Shadows and slices of polytopes

Nina Amenta*
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304, USA
amenta@parc.xerox.com

Günter M. Ziegler**
Dept. Mathematics 6-1
Technische Universität Berlin
10623 Berlin, Germany
ziegler@math.tu-berlin.de
http://www.math.tu-berlin.de/~ziegler

Abstract

We give a lower bound of $\Omega(f^{\lfloor d/2 \rfloor})$ for the number of vertices of a d-dimensional polytope with f facets which can appear on the outer boundary of a projection to any dimension $2 \le k < d$. By duality, this implies a lower bound of $\Omega(n^{\lfloor d/2 \rfloor})$ for the number of facets in a k-dimensional slice of a d-dimensional polytope with n vertices. At the same time, the Upper Bound Theorem provides an $O(n^{\lfloor d/2 \rfloor})$ upper bound for this quantity. For cyclic polytopes, however, we show an upper bound of $O(n)$ on this quantity in dimension four. We give a new algorithm for the construction of the boundary of the projection.

* Partially supported by National Science Foundation grant number CCR-9404113. Some of this work was done while the author was at The Geometry Center, supported by NSF/DMS-8920161, and during a visit to the Freie Universität Berlin, supported by the Deutscher Akademischer Austauschdienst.

** Supported by a DFG Gerhard-Hess-Forschungsförderungspreis (Zi 475/1-1) and by the German-Israeli Foundation grant I-0309-146.06/93.

1 Introduction

A d-dimensional polytope with f facets may have no more than

$$M(f,d) \;:=\; \binom{f - \lceil \frac{d}{2} \rceil}{\lfloor \frac{d}{2} \rfloor} + \binom{f - 1 - \lceil \frac{d-1}{2} \rceil}{\lfloor \frac{d-1}{2} \rfloor}$$

vertices, which is $O(f^{\lfloor \frac{d}{2} \rfloor})$; this is the dual statement of the Upper Bound Theorem for polytopes. This bound is achieved by the duals of the cyclic polytopes, defined below. The *shadow* of a d-dimensional polytope P is the set of points (x_1, x_2) such that some point $(x_1, x_2, \ldots, x_d)$ belongs to P, or, equivalently, the projection of P to the (x_1, x_2)-plane. The shadow is a convex polygon. How many vertices can the shadow of a d-polytope with f facets have? This question can be generalized in the obvious way, to k-dimensional shadows (k-shadows) for $2 \le k \le d-1$, where it also makes sense to ask for the number of i-faces, $0 \le i \le k-1$.

It is not inherently unreasonable to hope that the complexity of the shadow of P might be asymptotically less than that of P itself. We show, however, that the k-dimensional shadow of a polytope with f facets in d dimensions might, in the worst case, have complexity $\Omega(f^{\lfloor \frac{d}{2} \rfloor})$. Let $M_2(f,d)$ denote the maximal number of vertices in the shadow (i.e., of any 2-dimensional projection) of a simple d-polytope with f facets. Then $M_2(f,d) \le$

$M(f,d)$ is immediate. We show that both functions have the same order $\Theta(f^{\lfloor d/2 \rfloor})$ for fixed d. This is somewhat puzzling, since our results on projections of duals of cyclic polytopes suggest that $M_2(f,d)$ may not coincide with $M(f,d)$ even for $d=4$.

Shadows are natural objects in combinatorial geometry, and they have a number of algorithmic applications. Perhaps most importantly, the shadow vertex algorithm for linear programming chooses a simplex path by following a two-dimensional shadow [GS, Borg]. Lower bounds on the size of the two-dimensional shadow provide lower bounds for this algorithm. Exponential lower bounds in the special case $f = 2d$ were given by Murty [Mu] and Goldfarb [Gol1, Gol2]; our example handles the case where d is fixed and provides worst-case bounds of $\Theta(f^{\lfloor d/2 \rfloor})$. This is in contrast to the results of Borgwardt [Borg], who has established a polynomial bound of at most $O(d^4 f)$ for the expected number of steps of the shadow vertex algorithm on random linear programs. A randomized version of the shadow-vertex rule that may be polynomial on every linear program was suggested by Gärtner & Ziegler [**?**].

Our main result is a construction for a d-dimensional polytope, with at most f facets, such that $M_2(d,f) = O(f^{\lfloor d/2 \rfloor})$. In the full version of this paper we define a class of polytopes, called *deformed products*, which include these polytopes, the constructions of Goldfarb, above, the polytopes of Klee and Minty [KlMi], and a number of other bad examples for various simplex rules.

The case of fixed dimension is important also for other optimization problems involving shadows. One optimization problem is the maximization of a convex function in k variables $x_1, \ldots, x_k$ over a polytope P. The maximum is achieved at some vertex of the k-shadow. In [AAAS] this maximization is applied to finding the largest similar copy of one convex polygon contained in another, a problem with applications in vision and robotics. The computation of shadows was studied in [PoFa] and [PSSBM]. In [PoFa] a set of stable three-finger grasps, with friction, of a polygon in the plane, is represented by the three-shadow of a five-polytope. This result is extended in [PSSBM], where a subset of the stabe four-finger grasps, with friction, of a polyhedron in R^3 is computed by taking the 8-shadow of an 11-dimensional polytope. Their algorithms and experimental results are reviewed in section 6.

This question about shadows is only interesting in dimensions four and higher. In three dimensions, it is not too difficult to construct simple polytopes in which every vertex appears on the shadow. Other questions concerning the shadows of three-dimensional polytopes are considered in [CEG].

2 Results

We relay on the following observation.

Observation 1 *A lower bound on the complexity of the two-dimensional shadow of a polytope also is a lower bound on the complexity of any k-shadow.*

This becomes obvious when we imagine doing the projection to two dimensions by projecting first to dimension k and then to the plane. Any vertex which shows up on the planar shadow has to correspond to at least one vertex on the k-shadow.

To prove the theorem, then, it is sufficient to exhibit a d-dimensional polytope with f facets which has a two-dimensional shadow with $O(f^{\lfloor d/2 \rfloor})$ vertices. Our construction of such a polytope essentially follows an example by Klee and Minty [KlMi] of a polytope with a long monotone path. We do the construction so as to be able to project the whole path to the plane. This gives us our main theorem.

Theorem 2 *For all $d \geq 4, f \geq 2d$, there is a d-dimensional polytope P with at most f facets, such that the k-shadow of P has $O(f^{\lfloor d/2 \rfloor})$ vertices.*

This result answers an equally natural question in the dual setting. The dual of a polytope P with f facets is a polytope with f vertices. But what is the dual of its k-shadow? The $(d-1)$-dimensional shadow of P is the intersection of the linear halfspaces parallel to the x_d axis and containing P. In the dual, this is the intersection of the dual of P with the hyperplane $x_d = 0$. So the dual of the k-shadow of a polytope is the intersection of the dual polytope with a k-dimensional hyperplane.

Corollary 3 *For all $d \geq 4, n \geq 2d$, there is a d-dimensional polytope P with at most n vertices, and a k-dimensional plane p in $\mathbb{R}^d$, such that the intersection $p \bigcap P$ is a k-polytope with $O(n^{\lfloor d/2 \rfloor})$ facets.*

Recall that the projection of the lower envelope of a polytope to $\mathbb{R}^{d-1}$ is a regular $(d-1)$-dimensional triangulation of the projected vertices of the lower envelope. If, in the example above, we consider a projection to $\mathbb{R}^{d-1}$ in any direction parallel to p, we get the following configuration.

Corollary 4 *There is a regular triangulation T of a set of n points in $\mathbb{R}^{d-1}$, and a $(k-1)$-plane p, which intersects $O(n^{\lceil d/2 \rceil})$ of the simplices of T.*

The cyclic polytope duals maximize the number of vertices over all d-dimensional polytopes with f facets, providing the lower bound example matching the Upper Bound Theorem. Somewhat surprisingly, we show the following upper bound on the complexity of the shadow of a cyclic polytope dual.

Theorem 5 *The projection of the dual of a 4-dimensional cyclic polytope with f facets to the plane can have at most $3f$ vertices.*

We also give an example of a projection of a 4-dimensional cyclic polytope dual which achieves $3f - 10$ vertices on the boundary of the projection.

Finally, although the k-shadow of a polytope P may have asyptotically as many vertices as P itself, in many cases it has far fewer. We survey existing algorithms for the output-sensitive construction of k-shadows, and present a new algorithm which is efficient for large k.

3 The fourth dimension

In this section we develop the 4-dimensional case in detail. In the following section, we give the higher dimensional generalization.

Theorem 6 *There is a 4-dimensional polytope P with $2m$ facets, such that the shadow of P has $m(m+1)/2$ vertices.*

Proof: We construct P in three steps: first, we take the cross-product of two m-gons to get a 4-dimensional polytope P'. Then we deform P', without changing its combinatorial structure, to make a new polytope P''. Finally we perform a projective transformation of P'' to get P.

Let A be an m-gon, m divisible by 4, in the (x_1, x_2) coordinate plane, with vertices evenly spaced on the unit circle, so that vertex $v_i(A) = (\cos(i\alpha), \sin(i\alpha))$ for $\alpha = 2\pi/m$ and $i = 0 \ldots m-1$, and edge $e_i(A) = v_i v_{i+1}$ ($i+1$ is taken mod m, here and throughout; that is, the edge $v_i v_{i+1}$ might be $v_m v_1$). Let B be the same m-gon in the x_3, x_4 coordinate plane. The cross-product of A and B (the set of all points with x_1, x_2 in A and x_3, x_4 in B) is a 4-dimensional polytope P'.

A facet of P' is the cross-product either of A with an edge of B (an A-facet), or of B with an edge of A (a B-facet), so P' has $2m$ facets, each a cylinder over an m-gon. There are two kinds of two-faces.

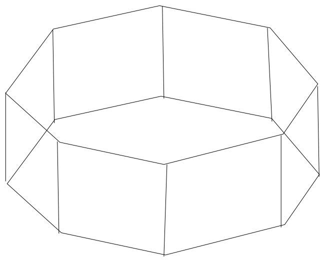

Figure 1: A facet of P'

The cross-product of an edge of A with an edge of B is a square, a side of a cylindrical three-face. (The topologically inclined may notice that these square faces are a polygonalization of the flat torus.) The cross-product of B with a vertex $v_i(A)$ is a copy of B in the two-flat $(x_1, x_2) = v_i(A)$. We will call these m-gonal faces B-ridges; the A-ridges are defined similarly. There are m B-ridges, each containing m distinct vertices, so P' has m^2 vertices. We write $u_{i,j}$ for the vertex which is the cross product of $v_i(A)$ with $v_j(B)$.

For a fixed i, the orthographic projection to the (x_1, x_2)-plane takes all the vertices in a B-ridge B_i to a single point (see Figure 3). We now deform P' into P'', so that these vertices are distributed along a line segment in the projection, without disturbing the combinatorial structure of the polytope.

We do this by tilting each of the B-facets of P'. All the B-facets are parallel to the x_3 and x_4 axes. For i even, we tilt the supporting three-plane of the B-facet containing edge $e_i(A)$ in towards the positive x_3 axis, maintaining the incidence with $e_i(A)$ and keeping it parallel to the x_4 axis. For i odd, we tilt it towards the negative x_3 axis in the same way and by the same amount. We use a gentle enough angle, defined precisely in a moment, so that the combinatorial structure of P'' remains the same as that of P'.

After the tilting, the three-planes supporting the B-facets are defined by linear equations in (x_1, x_2, x_3). Each B-ridge lies in the intersection of two B-facets. We can use this intersection to eliminate the x_3 variable, which means that a B-ridge lies in a three-plane determined by a linear equation in (x_1, x_2), so that it does indeed project to a line segment in the x_1, x_2 plane.

In order to verify that we can accomplish this tilting without changing the combinatorial structure, we consider the motion of the vertices induced by the tilting. A vertex is the intersection of two adjacent A-facets and two adjacent B-facets. The intersection of two adjacent A-facets is a two-plane with constant x_3, x_4 coordinates, so tilting the B-facets will not affect the x_3, x_4 coordinates of the vertices.

We consider the extremal two-plane p_{max} in the positive x_3 direction containing A-ridge A_0. The intersection of p with any plane supporting a B-facet is a line, and the intersections of all of the positive halfspaces of these lines is A_0, an m-gon. For even i, the tilting causes the line corresponding to the B-facet through $e_i(A)$ to move towards the origin, without changing its slope. For odd i, the line moves away from the origin. So the edges of A_0 corresponding to even i get longer, and the edges corresponding to odd i get shorter. The behavior of $A_{m/2}$ on the the extremal two-plane in the negative x_3 direction is just the opposite; even edges get shorter and odd edges get longer. See Figure 2.

In the central two-planes containing ridges $A_{m/4}$ and $A_{3m/4}$, there will be no motion at all, and the intermediate two-planes will exhibit more moderate behavior than the extremal ones.

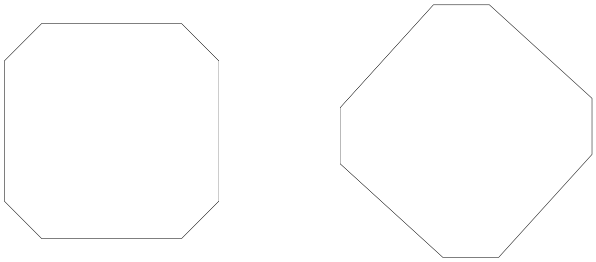

Figure 2: A_0 and $A_{m/2}$ after tilting

We claim that so long as we tilt the the three-planes supporting the B-facets by a small enough amount so that none of the shrinking edges of the A-ridges disappear entirely, the combinatorial structure of the resulting polytope P'' remains the same as that of P'.

A combinatorial change would occur if the orientation between a vertex and one of the three-planes supporting a facet changed. This cannot happen in the case of the A-facets, since they are supported by planes defined by a single linear equation in x_3, x_4, and the x_3, x_4 coordinates of all the vertices do not change. In the case of the B-facets, the orientation of all the vertices in each A-ridge remains the same with respect to each three-plane supporting a B-facet, since each A-ridge remains a convex m-gon. This establishes the claim.

The final step in the construction is just a projective transformation. For some very small constant ϵ, we multiply every point in the space by the matrix

$$\begin{bmatrix} 1 & 0 & 0 & \epsilon & 0 \\ 0 & 1 & 0 & \epsilon & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

The resulting polytope is P. The projection of P'' to the x_1, x_2 plane took every B-ridge B_i to a line segment. This transformation "adds back" some of the x_4 coordinate of each 4-dimensional vertex to its x_1, x_2 coordinates, so that the vertices of B_i lie on an ellipse in the projection, with the arc containing vertices $u_{i,1}, \ldots, u_{i,m/2-1}$, where x_4 is positive, curving away from the origin and onto the convex hull. See Figure 3 once again. The

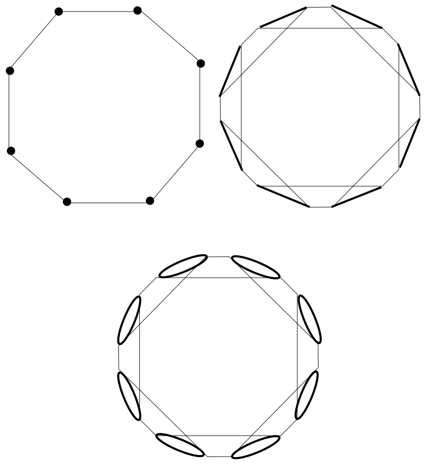

Figure 3: The projections of the ridges B_i in P', P'', and, below, P

constant ϵ can be chosen small enough so that the angle formed, in the projection, between $u_{i-1,m}, u_{i,0}, u_{i,1}$ remains smaller than π. In that case, the vertices $u_{i,0}, \ldots, u_{i,m/2}$ appear on the boundary of the projection, for every B_i, giving the shadow a total of $m(m+1)/2$ vertices. $\square$

Note that B might be replaced with a roughly semi-circular $(m/2+1)$-gon, the convex hull of vertices $v_0, \ldots, v_{m/2}$, giving a polytope with fewer facets but the same number of vertices on the shadow. This improves the constants in the construction but makes it uglier.

4 Main theorem

We now generalize the 4-dimensional construction to any higher dimension d. Basically, we replace B in the construction above with a $(d-2)$-dimensional polytope whose shadow has $O(f^{\lfloor d/2-1 \rfloor})$ vertices.

Theorem 7 *For all d, there is a d-dimensional polytope P with f facets, such that the shadow of P has $O(f^{\lfloor d/2 \rfloor})$ vertices.*

Proof: Let A be a planar $f/2$-gon in the (x_1, x_2) coordinate plane. We recursively construct a $(d-2)$-dimensional polytope $\mathcal{B}$ in $x_3, \ldots, x_d$, with $f/2$ facets, such that the shadow of $\mathcal{B}$ in the (x_3, x_4)-plane has $O(f^{\lfloor d/2-1 \rfloor})$ vertices. For simplicity, we also stipulate that the shadow is a polygon symmetrical about the x_3 and x_4 axes, with unique maximal and minimal vertices in x_3, properties that this construction recursively ensures.

We take the cross-product $\mathcal{P}'$ of A and $\mathcal{B}$. The A-facets $\mathcal{P}'$ are the cross-products of A with the facets of $\mathcal{B}$, and the $\mathcal{B}$-facets are the cross-products of $\mathcal{B}$ with the edges of A, so $\mathcal{P}'$ has f facets. A $\mathcal{B}$-ridge is again the cross-product of $\mathcal{B}$ with a vertex of A.

We now deform $\mathcal{P}'$ into $\mathcal{P}''$. For even edges i of A, we tilt the corresponding $\mathcal{B}$-facet slightly towards the positive x_3 axis, maintaining its contact with the edge i and keeping it parallel to the $x_4, \ldots, x_d$ axes. For odd edges i, we similarly tilt the corresponding $\mathcal{B}$-facet towards the negative x_3 axis.

We argue again that a small enough tilting leaves the combinatorial structure of $\mathcal{P}''$ the same as that of $\mathcal{P}'$. An A-facet is supported by a half-space orthogonal to a facet of $\mathcal{B}$, a linear equation in $x_3, \ldots, x_d$. The A-facets are unmoved, and the intersection of $d-2$ of them is a two-plane with constant $x_3, \ldots, x_d$ coordinates. A vertex is the intersection of $d-2$ A-facets and two $\mathcal{B}$-facets, so the $x_3, \ldots, x_d$ coordinates of each vertex are unchanged by the tilting, and the relationship of each A-facet with the vertices is undisturbed.

Now consider the two-spaces formed by the intersection of the hyperplanes supporting $d-2$ adjacent A-facets (the facets are adjacent if they are the cross-product of a single edge of A with the $d-2$ facets of $\mathcal{B}$ meeting at a vertex). The intersection of the halfspace supporting a $\mathcal{B}$-facet with such a two-plane is a half-plane. Before the tilting, the intersections of these $f/2$ halfplanes form an $f/2$-gon identical to A. Consider the minimal and maximal two-faces with respect to x_3 coordinate, p_{min} and p_{max}. In p_{max}, the tilting again moves the even edges of these $f/2$-gons in towards the origin, and the odd ones outwards, and it does the opposite in p_{min}. Again, if we tilt the $\mathcal{B}$ facets gently enough so that all of these two-faces remain convex $f/2$-gons, no combinatorial change can occur between a $\mathcal{B}$-facet and a vertex as a result of the tilting.

Finally, we apply the projective transformation

$$\begin{bmatrix} 1 & 0 & 0 & \epsilon & 0 & \cdots & 0 & 0 \\ 0 & 1 & 0 & \epsilon & 0 & \cdots & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & \cdots & 0 & 0 \\ & & & \cdots & & & & \\ 0 & 0 & 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 1 \end{bmatrix}$$

to $\mathcal{P}''$ to produce $\mathcal{P}$. This "adds back" some of the x_4 coordinate of each vertex to the (x_1, x_2)-coordinates, causing every $\mathcal{B}$-ridge to project to a convex polygon in the (x_1, x_2)-plane. Again, ϵ can be chosen small enough so that half of the vertices of every $\mathcal{B}$-ridge end up on the boundary of the shadow. $\square$

5 The shadow of a cyclic polytope

The dual of a cyclic polytope has the maximum number of faces of all dimensions among polytopes with f facets. We show, however, that the duals of cyclic polytopes do not maximize the complexity of the shadow among all polytopes with f facets. In fact we show that the shadow of a 4-dimensional cyclic polytope dual must have asymptotically fewer vertices than the polytope itself.

Let us review the definition and properties of a cyclic polytope; for more details, see [Z]. Let C be a *curve of order* d in $\mathbb{R}^d$, meaning that any $(d-1)$-plane intersects C in at most d points. The convex hull of any set of n points on C is a d-dimensional cyclic polytope P_d. Sturmfels [Stur] has shown that any polytope that is combinatorially isomorphic to a cyclic polytope has its vertices on some curve of order d. In four and higher dimensions, every pair of vertices in P_d is connected by an edge, in dimension six or higher every triple form a two-face, and so on.

A facet of P_d is supported by a $(d-1)$-plane which passes through d vertices. Let us index the vertices $v_1, \ldots, v_n$ along C. If C passes outside a facet at a vertex v_i, it must come back inside at v_{i+1}, since otherwise v_{i+1} would be outside P_d. So the set of vertices determining a facet is made up of adjacent pairs v_i, v_{i+1}. Every face of smaller dimension is determined by a subset of the set of vertices determining a facet. In four dimensions this means that every two-face of P_4 is the convex hull of three vertices $\{v_i, v_{i+1}, v_j\}$, with v_j distinct from both v_i and v_{i+1} (recall that $i+1$ is taken mod n, as above).

Theorem 8 *Let P_4 be any cyclic polytope in $\mathbb{R}^4$ with n vertices, and let P_4' be the dual of P_4. The shadow of P_4' in $\mathbb{R}^2$ may have at most $3n$, and might have at many as $3n-10$, vertices.*

Proof: Any projection of the dual P_4' into $\mathbb{R}^2$ is the dual of the the intersection of the cyclic polytope P_4 with some two-plane F_2. This intersection is a convex polygon P_2 in F_2. We show that P_2 has at most $3n$, and might have as many as $3n-10$, vertices.

Consider each possible vertex v_j in turn. There are at most $n-2$ triangular faces T_{ij} of P_4 involving v_j, one for each possible v_i. Construct a three-plane F_{v_j} through F_2 and v_j. If F_2 hits a triangular face T_{ij}, then the edge $\{v_i, v_{i+1}\}$ is cut, in P_4, by the halfspace of F_{v_j} bounded by F_2 and *not* containing v_j. The vertices v_i and v_{i+1} are

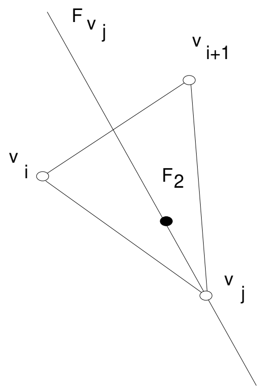

Figure 4: Projection to $\mathbb{R}^2$ along F_2. F_2 projects to a point.

connected by a segment of C; if one of them lies below F_{v_j} and the other lies above it, this segment of C must also cross F_{v_j} at least once. But any three-plane intersects C at most four times; and for F_{v_j}, one of those intersections is v_j. Hence there are at most three pairs v_i, v_{i+1} separated by F_{v_j}, and there are at most three vertices of P_2 for every vertex v_j. This gives the upper bound of $3n$.

Now we construct a cyclic polytope that realizes the lower bound. Select two three-planes F_{v_n} and F_{v_5}, intersecting in a common two-plane F_2. Figure 5 again represents the projection from $\mathbb{R}^4$ to $\mathbb{R}^2$ along F_2. Select an order four curve C such that intersections of the 3-flats with C occur in the order indicated.

Let v_n be the intersection of F_{v_n} with C, v_5 be the intersection of F_{v_5} with C, and put vertices $v_6, \ldots, v_{n-1}$ along C, between v_5 and v_n. Finally, position vertices v_1 through v_4 on C as shown, so that the segments connecting the adjacent pairs $\{v_1, v_2\}$, $\{v_2, v_3\}$ and $\{v_3, v_4\}$ all cross both F_{v_5} and F_{v_n}.

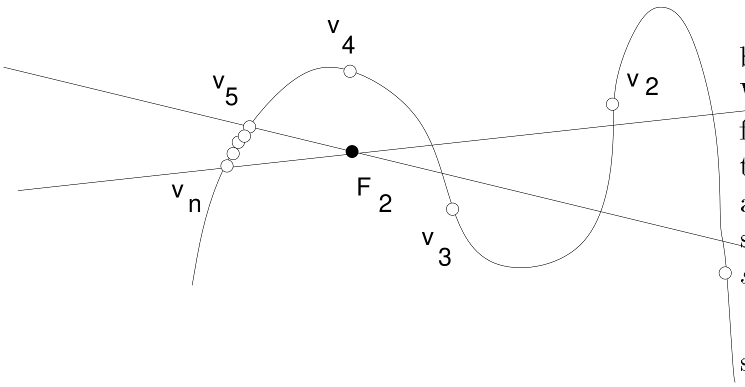

Figure 5: Projection to $\mathbb{R}^2$ along F_2.

Every 2-face of P_4 formed by $\{v_1, v_2\}$, $\{v_2, v_3\}$ or $\{v_3, v_4\}$, together with any one of the vertices $v_5, \ldots, v_n$, crosses F_2. This gives $3(n-4)$ vertices in P_2. In addition, the two-faces $\{v_1, v_4, v_5\}$ and $\{v_4, v_n, v_1\}$ also cross F_2, for a total of $3n - 10$. $\square$

6 Computation of shadows and slices

In this section we review the known algorithms for computing k-shadows, and give a new one which should be more efficient in some situations.

Since we have shown that the asymptotic complexity of a polytope shadow can be as great as that of the polytope itself, a worst-case optimal algorithm for computing the k-shadow of a polytope P given as the intersection of a family H of halfspaces is to compute P, using a worst-case optimal algorithm, and then test each of its faces to determine which of them fall on the boundary of the projection into the k-dimensional subspace. So the interesting problem is to develop an *output-sensitive* algorithm, where the running time is bounded by some function of the size of the shadow. This problem is closely related to the more basic problem of finding an output-sensitive algorithm for convex hulls, which is still not entirely solved in higher dimensions.

We might measure the size of a k-shadow by counting either its facets or its vertices. We will let s be the number of shadow facets and v be the number of shadow vertices. Since a vertex in the polytope P is adjacent to at most $\binom{d}{k-1}$ faces of dimension $(k-1)$, a shadow vertex is as well, so $s = O(v)$. But v can be $O(s^{\lfloor k/2 \rfloor})$.

In the following discussion we will assume that P is simple, that is, that H is in general position. This can be ensured using standard perturbation techniques, although the size of P, and of its shadow, might increase as result. We also assume, without loss of generality, that no face is parallel to the x_1 axis.

We define a *vertical* flat as one which is parallel to the direction of projection, that is, to the $x_{k+1}, \ldots, x_d$ axes. Notice that there is exactly one vertical hyperplane through every non-vertical $(k-1)$-flat.

Significant work on the computation of k-shadows appears in Ponce, et al. [PSSBM]. They analyze two algorithms and report on experiments with implementations of both on their grasp polytopes. The first algorithm is a refinement of the Fourier-Motzkin elimination algoithm for convex hulls. We test each $(d-k+1)$-tuple of input halfspaces to determine whether their intersection supports a vertical hyperplane. Those that do are potential shadow facets. Each of these is then tested, by linear programming, to see if it is in fact a shadow facet. The running time is $O(f^{d-k+2})$. [1] This algorithm only finds the shadow facets; it does not compute their convex hull to find the shadow vertices. This is sufficient for their application.

The second algorithm is a pivoting algorithm. We find the minimum vertex in the x_1 direction, which is a shadow vertex, by linear programming. We then trace the

[1] Note that we use the notation $(d-k)$ where [PSSBM] uses k, and visa versa.

the shadow by following the edges that are supported by vertical hyperplanes. At each vertex p, we have to find the vertex at the other end of each new shadow edge e. We intersect the ray supporting e anchored at p with each of the input halfspaces. The intersection nearest p is the next vertex. The running time of this second algorithm is $O(vf)$.

They implemented both algorithms, and tested them on a sequence of 11-dimensional grasp polytopes. The running time pivoting algorithm was a slow growing linear function in the size of the polytope, while the running time of the elimination algorithm grew quickly with f, as expected. Theoretically, they note that the running time of the pivoting algorithm can be improved using the ray-shooting data structure of Matoušek and Schwazkopf [Mat] (a similar observation is found in [Chan]).

There is another approach to output-sensitive convex hulls, due to Seidel [Sei], which is more efficient for computing intersections of halfsapces when the number of vertices greatly exceeds the number of input halfspaces. This is a plane-sweep algorithm; a sweep-hyperplane moves across P from vertex to vertex, maintaining the $(d-1)$-dimensional slice of P.

The algorithm takes advantage of the fact that, in higher dimensions, most of the events in the sweep are vertices for which every hyperplane which participates in an outgoing edge also participates in an incoming edge (which is certainly not the case in dimensions two and three!). Such events can be scheduled dynamically, once the sweep-hyperplane intersects all the incoming edges. Scheduling and processing all these events requires $O(v \lg f)$ time.

The remaining events are those vertices at which a new facet makes its first appearance. These events are scheduled before the start of the sweep, by doing a $(d-1)$-dimensional linear program in the bounding hyperplane of each input halfspace, to determine where, if ever, it first intersects P. This requires $O(f^2)$ time, for an overall running time of $O(f^2 + v \lg f)$. The $O(f^2)$ term can be improved to $O(f^{2-2/(\lfloor \frac{d}{2} \rfloor)} + \epsilon)$ by using the data structure for linear prgramming queries due to Matoušek and Schwarzkopf.

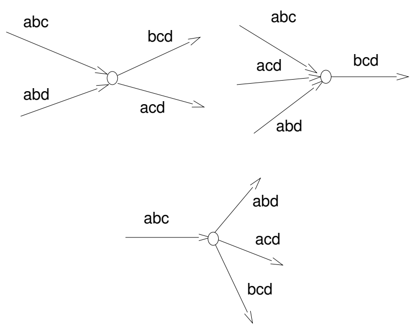

Figure 6: Two vertices of a four-polytope which can be scheduled dynamically, and, below, one which cannot. The labels on the edges indicate incedent facets.

We adapt this convex hull algorithm to get an algorithm for higher dimensional k-shadows that is more efficent when the number of shadow vertices greatly exceeds the number of shadow facets. Here is the basic algorithm.

Theorem 9 *A k-shadow of a polytope P, given as the intersection of an input set H of halfspaces, can be computed in time $O(sf^2 + v \lg f)$, where $f = |H|$, s is the number of facets in the shadow, and v is the total number of faces of the shadow.*

Proof: The idea is to trace, by pivoting, the $(k-1)$-faces of P which project to shadow facets, and find the minimal vertex of each one with respect to the sweep direction, and finally to run the sweep, exactly as in the convex hull algorithm, to compute the k-dimensional convex hull of the projected shadow facets.

We will sweep in the increasing x_1 direction. We find the vertex p of P which minimzes x_1 by linear programming. The vertex p is adjacent to $\binom{d}{k-1}$ faces of dimension $(k-1)$, some of which will be supported by vertical hyperplanes, and hence project to shadow facets.

A $(k-1)$-face f of P lies in the intersection of some set H_f of $(d-k+1)$ input hyperplanes. An adjacent $(k-1)$-face f' lies in the intersetion of some some set H'_f, which differs from H_f in only one element. For each known shadow facet f, we find all the adjacent shadow facets as follows. For each input hyperplane h not in H_f, we substitute h for each of the $(d-k+1)$ hyperplanes of H_f in turn to get a different $(k-1)$ flat f'. For efficiency, we can avoid checking any $(d+-k+1)$-tuple more than once by using a dictionary. If f' does not have a vertical supporting hyperplane, we discard it. Otherwise we then run a $(k-1)$-dimensional linear program f', again with the sweep direction as the objective function and with the intersections of the input halfspaces with f' as constraints, and find the first point p in which f' first intersects P. If f' fails to intersect P, then it does not support a shadow facet; otherwise the point p found by the linear program projects to the first point on the shadow facet encountered by the sweep plane.

A shadow facet might be adjacent to as many as f others. For each shadow facet, we find all adjacet facets by running at most f linear programs in dimension $(k-1)$, each requiring $O(f)$ time. This requires $O(sf^2)$ time.

With the first vertex of every shadow facet known, we can compute their convex hull using the sweep phase of Seidel's convex hull algorithm, in $O(v \lg f)$ time. □

Notice that the first phase of the algorithm consists of fs linear programming queries on the set H of input halfspaces, so we can apply the linear programming query data structure. To optimize the overall running time, the preprocessing time for the data structure must be balanced against the time required to answer all the queries. Since we do not know s in advance, we need to apply the (fairly standard) trick of answering queries $2^i f$ through $2^{i+1}-1f$ using a data structure designed for 2^i queries, as in the remarks after Corollaries 2.4 and 3.4 in [Chan]. This gives us the following theorem.

Theorem 10 *A k-shadow of a polytope P, given as the intersection of an input set H of f halfspaces, can be computed in time $O((sf^2)^{1-1/(\lfloor\frac{d}{2}\rfloor)+\epsilon} + v \lg f)$, where s is the number of shadow facets and v is the number of shadow vertices.*

The first stage of our algorithm finds the shadow facets at least as efficiently as the elimination algorithm of [PSSBM], and, in addition, it sets up the computation of the convex hull. Our algorithm is more efficient than their pivoting algorithm when $v > O(sf)$, which is possible only for $k \geq 4$.

Acknowledgments

We are grateful to Bernd Gärtner for helpful discussions.

References

[AAAS] Pankaj Agarwal, Nina Amenta, Boris Aronov and Micha Sharir. Largest placements and motion planning of a convex polygon, submitted to *Discrete and Computational Geometry.*

[Borg] Karl Heinz Borgwardt. The Simplex Method. A Probabilistic Analysis, *Algorithms and Combinatorics* **1**, Springer 1987.

[Chan] Timothy Chan. Output-sensitive results on convex hulls, extreme points and related problems, *Proceedings of the 11th Annual Symposium on Computational Geometry* (1995), 10–19.

[CEG] Bernard Chazelle, Herbert Edelsbrunner and Leonidas J. Guibas. The complexity of cutting complexes, *Discrete and Computational Geometry* **4** (1989), 139–182.

[GZ] Bernd Gärtner and Günter M. Ziegler. Randomized simplex algorithms on Klee-Minty cubes, *Proceedings of the 35th Annual Symposium on Foundations of Computer Science* (1994) 502-510.

[GS] Saul I. Gass and Thomas Saaty. The computational algorithm for the parametric objective function, *Naval Research Logistics Quarterly* **2** (1955), 39-45.

[Gol1] Donald Goldfarb. Worst case complexity of the shadow vertex simplex algorithm, *preprint*, Columbia University (1983), 11 pages.

[Gol2] Donald Goldfarb. On the complexity of the simplex algorithm, in: *Advances in optimization and numerical analysis*, Proc. 6th Workshop on Optimization and Numerical Analysis, Oaxaca, Mexico, January 1992; Kluwer, Dordrecht 1994, 25-38.

[KlMi] Victor Klee and George J. Minty. How good is the simplex algorithm?, *Inequalities III*, (O. Shisha, ed.), Academic Press, New York, (1972), 159–175.

[Mat] Jiří Matoušek. Linear optimaization queries, *Journal of Algorithms* **14** (1993), 432-448. Also, with Otfried Schwarzkopf in *Proceedings of the 8th Annual Symposium on Computational Geometry* (1992) 16-25.

[Mu] Katta G. Murty. Computational complexity of parametric linear programming, *Mathematical Programming* **19** (1980), 213-219.

[PoFa] Jean Ponce and Bernard Faverjon. On computing three-finger force-closure grasps of polygonal objects, to appear in *IEEE Transactions on Robotics and Automation*, **11**:6, (1995).

[PSSBM] Jean Ponce, Steve Sullivan, Attawith Sudsang, Jean-Daniel Boissonnat and Jean-Pierre Merlet. On computing four-finger equilibrium and force-closure grasps of polyhedral objects, to appear in the *International Journal of Robotics Research*, (1996).

[Sei] Raimund Seidel. Constructing higher dimensional convex hulls at logarithmic cost per face, *Proceedings of the 18th Annual Symposium on the Theory of Computing*, (1986), 404-413.

[Stur] Bernd Sturmfels. Cyclic polytopes and d-order curves, *Geometriae Dedicata* **24** (1987) 103-107.

[Z] Günter M. Ziegler. *Lectures on Polytopes*, Graduate Texts in Mathematics **152**, Springer-Verlag, New York 1995.