies that every vertex of p_F is in F.

Remark 3.5. This non-revisiting property for facets immediately implies the non-revisiting property for all lower dimensional faces as well, since each face is an intersection of facets.

Definition 3.6. A directed graph $G(P,\mathbf{c})$ on the 1-skeleton of a polytope has the **Hasse diagram property** if it is the Hasse diagram of a poset.

Lemma 3.7. *The non-revisiting property for 1-dimensional faces implies that the directed graph $G(P,\mathbf{c})$ will be the Hasse diagram of a poset. Specifically, $G(P,\mathbf{c})$ is the Hasse diagram for the poset having $u\le v$ if and only if there is a directed path from u to v in $G(P,\mathbf{c})$.*

Proof. Acyclicity of $G(P,\mathbf{c})$ ensures that the directed paths indeed specify comparabilities in a partially ordered set. The non-revisiting property for 1-dimensional faces ensures a directed path cannot visit the sink of a directed edge after departing from the source of that directed edge in a way that departs from the edge itself. This shows that each directed edge gives rise to a cover relation. ☐

Example 3.8. An especially instructive and significant family of simple polytopes failing the Hasse diagram property are the Klee-Minty cubes (cf. [24]). These were the first known polytopes exhibiting that the simplex method from optimization is not always efficient. See [24], or e.g. see [14] where (among other things) a particularly helpful illustration of one of these polytopes appears.

The Klee-Minty cubes are hypercubes with carefully chosen (non-standard) locations for the vertices. They were designed so that there exists a cost vector $\mathbf{c}$ such that $G(P,\mathbf{c})$ has a directed path from source to sink that visits all of the vertices. Thus, these are d-polytopes with a directed path in $G(P,\mathbf{c})$ visiting 2^d vertices, exhibiting that a pivot rule exists which makes the simplex method highly non-efficient.

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

In this section, we develop a series of general results about directed paths in posets that are derived from 1-skeleta of simple polytopes with respect to a choice of a generic cost vector. This also will include results on the topological structure of associated poset order complexes.

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, then $G(P, \mathbf{c})$ cannot have an edge directed from $v \in F$ outward to some vertex $v' \in G \setminus F$ and an edge directed from $w' \in G \setminus F$ inward to $w \in F$.*

Proof. 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. The fact that $v \in p_F$ has $o(v) = +1$ while $w \in p_F$ has $o(w) = -1$ allows us to apply the discrete version of the intermediate value theorem to deduce the existence of two consecutive vertices $v_1 \to v_2$ in p_F with $o(v_1) = +1$ and $o(v_2) = -1$.

Next we show that the edges e_{v_1} and e_{v_2} must both be contained in a single 2-dimensional face $F(e_{v_1}, e_{v_2})$ of P. 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_{v_2,x_2}, and we have $F(v_1, v_2, x_2) \not\subseteq F$ 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 in G such that this 2-face is not contained in F, by virtue of each upper interval in the face poset of a simple polytope being a Boolean lattice. Applying this observation specifically 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 same 2-face containing d_{v_1,v_2} and not contained in F. This shows that there exists a unique 2-face $F(v_1, v_2, x_1) = F(v_1, v_2, x_2)$ which contains all three edges e_{v_1,v_2}, e_{v_1,x_1}, and e_{v_2,x_2}.

Applying Remark 3.2 to this face $F(v_1, v_2, x_1)$, we note 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 the edge e_{v_1,v_2} in the path p_F. This directed edge by itself must be one of the two directed paths from the source to the sink in the boundary of $F(v_1, v_2, x_1)$. For $G(P, \mathbf{c})$ to be the Hasse diagram of a poset Q, the edge 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 $F(v_1, v_2, x_1)$ includes x_1 as an intermediate element and gives rise to a saturated chain from v_1 to v_2 in Q having x_1 as an intermediate element. This contradicts $v_1 \prec v_2$ being a cover relation in Q, giving a contradiction to $G(P, \mathbf{c})$ being a Hasse diagram. This completes the proof. ☐

Lemma 4.1 together with reasoning as in its proof yields the following further result.

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. Then any face F of P containing the source vertex of P has the property that each edge $e_{v,w}$ for $v \in F$ and $w \notin F$ must be oriented from v to w. Likewise, for each face F' containing the sink vertex of P and each edge $e_{x,y}$ for $x \in F'$ and $y \notin F'$ must be oriented from y to x.*

Proof. Let F be a face of P 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, 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 with $\dim(G) = \dim(F) + 1$.

The proof of the statement for faces F' containing the sink of P is entirely analogous. ☐

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 depart F and later revisit F. Likewise for each face G containing $\hat{1}$, there cannot be any directed paths that depart G and later revisit G.*

Corollary 4.3 has the following important special case:

Corollary 4.4. *For each set S of atoms, each directed path from $\hat{0}$ to $psj(S)$ stays within the unique smallest face F_S containing $\hat{0}$ and all of the atoms in 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 $J(S)$ of any collection S of atoms in L is contained in the unique smallest face F_S containing all of the atoms in S. Likewise the meet of any collection T of coatoms in L is contained in the unique smallest face G_T containing all of the coatoms in T.*

Proof. There is a directed path from any atom $a \in S$ to $J(S)$ by definition of upper bound. There is also a directed path from $J(S)$ to $psj(S)$ 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 $J(S)$ for the elements of S. Concatenating these directed paths yields a directed path p from a to $psj(S)$. By Lemma 4.2, this directed path p must stay within F_S. In particular, this implies $J(S) \in F_S$. The proof for coatoms is entirely analogous, by dualizing everything. □

One might be tempted to try to generalize Lemma 4.2 by replacing $\hat{0}$ by an arbitrary element $u \in L$ while replacing the atoms of L by the atoms of a closed interval $[u,v]$ in L. However, our proof of Lemma 4.2 does not apply for $u \neq \hat{0}$ due to $G(P,\mathbf{c})$ having edges directed towards any $u \neq \hat{0}$. Nonetheless, with a good bit of effort we will obtain weaker results seeminingly heading towards an analogue of Lemma 4.2 for arbitrary intervals $[u,v]$.

Next is a result regarding the pseudo-join of a pair of elements x, y, denoted $psj(x,y)$, together with a dual result regarding the pseudo-meet of elements x', y', denoted $psm(x',y')$. This will serve as the base case for the induction in the upcoming proof of Theorem 4.7.

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 in F, and let $x, y \in F$ be vertices both covering u in L. Then $psj(x,y) = x \vee y$. Likewise for any 2-face F' with sink v and elements $x', y' \in F'$ both covered by v in L, then $psm(x',y') = x' \wedge y'$.*

Proof. Suppose that either there exists a 2-face F in P with source u having $x, y \in F$ satisfying $u \prec x$ and $u \prec y$ in L with $psj(x,y) \neq x \vee y$ or that there exists a 2-face F' in P with sink v having $x', y' \in F'$ satisfying $x' \prec v$ and $y' \prec v$ in L with $psm(x',y') \neq x' \wedge y'$. For u the source of a face F, let $d(F,u)$ be the length of the longest saturated chain in L from $\hat{0}$ to u. For v the sink of a face F', let $d(F',v)$ be the length of the longest saturated chain in L from v to $\hat{1}$. A critical property in what follows will be the following sort of monotonicity: for $u_1 < u_2$ (resp. $v_1 < v_2$) with u_1 (resp. v_1) the source (resp. sink) of a 2-face F_1 and

u_2 (resp. v_2) the source (resp. sink) of a 2-face F_2 that we have $d(F_1,u_1) < d(F_2,u_2)$ (resp. $d(F_1,v_1) > d(F_2,v_2)$).

Consider a 2-face of one of the two types above, namely either having the join of its atoms not equalling the pseudo-join of its atoms or having the meet of its coatoms not equalling the pseudo-meet of these coatoms; more specifically, choose such a 2-face F in a way that achieves the overall largest possible value $d^{\max}$ achieved anywhere within P by either of the two quantities $d(F,u)$ and $d(F',v)$. For convenience in notation, let us assume that $d^{\max} = d(F,u)$ for a 2-face F having elements x,y covering the source u in F with $x \vee y \neq psj(x,y)$; the case of $d^{\max} = d(F',v)$ for v a sink in F' is entirely analogous with everything dualized.

Denote the vertices proceeding upward along one of the two directed paths from u to $psj(x,y)$ in the boundary of F as $u, x_1, x_2, \dots, x_r$. Denote the vertices proceeding upward along the other directed path in the boundary of F from u to $psj(x,y)$ as $u, y_1, y_2, \dots, y_s$. Do this in such a way that we have $x = x_1, y = y_1$ and $x_r = y_s = psj(x,y)$. Observe that the coatoms $c_1 = x_{r-1}$ and $c_2 = y_{s-1}$ in $[u, psj(x,y)] \cap F$ satisfy $c_1 \vee c_2 = psj(x,y)$ by virtue of both being covered by $psj(x,y)$. Our choice of F guarantees on the other hand the existence of $x' = x_l$ and $y' = y_m$ both in F with $1 \leq l \leq r-1$ and $1 \leq m \leq s-!$ such that $x' \vee y' < psj(x,y)$. Choose such x_l and y_m with l and m as large as possible in the sense that $x_{l+1} \vee y_m = psj(x,y) = x_l \vee y_{m+1}$. Notice for such x', y' that we must either have (a) $x' \vee y' \notin F$ or (b) the existence of a directed path from either x' or y' to $x' \vee y' \in F$ that includes at least one element outside of F. In what follows, we will rule out both (a) and (b) to get our desired contradiction.

Suppose we are in case (a), namely that we have $x' \vee y' \notin F$. Since $psj(x,y)$ is an upper bound for x' and y' by virtue of being the sink of a face containing both x' and y', there must be a directed path from $x' \vee y'$ to $psj(x,y)$ in $G(P,\mathbf{c})$ by definition of join. There also must be a directed path from y' to $x' \vee y'$. Since $x' \vee y' \notin F$, concatenating these directed paths exhibits the existence of a directed path from y' to $psj(x,y)$ that departs F and later revisits it. Consider a directed path from y' to $psj(x,y)$ that exits from F via an edge from $y_i \in F$ to some $z \notin F$ and re-enters F via an edge from $z' \notin F$ to $y_k \in F$ with $k-i$ chosen as small as possible. By our set-up, y_k is an upper bound for y_{i+1} and z, which implies $y_{i+1} \vee z \leq y_k$. But now the fact that we chose F with $d(F,u) = d^{\max}$ ensures that $y_{i+1} \vee z = psj(y_{i+1}, z)$, since otherwise y_i, y_{i+1}, z would all belong to a face F_1 with source y_i satisfying $d(F_1, y_i) > d(F,u)$. But $y_{i+1} \vee z = psj(y_{i+1}, z)$ 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. $G \neq F$, since $z \in G$ and $z \notin F$. Distinctness of the 2-faces F and G implies that F and G intersect in at most an edge, hence share at most two vertices. By definition, F and G do share y_i and y_{i+1}; we also have $y_{i+1} \vee z \in G$ and $y_k \in F$, implying $y_{i+1} \vee z \neq y_k$. Since we have exhibited directed paths from y_{i+1} to y_k and from z to y_k, the definition of least upper bound combined with $y_{i+1} \vee z \neq y_k$ implies $y_{i+1} \vee z < y_k$. But then we will have a directed path p_{i+1} from y_{i+1} to $y_{i+1} \vee z \notin F$ with p_{i+1} departing from F either at y_{i+1} or at some $y_{i'}$ with $i' > i+1$. We will also have a directed path from $y_{i+1} \vee z$ to y_k re-entering F at some $y_{k'}$ with $k' \leq k$. Since $k' - i' \leq k - (i+1) < k - i$, this contradicts $k-i$ having been chosen as small as possible. Thus, (a) is ruled out.

For (b), a completely analogous argument to the one just given above for (a) will apply. The point is now to choose a directed path exiting F at some y_i and re-entering F at some x_k (or exiting F at some x_i and re-entering F at some y_k). Again do this with $k-i$ as small as possible, and again use the fact that whenever three vertices belong to a 2-dimensional face, the three vertices uniquely determine that 2-dimensional face. ☐

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 $a_1, \dots, a_j \in L$ all covering u in L, then the pseudo-join of $a_1, \dots, a_j$ equals $a_1 \vee \cdots \vee a_j$. Likewise for any $v \in L$ and any $c_1, \dots, c_j \in L$ all covered by v in L, the pseudo-meet of $c_1, \dots, c_j$ equals $c_1 \wedge \cdots \wedge c_j$.*

Proof. The proof is by induction on j. The $j = 1$ case holds tautologically. The $j = 2$ case is proven in Theorem 4.6. In what follows, let $J(a_1, \dots, a_j)$ denote $a_1 \vee \cdots \vee a_j$. Let $J(S)$ denote the join of the elements in a set S. Let F_S be the unique smallest face containing all of the elements in $S \cup \{u\}$ for S a set of elements all covering some element $u \in L$.

To prove the result for $j \geq 3$, we inductively assume for all $u \in L$ and all $T \subseteq \{a_i \in L | u \prec a_i\}$ satisfying $|T| < j$ that $J(T) = psj(T)$. To prove $J(a_1, \dots, a_j) = psj(a_1, \dots, a_j)$, our plan is to prove $psj(a_1, \dots, a_j) \leq J(a_1, \dots, a_j)$ and $J(a_1, \dots, a_j) \leq psj(a_1, \dots, a_j)$. In order to prove $psj(a_1, \dots, a_j) \leq J(a_1, \dots, a_j)$, our plan is to prove for any $x \in F_{\{a_1,\dots,a_j\}}$ that $x \leq J(a_1, \dots, a_j)$ by first assuming that we have already shown $x' \leq J(a_1, \dots, a_j)$ for all $x' \in F_{\{a_1,\dots,a_j\}}$ with $x' < x$. In this manner, we will progressively deduce this upper bound of $J(a_1, \dots, a_j)$ for all of the finitely many elements of $F_{\{a_1,\dots,a_j\}}$. In fact, it will be convenient to do more: we will prove not only that $x \in F_{\{a_1,\dots,a_j\}}$ satisfies $x \leq J(a_1, \dots, a_j)$ but also that all elements y covering x that are also in $F_{\{a_1,\dots,a_j\}}$ also satisfy $y \leq J(a_1, \dots, a_j)$, using the assumption that we have already shown this same pair of claims for each $x' \in F_{\{a_1,\dots,a_j\}}$ that has $x' < x$. That is, we assume not only that all $x' \in F_{\{a_1,\dots,a_j\}}$ with $x' < x$ satisfy $x' \leq J(a_1, \dots, a_j)$ but also that all elements $y' \in F_{\{a_1,\dots,a_j\}}$ covering such x' also satisfy $y' \leq J(a_1, \dots, a_j)$, and we use these claims to prove the same statements for x. To get started in checking this condition progressively as one works one way upward from u, note first that this condition holds for u itself due to each a_i which covers u in $F_{\{a_1,\dots,a_j\}}$ by definition satisfying $a_i \leq J(a_1, \dots, a_j)$.

Let i be the number of cover relations upward from elements of $F_{\{a_1,\dots,a_j\}} \cap [u, x)$ to a given $x \in F_{\{a_1,\dots,a_j\}}$ we wish to handle next. First consider the case with $2 \leq i \leq j - 1$. Since P is a simple polytope, these i edges incident to x oriented towards x all must belong to an i-face F in $F_{\{a_1,\dots,a_j\}}$ having x as its sink. By definition, x is the pseudo-join of the set T of elements of F all covering the source of F. Since $|T| \leq i < j$, we have $psj(T) = J(T)$ by our inductive hypothesis on j, and hence we have $x = J(T)$. But each element of T would have already been proven to be less than or equal to $J(a_1, \dots, a_j)$ by virtue of being less than x and belonging to $F_{\{a_1,\dots,a_j\}}$. Since x is the join of the elements of T, this implies $x \leq J(a_1, \dots, a_j)$. Now consider any $y \in F_{\{a_1,\dots,a_j\}}$ which covers x. There must also be some $w \in F_{\{a_1,\dots,a_j\}}$ covered by x, since we already handled $x = u$ so may assume $x > u$. Consider the unique 2-face $F(x, y, w)$ containing x, y and w. Let u be its source. Since $u < x$, we will have already shown both $u \leq J(a_1, \dots, a_j)$ and the same for all elements of $F_{\{a_1,\dots,a_j\}}$ covering u, so in particular we will have already shown that the two atoms of $F(x, y, w)$ are both bounded above by $J(a_1, \dots, a_j)$. But the pseudo-join of these two atoms equals the join of these two atoms, hence also will be less than or equal to $J(a_1, \dots, a_j)$. This implies that every element of $F(x, y, w)$ will be bounded above by $J(a_1, \dots, a_j)$, in particular yielding $y \leq J(a_1, \dots, a_j)$.

Next suppose the element $x \in F_{\{a_1,\dots,a_j\}}$ we wish to handle next has $i = 1$ cover relations upward to it from elements in $F_{\{a_1,\dots,a_j\}}$. Let w' be the unique element in $F_{\{a_1,\dots,a_j\}}$ covered by x. If $u = w'$, then $x \in \{a_1, \dots, a_j\}$, implying $x \le J(a_1, \dots, a_j)$. If $u < w'$, then there exists $v' \in F_{\{a_1,\dots,a_j\}}$ with $u \le v' \prec w' \prec x$. Since P is simple, there exists a unique 2-face $F(v', w', x)$ in $F_{\{a_1,\dots,a_j\}}$ containing both of the edges $e_{v',w'}$ and $e_{w',x}$. Let u' be the source of $F(v', w', x)$, and let a_1' and a_2' be the elements of $F(v', w', x)$ covering u'. At the step when we are seeking to prove the desired claim for x, we would have already proven $u' \le J(a_1, \dots, a_j)$ and also that all elements covering u' are bounded above by $J(a_1, \dots, a_j)$, so in particular $a_1' \le J(a_1, \dots, a_j)$ and $a_2' \le J(a_1, \dots, a_j)$. But $a_1' \le J(a_1, \dots, a_j)$ and $a_2' \le J(a_1, \dots, a_j)$ implies $a_1' \vee a_2' \le J(a_1, \dots, a_j)$. Theorem 4.6 yields $psj(a_1', a_2') = a_1' \vee a_2'$, implying $psj(a_1', a_2') \le J(a_1, \dots, a_j)$. Since $x \le psj(a_1', a_2')$, we may conclude $x \le psj(a_1', a_2') \le J(a_1, \dots, a_j)$, giving $x \le J(a_1, \dots, a_j)$ as desired.

Now we turn to the case with $x = psj(a_1, \dots, a_j)$, namely the $i = j$ case. Let $x_1 \prec x$ and $x_2 \prec x$ be distinct cover relations upward to x within $F_{\{a_1,\dots,a_j\}}$. We already have $x_1 \le J(a_1, \dots, a_j)$ and $x_2 \le J(a_1, \dots, a_j)$. But $x = x_1 \vee x_2$ due to x covering both these elements, implying $x \le J(a_1, \dots, a_j)$. Thus we have proven $psj(a_1, \dots, a_j) \le J(a_1, \dots, a_j)$. But we also have $J(a_1, \dots, a_j) \le psj(a_1, \dots, a_j)$ by virtue of $psj(a_1, \dots, a_j)$ being an upper bound for all of the elements $a_1, \dots, a_j$ and $J(a_1, \dots, a_d)$ being the least upper bound for this same set of lattice elements. Thus, $psj(a_1, \dots, a_j) \le J(a_1, \dots, a_j) \le psj(a_1, \dots, a_j)$, implying $psj(a_1, \dots, a_j) = J(a_1, \dots, a_j)$, as desired.

The desired equivalence of meet and pseudo-meet for any set $\{c_1, \dots, c_j\}$ of elements all covered by a single element v follows by an entirely analogous argument. Specifically, one works instead with the dual poset by negating the cost vector and exchanging the roles of meets and joins, allowing precisely the same proof to go through. ☐

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 $a_1, a_2, \dots, a_j$. Likewise 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 $c_1, c_2, \dots, c_j$.*

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

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 L. Let S and T be distinct sets of atoms (resp. coatoms) in L. Then psj(S) $\ne$ psj(T) (resp. psm(S) $\ne$ psm(T)). If $G(P, \mathbf{c})$ is the Hasse diagram of a lattice, then the same pair of statements holds for each interval $[u, v]$ in L.*

Proof. Again we let F_S be the smallest face containing a collection S of atoms. First note that $S \ne T$ implies $F_S \ne F_T$ since each face F_S is simple with $dim(F_S) = |S|$ and with the neighbors of $\hat{0}$ in F_S being exactly the elements of S.

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 exactly one edge directed from some vertex $v_S \in F_S \setminus F_T$ to $v \in F_T$, by virtue of v being the sink of F_S. 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. There is also a directed path from $\hat{0}$ to v that stays within F_T, again by virtue of $\hat{0}$ being the source. Thus, Lemma 4.1 applies in this case, giving a contradiction to such faces F_S and F_T having the same sink.

Next turn to $F_T \subsetneq F_S$ both containing $\hat{0}$ with F_T of codimension higher than one in F_S. 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. If F_T and F_S have the same source and sink, then $F_{T'}$ 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 show the nonexistence of such $F_T \subseteq F_S$ of higher codimension having the same sink and both having source $\hat{0}$.

The case of any two faces F_S and $F_{S'}$ both containing $\hat{0}$ and having the same sink as each other also is ruled out by reduction to cases above as follows. 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 vertex as its sink. Thus, we may reduce the case of any F_S and $F_{S'}$ both containing $\hat{0}$ and both having the same sink as each other to the case of $F_T \subsetneq F_S$ already handled above by letting $T = S \cap S'$.

Turning now to arbitrary intervals $[u, v]$ of L, we use the fact that v is an upper bound for the set S all of the atoms of $[u, v]$, implying that the join of all of the elements of S is contained in $[u, v]$. Theorem 4.7 ensures that this join equals the pseudo-join of these same elements. This allows the above argument to be applied more generally to arbitrary intervals $[u, v]$ by applying our above argument for the case with $u = \hat{0}$ now instead to the subposet $[u, v] \cap F_S$; by definition of F_S, the face F_S also will include the pseudo-join of any set of atoms of $[u, v]$, just as needed.

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

Corollary 4.10. *Let P be a simple polytope and let $\mathbf{c}$ a generic cost vector such that $G(P, \mathbf{c})$ is the Hasse diagram of a poset L. Then the subposet of L 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). If L is a lattice, then this same property holds for each interval $[u, v]$ in L.*

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 poset L. Then each open interval (u, v) has order complex that is homotopy equivalent to a ball or a sphere. Thus, $\mu_L(u, v)$ equals $0, 1$, or -1 for each $u \le v$.*

Proof. Given any open interval (u, v) in L, the point is to define a surjective poset map $f : (u, v) \to Q$ for $Q = B_n \setminus \{\hat{0}, \hat{1}\}$ or $Q = B_n \setminus \{\hat{0}\}$ for B_n a Boolean lattice and to check the requisite contractibility of fibers of f needed to use the Quillen Fiber Lemma (see Lemma 2.2). The map f that we will use sends each $z \in (u, v)$ to the pseudo-join of the set of elements $a \in [u, v]$ satisfying $u \prec a \le z$. $B_n \setminus \{\hat{0}, \hat{1}\}$ has order complex homeomorphic to a sphere, by virtue of being the barycentric subdivision of the boundary of a simplex, while $B_n \setminus \{\hat{0}\}$ has order complex with a cone point at $\hat{1}$ which is therefore is contractible.

Under our hypotheses, the join of a set of atoms equals the pseudo-join of this same set of atoms, by Theorem 4.7, and likewise for the restriction to any closed interval $[u, v]$. We also 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 in Corollary 4.10. Our map f is a poset map, by the result from Theorem 4.7 that each pseudo-join of atoms of an interval equals the join of the same set of atoms of that interval. The fibers of f meet the contractibility requirement for the Quillen Fiber Lemma by virtue of each having a unique highest element, hence a cone point in the order complex of the fiber.

The claim that $\mu_L(u,v) \in \{0, \pm 1\}$ for each $u \leq v$ now follows directly from the well-known interpretation for $\mu_L(u,v)$ 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 considered in Theorem 4.11 typically are not shellable. A shelling would force every 2-dimensional face to have one of the two directed paths comprising its boundary to be of length exactly 2. See e.g. [9] for background on shellability.

5. Applications to large 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 as possible, we focus on some important families where our theory applies particularly naturally.

5.1. **Applications to 3-polytopes and to spindles.** We begin with two important classes of polytopes for which the face nonrevisiting property will follow from our earlier results in particularly direct and natural seeming ways.

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. However, the fact that there is no directed path from u to v that stays entirely in the 2-dimensional face F, implies that v is an upper bound in the lattice L for the two elements a_1, a_2 of F that cover the source vertex of F, since (without loss of generality) there is a directed path within F from a_1 to u as well as a directed path in $G(P, \mathbf{c})$ from u to v, and there is also a directed path within F from a_2 to v. 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 2.1 for a review of the notion of 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 there are no directed paths that depart from F and later revisit F for such F. 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 being greater than $n-d$. □

This result may be rephrased as saying 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. **Polytopes with well-known posets arising as their 1-skeleta.** We now turn to two well-known families of posets, namely the weak order and the Tamari lattice. Their Hasse diagrams will arise as 1-skeleta of the permutahedron and the associahedron, 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

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

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 [26]. 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. $M(t) = (a_1b_1, \ldots, a_ib_i, \ldots, a_{n-1}b_{n-1})$ 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 $a_i b_i$ replaced by $a_i b_i + a_i b_{i+r}$ and $a_{i+r} b_{i+r}$ replaced by $a_{i+r} b_{i+r} - a_i b_{i+r}$. We may again use any cost vector $\mathbf{c} = (c_1, \ldots, c_n) \in \mathbb{R}^n$ 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], [20] 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 [12], [13], 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 [29]. 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 [30].

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 [18] for the fact that the Hasse diagram of the c-Cambrian lattice is obtained from the polytope $Asso_c^a(W)$ by choosing a suitable cost vector and taking the directed graph it induces on the 1-skeleton of the polytope. Just before Example 3.5 in [18], it is asserted that all of these polytopes are simple. This is proven as Theorem 3.4 in [19]. 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 non-revisiting 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. [39] or [17]. 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. One may easily check that the proof of each statement above that has already been given earlier in the paper for $\mathcal{O}$ induced by a cost vector will indeed also apply more generally for facial orientations without need for any modification. These earlier results to be generalized are Theorem 4.7, Lemma 4.9 and Theorem 4.11, respectively. Checking that the proofs still hold unchanged is left as a completely straightforward exercise for the 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.

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 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 $i \ne j$ other than the pair $i = 1, j = 3$. Likewise introduce edges e_{w_i,w_j} for each $i \ne j$ 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 [15], 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 [23] for interesting, related work that takes a somewhat similar perspective to ours, work which provided some inspiration for parts of our work. In [23], 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 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 replace 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 join 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 $\mathbf{c} = (\mathbf{100}, \mathbf{2}, \mathbf{1})$, 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 join of the atoms in the original cube still remains as the join 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 question, 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 non-revisiting 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 $G(P, \mathbf{c})$ 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 [37] of a polytope violating the Hirsch Conjecture was a spindle (see Definition 2.1) 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 [37] 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 one would only need to consider the "good orientations" in the sense of Kalai from [23]. 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 [27] are simple polytopes that are spindles. All of these known violations to 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 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.

Another natural seeming class of polytopes to try, the generalized permutahedra (see [31]), also will not always satisfy all of our hypotheses. One sees this by noting that generalized permutahedra sometimes have triangular faces, which forces the Hasse diagram property to fail for all choices of cost vector.

A question we have not yet considered that seems worthwhile is as follows:

Question 6.9. Do some of the important classes of polytopes coming from real-world problems and operations research fit into our framework?

An affirmative answer could help explain (in such cases) the widely observed phenomenon that the simplex method typically is much more efficient in practice in real-world applications than theoretical results would predict.

Acknowledgments

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 helpful discussions, references, and in some cases counterexamples to questions she raised while working on this project. She thanks the anonymous referee for lots of very helpful feedback. 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.

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**, no. 2 (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] F. Chapoton, S. Fomin, and A. Zelevinsky, Polytopal realizations of the generalized associahedra, *Canadian Mathematical Bulletin* **45** (2002), 537–566.

[11] W. Cook, In pursuit of the Traveling Salesman: Mathematics at the Limits of Computation, Princeton University Press, 2011.

[12] P. Edelman, A partial order on the regions of $\mathbb{R}^n$ dissected by hyperplanes, *Trans. Amer. Math. Soc.* **283**, no. 2 (1984), 617–631.

[13] P. Edelman and J. Walker, The homotopy type of hyperplane posets, *Proc. Amer. Math. Soc.* **94** (1985), 221–225.

[14] B. Gärtner, M. Henk, and G. Ziegler, Randomized simplex algorithms on Klee-Minty cubes, *Combinatorica* **18** (3) (1998), 349-372.

[15] C. Greene, A class of lattices with Möbius function $\pm 1, 0$, *Europ. J. Combinatorics* (1988), **9**, 225–240.

[16] 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.

[17] P. Hersh and E. Swartz, Coloring complexes and arrangements, *J. Algebraic Combinatorics* **27** (2008), no. 2, 205–214.

[18] C. Hohlweg, Permutahedra and associahedra: generalized associahedra from the geometry of finite reflection groups, 129–159, *Prog. Math. Phys.*, **299**, Birkhäuser/Springer, Basel, 2012.

[19] C. Hohlweg, C. Lange, H. Thomas, Permutahedra and generalized associahedra, *Advances in Mathematics* **226** (2011), no. 1, 608–640.

[20] 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.

[21] V. Kaibel and M. Pfetsch, Computing the face lattice of a polytope from its vertex-facet incidences, *Comput. Geom.* **23** (2002), 281–290.

[22] V. Kaibel and M. Pfetsch, Some algorithmic problems in polytope theory, *Algebra, geometry, and its software systems*, 23–47, Springer, Berlin, 2003.

[23] G. Kalai, A simple way to tell a simple polytope from its graph, *J. Combinatorial Theory, Ser. A* **49** (1988), 381–383.

[24] V. Klee and G. Minty, How good is the simplex algorithm?, in: *Inequalities III* (O. Sisha ed.), Academic Press, New York, 1972, 159–175.

[25] V. Klee and D. Walkup, The d-step conjecture for polyhedra of dimension $d < 6$, *Acta Math.* **133** (1967), 53–78.
[26] J.-L. Loday, Realization of the Stasheff polytope, *Arch. Math.* **83** (2004), 267–278.
[27] B. Matschke, F. Santos, and C. Weibel, The width of a 5-dimensional prismatoid, *Proc. London Math. Soc.* **110** (2015, 647–672.
[28] P. McMullen, The maximum numbers of faces of a convex polytope, *Mathematika* **17** (1970), 179–184.
[29] J. M. Pallo, Enumerating, ranking and unranking binary trees, *Computer J.* **29** (1986), 171–175.
[30] 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.
[31] A. Postnikov, Permutahedra, associahedra, and beyond, *IMRN* **2009** (2009), 1026–1106.
[32] D. Quillen, Homotopy properties of the poset of non-trivial p-subgroups of a group, *Advances in Mathematics*, **28** (1978), 101–128.
[33] N. Reading, Lattices congruences of the weak order, *Order*, **21** (2004) no. 4, 315–344.
[34] N. Reading, Cambrian lattices, *Adv. Math.* **205** (2006), no. 2, 313–353.
[35] N. Reading, Lattice congruences, fans, and Hopf algebra, *J. Combin. Theory, Ser. A* **110** (2005), no. 2, 237–273.
[36] V. Reiner, The generalized Baues problem, chapter in: New Perspectives in Algebraic Combinatorics, *Mathematical Sciences Research Institute Publications* **38**, 293–336.
[37] F. Santos, A counterexample to the Hirsch conjecture, *Annals of Mathematics (2)* **176** (2012), 383–412.
[38] R. P. Stanley, *Enumerative Combinatorics, vol. 1*, Cambridge Studies in Advanced Mathematics, no. 49, Cambridge University Press, Cambridge, 1997.
[39] G. Ziegler, Lectures on Polytopes, Graduate Texts in Mathematics **152**, Berlin, New York, Springer, 1995.

Department of Mathematics, North Carolina State University, Raleigh, NC 27695 and Department of Mathematics, University of Oregon, Eugene, OR 97403

E-mail address: `plhersh@ncsu.edu and plhersh@uoregon.edu`