Realizability of Polytopes as a Low Rank Matrix Completion Problem

Michael Gene Dobbins

Abstract This article gives necessary and sufficient conditions for a given relation to be the incidence relation between the facets and vertices of a polytope, or equivalently for a given poset to be the face lattice of a polytope. These conditions pair non-emptiness of the real algebraic variety that parameterizes the combinatorial polytope's realization space with simple combinatorial conditions. Similar realizability conditions are also given for polytopes in geometric spaces (spherical and hyperbolic) in terms of their dihedral angles.

Keywords Polytope · Realizability · Steinitz's Theorem · Matrix completion

1 Introduction

Given a polytope, there are various kinds of combinatorial data we could use to describe it. The purpose of this article is to go in the other direction. Specifically, we would like to know if a given relation can be realized as the incidence relation between a polytope's facets and vertices. As a combinatorial invariant, this facet–vertex incidence relation is equivalent to a polytope's face lattice, the partial ordering of a its faces by inclusion, and defines the polytope's combinatorial type. When such a realization exists for a given poset, we call the poset a combinatorial polytope.

An early realizability result for polytopes was given by Steinitz, who characterized 3-dimensional combinatorial polytopes by showing that a graph can be realized as the 1-skeleton of some 3-dimensional polytope if and only if the graph is planer and 3-connected [17]. Note that the combinatorial type of a 3-dimensional polytope is determined by its 1-skeleton, but this is not the case in higher dimensions. While there

M. G. Dobbins
GAIA, POSTECH, Daejeon and Pohang, South Korea
e-mail: michaelgenedobbins@gmail.com

are impediments to such a purely combinatorial characterization in higher dimensions [14], algebraic characterizations are known [1,3,8]. Here, the algebraic conditions characterizing combinatorial polytopes are simplified to the non-emptiness of a natural parameterization of the combinatorial polytope's realization space by imposing additional combinatorial conditions.

The main result of this article, Theorem 1 in Sect. 3, gives necessary and sufficient conditions for a given relation to be the facet–vertex incidence relation of a polytope in two parts, a combinatorial part and an algebraic part. The combinatorial part gives conditions on the Dedekind–MacNeille completion of the given relation, which we call the maxbiclique lattice. For a facet–vertex incidence relation of a polytope, this lattice is isomorphic to the polytope's face lattice. For the algebraic part, we associate a matrix to any centered polytope, called the facet–vertex matrix. The entries of this matrix are the inner products of the vertices of the polytope with the vertices of its polar polytope. Theorem 1 gives conditions for a matrix to be the facet–vertex matrix of a polytope of the specified combinatorial type, namely that the matrix has the expected rank and entries satisfying the obvious bounds. Theorem 1 also shows that facet–vertex matrices parameterize the space of centered realizations of a combinatorial type of polytope modulo linear transformations, its linear realization space.

Homogenizing a polytope gives a cone one dimension higher, and intersecting this cone with an appropriate affine hyperplane returns a projectively equivalent polytope. This makes the problem of realizing a combinatorial type of polytope equivalent to the problem of realizing the corresponding combinatorial type of cone. We associate a similarly defined facet–ray matrix to any pointed polytopal cone. Theorem 5 in Sect. 3, gives similar conditions for a relation to be realizable in terms of facet–ray matrices, and shows that such matrices (up to positive scaling of rows and columns) parameterize the space of realizations of a combinatorial polytope modulo projective transformations, its projective realization space.

For a polytope defined by linear constraints, these matrices are related to the polytope's slack matrix, the matrix of slack in each linear constraint on each vertex. Slack matrices of polytopes have applications in linear programming and combinatorial optimization. Yannakakis showed that low non-negative rank of a given polytope's slack matrix is equivalent to the polytope being the projection of some higher dimensional polytope with few facets. This can be used to determine if an optimization problem can be expressed as a small linear programming problem [4,7,19].

The algebraic conditions of Theorem 1 are the same as the conditions Robertson gave for a small perturbation of a polytope's facet–vertex matrix to be the facet–vertex matrix of a perturbed polytope of the same combinatorial type [15]. The algebraic conditions of Theorem 5 are the same as those Díaz [6] gave to characterize the facet–ray matrices of a given combinatorial polytope. The results of Robertson and Díaz, however, both assume we start with a combinatorial polytope. If a given cell complex is homeomorphic to a sphere and the inclusion relation of its cells is a lattice, then the algebraic conditions imply the lattice is realizable, but these algebraic conditions alone are not enough to decide if an arbitrary relation or poset is realizable. In this sense, the combinatorial conditions are an alternative to the topological condition of being a CW-sphere, which is algorithmically undecidable in general [18].

The combinatorial conditions of Theorem 1 are closely related to the axioms of abstract polytopes. Abstract polytopes are a larger class of posets that includes combinatorial polytopes, but are defined in a purely combinatorial way [10]. The conditions used here are properties of the poset's flag graph [12]. If we want to consider purely combinatorial objects resembling combinatorial polytopes, more faithful objects have been studied, such as shellable lattices [20] or Eulerian lattices [16]. Even among posets that satisfy both of these properties, only a vanishingly small portion are actually combinatorial polytopes [2,11]. These stronger properties could provide alternative combinatorial conditions for Theorem 1, but here we are interested in weakening the conditions needed for realizability.

The motivation Díaz gave for characterizing the facet–ray matrix of a combinatorial type of cone was for use as a lemma to characterize the Gramian matrix of combinatorial polytopes in other geometric spaces, such as spherical or hyperbolic space. In Sect. 4, Theorem 5 is used instead to give realizability conditions in terms of a polytope's Gramian. The Gramian has the geometric appeal that its entries are the cosine of the dihedral angles of the polytope.

Section 2 gives basic definitions with some terms in *italic*. Section 3 gives necessary and sufficient conditions for a relation to be realizable and shows that the facet–vertex (facet–ray) matrices parameterize the linear (projective) realization spaces of polytopes. Section 4 gives realizability conditions in terms of the Gramian of a polytope in more general geometric spaces. Finally, Section 5 ends with some further remarks about related results, complexity, and Gale duality.

2 Definitions

Here we consider a relation to be a subset of the product of two sets and use terminology for bipartite graphs where convenient. An *induced biclique* (I_0, J_0) of a relation $\mathcal{R} \subset I \times J$ is a pair of sets $I_0 \subset I$, $J_0 \subset J$ such that every element of one part is incident to every element of the other part $I_0 \times J_0 \subset \mathcal{R}$. A *maxbiclique* of a relation is a maximal induced biclique. The *maxbiclique lattice* of a relation on ordered parts (here J is specified as the lower part) is the poset consisting of maxbicliques ordered by containment of elements in the lower part, $(I_0, J_0) \le (I_1, J_1) := J_0 \subset J_1$.

For a pair of elements a, b of a lattice, we denote the meet by $a \wedge b$ and the join by $a \vee b$. We denote the greatest element by $\top$ and least element by $\bot$. A meet (join) irreducible is an element of a lattice that cannot be expressed as the meet (join) of other elements. Lemma 2 below says that sending a relation to its maxbiclique lattice and sending a lattice to its irreducibles are inverse operations. A *flag* of a lattice, or more generally of a poset, is a maximal totally ordered subset. Here we will generally assume posets are graded, all flags have the same length. The *rank* of an element of a graded poset is the number of elements below it in a flag, and the rank of the poset is that of its flags. Note that we begin counting at 0, so the rank of the singleton poset is 0. The *flag graph* of a poset is a graph consisting of a node for each flag $\mathcal{F}$, and edges between pairs of flags that differ in only one element. That is, $\mathcal{F}$ and $\mathcal{F}'$ are neighbors when $|\mathcal{F} \setminus \mathcal{F}'| = 1$. Figure 1 shows the flag graph of a polytope, which is the flag graph of its face lattice. When the flag graph of a graded poset is connected we
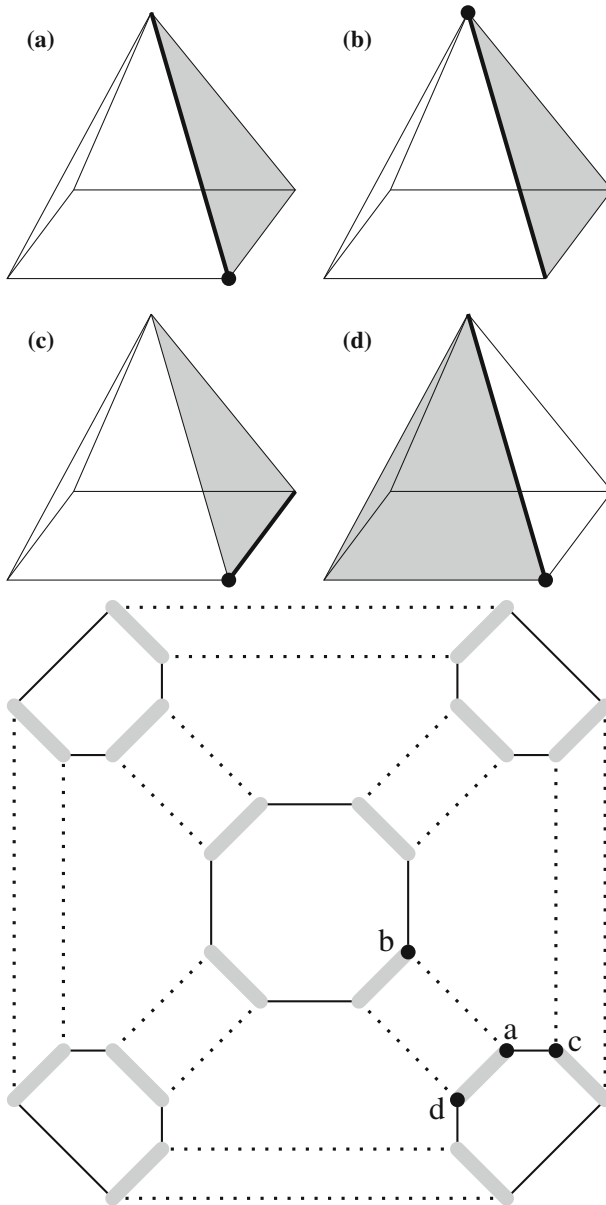
Fig. 1 *Top* four flags of a polytope. *Bottom* the polytope's flag graph with edge between flags that differ by a vertex, edge, and facet shown *dotted*, *thin*, and *thick gray* respectively with vertices corresponding to the 4 flags indicated

say the poset is *flag-connected*. An interval $[a, b]$ of a poset is the set of all elements between two particular elements a, b of the poset, $[a, b] = \{c : a \leq c \leq b\}$. Finally, a *diamond* poset is a graded poset where every rank 2 interval has 4 elements. We will use the same terminology as for face lattices of polytopes to refer to elements of a

lattice. That is, elements of a lattice are called faces, atoms are vertices, and coatoms are facets.
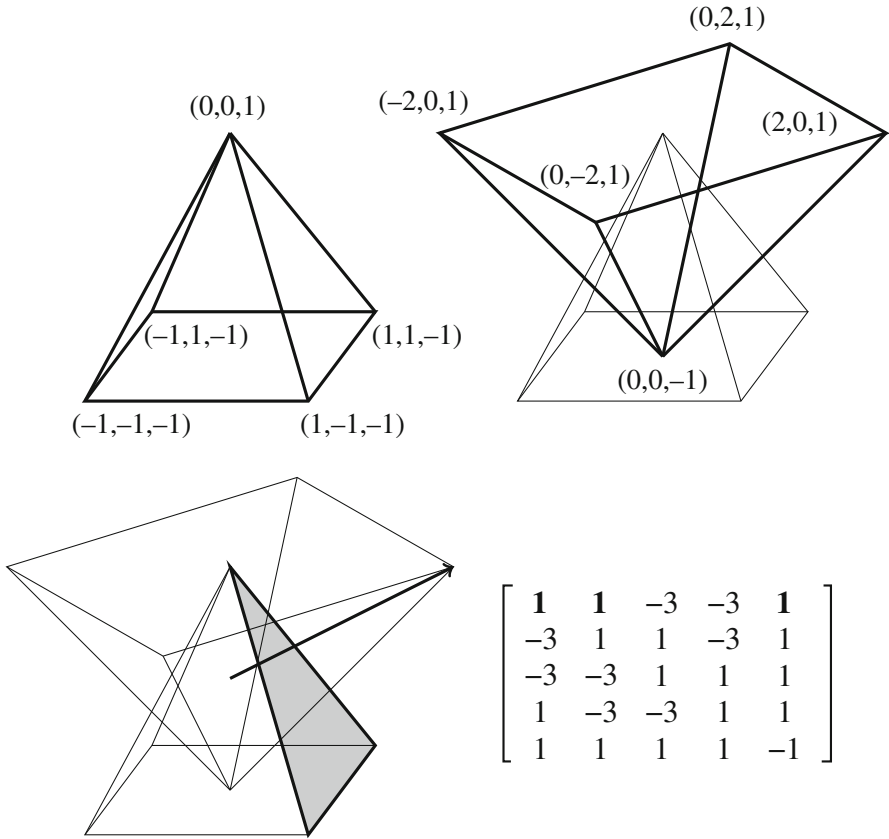
We denote the ith entry of a vector v by $[v]_i$, and the (i, j)th entry of a matrix A by $[A]_{i,j}$. We denote the origin by $\mathbf{0}$, the ith standard basis vector by e_i, the vector of ones by $\mathbf{1}$, the identity matrix by I, and the matrix of ones by E, and for a vector v, denote the diagonal matrix with diagonal entries $[v]_1, [v]_2, \ldots$ by $\operatorname{diag}(v)$. A *filled x-incidence matrix A* of a relation $\mathcal{R} \subset I \times J$ between row indices I and column indices J is a matrix with entries $[A]_{i,j} = x$ for $(i, j) \in \mathcal{R}$ and $[A]_{i,j} < x$ for $(i, j) \notin \mathcal{R}$. For a matrix with compact singular value decomposition $A = U \Sigma V^*$, its pseudoinverse is $A^+ = V \Sigma^{-1} U^*$. The square root is defined for any positive symmetric semi-definite matrix by taking the square root of its singular values. That is, for a non-negative diagonal matrix D, $D^{1/2} = \operatorname{diag}(([D]_{1,1}^{1/2}, [D]_{2,2}^{1/2}, \ldots))$, and for a matrix with compact singular value decomposition $A = U \Sigma U^*$, $A^{1/2} = U \Sigma^{1/2} U^*$.

For now, a polytope is the convex hull of finitely many points in $\mathbb{R}^d$. In Sect. 4 we will consider a more general notion of polytope. A polytope is centered when it contains the origin, and the *covertices* of a full-dimensional centered polytope P are the vertices of its polar. The *facet–vertex matrix M* of a centered polytope P with vertices w_j and covertices h_i has entries $[M]_{i,j} = \langle h_i, w_j \rangle$. See Fig. 2 for an example. A polytope is type $\mathcal{R}$ when its facets and vertices are indexed by the respective parts of $\mathcal{R}$ and its facet–vertex incidence is $\mathcal{R}$. Clearly the facet–vertex matrix of a type $\mathcal{R}$ polytope is a rank d filled 1-incidence matrix of $\mathcal{R}$. Theorem 1 will give conditions on $\mathcal{R}$ for any rank d filled 1-incidence matrix to be a facet–vertex matrix. Note that the transpose of the facet–vertex matrix of a centered polytope is the facet–vertex matrix of its polar polytope. We use the convention that rows correspond to covertices and columns to vertices, since these are respectively covariant and contravariant under change of coordinates. The realization space of a d-dimensional type $\mathcal{R}$ polytope is the subspace of $\mathbb{R}^{d|J|}$ corresponding to all families of points $\{x_j\}_{j \in J} \subset \mathbb{R}^d$ that are the respective vertices of a type $\mathcal{R}$ polytope. The centered realization space further restricts this to centered realizations. The linear realization space of a polytope is the quotient of its centered realization space by linear transformations on $\mathbb{R}^d$, and the projective realization space is the quotient of its realization space by projective transformations.

A *cone C* is the set of positive linear combinations of vectors in $\mathbb{R}^{d+1}$ that does not contain a line. That is, by 'cone' we mean what may be called a pointed polytopal cone with vertex at the origin. A set of *generators* of C is a minimal set of vectors such that C consists of positive linear combinations of these vectors, and a set of *cogenerators* is a set of generators of its polar. The *facet–ray matrix N* of a cone C with generators w_j and cogenerators h_i has entries $[N]_{i,j} = \langle h_i, w_j \rangle$. Note that the generators and cogenerators of a cone are defined up to positive scaling of each vector. As such, a cone does not have a unique facet–ray matrix. We will, however, often assume that a minimal set of generators and cogenerators are given.

3 Realizability

Here is the main theorem for centered polytopes and their linear realization spaces.

$$\begin{bmatrix} \mathbf{1} & \mathbf{1} & \mathbf{-3} & \mathbf{-3} & \mathbf{1} \\ -3 & 1 & 1 & -3 & 1 \\ -3 & -3 & 1 & 1 & 1 \\ 1 & -3 & -3 & 1 & 1 \\ 1 & 1 & 1 & 1 & -1 \end{bmatrix}$$

Fig. 2 *Top from left* A 3-polytope and its polar with vertex coordinates. *Bottom* A covertex with the corresponding facet shaded and the facet–vertex matrix with the corresponding row in *bold*

Theorem 1 *Given a relation $\mathcal{R}$, there is a polytope with facet–vertex incidence $\mathcal{R}$ if and only if both the maxbiclique lattice of $\mathcal{R}$ is flag-connected, rank $d + 1$, and diamond, and $\mathcal{R}$ has a rank d filled 1-incidence matrix.*

Moreover, for such $\mathcal{R}$, M is the facet–vertex matrix of a centered type $\mathcal{R}$ polytope if and only if M is a rank d filled 1-incidence matrix of $\mathcal{R}$, and these matrices parameterize the linear realization space of centered type $\mathcal{R}$ polytopes.

While the face lattice of a polytope includes the facet–vertex incidence relation $\mathcal{R}$, the following lemma shows that these objects carry the same information. A proof of this lemma can be found in [5].

Lemma 2 *Any lattice where all flags are finite is isomorphic to the maxbiclique lattice of the comparability relation between its meet and join irreducibles.*

For polytopes, this means that each face of a polytope can be uniquely determined from the vertices it contains and the facets it is contained in, and a set of facets and vertices are those of a face if and only if the following two conditions hold: (biclique)

all the vertices of the face are contained in all its facets, and (maximal) this first condition fails if any other vertex or facet is added to those of the face.

The proof of Theorem 1 uses ideas that are similar to what Díaz used in [6], but gives combinatorial rather than geometric arguments. The point of this is to make explicit the purely combinatorial conditions that are needed. The proof works by constructing a polytope from a filled 1-incidence matrix of a given relation and showing that the maxbiclique lattice of this relation can be embedded in the face lattice of this polytope. The following lemma shows that such an embedding is actually an isomorphism, so exhibiting this embedding is enough to show that the polytope realizes the given relation.

Lemma 3 *An order embedding between flag-connected diamond lattices of the same finite rank is an isomorphism.*

Proof Suppose the lemma fails, then there are flag-connected diamond lattices $\mathcal{P}$ and $\mathcal{Q}$ of rank $d+1$ with an embedding $\varphi : \mathcal{P} \to \mathcal{Q}$ such that there is some element $f \in \mathcal{Q}$ that is not in $\varphi(\mathcal{P})$. Without loss of generality we assume that $\mathcal{P} \subset \mathcal{Q}$ and the embedding is the identity, otherwise just replace $\mathcal{P}$ with its image. Consider now the respective flag graphs G, H of $\mathcal{P}$, $\mathcal{Q}$. Every flag of $\mathcal{P}$ is a flag of $\mathcal{Q}$, and whether two flags of $\mathcal{P}$ differ by one element is independent of whether they are regarded as flags of $\mathcal{P}$ or as flags of $\mathcal{Q}$, so G is an induced subgraph of H. The lattice $\mathcal{P}$ is non-empty, and f must belong to some flag of $\mathcal{Q}$, so $\emptyset \neq G \subsetneq H$. The diamond condition is equivalent to each flag having 1 neighbor in the flag graph for each element except $\top$ and $\bot$, and no other neighbors. This means G is a d-regular proper induced subgraph of a connected d-regular graph, H, which is impossible. To see why, consider a path of H from a node that is in G to a node that is not in G. The last node of this path that is in G must have d neighbors total, of which at least 1 neighbor must not be in G and d neighbors must be in G. □

To show that the conditions for realizability given in Theorem 1 are sufficient, we construct a d-polytope P from a decomposition of a rank d filled 1-incidence matrix $M = H^*W$ of a relation $\mathcal{R}$ satisfying the specified combinatorial conditions into a $n \times k$ matrix H^* and a $k \times m$ rank d matrix W. Specifically, P's vertices are the columns of W, and if $k = d$ its covertices are the columns of H. We then see that the facet–vertex incidence of the constructed polytope P is $\mathcal{R}$. To make this construction more explicit, we use the compact singular value decomposition to get H, W and their pseudoinverses H^+, W^+.

If we start with a polytope, find its facet–vertex matrix M, and do the above construction, we get back a linear copy of the polytope we started with. This is because, if the columns of a matrix X are the vertices of a polytope P and W is any matrix with the same null space as X, then $X = AW$ factors though W by a matrix $A = XW^+$ that does not change the rank, so the columns of W are the vertices of some linear copy of P. We can do the above construction by finding a bases for $\text{null}(M)^\perp = \text{null}(X)^\perp$, which are then the rows of W, to get a linear copy of P in $\mathbb{R}^d$. Alternatively, the trivial decomposition $M = I^*M$ gives the following corollary, which is seen by the proof of Theorem 1.

Corollary 4 *The columns of the facet–vertex matrix M of a polytope P are the vertices of a linear copy of P.*

We would, however, like to start with just a matrix M instead of a polytope.

Proof of Theorem 1 We have the 'only if' direction of the first part of the theorem immediately. To show that any relation $\mathcal{R}$ satisfying the combinatorial conditions with a rank d filled 1-incidence matrix M can be realized as the facet–vertex incidence of a polytope, let $\mathcal{P}$ be the maxbiclique lattice of $\mathcal{R}$, which by hypothesis is rank $d + 1$, flag-connected, and satisfies the diamond condition, let H and W be matrices such that W has rank d and $M = H^*W$, let $\{h_i\}_{i \in [n]}$ and $\{w_j\}_{j \in [m]}$ be the columns of H and W respectively, and let $P = \text{conv}(\{w_j\}_{j \in [m]})$.

We will show that P is a realization of $\mathcal{R}$. Since W has rank d, P has dimension at most d. Furthermore, w_j is in the hyperplane $h_i^{=1} := \{x : \langle h_i, x \rangle = 1\}$ for $(i, j) \in \mathcal{R}$, but is in the open half-space $h_i^{<1} := \{x : \langle h_i, x \rangle < 1\}$ for $(i, j) \notin \mathcal{R}$. To be more explicit we could let $H = I$ and $W = M$ for example, in which case $h_i^{=1}$ becomes $[x]_i = 1$ and $h_j^{<1}$ becomes $[x]_i < 1$.

We will construct a map from $\mathcal{P}$ to the face lattice of P and show that this map is an isomorphism. For $a = (I_a, J_a) \in \mathcal{P}$, let $\varphi(a) = \bigcap_{i \in I_a} h_i^{=1} \cap P$ be the face of P we get by intersecting P with the hyperplanes corresponding to covertices of a. This must be a face of P since $h_i^{=1}$ are all supporting hyperplanes of P.

First we will see that φ preserves order. That is, we will show the very strong condition that $\varphi(a) \subseteq \varphi(b)$ if and only if $a \leq b$. Suppose $a \leq b$, then $I_a \subseteq I_b$ and $\bigcap_{i \in I_a} h_i^{=1} \subseteq \bigcap_{i \in I_b} h_i^{=1}$ so $\varphi(a) \subseteq \varphi(b)$. For the other direction suppose $a \nleq b$, then there is some $j \in J_a$ but $j \notin J_b$, so $w_j \in \varphi(a)$ but $w_j \notin \varphi(b)$ and $\varphi(a) \nsubseteq \varphi(b)$. Thus order is preserved.

We will also see that φ is an injection, and therefore an embedding. To see this consider a pair of faces a, b of $\mathcal{P}$ that map to the same face $\varphi(a) = \varphi(b) = f$ of P. With this, $w_j \in f \subset h_i^{=1}$ for any $i \in I_a \cup I_b$ and $j \in J_a \cup J_b$, so $[M]_{ij} = 1$ and $(i, j) \in \mathcal{R}$. This means $a = (I_a, J_a)$ and $b = (I_b, J_b)$ are maxbicliques that are subsets of the same biclique $c = (I_a \cup I_b, J_a \cup J_b)$, so they must be equal $a = b = c$, which makes φ an embedding.

This induces an injection from a flag $\perp < a_1 < \cdots < a_d < \top \in \mathcal{P}$ to a totally ordered set of P's faces $\varphi(\perp) = \emptyset < \varphi(a_1) < \cdots < \varphi(a_d) < \varphi(\top) = P$, which must be of the same size or less. Since a larger set cannot be injected into a smaller one, these flags must be the same size, so P must have dimension d. Now φ is an embedding between flag-connected diamond lattices of the same rank, and by Lemma 3 is therefore an isomorphism. Thus, P is a realization of $\mathcal{R}$ and the 'if' direction of the first part of the theorem holds.

For the second part of the theorem, we consider realizations in $\mathbb{R}^d$. For this, let $M = U\Sigma V^*$ be the compact singular value decomposition of M and let $H = \Sigma^{1/2}U^*$ and $W = \Sigma^{1/2}V^*$. From what we have just seen, M is the facet–vertex matrix of the polytope P with vertices given by the columns of W. Suppose we have some other centered polytope in $\mathbb{R}^d$ with vertices and covertices respectively given by the columns of matrices W_0 and H_0 such that $M = H_0^*W_0$. The rows of both W and W_0 give a basis of the same space, $\text{null}(M)^\perp$, so there is an invertible linear transformation between them $W_0 = AW$, $A = W_0W^+ = W_0V\Sigma^{-1/2}$. Similarly we have $A^{-*} = H_0H^+ = H_0U\Sigma^{-1/2}$ and $H_0 = A^{-*}H$, and for any nonsingular matrix A we get such a decomposition $M = (A^{-*}H)^*AW$. This means that any centered type $\mathcal{R}$

polytope with facet–vertex matrix M is a linear copy of P. Applying an invertible linear transformation to P will not change the facet–vertex matrix of the resulting polytope, so M parameterizes centered realizations of $\mathcal{R}$ modulo linear transformations.   □

Here is the main theorem for cones and the projective realization spaces of polytopes.

Theorem 5 *If the maxbiclique lattice of $\mathcal{R}$ is flag-connected, rank $d+1$, and diamond, then $\mathcal{R}$ has a rank d filled 1-incidence matrix if and only if $\mathcal{R}$ has a rank $d + 1$ filled 0-incidence matrix N.*

Moreover, for such $\mathcal{R}$, N is the facet–ray matrix of a type $\mathcal{R}$ cone if and only if N is a rank $d + 1$ filled 0-incidence matrix, and equivalence classes of these matrices modulo positive scaling of rows and columns parameterize the projective realization space of type $\mathcal{R}$ polytopes.

Proof For the 'only if' direction, Theorem 1 gives us a polytope realizing $\mathcal{R}$ and a facet–ray matrix of the cone over this polytope in homogeneous coordinates gives us a rank $d + 1$ 0-incidence matrix for $\mathcal{R}$. For the rest of the proof, we exclude the case where $\mathcal{R}$ is a singleton or empty. We will show that we can use a rank $d + 1$ filled 0-incidence matrix N for an appropriate relation $\mathcal{R}$ to find a rank d filled 1-incidence matrix M for $\mathcal{R}$. We first note that $N + E$ (where E is the matrix of ones) is a filled 1-incidence matrix that and can differ from N in rank by at most 1. If the rank of $N + E$ were 1 less than that of N we would be done, but that might not be the case. Instead we will find full diagonal matrices D_1, D_2 such that $M = D_1 N D_2 + E$ has rank d. We start by letting $N = U \Sigma V^*$ be the compact singular value decomposition and U_0 and V_0 be respective completions of U and V to orthogonal matrices.

Choose some vector $y \in \mathbb{R}_{>0}^{d+1}$ in the positive orthant such that the entries of $V y$ are all non-zero. That is, so y is not orthogonal to any row of V. Such a vector must exist unless V has a row of all zeros. In that case N would have a column of all zeros, which means some meet irreducible of $\mathcal{R}$ would be incident to all join irreducibles, and therefore the only meet irreducible, but $\mathcal{R}$ is not a singleton so the diamond condition implies there is more than one meet irreducible. Thus, we can find a vector y with all positive entries such that $V y$ has all non-zero entries, and similarly we can find such a vector x for U. Furthermore, we can choose these so that $\langle x, \Sigma^{-1} y \rangle = 1$.

With this in mind, let $D_1 = \mathrm{diag}(-Ux)^{-1}$, $D_2 = \mathrm{diag}(V y)^{-1}$, and $M = D_1 N D_2 + E$. We preform a change of bases to make the rank of the resulting matrix more apparent.

$$
\begin{aligned}
(U_0^* D_1^{-1}) M (D_2^{-1} V_0) &= U_0^* N V_0 + U_0^* D_1^{-1} E D_2^{-1} V_0 \\
&= U_0^* U \Sigma V^* V_0 + U_0^* D_1^{-1} \mathbf{1}\mathbf{1}^* D_2^{-1} V_0 \\
&= U_0^* U (\Sigma - x y^*) V^* V_0 \\
&= \begin{bmatrix} (\Sigma - x y^*) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}.
\end{aligned}
$$

Under this change of bases we see e_i is sent to $\mathbf{0}$ for $i > d + 1$, so the rank is at most $d + 1$. To observe that the rank is in fact d, consider only the first $d + 1$ coordinates. We show that the columns are orthogonal to $\Sigma^{-1} y$.

$$\langle (\Sigma - xy^*)e_i, \Sigma^{-1}y \rangle = \langle \Sigma e_i - [y]_i x, \Sigma^{-1}y \rangle$$
$$= [y]_i - [y]_i = 0.$$

Thus M has rank d and the 'if' direction holds.

For the second part of the theorem, recall that two polytopes P and Q are projectively equivalent if and only if the cones over these polytopes in homogeneous coordinates are linearly equivalent. That is for some A invertible,

$$A \left(\mathbb{R}_{\geq 0} \begin{bmatrix} P \\ 1 \end{bmatrix} \right) = \mathbb{R}_{\geq 0} \begin{bmatrix} Q \\ 1 \end{bmatrix}.$$

As we have seen, full rank matrices H and W of a decomposition $N = H^*W$ are determined up to linear transformation by N, so N determines a set of generators and cogenerators of C up to linear transformation. Also, a set of generators and cogenerators determine the same cone if and only if they are equivalent up to positive scaling of each, which amounts to scaling the rows and columns of N. In this way N modulo positive scaling of rows and columns uniquely determines a polytope up to projective transformation.

As a final remark, note that the diagonal matrices D_i above scale the generators and cogenerators of a cone to be the vertices of a polytope and the vertices of a dual polytope respectively. These polytopes are the intersection of the cone and its polar each with an affine hyperplane, and the choice of x and y amount to choosing these hyperplanes, which in turn amounts to choosing the origin and infinite horizon defining a vector space in projective space.   □

The flag graph of a polytope may be much larger than its face lattice, but is highly structured. The following theorem shows that we do not have to consider the full flag graph when deciding realizability. For diamond lattices, the *upper graph* of an element a of rank k is the graph having a node for each rank $k + 1$ element and an edge for each rank $k + 2$ element with the same incidences as those of the lattice. Note that for a face lattice, the upper graph of a is the 1-skeleton of the face figure of a.

Theorem 6 *For a diamond lattice, if the upper graph of every element is connected, then the lattice is flag-connected.*

Proof Consider two flags $\mathcal{F} = [\perp, f_0, \ldots, f_{d-1}, \top], \mathcal{G} = [\perp, g_0, \ldots, g_{d-1}, \top]$ of a diamond lattice satisfying the hypothesis. We will build a path from $\mathcal{F}$ to $\mathcal{G}$ in the flag graph inductively from vertices up to facets by first walking along the upper graph of $\perp$ to get a sequence of vertices, then walking along the upper graph of each vertex to insert a sequence of edges, etc. Then, we generate a sequence of flags from these sequences of faces. At each step of the induction we get the following. For $c \in [0, d-1]$ we claim there is a sequence of faces $H_c = [h_{c,1}, \ldots, h_{c,t_c}]$ such that for every $k \in [0, c]$ the subsequence of rank k faces begins with f_k, ends with g_k, and for any two consecutive faces $h_{c,i}, h_{c,j}$ in the subsequence the following holds. $h_{c,i} \wedge h_{c,j}$ is last face of rank $k - 1$ in the sequence before $h_{c,j}$ if $k > 0$, and $h_{c,i} \vee h_{c,j}$ is the last face of rank $k + 1$ in the sequence before $h_{c,j}$ if $k < c$ and is some face of rank $c + 1$ if $k = c$.

Before seeing why the claim holds, we will see why it implies the lattice is flag-connected. Consider the sequence H_{d-1} after the initial faces of $\mathcal{F}$ in this sequence, and set $h_i = h_{d-1,i} \in H_{d-1}$. Starting from $\mathcal{H}_0 = \mathcal{F}$ we get a sequence of flags $\mathcal{H}_i = \mathcal{H}_{i-1} \cup h_i \setminus h_j$ where h_j is the face of $\mathcal{H}_{i-1}$ having the same rank as h_i. Note that h_j is also the last face in H_{d-1} before h_i of the same rank. By the claim, $h_i \wedge h_j$ is the face of $\mathcal{H}_i$ with rank $k-1$ and $h_i \vee h_j$ is the face of $\mathcal{H}_i$ with rank $k+1$, so this is a path in the flag graph from $\mathcal{F}$ to $\mathcal{G}$.

Now we will see why the claim holds. The case $c = 0$ is exactly the the upper graph of $\perp$ being connected. Assume by induction that the claim hold for the case c and let $h_{c,x}, h_{c,y}, h_{c,z}$ be consecutive faces in the subsequence of faces with rank c. Since the upper graph of $h_{c,y}$ is connected, there is a sequence $E_y = [e_{y,1}, \ldots e_{y,s_y}]$ such that $e_{y,1} = h_{c,x} \vee h_{c,y}, e_{y,s_y} = h_{c,y} \vee h_{c,z}, e_{y,i} \wedge e_{y,i+1} = h_{c,y}$, and $e_{y,i} \vee e_{y,i+1}$ has rank $c + 1$. Inserting E_y after each face $h_{c,y}$ of rank c in the sequence H_c gives the desired sequence H_{c+1}. By induction we get a sequence H_{d+1} that determines a path in the flag graph between $\mathcal{F}$ and $\mathcal{G}$, and since this holds for any pair of flags, the lattice is flag-connected. $\square$

4 The Gramian

So far we have considered realizability in Euclidean space in terms of the facet–vertex or facet–ray matrices, now we will consider realizations in other geometric spaces in terms of the Gramian matrix of a polytope. A geometric space is a level set of a bilinear form ϕ, which is a function on two vectors of the form $\phi(x, y) = x^*\Phi y$ for some matrix Φ. If $\Phi = I$ is the identity matrix, for example, then $\phi(x, x) = \langle x, x \rangle$ is the standard inner product and $\phi(x) = \|x\|^2$ is the standard norm squared. A d dimensional affine space $\mathbb{A}^d$ may be represented in homogeneous coordinates as a non-zero level set of a linear functional on $\mathbb{R}^{d+1}$, and polyhedra in $\mathbb{A}^d$ are given by the intersection of $\mathbb{A}^d$ with a cone in $\mathbb{R}^{d+1}$, not necessarily pointed. Similarly, a geometric space $\mathbb{X}$ is given by a connected component of a non-zero level set of ϕ, and polyhedra in $\mathbb{X}$ are given by the intersection of $\mathbb{X}$ with a cone.

We will mostly be concerned with realizing cones. Recall the generators and cogenerators of a given cone are not unique. However, we can use ϕ to define a unique set of cogenerators for certain cones as follows. A hyperplane is called lightlike when ϕ vanishes on the hyperplane away from the origin. A half-space bounded by a hyperplane that is not lightlike has the form $h^{\phi, \leq 0} = \{x : \phi(h, x) \leq 0\}$ for a unique outward normal vector h, with $|\phi(h)| = 1$. Observe that if $\phi(h, h) = 0$ then h is in the hyperplane $h^{\phi, =0} = (\Phi h)^{=0}$. For a cone C such that no facet supporting hyperplane is lightlike, the outward normals h_i of the facet supporting half-spaces are cogenerators and the *Gramian* G of C with respect to ϕ has entries $[G]_{i,j} = \phi(h_i, h_j)$.

We will refer to elements of a lattice as the corresponding object in the face lattice of a polytope. A *cycle* is a sequence of facets $F_{i_1}, \ldots, F_{i_d}$ such that $\Lambda_{k=1}^t F_{i_k}$ is a $(d-t)$-face, and a *super cycle* is a cycle with an additional facet such that $\Lambda_{k=1}^{d+1} F_{i_k} = \perp$. Note that while Díaz called these maximal cycles, we will call these super cycles instead to emphasize that they are not cycles. A cycle or super cycle induces a flag with faces $\Lambda_{k=1}^t F_{i_k}$ for $1 \leq t \leq d$, $\top$, $\perp$. Note the flag graph of a combinatorial polytope $\mathcal{P}$ is

always bipartite, with bipartition given for each flag by the orientation of the outward normal vectors of the facets of a polytope realizing $\mathcal{P}$ in a cycle inducing that flag. With this in mind, we will require the flag graph to be bipartite and say two super cycles have the same *orientation* when the induced flags are in the same bipart. We denote the minor of a matrix G with rows $i_1 \cdots i_k$ and columns $j_1 \cdots j_l$ by $[G]_{j_1 \cdots j_l}^{i_1 \cdots i_k}$. The *signature* of a matrix, or bilinear form, is the pair (p, n) where p and n are the total ranks of eigenspaces with positive and negative eigenvalues respectively. For Φ with signature $(d+1, 0)$, such as with $\Phi = I$, a level set $\mathbb{X}$ has spherical geometry, and for signature $(d, 1)$, $\mathbb{X}$ has hyperbolic geometry. A linear transformation Q is an *orthogonal* transformation of ϕ if it preserves ϕ, $Q^*\Phi Q = \Phi$. In the case $\Phi = I$, these are the usual orthogonal transformations. The following theorem was proven by Díaz [6].

Theorem 7 *Given a combinatorial d-polytope $\mathcal{P}$ with n facets and a symmetric bilinear form ϕ, there is a type $\mathcal{P}$ cone with Gramian G if and only if G is an $n \times n$ symmetric matrix with diagonals ± 1 and the same signature as ϕ that satisfies the following:*

1. *For every vertex of $\mathcal{P}$ and all facets $F_{i_1} \cdots F_{i_k}$ incident to it, $\mathrm{rank}([G]_{i_1 \cdots i_k}^{i_1 \cdots i_k}) = d$.*
2. *For every pair of super cycles $F_{i_1}, \ldots, F_{i_{d+1}}$ and $F_{j_1}, \ldots, F_{j_{d+1}}$ with the same orientation, $\det([G]_{j_1 \cdots j_{d+1}}^{i_1 \cdots i_{d+1}})\det(\phi) > 0$.*

   Moreover, G parameterizes the realization space of type $\mathcal{P}$ cones modulo orthogonal transformations of ϕ.

We replace the combinatorial polytope in Díaz's theorem, with a purely combinatorial object.

Theorem 8 *Given an $n \times m$ relation $\mathcal{R}$ and symmetric bilinear form ϕ, there is a type $\mathcal{R}$ polytopal cone with Gramian G if and only if the maxbiclique lattice of $\mathcal{R}$ is a rank $d+1$ diamond lattice with connected bipartite flag graph, and G is an $n \times n$ symmetric matrix with diagonals ± 1 and the same signature as ϕ that satisfies the following:*

1. *For every vertex of $\mathcal{R}$ and all facets $F_{i_1} \cdots F_{i_k}$ incident to it, $\mathrm{rank}([G]_{i_1 \cdots i_k}^{i_1 \cdots i_k}) = d$.*
2. *For every pair of super cycles $F_{i_1}, \ldots, F_{i_{d+1}}$ and $F_{j_1}, \ldots, F_{j_{d+1}}$ with the same orientation, $\det([G]_{j_1 \cdots j_{d+1}}^{i_1 \cdots i_{d+1}})\det(\phi) > 0$.*

Moreover, G parameterizes the realization space of type $\mathcal{R}$ cones modulo orthogonal transformations of ϕ.

The proof of Theorem 8 will follow that given by Díaz [6], except Theorem 5 will be used in place of the corresponding result for combinatorial polytopes. Theorem 5 makes use of both the generators and cogenerators of a cone, but Theorem 8 only makes use of the cogenerators. To deal with this, we use the Hodge star operator to construct generators of a cone from its cogenerators in a way that is analogous to finding the generators of a polygonal cone in $\mathbb{R}^3$ with $\Phi = I$ by taking the cross product of outward normals of neighboring facets. Let $\Lambda^m X$ denote the m-fold exterior power of a real vector space X, and $(\wedge) : \Lambda^n X \times \Lambda^m X \to \Lambda^{n+m} X$ denote concatenation on exterior powers. Given a bilinear form ϕ on X we define a bilinear form on $\Lambda^m X$ by

$$(\Lambda^m \phi)(y_1 \wedge \cdots \wedge y_n, x_1 \wedge \cdots \wedge x_n) := \det(A),$$

where $[A]_{i,j} = \phi(y_i, x_j)$. Let $\star : \Lambda^m X \to \Lambda^{r-m} X$ denote the Hodge star operator. That is, $\star := \tau^{-1} \circ \psi$ where

$$\psi : \Lambda^m X \to (\Lambda^{r-m} X)^*, \quad \psi(u) := v \mapsto \star(u \wedge v)$$

$$\tau : \Lambda^{r-m} X \to (\Lambda^{r-m} X)^*, \quad \tau(x) := v \mapsto (\Lambda^{r-m} \phi)(u, v).$$

In the case of $m = 1$, $\Lambda^1 X = X$, $\Lambda^1 \phi = \phi$ this becomes

$$\phi(\star(x_1 \wedge \cdots \wedge x_{r-1}), x_r) = \star(x_1 \wedge \cdots \wedge x_{r-1} \wedge x_r).$$

Note that in $\mathbb{R}^3$ with $\Phi = I$ this is $\langle x_1 \times x_2, x_3 \rangle = \det([x_1 x_2 x_3])$.

Proof of Theorem 8 Since we are only weakening one side of a biconditional statement in the theorem given by Díaz, we only have to consider the argument showing one direction. That is, we will show that if a relation and a matrix satisfy these conditions, then they can be realized by such a cone. We only consider $d \geq 1$, since the theorem is trivial otherwise. We can represent ϕ by a real symmetric matrix Φ, and since G and Φ are real symmetric matrices with the same signature, there exists an $n \times (d+1)$ matrix H such that $G = H^* \Phi H$. We denote columns of H by h_j, and we claim $C := \bigcap_j h_j^{\phi, \leq 0}$ is a polytopal cone of type $\mathcal{P}$. Following the footsteps of [6], we show this by using the Hodge star operator to find generators of C, which we collect in a matrix W, and show that $N = (\Phi H)^* W = H^* \Phi W$ is a rank $d+1$ filled 0-incidence matrix of $\mathcal{P}$.

With this in mind, choose a cycle $F_{i_{j,1}}, \ldots, F_{i_{j,d}}$ incident to each vertex v_j of $\mathcal{P}$ such that these all have the same orientation, and let

$$w_j = \star(h_{i_{j,1}} \wedge \cdots \wedge h_{i_{j,d}}).$$

Since $\mathcal{P}$ is a diamond lattice with rank at least 2, we can choose another facet $F_{i_{j,d+1}}$ that is not incident to v_j giving a super cycle. By Condition 2 for all j we have $\det([G]_{i_{j,1} \cdots i_{j,d+1}}^{i_{j,1} \cdots i_{j,d+1}}) \neq 0$, so the vectors $\{h_{i_{j,k}}\}_{k \in [d+1]}$ are linearly independent and $h_{i_{j,1}} \wedge \cdots \wedge h_{i_{j,d}}$ is not the origin of $\Lambda^d \mathbb{R}^{d+1}$. Therefore, $w_j \neq \mathbf{0}$, since $\star$ is an isomorphism. We let W have columns w_j and $N = H^* \Phi W$.

We now show that N is a 0-incidence for the relation between the irreducibles of $\mathcal{P}$. For $v_j < F_i$ incident, by Condition 1 we have

$$[N]_{i,j} = \phi(h_i, w_j) = \phi(h_i, \star(h_{i_{j,1}} \wedge \cdots \wedge h_{i_{j,d}})) = \star(h_i \wedge h_{i_{j,1}} \wedge \cdots \wedge h_{i_{j,d}}) = 0,$$

since $\{h_i, h_{i_{j,1}}, \ldots, h_{i_{j,d}}\}$ has span at most d and is therefore linearly dependent. Alternatively consider a pair $v_j \not< F_i$ that are not incident. As we have seen, this

means $\{h_{i_{j,1}}, \ldots, h_{i_{j,d}}, h_i\}$ are linearly independent and therefore a basis of $\mathbb{R}^{d+1}$. We have already that $\phi(h_{i_{j,k}}, w_j) = 0$ for $k \in [d]$, and since ϕ has full rank we cannot have $\phi(x, w_j) = 0$ for every vector x of a basis, so we must have $\phi(h_i, w_j) \neq 0$. Thus $[N]_{i,j} = 0$ if and only if F_i and v_j are incident in $\mathcal{P}$.

We would like to show N is non-positive. We will actually show N is either that or non-negative, which can be fixed by redefining W to be $-W$, so this is enough. We do this by showing all nonzero entries have the same sign. Recall that we chose cycles that have the same orientation, but did not specify which orientation, so we should not expect to have chosen correctly. Let (i, j) and $(\iota, \jmath)$ both be indices of a pair of irreducible of $\mathcal{P}$ that are not incident. By Condition 2 we have the following.

$$\phi(h_i, w_j)\phi(h_\iota, w_\jmath) = \det\left([G]_{i_{\jmath,1}\cdots i_{\jmath,d}\ \iota}^{i_{j,1}\cdots i_{j,d}\ i}\right)\mathrm{sign}(\phi) > 0.$$

This gives us that $N = (\Phi H)^* W$ is a rank $d+1$ filled 0-incidence matrix of $\mathcal{P}$, adjusting W's sign if needed. By Theorem 5, C is a type $\mathcal{P}$ cone and the columns of ΦH are cogenerators of C. Thus, C has outward normals $\{h_i\}_{i\in[n]}$ and Gramian G with respect to ϕ. □

We now state versions the above theorem for spherical and hyperbolic polytopes. Díaz gave proofs of these for combinatorial polytopes, which work just as well in this context, so the proofs are not included here. Let $\mathbb{S}^d = \{x : \|x\| = 1\}$ denote the d-dimensional unit sphere in $\mathbb{R}^{d+1}$. A spherical d-polytope is the intersection of a pointed polytopal cone in $\mathbb{R}^{d+1}$ with $\mathbb{S}^d$, and the Gramian of this polytope is the Gramian of the cone.

Theorem 9 *Given a $n \times m$ relation $\mathcal{R}$, there is a type $\mathcal{R}$ spherical polytope with Gramian G if and only if the maxbiclique lattice of $\mathcal{R}$ is a rank $d+1$ diamond lattice with connected bipartite flag graph, and G is a $n \times n$ rank $d+1$ symmetric positive semi-definite matrix with all diagonals 1 that satisfies the following:*

1. *For every vertex of $\mathcal{R}$ and all facets $F_{i_1} \cdots F_{i_k}$ incident to it, $\mathrm{rank}([G]_{i_1\cdots i_k}^{i_1\cdots i_k}) = d$.*
2. *For every pair of super cycles $F_{i_1}, \ldots, F_{i_{d+1}}$ and $F_{j_1}, \ldots, F_{j_{d+1}}$ with the same orientation, $\det([G]_{j_1\cdots j_{d+1}}^{i_1\cdots i_{d+1}}) > 0$.*

Moreover, G parameterizes the realization space of type $\mathcal{R}$ spherical polytopes modulo isometries.

Let $\Phi = \mathrm{diag}(1, \ldots, 1, -1)$ so $\phi(x) = [x]_1^2 + \cdots + [x]_d^2 - [x]_{d+1}^2$, and let $\mathbb{H}^d = \{x : \phi(x) = -1, [x]_{d+1} > 0\}$ denote the d-dimensional upper hyperbola in $\mathbb{R}^{d+1}$. A finite-volume hyperbolic d-polytope $P = \mathbb{H}^d \cap C$ is the intersection of $\mathbb{H}^d$ with a pointed polytopal cone

$$C \subset \{x : \phi(x) \leq 0, [x]_{d+1} \geq 0\}.$$

Again the Gramian of P is the Gramian of C. When a ray of C does not intersect $\mathbb{H}^d$, we say the corresponding vertex, which does not appear in $\mathbb{H}^d$, is an ideal vertex of P. An isometry of $\mathbb{H}^d$ is an orthogonal transformation of ϕ.

Theorem 10 *Given a $n \times m$ relation $\mathcal{R}$ and vertices $\mathcal{W} \subset [m]$ of $\mathcal{R}$, there is a finite-volume type $\mathcal{R}$ hyperbolic polytope with all vertices except those of $\mathcal{W}$ in $\mathbb{H}^d$ and Gramian G if and only if the maxbiclique lattice $\mathcal{P}$ of $\mathcal{R}$ is a rank $d+1$ diamond lattice with connected bipartite flag graph, and G is a $n \times n$ symmetric matrix with all diagonals 1 that satisfies the following:*

1. *For every vertex of $\mathcal{R}$ and all facets $F_{i_1} \cdots F_{i_k}$ incident to it, $\mathrm{rank}([G]_{i_1 \cdots i_k}^{i_1 \cdots i_k}) = d$.*
2. *For every pair of super cycles $F_{i_1}, \ldots, F_{i_{d+1}}$ and $F_{j_1}, \ldots, F_{j_{d+1}}$ with the same orientation, $\det([G]_{j_1 \cdots j_{d+1}}^{i_1 \cdots i_{d+1}}) < 0$.*
3. *For $2 \le s \le d$ and every truncated cycle $F_{i_1}, \ldots, F_{i_s}$ incident to a $(d-s)$-face $f \in \mathcal{P}$, if $f \in \mathcal{W}$ then $\det(G_{i_1 \cdots i_s}) = 0$ otherwise $\det(G_{i_1 \cdots i_s}) > 0$*
4. *For every pair of super cycles $F_{i_1}, \ldots, F_{i_{d+1}}$ and $F_{j_1}, \ldots, F_{j_{d+1}}$ with the same orientation that are incident to a different vertex, $\det([G]_{j_1 \cdots j_{d+1}}^{i_1 \cdots i_{d+1}}) > 0$.*

Moreover, G parameterizes the isometric realization space of finite-volume type $\mathcal{R}$ hyperbolic polytopes with ideal vertices $\mathcal{W}$.

5 Concluding Remarks

Often when we are able to reduce a problem to linear algebra it becomes computationally tractable, but the results shown here do not provide a polynomial time algorithm to decide realizability. Indeed, the complexity of deciding realizability for d-polytopes for any fixed $d \ge 4$ is in the existential theory of the reals (ETR), the complexity of deciding a boolean sentence of polynomial equations and inequalities with only existential quantifies in prenex normal form. While no algorithm is known to be optimal, there is a $(p\delta)^{O(v)}$ time algorithm for p polynomial formulas in v variables of degree at most δ each [13]. The existence of a polynomial time algorithm for deciding realizability would imply that $P = NP = ETR$.

Richter-Gebert [14] showed ETR is a lower bound for deciding realizability for 4-polytopes by reduction to a realizability problem with a polynomial number of vertices. For the upper bound, Bokowski and Sturmfels [3] provided a method for deciding realizability by generating a partial chirotope from certain lattices in a way that always gives the partial chirotope of the vertices of a combinatorial polytope. For an incidence relation on n vertices and m facets, the time it takes to generate the partial chirotope is linear in the number of flags of the maxbiclique lattice, and deciding its realizability is ETR in $O(dn)$ variables and $O(dn^{d+1})$ total degree, giving $n^{O(d^2 n)}$ time. Desciding realizability of a relation using Theorem 1 is polynomial in the number of faces plus ETR in $O(d(n+m))$ variables and $O(2nm)$ total degree, giving $(nm)^{O(d(n+m))}$ time. Note that deciding the combinatorial conditions of Theorem 1 for a lattice with f faces can be done in f^2 time using Therorem 6, and the maxbiclique lattice can be computed from a relation with i incidences and $n \le m$ in $O(nif)$ time [9]. If d is not fixed, the number of faces in the maxbiclique lattice of a relation may be exponential in the size of the relation, and the number of flags may be super-polynomial in the number of faces, so these factors may be significant.

We may understand the matrices associated with a polytope (or cone) as linear maps sending formal linear combinations of vertices (or generators) to formal linear

combinations of facets. For for a cone C with facet–ray matrix N and $x \in \mathbb{R}^m$ in the closed positive orthant, if $N = H^*W$ is a decomposition into unit facet normals h_i and generators w_j, the map $-N$ first sends x to a point Wx in the cone, and then sends that point to the vector $-Nx = -H^*Wx$ of distances from the supporting hyperplane of each facet, the slack of the point Wx in each linear constraint on C.

Such decompositions are used by Yannakakis's Theorem on the size of extended formulations [19]. This theorem relates non-negative decompositions of slack matrices to projective maps between polytopes, or equivalently linear maps between cones, in the following way. If we are given a linear map A sending a cone C_0 with r facets to another cone C, then we can use this map to decompose the slack matrix of C into non-negative matrices mapping through $\mathbb{R}^r$. To define the first matrix W_0, choose a preimage $v_j \in C_0$ for each generator $w_j \in C$, $w_j = Av_j$. Now let W_0 be defined by acting on formal linear combinations of generators of C by sending each generator w_j to the slack of v_j in the constraints on C_0. The map A^* sends each cogenerator of C to a non-negative linear combination of cogenerators of C_0. Let H_0 be the corresponding map from cogenerators of C to cogenerators of C_0. The map H_0^* now sends the slack of v_j in constraints on C_0 to the slack of w_j in constraints on C, so $-N = H_0^*W_0$ is the desired decomposition. Alternatively, if we start with such a non-negative decomposition, and we also have a full rank decomposition $N = H^*W$, then the map $H^{+*}H_0^*$ sends the closed positive orthant $\mathbb{R}^r_{\geq 0}$ to the cone C.

Gale duality can also be understood from these matrices. Each point in a cone C comes from an equivalence class of non-negative linear combinations of generators. These equivalence classes are defined by the linear relations (dependencies) of the generators, which are the same relations as those of the columns of N. That is, the points of the cone correspond to the classes $\mathbb{R}^m_{\geq 0}/\sim$ where $x \sim y$ when $x - y \in R = \text{null}(N)$. The orthogonal projection of the closed positive orthant $\mathbb{R}^m_{\geq 0}$ into the orthogonal complement of R gives one representative of each class, and hence a linear copy C' of the cone C. The generators of C' are the orthogonal projections of the standard basis vectors, $x_j = \text{proj}_{R^\perp} e_j = N^+Ne_j$. This projection expresses each point in C as the linear combination of generators with smallest norm (allowing negative coefficients). In particular, each generator of C is expressed as a linear combinations of itself with the other generators, $w_j = We_j = Wx_j$, which associates a linear relation to each generator $r_j = e_j - x_j$. The Gale dual of the cone is the vector configuration $\{r_j\}_{j \in [m]} \subset R$ consisting of the relation associated to each generator in this way. This configuration is the orthogonal projection of the standard basis vectors into $R = \text{null}(N)$ and is given by the columns of $I - N^+N$.

In the case of polytopes, the facet–vertex matrix M acts on formal weighted averages of the vertices, so we must impose one additional relation, that the weights sum to 1. This makes the Gale dual the orthogonal projection of the standard basis vectors into $\text{null}([M; \mathbf{1}])$. Or equivalently, the Gale dual of the cone $\mathbb{R}_{\geq 0}[P; \mathbf{1}]$.

Here we presented the Gale dual as a vector configuration in $R = \text{null}(N) \subset \mathbb{R}^m$. The Gale dual is more commonly presented, however, by extending $\mathbb{R}^{d+1} \supset C$ to $\mathbb{R}^m$. For this, we find some orthonormal set of vectors $q_j = Qe_j$, such that projecting q_j to the first $d + 1$ coordinates gives the jth generator of the cone. Then projecting q_j to the last $m - d - 1$ coordinates gives the jth vector of the Gale dual. This difference amounts to an orthogonal change of coordinates by Q.

If we choose unit length generators and cogenerators of a cone and find the full singular value decomposition of the facet–ray matrix $N = U \Sigma V^*$, we get the Gale dual of the cone along with the Gale dual of its polar in this way. That is, from the compact singular value decomposition, extend U and V to full orthogonal matrices and augment Σ with 0. The additional columns of V give an $(m - d - 1)$-vector for each generator, which is the Gale dual of C. Also, the additional columns of U give a $(n - d - 1)$-vector for each facet, which is the Gale dual of the polar C^*. We also get the Gramian of a cone's generators and cogenerators together from the facet–ray matrix N as follows.

$$\tilde{G} := \begin{bmatrix} \sqrt{NN^*} & N \\ N^* & \sqrt{N^*N} \end{bmatrix} = \begin{bmatrix} U\Sigma U^* & U\Sigma V^* \\ V\Sigma U^* & V\Sigma V^* \end{bmatrix} = \begin{bmatrix} U \\ V \end{bmatrix} \Sigma \begin{bmatrix} U^* & V^* \end{bmatrix}.$$

This perhaps innocuous looking formula relates several objects, a polytopal cone, its Gale dual, its dihedral angles, and its polar cone along with the Gale dual and dihedral angles of the polar. the Gramian of the cone is $\sqrt{NN^*}$, and the Gramian of its polar is $\sqrt{N^*N}$. That is, the entries of the first diagonal block are the inner products of the outward normals of the cone's facets, which are cosine of the dihedral angles. Similarly the second diagonal block is that for the polar cone. It would be nice to add to this a geometric interpretation of the singular values Σ.

References

1. Björner, A., Las Vergnas, M., Sturmfels, B., White, N., Ziegler, G.M.: Oriented matroids. In: Encyclopedia of Mathematics and its Applications, vol. 46, 2nd edn. Cambridge University Press, Cambridge (1999)
2. Björner, A., Paffenholz, A., Sjöstrand, J., Ziegler, G.M.: Bier spheres and posets. Discrete Comput. Geom. **34**(1), 71–86 (2005)
3. Bokowski, J., Sturmfels, B.: Computational Synthetic Geometry. Lecture Notes in Mathematics, vol. 1355, Springer (1989)
4. Conforti, M., Cornuéjols, G., Zambelli, G.: Extended formulations in combinatorial optimization. Ann. Oper. Res. **204**, 97–143 (2013)
5. Davey, B.A., Priestley, H.A.: Introduction to Lattices and Order. Cambridge Math. Textbooks. Cambridge University Press, Cambridge (1990)
6. Díaz, R.: A characterization of Gram matrices of polytopes. Discrete Comput. Geom. **21**(4), 581–601 (1999)
7. Fiorini, S., Rothvoß, T., Tiwary, H.R.: Extended formulations for polygons. Discrete Comput. Geom. **48**(3), 658–668 (2012)
8. Grünbaum, B.: Convex Polytopes, vol. 221, 2nd edn. Springer, New York (2003)
9. Kaibel, V., Pfetsch, M.E.: Computing the face lattice of a polytope from its vertex–facet incidences. Comput. Geom. **23**(3), 281–290 (2002)
10. McMullen, P., Schulte, E.: Abstract Regular Polytopes. Cambridge University Press, Cambridge (2002)
11. Paffenholz, A.: Constructions for Posets, Lattices, and Polytopes. PhD thesis, Technischen Universität Berlin (2005)
12. Peterin, I.: Characterizing flag graphs and induced subgraphs of cartesian product graphs. Order **21**(4), 283–292 (2004)

13. Renegar, J.: On the computational complexity and geometry of the first-order theory of the reals. Part I: Introduction. Preliminaries. The geometry of semi-algebraic sets. The decision problem for the existential theory of the reals. J. Symb. Comput. **13**(3), 255–299 (1992)
14. Richter-Gebert, J.: Realization Spaces of Polytopes, Lecture Notes in Mathematics, vol. 1643. Springer, Berlin (1996)
15. Robertson, S.A.: Polytopes and Symmetry, vol. 90. Cambridge University Press, Cambridge (1984)
16. Stanley, R.P.: Enumerative Combinatorics, vol. 49. Cambridge University Press, Cambridge (2011)
17. Steinitz, E.: Polyeder und raumeinteilungen. Encyclopädie der Mathematischen Wissenschaften **3**, 1–139 (1922)
18. Volodin, I.A., Kuznetsov, V.E., Fomenko, A.T.: The problem of discriminating algorithmically the standard three-dimensional sphere. Uspekhi Matematicheskikh Nauk **29**(5), 71–168 (1974)
19. Yannakakis, M.: Expressing combinatorial optimization problems by linear programs. J. Comput. Syst. Sci. **43**(3), 441–466 (1991)
20. Ziegler, G.M.: Lectures on Polytopes, Graduate Texts in Mathematics, vol. 52. Springer, New York (2007)