Bounds on the Complexity of Halfspace Intersections when the Bounded Faces have Small Dimension

David Eppstein · Maarten Löffler

Received: 9 September 2011 / Revised: 16 April 2013 / Accepted: 22 April 2013 /
Published online: 14 May 2013
© Springer Science+Business Media New York 2013

Abstract For a polyhedron P let $B(P)$ denote the polytopal complex that is formed by all bounded faces of P . If P is the intersection of n halfspaces in $\mathbb{R}^D$, but the maximum dimension d of any face in $B(P)$ is much smaller, we show that the combinatorial complexity of P cannot be too high; in particular, that it is independent of D . We show that the number of vertices of P is $O(n^d)$ and the total number of bounded faces of the polyhedron is $O(n^{d^2})$. For inputs in general position the number of bounded faces is $O(n^d)$. We show that for certain specific values of d and D , our bounds are tight. For any fixed d , we show how to compute the set of all vertices, how to determine the maximum dimension of a bounded face of the polyhedron, and how to compute the set of bounded faces in polynomial time, by solving a number of linear programs that is polynomial in n .

Keywords Polyhedral combinatorics · Intersection of halfspaces · Bounded subcomplex · Dimension · Vertex enumeration

1 Introduction

Bounds on the complexity of halfspace intersections and convex hulls [31, 36], and algorithms for constructing halfspace intersections and convex hulls [1, 3, 6, 8, 15, 29,

D. Eppstein
Department of Computer Science, University of California, Irvine,
Irvine, CA, USA
e-mail: eppstein@uci.edu

M. Löffler (✉)
Department of Information and Computing Sciences, Utrecht University,
Utrecht, Netherlands
e-mail: m.loffler@uu.nl

Fig. 1 An unbounded 3-dimensional polyhedron, with a bounded subcomplex consisting of four 2-dimensional and three 1-dimensional maximal faces. On the *right* we show the projection onto a plane that is perpendicular to an unbounded direction

[35, 38], have long been a mainstay in discrete and computational geometry. However, because of their inherent exponential dependence on the dimension of the input, these worst case bounds are useful primarily for polytopes of low to moderate dimension. In contrast, linear programming allows single vertices of halfspace intersections to be found numerically in polynomial time even for high-dimensional inputs [25, 26, 37]. Is there an intermediate family of polytope construction problems, with low complexity even for high-dimensional input, but capable of describing more complex sets of features than a single polytope vertex?

Given the obstacles that have been identified to finding a truly output-sensitive halfspace intersection algorithm [2], it may not be reasonable to ask for an output-sensitive algorithm that works efficiently on all input sets of halfspaces that define output polytopes with few vertices. Instead, in this paper, we identify a different parameter of instance complexity, the dimension of the bounded faces of a halfspace intersection (that is, the faces that do not extend to infinity). When this dimension is bounded, we obtain both low numbers of vertices and efficient algorithms for listing them, despite the possibility that the input may have very large numbers of unbounded faces. Figure 1 shows an example polyhedron of which the bounded faces have been marked. These faces form a combinatorial structure that can be described equivalently in several ways:

- The *bounded subcomplex* of an intersection of n halfspaces is the cell complex formed by the set of bounded faces of the intersection.
- The subset of faces of a convex polytope that lie above (or below) some linear threshold is combinatorially equivalent to a bounded subcomplex, under a projective transformation that takes the threshold hyperplane to the hyperplane at infinity. A transformation of this type necessarily transforms the faces of the polytope that are entirely above the threshold into bounded faces of the transformed polyhedron, but the faces that cross the threshold (if any exist) become unbounded.
- The subset of faces of a convex polytope that are disjoint from some specified facet is again equivalent to a thresholded intersection or a bounded subcomplex.

To see this, one may consider a linear threshold determined by a linear function that is zero on the specified facet, positive on the points of the polytope disjoint from the facet, and negative on the other side of the facet.

- The complexes considered above are equivalent under projective duality to the subset of faces of a convex polytope that are nonadjacent to some specified vertex.

Our interest in the work presented here is in characterizing structures of this type, bounding their combinatorial complexity, and finding efficient algorithms to generate these structures from the halfspaces or points that define them.

1.1 Motivating Applications

One motivation for studying bounded subcomplexes comes from the study of algorithms on arbitrary finite metric spaces, and more specifically from the *tight span* construction, a canonical method of embedding any metric space into a continuous space with properties similar to those of L^∞ spaces [9, 14, 24]. One way of defining the tight span, for a finite metric space with n points p_i , is to coordinatize n -dimensional L^∞ space by n variables x_i , and to define a polyhedral subset of the space by the $\binom{n}{2}$ linear inequalities

$$x_i + x_j \geq \text{dist}_{L^\infty}(p_i, p_j)$$

for each possible pair (i, j) . Then, the tight span is the bounded subcomplex of this polyhedron. For metric spaces satisfying an appropriate general position assumption, the dimension of the bounded subcomplex is between $\lceil n/3 \rceil$ and $\lfloor n/2 \rfloor$ [11], but certain combinatorially defined metrics, such as the metrics of distances on certain classes of graphs, can have tight spans of much lower dimension [17]. Our results bound the complexity of these low-dimensional tight spans and allow them to be constructed efficiently, generalizing our previous algorithms for constructing tight spans when they are homeomorphic to subsets of the plane [18].

Our results can also be interpreted as statements about the complexity of Delaunay triangulations for inputs satisfying strong convex position assumptions. Delaunay triangulations are closely related to convex hulls: the Delaunay triangulation is combinatorially equivalent to the convex hull of a point set lifted to a sphere in one higher dimension, augmented by an extra point at the pole of the sphere, followed by the removal of all faces incident to that pole [4]. If a D -dimensional set of n points has the property that every interior point of the convex hull of the set belongs to a Delaunay triangulation feature of dimension at least $D - d$, then our results imply via this lifting relation that the Delaunay triangulation has $O(n^d)D$ -dimensional simplices. For instance, if the convex hull of a point set is a stacked polytope (a convex figure formed by gluing simplices facet-to-facet) and the Delaunay triangulation coincides with the gluing pattern of the simplices, then every interior point of the convex hull belongs either to one of the simplices or to one of the glued facets, so $d = 1$; in this case, the number of simplices in the Delaunay triangulation is exactly $n - D$.

Additionally, many combinatorial optimization problems such as shortest path trees, minimum spanning trees, bipartite minimum weight perfect matchings, and network flows can be expressed as linear programs. The problems of finding the k best solutions to optimization problems of these types have been much studied (e.g., see [16, 19, 21]), and from the polyhedral point of view these sets of the k best solutions can be represented by the vertices of the feasible polyhedron that have the k smallest values of the linear objective function used for the problem. Therefore, it is of interest to understand the combinatorial behavior of the subsets of polyhedron vertices and faces satisfying a linear threshold constraint, or equivalently of the bounded faces of polyhedra. As a somewhat trivial example (easy to prove directly without resorting to our polyhedral point of view), consider the bin packing problem in which the input consists of a set of items with weights w_i and a capacity W , and a solution consists of a subset of the items with total weight at most W . If x_i is an indicator variable for the inclusion of item i , then the solutions are the vertices of the hypercube $0 \leq x_i \leq 1$ that satisfy the linear constraint

$$\sum w_i x_i \leq W.$$

The number of halfspaces defining the hypercube is $2n$, and the dimension d of the bounded subcomplex of the hypercube determined by the linear constraint is the maximum number of items that can be packed into a single solution. Therefore, our results imply that the total number of solutions is $O(n^d)$.

1.2 Related Work

Herrmann et al. [22] consider the problem of computing the bounded subcomplex of a polyhedron, as we do here, but they do not provide combinatorial bounds on the complexity of the bounded subcomplex. Their algorithms assume that the vertices and facets of the polyhedron are both already known, and are output-sensitive given this information.

Bounded subcomplexes have also been investigated in other contexts. In tropical geometry, the *tropical complex* defined by a set of points is also equivalent to the set of bounded faces of a polytope defined from the points [12]. Develin [10] studied the bounded faces of a certain polytope arising from a problem in algebra, and showed that they are all isomorphic to subpolytopes of permutahedra. Queyranne [33] used polytopes to model scheduling problems; the polytope defined by Queyranne has a unique bounded facet. In connection with the tight span application, Hirai [23] showed that the bounded faces of a halfspace intersection form a contractable complex.

The view of a bounded subcomplex as the subset of a polytope that is not adjacent to a particular facet is reminiscent of a *polytope pair*, the subset of a polytope that is adjacent to a particular facet [27]. However, known results on the complexity of polytope pairs do not appear to apply directly to the complexity of bounded subcomplexes.

1.3 New Results

We show that, when the bounded subcomplex consists only of faces of dimension at most d , where d is significantly smaller than the ambient dimension, then it has smaller combinatorial complexity than the bounds given by the upper bound theorem. Specifically, it has $O(n^d)$ vertices, and $O(n^{d^2})$ faces. More strongly, when the half-spaces forming the intersection are in general position, we obtain a tighter bound of $O(n^d)$ on the number of faces as well. We also provide several examples showing that our bounds are tight.

Based on these combinatorial bounds, we provide algorithms for listing all the vertices or faces of the bounded subcomplex in polynomial time whenever d is a fixed constant.

1.4 Organization

The rest of this document is organized as follows. In Sect. 2, we review some standard terminology and facts about polytopes and polyhedra. In Sect. 3, we provide upper bounds on the complexity of the bounded subcomplex of unbounded polyhedra. In Sect. 4, we provide lower bounds, matching the upper bounds in some cases. In Sect. 5, we address the algorithmic question of reporting the bounded subcomplex efficiently. Finally, we conclude in Sect. 6.

2 Preliminaries

2.1 Standard Definitions

We begin with some standard material in polyhedral combinatorics. For more background on this material, see [20] or [39].

Define a *polytope* to be the convex hull $\text{CH}(S)$ of a finite set S of points, and a *polyhedron* to be the intersection $\bigcap T$ of a finite set T of closed halfspaces. Every polytope is a polyhedron; a polyhedron is a polytope if and only if it is bounded.

The *faces* of a polytope or polyhedron P are the sets of the form $P \cap H$ where H is a closed halfspace whose boundary is disjoint from the relative interior of P . Every face of a polytope is also a polytope, and every face of a polyhedron is also a polyhedron. The *dimension* $\dim f$ of a face f is the dimension of its affine hull $\text{aff } f$. With these definitions, P itself is a face, as is the empty set. By convention the dimension of the empty set is -1 . Two faces are *incident* if one is a subset of the other.

We may label the faces of a polytope or polyhedron combinatorially, by subsets of the generating objects for the polytope or polyhedron. That is, a face f of a polytope $\text{CH}(S)$ may be identified with the set $f \cap S$, and a face f of a polyhedron $\bigcap T$ may be identified with the set of halfspaces of T whose boundary hyperplanes contain f . This identification is one-to-one: each face is identified with a unique set. The subset relation between faces is either the same as or the reverse of the subset relation between the sets the faces are identified with, for polytopes and polyhedra, respectively.

Fig. 2 The link of a vertex of P is a polytope of dimension one lower than P

A *vertex* of a polytope or polyhedron is a 0-dimensional face; it must be a singleton set, containing a single point. An *edge* of a polytope or polyhedron is a bounded 1-dimensional face, consisting of a line segment connecting two vertices. Polyhedra may also have *rays* and *lines*, unbounded 1-dimensional faces with one or zero vertices on them, respectively. (In many sources, all one-dimensional faces are called edges, but we prefer this more specific terminology.) If a polytope or polyhedron P has dimension D , then a *facet* of P is a $(D - 1)$ -dimensional face. If D is also the dimension of the space containing the points or halfspaces from which P was defined, then P is said to be *full-dimensional*. More generally we define a d -*face* of P to be a face of dimension d , so for instance a 1-face may be an edge, a ray, or a line.

Let v be a vertex of a polytope or polyhedron P , and let h be a hyperplane that separates v from all other vertices of P such that $P \cap h$ is bounded. We define the *link* of v to be the intersection $P \cap h$. Figure 2 shows an example. Different choices of h will give geometrically different intersections, but they are all related to each other by affine transformations, so as a combinatorial polytope the link is well defined; however, in our proofs we will need to refer to the link as a geometric object, defined as above, rather than as a combinatorial object, so in each such usage we will need to verify that the choice of h does not affect the result. The faces of the link correspond one-for-one, in an incidence-preserving way, with the faces of P that are incident to v .

A full-dimensional polytope or polyhedron is *simple* if every vertex is incident to at most D edges and rays. Equivalently, a polytope or polyhedron is simple if the link of every vertex is a simplex. A polyhedron defined as the intersection of a set of halfspaces in general position (meaning that a small perturbation of any of the halfspaces does not change the combinatorial structure of the polyhedron) is necessarily a simple polyhedron. Dually, a polytope is *simplicial* if all of its faces except for the polytope itself are simplices. A polytope defined as the convex hull of a set of points in general position is necessarily simplicial.

If P is a polytope containing the origin, the *polar polytope* of P is the intersection of a system of halfspaces, one for each vertex of P : if v is a vertex of P then the corresponding halfplane is the set of points w such that $v \cdot w \leq 1$. Although defined as an intersection of halfspaces, the polar polytope is bounded and is therefore a polytope. The faces of the polar polytope correspond one-for-one with the faces of P but in a dimension-preserving way: an i -dimensional face of P corresponds to a $(D - i - 1)$ -dimensional face of the polar polytope. The polar polytope of the polar polytope is P .

again. If P is simple, its polar polytope is simplicial, and if P is simplicial, its polar polytope is simple.

2.2 New Notation

If P is a polyhedron or polytope and ℓ is a linear function, we define $\max_{\ell} P$ and $\min_{\ell} P$ to be the face on which ℓ takes its minimum or maximum. If the minimum or maximum is unbounded, then we define $\max_{\ell} P$ or $\min_{\ell} P$ to be the empty set. A polyhedron is *pointed* if it has at least one vertex. If ℓ^0 is a tangent hyperplane to a polyhedron P at a given vertex v , that is not tangent to any higher-dimensional face, and ℓ is a linear function that is positive on the side of ℓ^0 containing P and negative on the other side, then $v = \min_{\ell} P$, so a polyhedron P is pointed if and only if there exists a linear function ℓ for which $\min_{\ell} P$ is a single vertex.

A *polyhedral complex* is a finite set C of polyhedra, all in the same ambient space, such that C contains each face of each polyhedron in C and such that the intersection of any two polyhedra in C is a face of both. We define a *polytopal complex* to be a polyhedral complex in which each member of C is a polytope.

Let P be a polytope or polyhedron, ℓ a linear function, and B any real number. We define $P^{\ell < B}$ as the polyhedral complex formed by the set of faces of P such that, for every point p of a face in $P^{\ell < B}$, $\ell(p) < B$. Analogously, for $B = \infty$, we define $P^{\ell < \infty}$ to be the set of faces of P on which ℓ is bounded. (This notation directly implies only that ℓ is bounded from above, but we define it to mean that it is bounded from below as well.) We define $\dim P^{\ell < B}$ to be the maximum dimension of a face in $P^{\ell < B}$.

3 Upper Bounds

3.1 The Slicing Lemma

In order to help prove our main result, we begin with a technical lemma stating that if we slice a polytope by a hyperplane in such a way that one side of the slice contains only faces of low dimension, then the other side of the slice contains a face of high dimension. Figure 3 illustrates the lemma.

Lemma 1 *Let P be a polytope, and λ be a hyperplane that does not pass through any vertex of P . Then there exist faces f^+ and f^- of P , on opposite sides of λ , such that $\text{aff}(f^+ \cup f^-) = \text{aff } P$.*

Fig. 3 No matter how we slice a tetrahedron with a plane, there is always a pair of faces (one of which could be the empty set) on both sides of the plane such that the affine hull of these faces is the full space. Lemma 1 claims that this is still true for any polytope, in any dimension

Proof We use induction on dimension. Note that the empty set is also a face of P , so λ need not intersect P . Let ℓ be a linear function that is zero on λ , let $P^+ = P^{\ell > 0}$ be the complex of faces of P on which ℓ is entirely positive, and let $P^- = P^{\ell < 0}$ be the complex of faces of P on which ℓ is entirely negative. As a base case, if P^+ is empty, then we may take $f^+ = \emptyset$ and $f^- = P$. As a second base case, if P^+ is a single vertex v , then we may take $f^+ = v$ and f^- to be any facet of P disjoint from v ; such a facet must exist, for otherwise P would be unbounded.

We may assume without loss of generality (by perturbing λ , if necessary) that no hyperplane parallel to λ passes through two or more vertices of P . By the assumption of general position on λ , ℓ does not take the same value on any two vertices of P . Let v be the vertex in P^+ minimizing ℓ , and let L be the link of v (with an arbitrary choice of the intersecting hyperplane μ defining the link). Let $L^+ = L^{\ell > \ell(v)}$ and $L^- = L^{\ell < \ell(v)}$. By the induction hypothesis, applied to the polytope L , we may find faces g^+ and g^- in L^+ and L^- , respectively, such that $\text{aff}(g^+ \cup g^-) = \text{aff} L$. (When P is full-dimensional, this affine hull is just the hyperplane μ .) The face g^+ of L^+ corresponds to a face f^+ of P in which all points have a value of ℓ larger than or equal to $\ell(v)$; since $\ell(v)$ is strictly positive, f^+ belongs to P^+ . However, the face g^- of L^- corresponds to a face h of P such that v lies in P^+ and the remaining vertices of h lie in P^- . Let f^- be any facet of h disjoint from v , as in the second base case.

Then

$$g^+ = (f^+ \cap \mu) \subset f^+ \subset \text{aff}(f^+ \cup f^-).$$

Additionally, $v \in f^+$, so

$$g^- = (h \cap \mu) \subset h \subset \text{aff}(v \cup f^-) \subset \text{aff}(f^+ \cup f^-),$$

and

$$\text{aff} P = \text{aff}(L \cup v) = \text{aff}(g^+ \cup g^- \cup v) \subset \text{aff}(f^+ \cup f^-).$$

But $f^+ \cup f^-$ is a subset of P , so its affine hull cannot be a proper superset of the affine hull of P , and the two affine hulls are equal, as desired. $\square$

3.2 Labeling Low-Dimensional Faces by Small Sets

As we now show, the vertices of a polyhedron with low-dimensional bounded faces can be associated to small sets of facets of the polyhedron, leading to polynomial bounds on the number of vertices of the polyhedron. Figure 4 illustrates the lemma for 3-dimensional polyhedra.

Lemma 2 *Let P be a full-dimensional pointed polyhedron, defined as the intersection of n halfspaces in $\mathbb{R}^D$. Let ℓ be a linear function with a vertex of P as its minimum, and suppose that ℓ does not take the same value at any two different vertices of P . Let B be a given real number, and let $d = \dim P^{\ell < B}$. Let v be any vertex of $P^{\ell < B}$. Then*

Fig. 4 A vertex v of a 3-dimensional polyhedron has at least three incident edges. Let S be the set of halfspaces as in Lemma 2. On the left, two edges point “up” (in the direction of ℓ), so v is the minimum of the face between them: $|S| = 1$. In this case, the only face that is necessarily bounded is the edge that points down (which has dimension 1). On the right, only one edge points up, so v is the minimum of the intersection of the two faces on both sides of it: $|S| = 2$. In this case, however, at least two edges point down, so the face between them is necessarily bounded (it has dimension 2)

there exists a set S of d of the halfspaces defining P , such that $v = \min_{\ell}(P \cap A)$, where A is the intersection of the boundary hyperplanes of S .

Proof Define a linear function ℓ_v by $\ell_v(x) = \ell(x) - \ell(v)$. Then ℓ_v differs from ℓ by a simple translation; therefore, it has the same maxima and minima as ℓ on each face of P . Let V be the link of v , let V^+ be the complex of faces of V on which ℓ_v is positive, and let V^- be the complex of faces of V on which ℓ_v is negative; this definition of V^+ and V^- does not depend on the intersecting hyperplane used to define the link geometrically.

By Lemma 1, we can find faces f^+ and f^- in V^+ and V^- respectively, such that $\text{aff}(f^+ \cup f^-) = \text{aff} V$. Since $\dim V = \dim P - 1 = D - 1$, we have $\dim f^+ + \dim f^- \geq D - 2$. Then these two faces f^+ and f^- correspond to faces g^+ and g^- in P , such that $v = \min_{\ell} g^+ = \max_{\ell} g^-$, and such that $\dim g^+ + \dim g^- \geq D$. But we know that $\dim g^- \leq d$ by assumption, so $\dim g^+ \geq D - d$. Therefore, g^+ may be represented as the intersection of P with a set of the boundary hyperplanes of at most d of the defining halfspaces of P . If A is the intersection of these hyperplanes, then $v = \min_{\ell} g^+ = \min_{\ell}(P \cap A)$ as desired. $\square$

3.3 Bounding the Number of Vertices

We are now ready to state our main result.

Theorem 1 Let P be a full-dimensional pointed polyhedron, defined as the intersection of n halfspaces in $\mathbb{R}^D$, let ℓ be a linear function which attains a minimum on P , let B be given and greater than the minimum value of ℓ , and suppose that $\dim P^{\ell < B} = d < D$. Then the number of vertices of $P^{\ell < B}$ is at most

$$\binom{n}{d} - \binom{D}{d} + 1.$$

Proof Lemma 2 gives a formula $v = \min_{\ell}(P \cap A)$ by which each vertex v of P may be obtained from a subset of at most d of the defining halfspaces of P . We may augment

any set of halfspaces determining v in this way so that it has exactly d halfspaces, by making an arbitrary choice of a sufficient number of the other halfspaces whose boundaries are incident to v . There are $\binom{n}{d}$ sets of exactly d halfspaces, each one determines at most one vertex, and all vertices can be determined in this way, but some vertices may be duplicated. In particular the vertex $\min_{\ell} P$ belongs to at least D facets, and is the vertex that is determined by each of the $\binom{D}{d}$ different d -tuples of these D halfspaces, giving rise to the correction terms in the formula. $\square$

Corollary 1 *Let P be a polyhedron, defined as the intersection of n halfspaces in $\mathbb{R}^D$, let ℓ be a linear function which attains a minimum on P , let B be given, and suppose that $\dim P^{\ell < B} = d < D$. Then the number of vertices of $P^{\ell < B}$ is $O(n^d)$.*

Proof We may assume without loss of generality that P is full-dimensional, for otherwise we may restrict our attention to the affine hull of P and eliminate any halfspaces that contain the entire affine hull.

If P is not pointed or B does not exceed the minimum value of ℓ , then $P^{\ell < B}$ has zero vertices and the result is obvious. Otherwise, we may ignore the $-\binom{D}{d} + 1$ correction terms in Theorem 1, obtaining a formula which grows asymptotically as $O(n^d)$ and is independent of D . $\square$

3.4 Euler's Formula

Letting $\Phi(P^{\ell < B})$ denote the set of faces of $P^{\ell < B}$, define the *Euler characteristic*

$$\chi(P^{\ell < B}) = \sum_{f \in \Phi(P^{\ell < B})} (-1)^{\dim f}.$$

Recall that we include the empty set as a face, as well as P itself in the case that $\max_{\ell} P < B$. As is well known, with this definition, $\chi(P) = 0$ for any convex polytope P ; Ziegler [39] gives a nice proof of this fact based on shelling.

As we now show, the Euler characteristic of $P^{\ell < B}$ behaves similarly to the Euler characteristic of an entire polytope. This fact will lead to improved bounds on the number of faces of $P^{\ell < B}$, because it will imply that the d -dimensional faces can be charged against other faces of lower dimension. Our proof can be viewed as a dual form of the shelling proof of Euler's formula for polytopes.

Theorem 2 *Let P be a polytope or polyhedron, let ℓ be a linear function, and let B be either ∞ or a real number. Additionally, suppose that $P^{\ell < B}$ is nonempty. Then $\chi(P^{\ell < B}) = 0$.*

Proof We use induction, both on the dimension of $P^{\ell < B}$ and on the number of vertices of $P^{\ell < B}$. As a base case, the result is clearly true if $P^{\ell < B}$ consists only of a single vertex: for, in that case, $\Phi(P^{\ell < B})$ contains two faces, the empty set of dimension -1 and that vertex, and these two faces have opposite signs in the sum defining $\chi(P^{\ell < B})$, cancelling each other out.

Otherwise, $P^{\ell < B}$ has at least two vertices. We may assume without loss of generality (by perturbing ℓ if necessary) that no two vertices of $P^{\ell < B}$ have the same value of ℓ .

Let v be the vertex of $P^{\ell < B}$ with the maximum value of ℓ , and let $B' = \ell(v)$. Then, by induction, $\chi(P^{\ell < B'}) = 0$.

$P^{\ell < B}$ differs from $P^{\ell < B'}$ by the set of faces incident to v . Let L be the link of v in P ; then the face structure of $L^{\ell < B'}$ does not depend on the intersecting hyperplane used to define the geometry of L , and the faces incident to v in $P^{\ell < B}$ are in one-to-one correspondence with the faces of $L^{\ell < B'}$. For instance, the empty face of $L^{\ell < B'}$ corresponds to the vertex v of $P^{\ell < B}$, and each vertex u of $L^{\ell < B'}$ corresponds to an edge uv of $P^{\ell < B}$. In this correspondence, the dimension of a face in $L^{\ell < B'}$ is one less than the dimension of the corresponding face of $P^{\ell < B}$. Therefore, each face of $L^{\ell < B'}$ makes a contribution to the Euler characteristic of $L^{\ell < B'}$ with the opposite sign to the contribution of the corresponding face in $P^{\ell < B}$. Therefore, $\chi(P^{\ell < B}) = \chi(P^{\ell < B'}) - \chi(L^{\ell < B'})$. But $L^{\ell < B'}$ has lower dimension than $P^{\ell < B}$, and is non-empty (it includes at least one vertex corresponding to an edge that can be reached from v by a single step of the simplex method for minimizing the objective function ℓ), so by induction on dimension, $\chi(L^{\ell < B'}) = 0$. Therefore, $\chi(P^{\ell < B}) = \chi(P^{\ell < B'}) - \chi(L^{\ell < B'}) = 0 - 0 = 0$. $\square$

As a simpler proof for the special case in which P is a polytope and f is a facet of P on which $\ell = B$, let ε be sufficiently small that there are no vertices of P for which ℓ is between B and $B - \varepsilon$, and let λ be the hyperplane $\ell = B - \varepsilon$. Then

$$\chi(P^{\ell < B}) = \chi(P) + \chi(P \cap \lambda) - \chi(f) = 0 + 0 - 0 = 0.$$

3.5 Bounding the Number of Faces

Theorem 3 *Let P be a polyhedron, defined as the intersection of n halfspaces in $\mathbb{R}^D$, let ℓ be a linear function which attains a minimum on P , let B be given, and suppose that $\dim P^{\ell < B} = d < D$. Then the number of faces of $P^{\ell < B}$ is $O(n^{d^2})$.*

Proof By Theorem 2, the number of d -dimensional faces is no larger than the total number of faces of all lower dimensions. Therefore, we need only count faces of dimension up to $d - 1$. But each such face is the intersection of P with the affine hull of a set of at most d vertices of $P^{\ell < B}$. There are $O(n^d)$ vertices, so there are $O(n^{d^2})$ sets of at most d vertices, and therefore $O(n^{d^2})$ faces. $\square$

For instances in general position, a considerably sharper bound may be obtained.

Theorem 4 *Let P be a polyhedron, defined as the intersection of n halfspaces in $\mathbb{R}^D$, let ℓ be a linear function which attains a minimum on P , let B be given, and suppose that $\dim P^{\ell < B} = d < D$. Additionally, suppose that the halfspaces defining P are in general position and that P has N vertices. Then the number of i -dimensional faces of $P^{\ell < B}$ is at most $N \binom{d}{i}$.*

Proof We charge each i -dimensional face f of $P^{\ell < B}$ to the vertex $\max_{\ell} f$. Let v be any vertex, and let L be its link. Then the faces charged to v correspond one-to-one with the faces of $L^{\ell < \ell(v)}$. By the general position assumption, L is a simplex, and

therefore $L^{\ell < \ell(v)}$ is also a simplex. Because the faces of $P^{\ell < B}$ correspond to faces with dimension one less in $L^{\ell < \ell(v)}$, $\dim L^{\ell < \ell(v)} \leq d - 1$. Thus, the number of i -dimensional faces of $P^{\ell < B}$ that are charged to v is at most the number of $(i - 1)$ -dimensional faces of a $(d - 1)$ -dimensional simplex (counting the empty set as a -1 -dimensional face, in the case $i = 0$), which is $\binom{d}{i}$. $\square$

Corollary 2 *Let P be a polyhedron, defined as the intersection of n halfspaces in $\mathbb{R}^D$, let ℓ be a linear function which attains a minimum on P , let B be given, and suppose that $\dim P^{\ell < B} = d < D$. Additionally, suppose that the halfspaces defining P are in general position. Then the number of faces of $P^{\ell < B}$ is $O(n^d)$.*

Proof By Theorems 1 and 4, adding the bounds on the number of faces of $P^{\ell < B}$ over all dimensions up to d gives a total that is less than $\binom{n}{d} 2^d < n^d \frac{2^d}{d!} = O(n^d)$. $\square$

4 Lower Bounds

Because of the upper bound theorem, our bounds of $O(n^d)$ on the number of vertices in a bounded subcomplex can only be tight when $d \leq D/2$. However, as our main interest is for bounded d and unbounded D , this is not a significant limitation. In this section we investigate lower bounds for small d : examples of bounded subcomplexes that show that some of our upper bounds are tight, and that our bounds on the number of bounded faces do not extend to unbounded faces even when the bounded subcomplex has low dimension.

4.1 When $d = 0$

Let P be any $(D - 1)$ -dimensional polytope, having facets determined by the linear inequalities $\bar{a}_i \cdot x \geq b_i$. Then the D -dimensional polyhedron defined by the linear inequalities $\bar{a}_i \cdot x - b_i y \geq 0$ is an unbounded cone over P : it has P as its cross-section (say in the hyperplane $y = 1$) but has only a single vertex at the origin. Examples of this form can be used to show that, even when $\dim P^{\ell < B}$ is small, the number of unbounded faces can be large. For instance, if P is a hypercube, defined as the intersection of $n = 2(D - 1)$ facets, then the cone over P has $3^{n/2} + 2$ faces despite having only a single bounded vertex.

For any polyhedron, the simplex algorithm may be used to find a path of bounded edges from any vertex to any other vertex. However, if $d = 0$, there can be no bounded edges. Therefore, in this case, there can only be one vertex. The formula of Theorem 1 always gives the same value

$$\binom{n}{0} - \binom{D}{0} + 1 = 1$$

for this case, so it is tight.

Fig. 5 Trees that may be realized as the bounded faces of two- and three-dimensional polyhedra

4.2 When $d = 1$

For $d = 1$, the bounded subcomplexes may be described as graphs, as they consist only of vertices and edges. We may characterize more precisely exactly which graphs are possible:

- When $D = 1$, the bounded subcomplex must be a single line segment.
- When $D = 2$, an unbounded two-dimensional polyhedron with n facets always has exactly $n - 1$ vertices, connected into a path of $n - 2$ edges (Fig. 5, left).
- When $D > 2$, a bounded subcomplex of dimension one must be a tree (Fig. 5, right). For, the simplex algorithm finds a path between any two vertices, so it is a connected graph, but by Euler's formula (Theorem 2) it must have exactly one fewer edge than it has vertices. The only connected graphs with this number of edges are the trees.

It is known that every tree without degree-two vertices may be realized as the Voronoi diagram of a set of points in the plane, in such a way that the leaves of the tree correspond to infinite rays in the Voronoi diagram [28]; see [5] for related tree-drawing algorithms. By a standard lifting transformation [4] this Voronoi diagram in turn may be represented as a polyhedron in three dimensions, with one facet for each of its generating sites, having the same bounded subcomplex as the Voronoi diagram. If T is any tree, then let T' be the tree formed from T by adding two new leaf nodes adjacent to each leaf in T , and one new leaf node adjacent to each degree-two vertex in T ; the Voronoi realization of T' will have T as its bounded subcomplex. Therefore, every tree may be realized as a bounded subcomplex for $D > 2$.

The Voronoi diagram of n sites in the plane that are in convex position with no four points cocircular forms a tree with n leaves and all other vertices having degree three; therefore, it lifts to a polyhedron that has n facets and $n - 2$ bounded edges. More generally, embedding the generating set of a Voronoi diagram in D dimensions into a $(D + 1)$ -dimensional space and adding one more generating site far from the others

Fig. 6 The bounded subcomplex of the Voronoi diagram of n points on the three-dimensional moment curve, for $4 \leq n \leq 7$

will lead to a Voronoi diagram with a combinatorially equivalent bounded complex. By repeating this construction, starting from the two-dimensional case, and then lifting these Voronoi diagrams to polyhedra, it is possible to find a D -dimensional polyhedron with n facets whose bounded subcomplex is a tree with $n - D + 1$ edges, for any D . The formula of Theorem 1 simplifies to

$$\binom{n}{1} - \binom{D}{1} + 1 = n - D + 1,$$

and the Voronoi-based examples presented above show that this is tight. Since the bound in Theorem 3 on the number of faces in the bounded complex uses only Euler's formula, which gives the exact number of edges in any tree, that theorem is also tight for this case.

4.3 When $d = 2$

Any n points on the three-dimensional moment curve (t, t^2, t^3) , for $t \geq 0$, determine a Voronoi diagram with $\binom{n-2}{2}$ vertices [13]. The same lifting transformation used in the preceding section takes this Voronoi diagram to an unbounded four-dimensional polyhedron, with n defining halfspaces. In this polyhedron, all facets are unbounded (corresponding to the fact that all of the defining sites of the Voronoi diagram are in convex position and therefore define unbounded Voronoi cells), so the maximum dimension of a bounded face is two. There are quadratically many bounded faces (Fig. 6). As this example shows, a nontrivial bound on d does not necessarily exclude examples with superlinear complexity. Corollary 1 states that, for $d = 2$, the number of vertices in the bounded complex is $O(n^2)$, and this example shows the number is also $\Omega(n^2)$; i.e., the bound is tight in this case. More strongly, it shows that the formula of Theorem 1 is tight to within a $1 - O(1/n)$ factor in this case.

It is consistent with the bound in Theorem 3 for the number of edges and 2-faces of a two-dimensional bounded subcomplex to be quadratically larger than the number of vertices. To show that this is possible, we consider another example, the tight span of the metric space of distances in a complete bipartite graph $K_{n,n}$ with $n > 1$. That is, if U and V are the two sides of the bipartition, with n vertices each, then the distance between any two members of U or between any two members of V is defined to be two, while the distance from any member of U to any member of V is defined to be

Fig. 7 The tight span of $K_{n,n}$ has the combinatorial structure of a cone over $K_{n,n}$, with a triangle for each edge in the graph. This example (with $n = 3$) forms a two-dimensional bounded subcomplex of a six-dimensional polytope

one. The tight span of this space consists of all functions f from $U \cup V$ to the real numbers satisfying the following two properties:

- For any two vertices x and y (not necessarily distinct), $f(x) + f(y) \geq d(x, y)$.
- For any vertex x , there exists a vertex y with $f(x) + f(y) = d(x, y)$.

Distances between two different functions in the tight span are measured according to the sup norm, i.e., $\text{dist}(f, g) = \sup_{x \in U \cup V} |f(x) - g(x)|$.

Let f be a function in the tight span, and let $u \in U$ and $v \in V$ have the minimum value of f on each side of the bipartition. Then (by the second property of tight spans) $f(u)$ and $f(v)$ are both at most one, from which the defining properties of tight spans imply that, for every $u' \in U \setminus \{u\}$, $f(u') = 2 - f(u)$ and for every $v' \in V \setminus \{v\}$, $f(v') = 2 - f(v)$. Additionally, it must be the case that $f(u) + f(v) \geq 1$. For any choice of u and v , and any two numbers $f(u)$ and $f(v)$ with $f(u) \leq 1$, $f(v) \leq 1$, and $f(u) + f(v) \geq 1$, these properties define a unique function f in the tight span of $K_{n,n}$. Thus, geometrically, the tight span may be formed by gluing together n^2 triangles defined by the inequalities $f(u) \leq 1$, $f(v) \leq 1$, and $f(u) + f(v) \geq 1$, one triangle for each choice of u and v , forming a space with the structure of a cone over $K_{n,n}$ (Fig. 7). This space has $2n + 1$ vertices, $n^2 + 2n$ edges, and n^2 triangles. However, note that in this case the $2n$ -dimensional polyhedron for which this is the bounded subcomplex has $4n^2$ defining halfspaces, so it does not show that Theorem 3 itself is tight.

We do not know whether it is possible for a two-dimensional bounded subcomplex, defined from a polytope of n facets, to have simultaneously a number of vertices that is quadratic in n and a number of edges and faces that is quadratic in the number of vertices.

4.4 Additional examples

Let P be a four-dimensional pyramid over an octahedron, modified by folding its octahedral facet so that it is split into two square pyramids. Then the subset of faces of P that are nonadjacent to the apex of P are dual to a one-dimensional complex that takes the form of a single line segment: the two square pyramid facets are dual to the endpoints of the line segment, and the square ridge separating these two facets is dual to the line segment itself. However, if the vertices of P are perturbed into general position, then the square becomes a flat tetrahedral facet while each of the two square

pyramids is split into two tetrahedra; the dual complex becomes two triangles joined at a vertex. As this example shows, some instances cannot be perturbed into general position without increasing the dimension of the bounded subcomplex.

Observation 1 *There exists a polyhedron P for which the dimension of its bounded subcomplex is strictly smaller than the dimension of the bounded subcomplex of any general position perturbation of P .*

Let Q_D be the D -dimensional hypercube, formed by intersecting $n = 2D$ half-spaces determined by the inequalities $0 \leq x_i$ and $x_i \leq 1$, and let σ be the linear function $\sum x_i$. Then, for any integer $0 \leq d \leq D$, $\dim Q_D^{\sigma < d+1/2} = d$. The number of vertices of $Q_D^{\sigma < d+1/2}$ is $\sum_{i=0}^d \binom{D}{i}$. For any even integer n , and any constant value of d , this provides an example of a bounded subcomplex in general position with $\Theta(n^d)$ vertices.

Observation 2 *For every constant d , the bound of Corollary 1 is tight up to constant factors.*

5 Algorithms

We now discuss algorithms for finding the bounded subcomplex of a polyhedron. There are many alternative solutions available, depending on whether the dimension d of the subcomplex is known or unknown to the algorithm, on whether only the vertices of the polyhedron need to be found or whether the whole bounded subcomplex is to be constructed, and on whether we use as a subroutine the algorithm for computing bounded subcomplexes of Herrmann et al. [22] or whether we directly enumerate the faces of the bounded subcomplex, using linear programming to test whether each face is bounded.

Our algorithms will necessarily involve the solution of linear programs and linear feasibility problems. Especially in the case of inputs that are not in general position, it is important that these linear programs are solved in an exact model of computation that allows solution vertices to be compared for equality: it is possible that a small perturbation of an input problem, as might occur due to round-off error in a non-exact numerical linear programming algorithm, could significantly increase the dimension of the bounded subcomplex. However, this sort of exact computation model is standard in computational geometry algorithms, and solutions are available when the ambient dimension D is of moderate size [30]. Even when D is large, strongly polynomial algorithms are known for some special cases of linear programming, such as the case with two variables per inequality that arises in the tight span construction [32]. We let L denote the time to solve a linear program of ambient dimension D and n constraints, or to determine whether such a program is infeasible or unbounded.

It may be possible to extend our algorithms to work in non-exact computation models, provided that the input is described by limited precision rational numbers, using interior point methods and rounding techniques. See e.g., [34].

5.1 When the Dimension Bound is Known

When the input consists of a set of n halfspaces together with the dimension d of the bounded subcomplex, we have the following results:

- We can construct all vertices of the given polyhedron in time $O(n^d L)$. This follows immediately from Lemma 2 which gives a formula $v = \min_{\ell}(P \cap A)$ allowing each vertex v of P to be obtained as the solution to a linear program in a subspace of dimension $D - d$.
- If the input is in general position, we can construct the bounded subcomplex in time $O(n^d L + n^{5d})$. This method generates the vertices as above, exhaustively tests each vertex-facet pair to find all the vertex-facet incidences, and then uses the algorithm of Herrmann et al. [22] to construct the bounded subcomplex from the vertex-facet incidences. The algorithm of Herrmann et al. takes time $O(N^2 \beta \phi + N^3 \phi^2)$, where N is the number of vertices of the given polyhedron (here at most $O(n^d)$), β is the number of vertex-facet incidences (at most $O(n^{d+1})$), and ϕ is the number of faces of the bounded subcomplex (again, at most $O(n^d)$).
- If the input is not in general position, then N can be at most $O(n^d)$ and ϕ can be at most $O(n^{d^2})$, as our previous bounds show. Plugging these bounds into the analysis of the same method of Herrmann et al., in place of the tighter bounds for general position inputs, shows that this method takes time $O(n^d L + n^{2d^2+3d})$.
- If the input is in general position, we can alternatively construct the bounded subcomplex in time $O(n^{2d} L)$. The method is to maintain a list of faces of the bounded subcomplex, initially containing all the vertices. Then, for each face f added to the list, and each vertex v , we compute the affine hull of $f \cup v$, use linear programming subproblems to determine whether this affine hull lies on the boundary of the polyhedron and has a vertex of maximum value, and if so check that the face contained in this affine hull is distinct from the ones already discovered. Whenever we find a new face in this way we add it to the list. There are $O(n^{2d})$ face-vertex pairs, each taking time $O(L)$ to check, so the time bound is as given.
- If the input is not in general position, the same method takes time $O(n^{d^2+d} L)$.

5.2 When the Dimension Bound is Unknown

When d is not given as input (as seems more likely to occur in most applications of this problem), constructing the bounded subcomplex becomes more complicated. A naive approach would be to repeatedly apply one of the algorithms for known d , described above, for increasing values of d until it calculates a correct answer. It is possible to save some redundant work in this approach, by interleaving the construction of vertices (as the solutions of linear programming subproblems with decreasing dimension) with the construction of faces of increasing dimension; in this way, the total work for the whole sequence of d repetitions of the known- d algorithm may be bounded by the time for the repetition with the largest value of d . However, this does not provide a complete solution to the problem, because “until it calculates a correct answer” is not

Fig. 8 Lemma 3 states that if there exists a vertex of P that is the minimum on the intersection of d hyperplanes but no set of less than d hyperplanes, then there also exists a vertex of P that is the minimum on the intersection of $d - 1$ hyperplanes but no set of less than $d - 1$ hyperplanes. In the figure, $d = 2$

a well-specified condition for our algorithm to test. How can we tell, in an algorithm for computing a bounded subcomplex with unknown d , whether we have chosen d to be large enough? That is, how can we tell when our algorithm has found all the bounded faces and can stop? An answer is provided by the following lemma, which we illustrate in Fig. 8.

Lemma 3 *Let H be a given set of halfspaces, let P be the D -dimensional polyhedron formed as the intersection of the halfspaces in H , let $d \leq D$ be a number, and let ℓ be a linear function in general position with respect to P (in the same sense as used in Lemma 2, namely that no two vertices have the same value of ℓ). For any given set S of halfspaces in H , let P_S denote the intersection of P with the hyperplanes bounding the halfspaces in S . Let V be the set of all vertices of P of the form $\min_{\ell} P_S$ for sets S of at most d halfspaces in H , and F be a set of bounded faces of P that includes the empty face. Suppose additionally that the following three conditions are all true:*

1. *For each set S of at most $d + 1$ halfspaces in H , $\min_{\ell} P_S$ (if it exists) belongs to V .*
2. *Every face in F has dimension at most d .*
3. *For every pair (f, v) of a face f in F and a vertex v in V , with A being the affine hull of $f \cup v$, one of the following three possibilities is true:*
 - (a) *A contains an interior point of P and $d < D - 1$,*
 - (b) *Some face $f' \in F$ has affine hull A , or*
 - (c) *ℓ is unbounded on $A \cap P$.*

Then F is the bounded subcomplex of the given polyhedron.

Proof We assume for a contradiction that there is a bounded face f that does not belong to F ; among all such faces, let f have the minimum possible dimension, and let $v = \max_{\ell} f$. If f had dimension at most d , then by the same argument as in the proof of Lemma 2, v would belong to V , and the pair (f', v) would violate condition (3) of the lemma where f' is any facet of f nonincident to v . We can assume without loss of generality that the lower link of v (the subset of the link consisting of the faces whose values in ℓ are entirely less than that of the vertex itself) has dimension exactly d : it has dimension at least d , because it contains a face of dimension $\dim f - 1$ corresponding to f , and if it had a higher dimensional face then we could reduce the dimension by one unit per step by moving from f to a facet of this higher dimensional face; condition (2) of the lemma ensures that, at each step, we continue to have a face f that does not belong to F . But, if the lower link of v has dimension exactly d ,

then by Lemma 2, v would be representable as $\min_{\ell} P_S$ for a set S of at most $d + 1$ halfspaces of H , violating condition (1) of the lemma. This contradiction shows that no such face f exists. $\square$

Theorem 5 *Given a polyhedron whose bounded complex has dimension d (with d unknown to the algorithm) we can construct the bounded complex in time $O(n^{d^2+d}L)$. If the input is in general position, we can construct the bounded complex in time $O(n^{2d}L)$.*

Proof We choose a linear function ℓ in general position; this may be simulated symbolically by starting with an arbitrary linear function such as the one that maps each point to the value of its first coordinate and then applying a tie-breaking rule that compares tuples of coordinates lexicographically [7]. We then apply the algorithm for known d , for increasing values of d , until the condition of the lemma is met. Testing the condition takes time $O(n^{d^2+d}L)$, or $O(n^{2d}L)$ for inputs in general position, matching the time for the known- d algorithm on the correct value of d . $\square$

In all of the algorithms described in this section, the solution is found within an amount of time and a number of linear programming subproblems that is polynomial for any fixed value of d .

It would be of interest to determine whether the method of Herrmann et al. [22], which uses many fewer linear programs at the expense of greater running time in the other parts of the algorithm, can be adapted to the case where d is unknown.

6 Discussion

We have shown that, for any fixed d , bounded subcomplexes of polyhedra that have dimension at most d have polynomial complexity. Our bounds on the numbers of vertices of the bounded subcomplex are tight for small values of d . However, our bounds on the numbers of faces of higher dimensions do not appear to be tight. In other bounds on the complexity of polytopes, it is possible to considerably simplify the problem by assuming that the input is in general position; for our problem, too, such an assumption would be very helpful (it would lower the exponent of the polynomial from d^2 to d), but it is not always possible to perturb the input into general position without changing its dimension. Nevertheless, we would like to know whether instances that are not in general position have the same complexity as, or higher complexity than, instances that are in general position.

For the problem of constructing tight spans, the high exponents in our bounds become even more unsatisfactory, because the number of halfspaces of the polyhedron defining the tight span as its bounded subcomplex is already quadratic in the number of points of the input metric space. For instance, for tight spans of n -point metric spaces that have only two-dimensional faces (that is, when $d = 2$), we get a bound of $O(n^8)$ on the complexity of the tight span. In contrast, the first author [18] has proven that under a stronger two-dimensionality assumption, that the tight span is homeomorphic to a subset of the plane, it may be described as a complex of only $O(n)$ cells and may be constructed in time $O(n^2)$, optimal since the input distance matrix has size

$O(n^2)$. It would be of interest to obtain tighter bounds for d -dimensional tight spans, even for the case $d = 2$ without the additional requirement that the tight span be homeomorphic to a planar set.

Acknowledgments This work was supported in part by NSF Grant 0830403 and by the Office of Naval Research under Grant N00014-08-1-1015.

References

1. Avis, D., Fukuda, K.: A pivoting algorithm for convex hulls and vertex enumeration of arrangements and polyhedra. *Discrete Comput. Geom.* **8**(1), 295–313 (1992). doi:[10.1007/BF02293050](https://doi.org/10.1007/BF02293050)
2. Avis, D., Bremner, D., Seidel, R.: How good are convex hull algorithms? *Comput. Geom. Theory Appl.* **7**(5–6), 265–301 (1997). doi:[10.1016/S0925-7721\(96\)00023-5](https://doi.org/10.1016/S0925-7721(96)00023-5)
3. Barber, C.B., Dobkin, D.P., Huhdanpaa, H.L.: The quickhull algorithm for convex hulls. *ACM Trans. Math. Softw.* **22**(4), 469–483 (1996). doi:[10.1145/235815.235821](https://doi.org/10.1145/235815.235821)
4. Brown, K.Q.: Voronoi diagrams from convex hulls. *Inf. Process. Lett.* **9**(5), 223–228 (1979). doi:[10.1016/0020-0190\(79\)90074-7](https://doi.org/10.1016/0020-0190(79)90074-7)
5. Carlson, J., Eppstein, D.: Trees with convex faces and optimal angles. In: *Proceedings of the 14th International Symposium on Graph Drawing, Lecture Notes in Computer Science*, Vol. 4372, pp. 77–88. Springer, Berlin (2006). doi:[10.1007/978-3-540-70904-6_9](https://doi.org/10.1007/978-3-540-70904-6_9)
6. Chand, D.R., Kapur, S.S.: An algorithm for convex polytopes. *J. ACM* **17**, 78–86 (1970). doi:[10.1145/321556.321564](https://doi.org/10.1145/321556.321564)
7. Charnes, A.: Optimality and degeneracy in linear programming. *Econometrica* **20**, 160–170 (1952)
8. Chazelle, B.: An optimal convex hull algorithm in any fixed dimension. *Discrete Comput. Geom.* **10**(1), 377–409 (1993). doi:[10.1007/BF02573985](https://doi.org/10.1007/BF02573985)
9. Chrobak, M., Larmore, L.L.: Generosity helps or an 11-competitive algorithm for three servers. *J. Algorithms* **16**, 234–263 (1994). doi:[10.1006/jagm.1994.1011](https://doi.org/10.1006/jagm.1994.1011)
10. Develin, M.: A complexity bound on faces of the hull complex. *Discrete Comput. Geom.* **32**(4), 471–479 (2004). doi:[10.1007/s00454-004-1134-2](https://doi.org/10.1007/s00454-004-1134-2)
11. Develin, M.: Dimensions of tight spans. *Ann. Comb.* **10**(1), 53–61 (2006). doi:[10.1007/s00026-006-0273-y](https://doi.org/10.1007/s00026-006-0273-y)
12. Develin, M., Sturmfels, B.: Tropical convexity. *Doc. Math.* **9**, 1–27 (2004)
13. Dewdney, A.K., Vranich, J.K.: A convex partition of $\mathbb{R}^3$ with applications to Crum’s problem and Knuth’s post-office problem. *Utilitas Math.* **12**, 193–199 (1977)
14. Dress, A.W.M.: Trees, tight extensions of metric spaces, and the cohomological dimension of certain groups. *Adv. Math.* **53**, 321–402 (1984). doi:[10.1016/0001-8708\(84\)90029-X](https://doi.org/10.1016/0001-8708(84)90029-X)
15. Dyer, M.E.: The complexity of vertex enumeration methods. *Math. Oper. Res.* **8**(3), 381–402 (1983)
16. Eppstein, D.: Finding the k shortest paths. *SIAM J. Comput.* **28**(2), 652–673 (1998). doi:[10.1137/S0097539795290477](https://doi.org/10.1137/S0097539795290477)
17. Eppstein, D.: Manhattan orbifolds. *Topol. Appl.* **157**(2), 494–507 (2009). doi:[10.1016/j.topol.2009.10.008](https://doi.org/10.1016/j.topol.2009.10.008)
18. Eppstein, D.: Optimally fast incremental Manhattan plane embedding and planar tight span construction. *J. Comput. Geom.* **2**(1), 144–182 (2011)
19. Frederickson, G.N.: Ambivalent data structures for dynamic 2-edge-connectivity and k smallest spanning trees. In: *Proceedings of the 32nd IEEE Symposium on Foundations of Computer Science*, pp. 632–641, October 1991. doi:[10.1109/SFCS.1991.185429](https://doi.org/10.1109/SFCS.1991.185429)
20. Grünbaum, B.: *Convex Polytopes*. Graduate Texts in Mathematics, Vol. 221, 2nd edn. Springer, New York (2003)
21. Hamacher, H.W., Queyranne, M.: K best solutions to combinatorial optimization problems. *Ann. Oper. Res.* **4**(1), 123–143 (1985). doi:[10.1007/BF02022039](https://doi.org/10.1007/BF02022039)
22. Herrmann, S., Joswig, M., Pfetsch, M.E.: Computing the bounded subcomplex of an unbounded polyhedron. *Comput. Geom. Theory Appl.* **46**(5), 541–551 (2013). doi:[10.1016/j.comgeo.2011.11.002](https://doi.org/10.1016/j.comgeo.2011.11.002)
23. Hirai, H.: Characterization of the distance between subtrees of a tree by the associated tight span. *Ann. Comb.* **10**(1), 111–128 (2006). doi:[10.1007/s00026-006-0277-7](https://doi.org/10.1007/s00026-006-0277-7)

24. Isbell, J.R.: Six theorems about injective metric spaces. *Comment. Math. Helv.* **39**, 65–76 (1964). doi:[10.1007/BF02566944](https://doi.org/10.1007/BF02566944)
25. Karmarkar, N.: A new polynomial-time algorithm for linear programming. *Combinatorica* **4**(4), 373–395 (1984). doi:[10.1007/BF02579150](https://doi.org/10.1007/BF02579150)
26. Khachiyan, L.G.: A polynomial algorithm in linear programming. *Soviet Math. Dokl.* **20**(1), 191–194 (1979)
27. Klee, V.: Polytope pairs and their relationship to linear programming. *Acta Math.* **133**, 1–25 (1974)
28. Liotta, G., Meijer, H.: Voronoi drawings of trees. *Comput. Geom. Theory Appl.* **24**(3), 147–178 (2003). doi:[10.1016/S0925-7721\(02\)00137-2](https://doi.org/10.1016/S0925-7721(02)00137-2)
29. Matheiss, T.H., Rubin, D.S.: A survey and comparison of methods for finding all vertices of convex polyhedral sets. *Math. Oper. Res.* **5**(2), 167–185 (1980)
30. Matoušek, J., Sharir, M., Welzl, E.: A subexponential bound for linear programming. *Algorithmica* **16**(4), 498–516 (1996). doi:[10.1007/BF01940877](https://doi.org/10.1007/BF01940877)
31. McMullen, P.: The maximum numbers of faces of a convex polytope. *Mathematika* **17**, 179–184 (1970). doi:[10.1112/S0025579300002850](https://doi.org/10.1112/S0025579300002850)
32. Megiddo, N.: Towards a genuinely polynomial algorithm for linear programming. *SIAM J. Comput.* **12**(2), 347–353 (1983). doi:[10.1137/0212022](https://doi.org/10.1137/0212022)
33. Queyranne, M.: Structure of a simple scheduling polyhedron. *Math. Program.* **58**(1), 263–285 (1993). doi:[10.1007/BF01581271](https://doi.org/10.1007/BF01581271)
34. Roos, C., Terlaky, T., Vial, J.-P.: *Interior Point Methods for Linear Optimization*. Springer, New York (2006)
35. Seidel, R.: Constructing higher-dimensional convex hulls at logarithmic cost per face. In: *Proceedings of the 18th ACM Symposium on Theory of Computing (STOC '86)*, pp. 404–413 (1986). doi:[10.1145/12130.12172](https://doi.org/10.1145/12130.12172)
36. Seidel, R.: The upper bound theorem for polytopes: an easy proof of its asymptotic version. *Comput. Geom. Theory Appl.* **5**(2), 115–116 (1995). doi:[10.1016/0925-7721\(95\)00013-Y](https://doi.org/10.1016/0925-7721(95)00013-Y)
37. Spielman, D.A., Teng, S.-H.: Smoothed analysis of algorithms: why the simplex algorithm usually takes polynomial time. *J. ACM* **51**(3), 385–463 (2004). doi:[10.1145/990308.990310](https://doi.org/10.1145/990308.990310)
38. Swart, G.: Finding the convex hull facet by facet. *J. Algorithms* **6**(1), 17–48 (1985). doi:[10.1016/0196-6774\(85\)90017-3](https://doi.org/10.1016/0196-6774(85)90017-3)
39. Ziegler, G.M.: *Lectures on Polytopes*. Graduate Texts in Mathematics, Vol. 152. Springer, New York (1995)