

The Complexity of Cutting Complexes*

Bernard Chazelle,¹ Herbert Edelsbrunner,² and Leonidas J. Guibas³

¹ Department of Computer Science, Princeton University, Princeton, NJ 08544, USA

² Department of Computer Science, University of Illinois, Urbana, IL 61801, USA

³ Department of Computer Science, Stanford University/DEC-SRC, Palo Alto, CA 94301, USA

Abstract. This paper investigates the combinatorial and computational aspects of certain extremal geometric problems in two and three dimensions. Specifically, we examine the problem of intersecting a convex subdivision with a line in order to maximize the number of intersections. A similar problem is to maximize the number of intersected facets in a cross-section of a three-dimensional convex polytope. Related problems concern maximum chains in certain families of posets defined over the regions of a convex subdivision. In most cases we are able to prove sharp bounds on the asymptotic behavior of the corresponding extremal functions. We also describe polynomial algorithms for all the problems discussed.

1. Introduction

Given a convex subdivision of the plane, how should we place a straight line in order to intersect the most regions? Intuitively, the maximum number of intersections, the so-called *line span* of the subdivision, cannot be too small. Two companion problems thus arise: to establish the complexity of computing the line span, and to find sharp bounds on the minimum line span for subdivisions of a given size. For variants of this problem likely to occur in practice, we can place constraints on the number of allowed directions for edges of the subdivision, or set an upper bound on the number of edges adjacent to any region.

A related problem on subdivisions originates from motion-planning and computer graphics. Choose a direction ℓ and construct a directed graph as follows.

* Bernard Chazelle wishes to acknowledge the National Science Foundation for supporting this research in part under Grant No. MCS83-03925. Herbert Edelsbrunner is pleased to acknowledge the support of Amoco Fnd. Fac. Dev. Comput. Sci. 1-6-44862.

The vertices of the graph correspond to the regions of the subdivision. We place an edge from region r to region r' if, informally speaking, r can push r' in the direction ℓ . It is easy to see that this process always creates a dag. Taking the maximum size of the longest path over all directions ℓ gives the so-called *monotone span* of the subdivision. As before, we must ask how hard it is to compute the monotone span, and how small can the smallest monotone span get for subdivisions of a given size.

Related problems arise in three dimensions. Given a convex polytope, how many facets can a single plane intersect (the *cross-section-span* problem)? Given a source of light outside the polytope and a fixed screen, what is the largest number of edges on any projection of the polytope (the *silhouette-span* problem)? Note that if we restrict the source of light to be placed at infinity we then have the classical *shadow* problem of Moser [MP].

This paper investigates both the computational and combinatorial aspects of these problems. Our main contribution is to provide sharp bounds for most of the functions defined above. For example, we show that any n -facet convex polytope has a cross-section of span $\Omega(\log n / \log \log n)$; moreover this bound is tight.¹ This result solves an old open problem [MP] in a fairly unexpected manner: indeed, $O(\log n)$ was the prevalent conjecture. The extremal functions mentioned so far are of the form min-max. For completeness we also study the corresponding max-min questions. In all cases, we are able to give algorithms for computing these functions in time ranging from $O(n^2)$ to $O(n^3)$.

What is the significance of our results? From a combinatorial perspective, our work sheds light on the general study of extremal properties of planar subdivisions and convex polytopes. This involves a comprehensive investigation of several classes of problems and their natural variants. The proof techniques which we introduce enable us to derive optimal bounds for most of the problems considered. From a computational standpoint this paper introduces a unifying framework for solving various optimization problems efficiently. Heavy use is made of known algorithmic tools such as depth-first search, topological sweep, dual transforms, etc. The main appeal of our approach resides in its generality. We expect to see further applications of our techniques.

Our principal results, along with the necessary definitions, are given in the next section.

2. Summary of the Results

We begin by recalling the definition of a convex subdivision of the Euclidean plane E^2 . A set of open convex subsets (called *regions*) of E^2 defines a convex subdivision (or subdivision for short) if no two regions intersect and the union of their closures form E^2 . If the number of regions is finite, then each region is the intersection of a finite number of open half-planes. The relative interior of the intersection of the closures of two regions, if nonempty, is called a *vertex* if

¹ All logarithms are taken to the base 2, unless specified otherwise.

it is a point and an *edge* otherwise. A subdivision is said to be of size n if it is made of n convex regions, and it is of *degree* d if every region is adjacent to at most d other regions (two regions are *adjacent* if their closures contain a common edge). We define the *line span* of a subdivision as the maximum number of regions which can be intersected by a single line (Section 3).

Theorem 1 (Line Span). *The line span of a subdivision of size n can be computed in $O(n^2)$ time and $O(n)$ space. The minimum line span over all subdivisions of size n is in $\Theta(\log n / \log \log n)$. More generally, if only subdivisions of degree d or less are considered then the bound is in $\Theta(\log_d n + \log n / \log \log n)$. On the other hand, if the edges of the subdivisions assume at most a constant number of distinct slopes then the minimum line span is $\Theta(\log n)$.*

Given a subdivision S and a directed line ℓ , we say that a region r pushes an adjacent region r' if there exists a directed line parallel to ℓ which intersects r and r' in that order. (Similar incidence relations were studied by Guibas and Yao [GY].) We make the regions of S the vertices of a directed graph G and we put an edge from r to r' if and only if r pushes r' (Fig. 2.1). It is not difficult to show that G is acyclic and therefore that its longest path (or any of them if there are several) has finite length. This path (or any of them if there are several) is called the *contact chain* of S in direction ℓ . Let θ be the angular slope of ℓ ($-\pi/2 \leq \theta < \pi/2$). It is clear that the length of the contact chain depends only on the slope of ℓ and is the same whether ℓ is directed one way or the reverse. For this reason, we call G the *contact graph* of the subdivision in direction θ (and do not distinguish between θ and $\theta + \pi$). The maximum length of a contact chain taken over all angles θ in $[-\pi/2, \pi/2)$ is referred to as the *monotone span* of the subdivision. Our next result states that, surprisingly, line spans and monotone spans lead to similar bounds (Section 4).

Theorem 2 (Monotone Span). *The monotone span of a subdivision of size n can be computed in $O(n^2)$ time and $O(n)$ space. The minimum monotone span over all subdivisions of size n is in $\Theta(\log n / \log \log n)$. More generally, if only subdivisions*

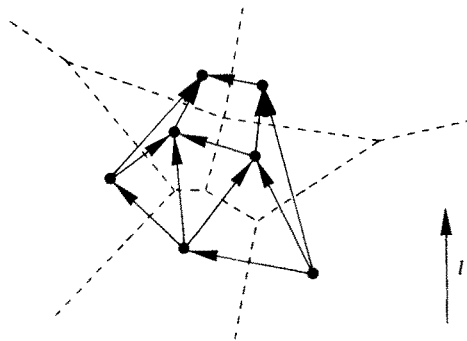


Fig. 2.1

of degree d or less are considered then the bound is $\Theta(\log_d n + \log n / \log \log n)$. If the number of distinct slopes in the subdivision is at most k then the minimum monotone span is in $\Omega(n^{1/k} + \log n / \log \log n)$.

Informally, a convex polytope is the bounded intersection of a finite number of closed half-spaces (see [G] for more details). We define the *cross-section span* of a convex polytope in E^3 as the maximum number of facets which can be intersected by a single plane. By analogy with subdivisions we say that a polytope is of degree d if no facet is adjacent to more than d others (Section 5).

Theorem 3 (Cross-Section Span). *The cross-section span of a convex polytope of n facets can be computed in $O(n^3)$ time and $O(n)$ space. The minimum cross-section span over all n -facet polytopes is $\Theta(\log n / \log \log n)$. If the degree of the polytope is d or less then the bound is $\Theta(\log_d n + \log n / \log \log n)$.*

Given a source of light and a fixed screen, the projection of a convex polytope on the screen is a convex polygon. The *silhouette span* of the polytope is the maximum number of edges of the projection polygon over all positions of the polytope. There is a one-to-one correspondence (actually, a dual mapping) between cross-sections and silhouettes. Thus, we can state an immediate corollary of Theorem 3. (Section 5).

Theorem 4 (Silhouette Span). *The silhouette span of a convex polytope of n facets can be computed in $O(n^3)$ time and $O(n)$ space. The minimum silhouette span over all n -facet polytopes is $\Theta(\log n / \log \log n)$. If no vertex of the polytope is incident to more than d edges then the bound is $\Theta(\log_d n + \log n / \log \log n)$.*

We have also obtained companion theorems concerning the max-min versions of all the problems mentioned above. Interestingly, when nontrivial, most of the combinatorial bounds are linear, while the complexity bounds remain pretty much the same (Section 6). The max-min version of the line span is called the *line width*; similarly we have the *monotone width*, *cross-section width*, and *silhouette width*: to give meaning to the latter quantities we require that the cutting plane (resp. source of light) passes through a fixed point (resp. be placed at infinity). In the following we call a silhouette whose light source is at infinity a *parallel view*.

Theorem 5 (Max-Min Problems). *The maximum line (monotone, cross-section, silhouette) width of a subdivision (polytope) of size n is $\Theta(n)$. In all cases, computing widths can be done within the same asymptotic complexity as computing spans.*

We have thus summarized the main results of this paper and included references to the appropriate sections. For completeness let us only add that Section 7 offers a short discussion of what we have solved and what we have left open.

3. The Line Span of a Convex Subdivision

This section consists of four subsections, the first three of which prove some of the combinatorial bounds claimed in Theorem 1. Section 3.1 establishes a lower bound on the line span of a convex subdivision of bounded degree. In Section 3.2 we relax the bounded-degree condition and prove a companion result. Section 3.3 deals with the issue of bounding the number of slopes. Finally, Section 3.4 addresses the problem of computing the line span of a given convex subdivision. All these partial results combine together to prove Theorem 1, as discussed in Section 3.5.

In the following, we use the notation $l_d(n)$ to designate the minimum line span over all subdivisions of size n and degree d . Since no degree need be defined above $n-1$, the minimum line span over all subdivisions of size n is $l_{n-1}(n)$, which we denote simply by $l(n)$.

3.1. Regions of Bounded Degree

In this section we establish a lower bound on $l_d(n)$ which is asymptotically tight unless d exceeds $\log n$. The proof is intimately based on the notion of wedges and subwedges. Two lines that intersect in a point p cut the plane into four open regions called wedges, each of which has p as its *apex*. A region W_2 is a *subwedge* of a wedge W if

$$W_2 = h \cap W_1,$$

where h is an open half-plane which does not contain p and W_1 is a wedge with apex p contained in W . The edge of W_2 contributed by h is called the *cap* of W_2 ; the other edges come from W_1 .

Lemma 3.1. *For any $n \geq 2$ and $d \geq 3$, we have $l_d(n) > \log_{d-1} n + \log_{d-1}(d-2)$.*

Proof. Let S be a convex subdivision of degree $d \geq 3$ with $n \geq 2$ regions, and let W_0 be a wedge with apex p that satisfies the following conditions:

- (i) W_0 intersects every region of S .
- (ii) W_0 contains at most $d-2$ vertices of the region that contains p .

If S contains a region with 0 or 1 vertices, then choose p inside this region. Otherwise, let v be a vertex of S that is also a vertex of the convex hull of all points that are vertices of S . Let r be an unbounded region of S that is incident to v , and let e be the edge of the convex hull that intersects r . Choose p in r but outside the convex hull and sufficiently close to the edge e . If one half-line bounding W_0 goes through vertex v and the other passes through the other endpoint of e , then W_0 satisfies (i) and (ii).

The wedge W_0 is the starting point of an iterative process. We show by induction that there exists a sequence of subwedges of W_0 ,

$$W_0 \supseteq W_1 \supseteq W_2 \cdots,$$

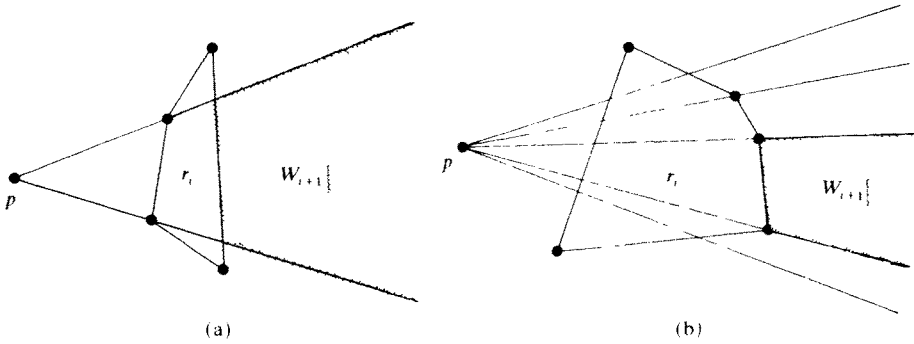


Fig. 3.1

such that W_i intersects w_i regions of S , with $w_0 = n$ and $w_{i+1} \geq (w_i - 1)/(d - 1)$, for $i \geq 0$. Let r_i be the region of S that intersects W_i and contributes its cap. Initially, r_0 is the region that contains p , and by convention, we say that r_0 contributes the (empty) cap of W_0 . If W_i contains no vertex of r_i , then we set

$$W_{i+1} = \text{int}(W_i \setminus r_i)$$

(Fig. 3.1(a)). Otherwise, r_i has at most $d - 2$ vertices within W_i . The lines passing through p and these vertices split W_i into at most $d - 1$ subwedges. Let W'_i be a subwedge that intersects the most regions of S . We set

$$W_{i+1} = \text{int}(W'_i \setminus r_i)$$

(Fig. 3.1(b)). We iterate on this process until W_{i+1} is empty.

It is also possible that $W_i \setminus r_i$ consists of two connected components. In this case, r_i must be unbounded and thus have at most $d - 1$ vertices. To embed this case in the general case (Fig. 3.1(b)) we act as if the two unbounded edges of r_i had a common vertex and we draw a line through p and this vertex (which separates the components of $W_i \setminus r_i$) when we recurse. Clearly, we have $w_0 = n$ and $w_{i+1} \geq (w_i - 1)/(d - 1)$. Thus, $w_i \geq u_i$, where $u_0 = n$ and $u_{i+1} = (u_i - 1)/(d - 1)$. Now,

$$u_i = \frac{n}{(d - 1)^i} - \left(\frac{1}{d - 1} + \frac{1}{(d - 1)^2} + \dots + \frac{1}{(d - 1)^i} \right)$$

for $i \geq 1$. We immediately derive

$$u_i > \frac{n}{(d - 1)^i} - \frac{1}{d - 2},$$

which implies that u_i , and therefore w_i , is positive if $i \leq \log_{d-1}((d - 2)n)$.

Notice that any line through p that intersects W_i also intersects regions r_0, \dots, r_{i-1} . Therefore, a line through p that intersects the last nonempty wedge

W_j intersects at least $j + 1$ regions, namely r_0 through r_{j-1} , but also r_j , a region which contains W_j . By the argument above we derive

$$j > \log_{d-1}((d-2)n) - 1,$$

which completes the proof of the lemma. □

It is interesting to notice that we get a lower bound of $\Omega(\log_{d-1} n)$ for almost any choice of the point p . The key is to avoid vertices of S . The proof is very similar to that of Lemma 3.1. Unfortunately, the lower bound becomes very weak as d grows, so a new line of attack is clearly needed. To begin with, we give up the liberty of choosing p anywhere we please.

3.2. General Convex Subdivisions

Recall that $l(n)$ denotes the minimum line span over all subdivisions of size n .

Lemma 3.2. For any $n > 2$, we have $l(n) \geq \log n / \log \log n$.

Proof. Let S be a subdivision of size n . There always exists a direction which we call *horizontal*, which is not parallel to any of the lines formed by joining pairs of vertices in S . Let K be the closed strip between two horizontal lines and let Σ be the subdivision of K induced by S . If L is an arbitrary nonhorizontal line, we say that a region of Σ is *critical* (with respect to L) if it intersects L . Let k be the number of critical regions. Through the topmost and bottommost vertex of each critical region, draw the longest horizontal line segment that does not intersect a critical region. (This is a well-defined operation since by definition regions are open.) The effect is to subdivide V , the complement of the critical regions in K , into a number of polygons with the following characteristics. Each one is bounded by two horizontal edges connected by two convex polygonal paths (Fig. 3.2). There might be degenerate cases where the two paths share an

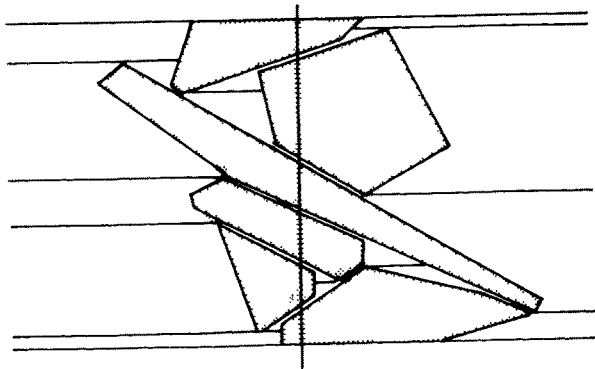


Fig. 3.2

edge, or edges are of null length, etc. In all cases, however, V gets to be subdivided into at most $3k - 1$ polygons. Why is that so? Initially, that is, before drawing the horizontal line segments, V consists of a single connected component (since critical regions are open). Let us count the polygons as they are encountered by a horizontal line h sweeping K bottom up. At the beginning h intersects two polygons. Then the line encounters one new polygon every time it sweeps over a topmost vertex of a critical region, and two new polygons every time it sweeps over a bottommost vertex (Fig. 3.2). This gives at most $3k - 1$ polygons since there is no contribution from the topmost vertex of the topmost critical region. This establishes our claim.

Among the polygons thus created, the one that intersects the most regions of Σ is called the *kernel* of (Σ, L) . If Σ has m regions, then S induces a subdivision of the kernel into at least $(m - k)/(3k - 1)$ regions.

We are now ready to go back to the original problem. Let k be any integer larger than some appropriate constant k_0 . We argue that if no nonhorizontal line intersects at least k regions of S then there is a horizontal line which intersects at least $\lceil \log_{3k-4}[(3k-5)n/(k-1)] \rceil$ regions. This yields a lower bound of

$$\max_{k \geq k_0} \min \left\{ k, \log_{3k-4} \frac{(3k-5)n}{k-1} \right\} \geq \frac{\log n}{\log \log n}$$

for any n large enough. To prove our claim we define an iterative construction, starting with K_0 equal to the entire plane and L_0 being an arbitrary nonhorizontal line. Let Σ_0 be the subdivision of K_0 induced by S . If L_0 intersects at least k regions of K_0 then we are done. Otherwise, let K_1 be the kernel of (K_0, L_0) and let Σ_1 be the subdivision of K_1 induced by S . A horizontal line which contains no vertex of S and intersects K_1 also intersects a critical region of K_0 . We iterate on this process inside K_1 . Because of the particular shape of K_1 we can find a nonhorizontal line L_1 which intersects K_1 in a single line segment s , such that the smallest horizontal strip that contains s also contains K_1 . This construction ensures that at stage i every horizontal line that contains no vertex of S and intersects the kernel K_i of (K_{i-1}, L_{i-1}) also intersects at least one critical region in each of K_0 through K_{i-1} . We can iterate on this process until either we encounter a nonhorizontal line crossing at least k regions of V_i , the complement of the critical regions in K_i , or until V_i is empty, whichever comes first.

The latter case can occur only if L_i intersects at most $k - 1$ regions of K_i . By the pigeonhole principle the kernel K_i intersects at least u_i regions of K_i , and hence of S , where $u_0 = n$ and

$$u_{i+1} = \frac{u_i - (k-1)}{3(k-1) - 1}.$$

If we fall in the second case then the index j of the last nonempty kernel must satisfy $u_{j+1} \leq 0$, and there exists a horizontal line that intersects at least $j + 1$ regions of S , one region per nonempty kernel. We have

$$u_i = \frac{n}{(3k-4)^i} - \left(\frac{k-1}{3k-4} + \frac{k-1}{(3k-4)^2} + \dots + \frac{k-1}{(3k-4)^i} \right)$$

for $i \geq 1$. Thus,

$$u_i > \frac{n}{(3k-4)^i} - \frac{k-1}{3k-5},$$

which implies that $u_{j+1} > 0$, unless $j+1 \geq \log_{3k-4}[(3k-5)n/(k-1)]$. □

3.3. Bounding the Number of Slopes

In this subsection we investigate the line span of a convex subdivision S under the condition that all edges of the subdivision have at most a constant number of slopes. We call such subdivisions *isothetic*. More specifically, we assume that there are k given directions d_1, d_2, \dots, d_k , such that each edge of S is parallel to one of these k directions. Without loss of generality we take the k directions as listed above to form a sorted circular sequence. As usual, we denote by n the size of S and we use $l^{(k)}(n)$ to represent the minimum, over all isothetic subdivisions of size n with sides parallel to the k given directions, of the maximum, over all lines, of the number of regions cut by a line in a given subdivision. We obtain matching upper and lower bounds for $l^{(k)}(n)$, but only under the assumption that k is a constant independent of n . Thus the results of this subsection are somewhat weaker than those of the previous two. When we prove the lower bound (Lemma 3.3), we consider only lines parallel to one of the k directions. Thus, the result in Lemma 3.3 is slightly stronger than the corresponding statement in Theorem 1. For the upper bound, however, we need to consider all lines.

Lemma 3.3. *For convex isothetic subdivisions of size n we have $l^{(k)}(n) \geq (\log n)/(2k(2 + \log k))$, where k is the number of directions assumed by the edges and the cutting lines.*

Proof. Let S be such a subdivision. Without loss of generality we assume that all bounded regions of S are contained within a large convex polygon P of $2k$ sides, two of which are parallel to each of the given directions d_1, d_2, \dots, d_k . Figure 3.3 gives an example for the case $k = 3$. We show that P can be broken up into $O(k^2)$ convex subpolygons, so that for each subpolygon there exists one of the given directions such that all lines in that direction cut the subpolygon in at least one fewer region than they cut P . To prove that such a construction implies the assertion, we associate each convex polygon Q with two quantities, $w(Q)$ and $a(Q)$, where $w(Q)$ denotes the number of regions of the subdivision of Q induced by S . To define $a(Q)$, we let $a_i(Q)$ be the largest number of regions of Q hit by a line with direction d_i , and let

$$a(Q) = a_1(Q) + a_2(Q) + \dots + a_k(Q).$$

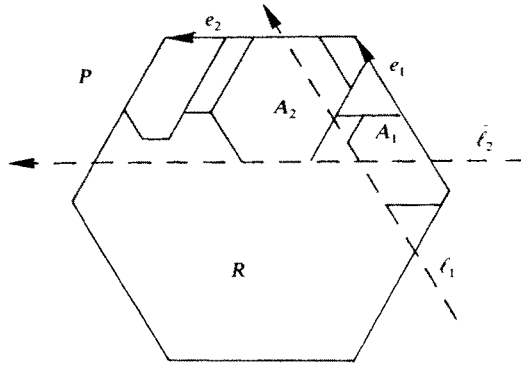


Fig. 3.3

The construction mentioned above allows us to generate a sequence of polygons with strictly decreasing values of a . The choice of polygons is guided by their w -values so that the sequence is guaranteed to be sufficiently long. This gives a lower bound on $a(P)$, where P is the first polygon of the sequence.

In order to produce such a partitioning of P , we arbitrarily choose a particular direction, say d_k . This direction occurs twice among the edges of P . Choose a particular way of getting around P from one of these edges to the other (say ccw) and denote by e_1, e_2, \dots, e_{k-1} the $k-1$ sides (of direction d_1, d_2, \dots, d_{k-1} , respectively) thus encountered. Choose an orientation of each edge e_i ($1 \leq i \leq k-1$) such that P lies to its left. Let ℓ_i be the rightmost line parallel to e_i , with the same orientation, such that all regions of P bounded by e_i lie between ℓ_i and e_i . The intersection of the left half-planes bounded by the ℓ_i 's and P defines a convex polygon R . Define $Q = \text{int}(P \setminus R)$ and let A_i denote a region of P that touches e_i and ℓ_i (Fig. 3.3).

We are finally ready to produce the desired subdivision of P . First, notice that any line in direction d_k must cut R in at least one region fewer than P , since all regions touching one of the e_i 's are fully outside R . Thus, we have $a(R) \leq a(P) - 1$. Any point x in $Q = \text{int}(P \setminus R)$ is to the right of some line ℓ_i , so that a line through it in direction d_i must cut A_i . This suggests subdividing Q (excluding the special regions A_i) into convex polygons in the following manner: first remove the regions A_i , which leaves at most $2k$ connected components. This is because we can draw a planar graph whose nodes are the A_i together with R and the unbounded region. Each connected component corresponds to a face of this graph. Next, we draw the infinite lines ℓ_i , which further subdivides Q into at most

$$2k^2 + \binom{k}{2}$$

polygons (the term $2k^2$ accounts for the fact that any one of the original $2k$ polygons may be cut into two by any one of the $k-1$ lines; also, any one of the at most $\binom{k}{2}$ intersections of two lines gives rise to one additional polygon). Some

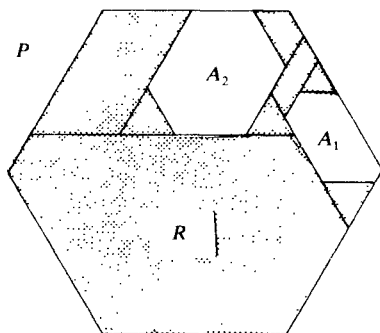


Fig. 3.4

of the remaining pieces may be nonconvex as the removal of the A_i 's leaves at most $2(k-1)^2$ reflex vertices (there are $k-1$ regions A_i with at most $2k$ vertices each and at least two of these vertices lie on the boundary of P). These nonconvex polygons can then be further subdivided into convex pieces, so that the total number of polygons is at most

$$2k^2 + \binom{k}{2} + 2(k-1)^2 = \frac{9}{2}k(k-1) + 2.$$

See region A_1 in Fig. 3.4: its removal from Q leaves one reflex vertex (one is resolved by drawing ℓ_1), which can be handled by extending an edge of A_1 . We have thus produced a convex partition of P into $O(k^2)$ polygons: the convex polygons R_1, R_2, \dots, R_{m-1} just mentioned, the polygon $R_0 = R$, and the regions A_1, A_2, \dots, A_{k-1} . By the above analysis we have $m \leq \frac{9}{2}k(k-1) + 3$. Figure 3.4 schematically depicts such a partitioning.

We already verified that $a(R_0) \leq a(P) - 1$. We have $a(R_j) \leq a(P) - 1$, for $1 \leq j \leq m-1$, since each polygon R_j lies between some line ℓ_j and edge e_j of P , which implies that each line parallel to direction d_j that intersects R_j , also meets the region A_j . Among the polygons R_j ($0 \leq j \leq m-1$), we call the one that maximizes $w(R_j)$, the number of regions of S that it overlaps, the *kernel* of P . We now recursively subdivide the kernel of P and thus generate a sequence of polygons

$$P = P_0, P_1, P_2, \dots$$

with the following properties: we have $w(P_{i+1}) \geq [w(P_i) - (k-1)]/m$, since P_{i+1} is the kernel of P_i , and $a(P_i) \geq a(P_{i+1}) + 1$, as shown above. From the first property, we derive

$$w(P_j) \geq \frac{n}{m^j} - \left((k-1) \sum_{1 \leq i < j} \frac{1}{m^i} \right) > \frac{n}{m^j} - \frac{k-1}{m-1},$$

which implies that $w(P_j)$ is positive, unless $j > \log_m [n(m-1)/(k-1)]$. If $w(P_j)$ is positive then $a(P_j) \geq k$, therefore

$$a(P) > k - 1 + \log_m \frac{n(m-1)}{k-1}$$

by virtue of the second property. By definition of $a(P)$ as the sum of the maxima of the numbers of regions meeting a line from any direction, we finally infer that there is a line that meets at least

$$\frac{k - 1 + \log_m [n(m - 1)/(k - 1)]}{k}$$

regions of P . Straightforward algebraic manipulations imply the assertion. \square

For the lower bound on $l^{(k)}(n)$ we considered only cutting lines parallel to the k directions, in order to get a result as strong as possible. For the upper bound, however, a stronger result will follow if we allow lines with arbitrary directions.

Lemma 3.4. *For convex isothetic subdivisions of size n we have $l^{(k)}(n) < 14 \log(n + 7) - 33$, where k is the number of distinct directions among the edges.*

Proof. We construct a specific subdivision of S of size n such that any line intersects at most $14 \log(n + 7) - 33$ regions. The edges of S will assume only two directions. This is sufficient to prove the lemma.

We start with a square and mark two tiny squares in it. These two squares should be small enough and sufficiently far apart such that a line can cut both of them only if its angular slope belongs to a small interval I_1 . Notice that we can make I_1 arbitrarily small by scaling down the two squares. Subdivide the complement of the two squares into seven rectangles, and then recurse within the two squares. When we do the latter, however, we choose the two tiny squares such that a line can cut both of them only if its angular slope lies in an interval I_2 disjoint from I_1 . In general, the i th recursive step uses an interval I_i disjoint from I_j , for all $j < i$. Figure 3.5 illustrates the construction. After t stages, the number of regions present will be

$$7 + 7 \times 2 + 7 \times 4 + \dots + 7 \times 2^{t-1} + 2^t = 2^{t+3} - 7.$$

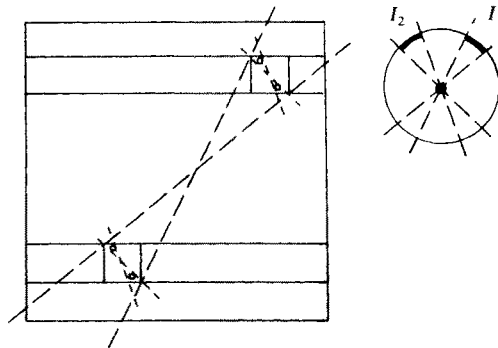


Fig. 3.5

In order to construct a subdivision with n regions, we choose the largest t such that $2^{t+3} - 7 < n$, and then further subdivide as many of the 2^t tiny squares as necessary to get n regions. To obtain an upper bound on the maximum number of regions we can cut with a single line, we assume that the subdivision is equal to the one obtained after $t+1$ stages. Any line cuts at most two tiny squares at each level of recursion (one at the first level), which implies that it meets at most

$$7 + 14 + 14 + \cdots + 14 + 16 = 14t + 9$$

regions. By the choice of t , we have

$$t < \log(n+7) - 3,$$

which implies the assertion. □

3.4. Computing the Line Span

The number of regions in which a line ℓ cuts a convex subdivision S is always one less than the number of edges of S intersected by the line. This assumes that ℓ is in general position with respect to S , i.e., that it does not pass through any of the vertices of S . A line ℓ that realizes the span of S can then be found simply by looking for a line that properly cuts as many edges of S as possible. This stabbing-line problem can be solved by the topological sweep techniques of Edelsbrunner and Guibas [EG, Section 5.2] in time $O(n^2)$ and linear space.

Lemma 3.5. *The line span of a subdivision of size n can be computed in $O(n^2)$ time and $O(n)$ space.*

3.5. Discussion

We are now in a position to prove Theorem 1. From Lemmas 3.1 and 3.2 we derive the fact that, among subdivisions of size n and degree d or less, the minimum line span is $\Omega(\log_d n + \log n / \log \log n)$. Setting d to the value $n-1$ is equivalent to relaxing the degree constraint. Since the line span of a subdivision cannot exceed its monotone span, Theorem 2 provides a matching upper bound. (The reader has our word that the argument is *not* circular!) Lemmas 3.3 and 3.4 cover the case of discrete slope domains. Finally, Lemma 3.5 provides the computational results of the theorem.

Theorem 1 (Line Span). *The line span of a subdivision of size n can be computed in $O(n^2)$ time and $O(n)$ space. The minimum line span over all subdivisions of size n is in $\Theta(\log n / \log \log n)$. More generally, if only subdivisions of degree d or less are considered then the bound is in $\Theta(\log_d n + \log n / \log \log n)$. On the other hand, if the edges of the subdivisions assume at most a constant number of distinct slopes then the minimum line span is $\Theta(\log n)$.*

4. The Monotone Span of a Convex Subdivision

Recall that the definition of the monotone span of a subdivision S is based on the notion of a contact chain of S in a given direction. Let ℓ be a directed line that defines a direction. We say that a region r *pushes* an adjacent region r' if there exists a directed line parallel to ℓ that intersects r and r' in this order. A *contact chain* of S with respect to ℓ is a longest sequence of regions such that each region pushes its immediate successor. The *monotone span* of S is the maximum length of a contact chain taken over all directions.

Since the monotone span of a convex subdivision is a proper generalization of its line span, we would expect its asymptotic value to be greater. Surprisingly, as we shall see below, the two quantities have identical growth rates. In Section 4.1 we establish an upper bound on the minimum monotone span. In Section 4.2 we specialize our discussion to treat the number of distinct slopes as a parameter, and we derive a corresponding lower bound. Section 4.3 addresses the computational aspect of the monotone-span problem, and, finally, Section 4.4 wraps up with a proof of Theorem 2.

4.1. Upper Bounds on the Minimum Monotone Span

Our objective is to construct a subdivision of size n and monotone span $O(\log n / \log \log n)$. To achieve this goal we need a slightly more general result which allows us to control the monotone span by a parameter λ . Let λ be an arbitrary integer greater than one. We define an infinite sequence of subdivisions: for any $k \geq 0$ the size of the k th subdivision in the sequence (which we call the subdivision of order k) will be roughly λ^k and its monotone span will be $O(k + \lambda)$. Setting λ to the value $\lfloor \log n / \log \log n \rfloor$ will naturally give us the desired result. To facilitate the recursive definition of these subdivisions, we introduce the concept of *orientation*. Given two angles α and β ($-\pi/2 \leq \alpha \leq \beta \leq \pi/2$) we say that a subdivision of order k has orientation $[\alpha, \beta]$ if not only its monotone span is $O(k + \lambda)$ but, when measured with respect to any direction in $[-\pi/2, \pi/2] \setminus [\alpha, \beta)$, the span is actually $O(k)$. It is not so easy to make much sense of this subtlety right now. Roughly, the idea is that a given subdivision has favorable directions (those falling outside of its *orientation*) and others less favorable. Unfortunately, to avoid this distinction seems difficult. Given this fact, however, it is important to take it into consideration. Indeed, by ignoring it we would run the risk of overlapping orientations at several levels of the recursion, and making the span unacceptably large.

The construction is based on the idea of a *unit* and a *frame*. Both units and frames have orientations. A unit of order 0 is simply the subdivision of the plane consisting of one region, the plane itself. By convention, its orientation is any interval in $[-\pi/2, \pi/2)$. For any $k > 0$, a unit of order k and orientation $[\alpha, \beta]$ is built by assembling λ units of order $k - 1$ within a frame. For convenience, let us call these units *subunits*. The key features of the construction is that the orientations of all subunits as well as the orientation of the frame are pairwise

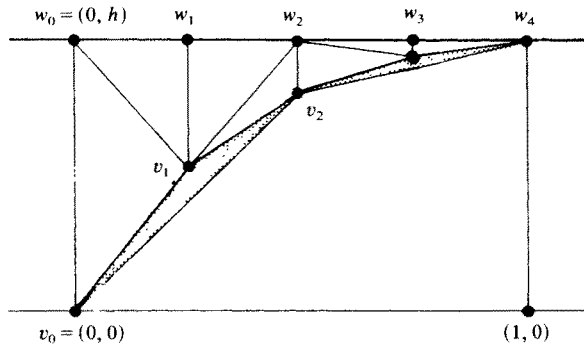


Fig. 4.1

disjoint and fall strictly within the interval $[\alpha, \beta]$. It appears from this discussion that the essential ingredient in the construction of a unit is the frame itself, so this is where we turn our attention next.

Let h be a real number ($0 < h < 1$). Given a Cartesian system of coordinates (Ox, Oy) , let R be the rectangle with the origin as its southwest corner and $(1, h)$ as its northeast corner. Since $h < 1$ there exists a unique circle passing through $(0, 0)$ and $(1, h)$ with its center on the line $x = 1$ (Fig. 4.1). On the circular arc lying inside R mark $2\lambda + 1$ points at regular intervals (along the x -axis), starting at $v_0 = (0, 0)$ and ending at $v_{2\lambda} = (1, h)$. For each i between 0 and $2\lambda - 1$ the difference in x -coordinates between v_{i+1} and v_i is $1/(2\lambda)$. For $i = 0, 1, \dots, 2\lambda$, let w_i be the vertical projection of the point v_i on the line $y = h$. To complete the construction of the frame we add the following segments:

- (i) $v_i v_{i+1}$ ($0 \leq i < 2\lambda$).
- (ii) $v_{2i} v_{2i+2}$ ($0 \leq i < \lambda$).
- (iii) $v_i w_i$ ($0 < i < 2\lambda$).
- (iv) $v_i w_{i-1}$ and $v_i w_{i+1}$ ($i = 1, 3, 5, \dots, 2\lambda - 1$).
- (v) The two lines $y = 0$ and $y = h$.

Since λ has a fixed value, a frame is completely characterized by its height h . For $i = 1, \dots, \lambda$, the triangle $v_{2i-2} v_{2i-1} v_{2i}$ is called the i th base of the frame. It is within each base that new frames will be added to form a unit. But before examining how frames are used to construct units, let us list a few simple properties relating to their monotone span.

Recall that the monotone span of a subdivision is obtained by maximizing the length of contact chains for all directions between $-\pi/2$ and $\pi/2$. Given a direction θ it is thus meaningful to define the θ -span of a frame as the maximum chain length in that direction. Last bit of terminology: the regions traversed by a contact chain are said to be hit. A question of interest is to find which regions of a frame can or cannot be hit in a certain direction. Of course, any single region can be hit in any direction. The interesting point is to know if, say, two or three given regions can be hit by a single chain.

Lemma 4.1. *There exist a function $\mu(h)$, defined for each h ($0 < h < 1$), and two constants (independent of λ) $c_1 \geq 1$ and $c_2 \geq 1$ such that:*

1. *The function μ is positive, continuous, and goes to 0 when h goes to 0.*
2. *For any θ in $[-\mu(h), \mu(h))$ no more than two bases (of a frame of height h) can be hit in direction θ . Moreover, the θ -span is always less than $c_1\lambda$.*
3. *For any θ outside of $[-\mu(h), \mu(h))$ no more than one base can be hit in direction θ . Moreover, the θ -span is always less than c_2 .*

Proof. Let δ be the largest absolute value amongst all the nonvertical directions defined by the edges of a frame. It is a simple exercise to verify that the function $\mu(h) = \delta$ not only goes to zero with h , but also satisfies all the claims of the lemma. \square

The interval $[-\mu(h), \mu(h)]$ is called the *orientation* of the frame. (Note that the function μ is not uniquely defined.) As it turns out, the essential properties of a frame will be found to be invariant under the group of motions and homothetic transformations (among the former we actually only concern ourselves with positive isometries). This means that we are free to move a frame around the plane anywhere we want as well as scale it to any desired magnitude. As a result, given any two angles α and β ($-\pi/2 \leq \alpha < \beta < \pi/2$), we can always define an arbitrarily small frame whose orientation falls within $[\alpha, \beta]$. We are now ready to construct a unit of order k and orientation $[\alpha, \beta]$. Because rotations “come for free,” we can assume without loss of generality that the orientation is of the form $[-\alpha, \alpha]$ for $0 < \alpha < \pi/2$.

A unit of order 0 is the subdivision consisting of the whole plane as its only region: its orientation is any closed interval in $[-\pi/2, \pi/2)$. A unit of order $k > 0$ and orientation $[-\alpha, \alpha]$ ($0 < \alpha < \pi/2$) is defined as follows. Let J be a frame of height h and orientation $[-\mu(h), \mu(h)]$. From Lemma 4.1 we can choose h small enough so that $\mu(h) < \alpha/2$. Inside each base $v_{2i}v_{2i+1}v_{2i+2}$ ($0 \leq i < \lambda$) we place a unit of order $k-1$ and orientation $[\alpha_i, \beta_i]$. We choose the values $\alpha_i = \alpha/2 + i\alpha/(4\lambda)$ and $\beta_i = \alpha/2 + (2i+1)\alpha/(8\lambda)$. By the induction hypothesis such a unit always exists (after appropriate scaling and isometric transformations). By “inside the base,” we mean that all the vertices of the unit should lie within the base. Of course, we clip the unit so that it fits entirely within the base $v_{2i}v_{2i+1}v_{2i+2}$. These λ units are called *subunits*, and the frame is referred to as the *master frame*. We easily verify that $[-\mu(h), \mu(h)]$ as well as the intervals of the form $[\alpha_i, \beta_i]$ are pairwise disjoint and lie strictly within $[-\alpha, \alpha]$. We are now in a position to prove the key property of a unit of order k .

Lemma 4.2. *Let c_1 and c_2 be the constants of Lemma 4.1, and let $d_2 = c_2 + 1$ and $d_1 = c_1 + 2d_2$. For any integer $k \geq 0$ and any reals α and β ($-\pi/2 \leq \alpha < \beta < \pi/2$), there exists a unit of order k and orientation $[\alpha, \beta]$. Moreover, for any θ in $[\alpha, \beta]$ the θ -span of the unit is less than $d_1(k + \lambda)$, and for any θ outside $[\alpha, \beta]$ the θ -span of the unit is less than $d_2k + 1$.*

Proof. Once again we assume without loss of generality that instead of $[\alpha, \beta]$ the orientation of the unit in question is $[-\alpha, \alpha]$. We prove the lemma by induction. The case $k = 0$ being easily resolved, we assume that $k > 0$. By construction, the orientation of the master frame used in the definition of the unit falls entirely within $[-\alpha, \alpha]$. From Lemma 4.1 it then follows that any contact chain whose direction falls outside $[-\alpha, \alpha]$ can hit at most one base and at most c_2 regions of the frame outside the bases. Now, we observe that the orientation of each subunit also falls within $[-\alpha, \alpha]$, therefore by induction hypothesis at most $d_2(k - 1) + 1$ regions of a subunit can be hit by a single contact chain. This gives us an upper bound of $c_2 + d_2(k - 1) + 1 < d_2k + 1$ on the θ -span of the unit, for $\theta \notin [-\alpha, \alpha]$. Let us now consider the case $\theta \in [-\alpha, \alpha]$. The argument is very similar, only slightly more complicated. If θ lies in the orientation of the master frame, then any contact chain in that direction will hit fewer than $c_1\lambda$ regions, not counting the one or two bases which may also be hit (Lemma 4.1). Fortunately, θ will then lie outside the orientation of any subunit, therefore at most $2d_2(k - 1) + 2$ subunit regions can be hit simultaneously. This gives an upper bound of $c_1\lambda + 2d_2(k - 1) + 2 < d_1(k + \lambda)$ on the θ -span. Next, if θ falls within the orientation of a subunit then it avoids the orientation of the master frame as well as of the other subunits. This leads to an upper bound of $c_2 + d_1(k - 1 + \lambda) < d_1(k + \lambda)$ on the θ -span. Finally, if θ lies in $[-\alpha, \alpha]$ but somehow fails to be in the orientation of the master frame or any of the subunits, we obtain an upper bound of $c_2 + d_2(k - 1) + 1 < d_1(k + \lambda)$. This completes our proof. \square

Lemma 4.3. *A unit of order $k > 0$ consists of precisely $5\lambda^k + 8[(\lambda^k - 1)/(\lambda - 1)] - 4$ regions.*

Proof. Let $N(k)$ be the number of regions in a unit of order k . (Clearly, this number is independent of the orientation.) We derive the recurrence relation $N(0) = 1$ and $N(k) = \lambda N(k - 1) + 4(\lambda + 1)$ for $k > 0$. \square

Given a fixed value of λ , it is clear from Lemma 4.3 that it is not always possible to construct a unit with a preassigned number of regions. However, we can approximate this goal quite well. Given any n large enough, we build a unit of highest order with at most n regions. Then we complete the desired subdivision by taking each base in turn, and adding a frame into it as long as the number of regions does not exceed n . When this process terminates we will be short of at most $O(\lambda)$ regions, so we can basically complete the addition of regions any way we like. From Lemma 4.2 and Fig. 4.1, it follows that the resulting subdivision will have degree $O(\lambda)$ and that its monotone span will be in $O(\lambda + \log_\lambda n)$. Observing that the degree can always be reduced by a constant factor at the cost of a multiplicative factor in the monotone span, we can state our results for all degrees $d \geq 3$.

Lemma 4.4. *For any positive integer n and any $d \geq 3$ there exists a subdivision of size n and degree d whose monotone span is $O(d + \log_d n)$.*

We might observe certain singularities in the definition of a unit which may raise the suspicion that these are actually needed in order to achieve the upper bound of Lemma 4.4. For example, many vertices are adjacent to more than three other vertices. More ominously perhaps, many edges are collinear, so minute perturbations might actually break the convexity of the subdivision. It is not difficult to see, however, that without violating the structural properties of a unit, each of its angles can be made strictly less than 180° and the degree of its vertices can be reduced to three, if so desired.

4.2. Bounding the Number of Slopes: A Lower Bound

As we did previously in the case of line spans, we investigate the monotone span of isothetic subdivisions. We are able to prove slightly stronger results if we assume that the number of distinct slopes is bounded above by a parameter k , not necessarily taken to be constant.

Lemma 4.5. *Over all subdivisions of size n using k or fewer distinct slopes the minimum monotone span is $\Omega(n^{1/k} + \log n / \log \log n)$.*

Proof. Let S be a subdivision of size n whose edges have slopes in the set $\{s_1, \dots, s_k\}$. Assuming that these slopes are in increasing order, let ℓ_i be an arbitrary direction with slope strictly between s_i and s_{i+1} ($0 < i < k$). We also add the vertical direction ℓ_k . We say that two regions r and r' are comparable in direction ℓ if there exists a sequence of regions $r = r_0, r_1, \dots, r_l = r'$ such that for all i between 0 and $l-1$, r_i pushes r_{i+1} in direction ℓ . If a line parallel to ℓ meets r and r' then r and r' are comparable in direction ℓ . It follows that r and r' are comparable in at least one of the directions ℓ_i . (Indeed, there is a whole interval of slopes $[s_a, s_b]$, or $(-\infty, s_b]$, or $[s_a, +\infty)$ within which r and r' are comparable.) Completing the proof is then easy. If for some parameter $m > 0$, fewer than m regions are mutually comparable in direction ℓ_1 then by Dilworth's theorem [H] at least n/m regions are pairwise incomparable in direction ℓ_1 . If of these n/m regions no m of them are pairwise comparable in direction ℓ_2 , then at least n/m^2 regions are pairwise incomparable in ℓ_1 and in ℓ_2 . Iterating on this process shows that $\max(m, n/m^{k-1}) \geq n^{1/k}$ regions are mutually comparable in some direction. Obviously, the lower bound on the minimum line span can always be used: in particular, when no restriction is placed on the slopes. \square

4.3. Computing the Monotone Span

Given a convex subdivision \mathcal{C} of n regions, it is relatively easy to compute its monotone span in $O(n^2)$ time. To do better appears to be an interesting open problem. We sketch the basic steps of a quadratic algorithm. The basic idea is to generate all possible contact graphs G_θ and to find longest paths in these graphs. Recall that the regions of the subdivision are the nodes of G_θ and that

we put a directed edge between regions r and r' if r pushes r' along direction θ (see Fig. 2.1). For simplicity, we may assume that the subdivision is represented with a quad-edge data structure [GS] (although, given the asymptotic complexity that we are aiming for, almost any type of representation would do just as well). In $O(n \log n)$ time we sort the slopes of all the edges of the subdivision. This partitions the set of directions $[-\pi/2, \pi/2)$ into a collection of at most $n + 1$ intervals, within each of which the contact graph of the subdivision remains invariant. Setting up the graph itself is a linear operation, while computing the longest path is a straightforward application of topological sorting, another linear procedure. This summarizes the description of the algorithm.

Lemma 4.6. *The monotone span of an n -region convex subdivision can be computed in $O(n^2)$ time and $O(n)$ storage.*

Notice that this algorithm takes only $O(kn)$ time if the edges assume only k different slopes. This is because there are only k different contact graphs (if we identify those for opposite directions). In the general case, it is trivial to construct all contact graphs in linear time, after sorting the directions of the edges. The challenging task is to maintain a longest path efficiently when the contact graph changes.

4.4. Discussion

It is easy to put the pieces together to complete the proof of Theorem 2. Let $\lambda = \lceil \log n / \log \log n \rceil$; setting $d = \lambda$ in Lemma 4.4 shows that the minimum monotone span over all subdivisions of size n and degree greater than λ is $O(\log n / \log \log n)$. For subdivisions of degree $d \leq \lambda$ we observe that, for n large enough, we have $\log_d n > d$ therefore the minimum monotone span is $O(\log_d n)$. Matching lower bounds are obtained directly from Theorem 1. The result of Theorem 2 concerning subdivisions with a bounded number of slopes is a restatement of Lemma 4.5. Similarly, the complexity of the relevant algorithms is given in Lemma 4.6.

To tidy things up it is good to mention that the proof of Theorem 2 establishes the validity of the upper bounds of Theorem 1 which had been left unproven.

Theorem 2 (Monotone Span). *The monotone span of a subdivision of size n can be computed in $O(n^2)$ time and $O(n)$ space. The minimum monotone span over all subdivisions of size n is in $\Theta(\log n / \log \log n)$. More generally, if only subdivisions of degree d or less are considered then the bound is $\Theta(\log_d n + \log n / \log \log n)$. If the number of distinct slopes in the subdivision is at most k then the minimum monotone span is in $\Omega(\max(n^{1/k}, \log n / \log \log n))$.*

We mention in passing that the technique of Lemma 4.4 can be used to derive an upper bound of $O(kn^{1/(c+\log k)})$, for some $c > 0$, on the minimum monotone

span of subdivisions of size n with at most k distinct slopes. This result is admittedly weak and, furthermore, follows the construction given above very closely, so we do not include a proof of it.

5. The Cross-Section and the Silhouette of a Convex Polytope

Our discussion of cross-sections and silhouettes is organized as follows. In Sections 5.1 and 5.2 we provide lower and upper bounds on the minimum cross-section span of a convex polytope. The proof of the upper bound is somewhat complicated, so the section in question is subdivided into Sections 5.2.1–5.2.4. In Section 5.3 we study the equivalence between cross-sections and silhouettes. Section 5.4 is concerned with the computational aspect of the problems. Finally, Section 5.5 assembles all the results above to prove Theorems 3 and 4.

5.1. Lower Bounds on the Minimum Cross-Section Span

Let $c_d(n)$ be the minimum cross-section span of any convex polytope of n facets of degree d (each facet has at most d edges). If we consider only polytopes containing the origin in their interior and only cutting planes passing through the origin then we obtain a restricted function, denoted $c_d^*(n)$. Clearly, we have $c_d^*(n) \leq c_d(n)$. Recall that $l_d(n)$ is the equivalent quantity with respect to the line span of a subdivision.

Lemma 5.1. *For any $n > 3$ and $d \geq 3$ we have the inequalities, $l_d(\lceil n/2 \rceil) \leq c_d^*(n) \leq c_d(n)$.*

Proof. Let P be a convex polytope of n facets, of degree d , containing O in its interior, and let P_1 (resp. P_2) be the polytope formed by the intersection of P with $(z \geq 0)$ (resp. $(z \leq 0)$). The central projection of P_1 (resp. P_2) about O on the plane $(z = 1)$ forms a convex subdivision of the plane, which we denote S_1 (resp. S_2). Clearly, both S_1 and S_2 are of degree d . Without loss of generality assume that S_1 contains at least as many regions as S_2 ; then it has at least $\lceil n/2 \rceil$ of them. Let $c_1(P, \pi)$ be the number of facets of P_1 (discounting $P \cap (z = 0)$) intersected by a plane π passing through the origin. Similarly, let $c(P, \pi)$ be the number of facets of P intersected by π . Let now π^* (resp. P) be the plane (resp. polytope) that achieves $c_d^*(n)$ and let π be the plane through O that maximizes $c_1(P, \pi)$. We have

$$c_d^*(n) = c(P, \pi^*) \geq c(P, \pi) \geq c_1(P, \pi).$$

But $c_1(P, \pi)$ is precisely the line span of S_1 . Since l_d is obviously nondecreasing, the proof is complete. \square

It is worthwhile to mention that there exist convex subdivisions of any degree which cannot be obtained as projections of convex polytopes [CH].

5.2. Upper Bounds on the Minimum Cross-Section Span

This section consists of four parts. Section 5.2.1 introduces the main concepts and gives the necessary definitions. Section 5.2.2 describes the construction of a polytope with small cross-section span. This polytope is used as a building block to create more complex polytopes recursively. The conditions necessary to carry out the recursion are discussed in Section 5.2.3. Finally, the desired upper bound is established in Section 5.2.4.

5.2.1. *Preliminaries.* The proof of the upper bound is somewhat similar to the proof of Lemma 4.4. It is more complicated, however, because cutting planes in E^3 have three degrees of freedom. The basic idea is to construct a convex polytope by truncating some of the vertices of a tetrahedron and attaching to them small polytopes defined recursively. Let λ be a parameter (the branching factor of the recursion tree), which for the time being we simply assume to be a fixed positive integer. As usual, we assume the existence of a Cartesian system of reference ($Oxyz$). A plane parallel to Oxy is called *horizontal* and a plane or a line parallel to the axis Oz is called *vertical*. We say that a tetrahedron is a *pyramid* if three of its facets are congruent to each other. The vertex common to them is called the *apex* of the pyramid and the opposite facet, an equilateral triangle, is referred to as its *base* (apex and base are not necessarily unique). When one apex is understood, we say that a pyramid is *well-grounded* if the vertical line passing through its apex intersects the interior of the base. Let π be a plane of equation $ax + by + z = \alpha$. If π is not vertical then we can define the transform $\Delta(\pi) : (a, b)$. These are the coordinates in the plane ($z = 1$) of the unique point of that plane which lies on the normal to π through O . Therefore any plane parallel to π is mapped into the same point via Δ . A more interesting property is that if A, B , and C are the transforms of the planes supporting the three nonbase facets of a well-grounded pyramid, then the set of planes that pass through the apex without intersecting the interior of the pyramid maps exactly into the convex hull of A, B, C .

The sphere centered at p of radius r is denoted $S(p, r)$. Similarly, when a plane of reference is understood, we let $W(A, L, \theta)$ denote the *double wedge* with parameters A, L, θ , where A is a point, L is a line, and θ is an angle: this is the locus of points p such that $|\angle(Ap, L)| \leq \theta$ (Fig. 5.1). The union of two double

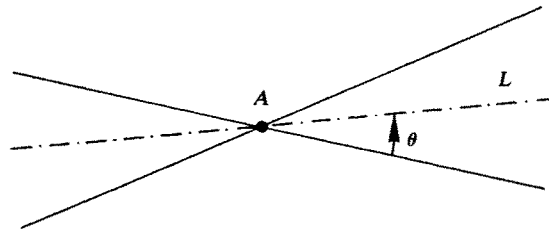


Fig. 5.1

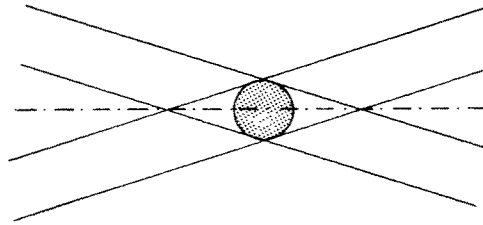


Fig. 5.2

wedges of the form $W(A, L, \theta)$ and $W(B, L, \theta)$ is called a *bow*. There exists a unique disk tangent to all four lines bounding the bow: it is called the *articulation* of the bow and the line L is called the *axis* (Fig. 5.2).

The last piece of terminology at this point concerns an operation on the convex hull of points p_1, \dots, p_k in E^3 . The hull can always be written as $\bigcap_{1 \leq i \leq k} h_i$, where each h_i is a closed half-space whose bounding plane contains a facet of the hull. Let f be a facet and assume, without loss of generality, that h_1 is the corresponding half-space. If h denotes the closure of the complement of h_1 , then we let $CH(p_1, \dots, p_k | f)$ designate the convex polyhedron $h \cap (\bigcap_{2 \leq i \leq k} h_i)$. An important observation is that $CH(p_1, \dots, p_k | f)$ is *not* the empty set.

5.2.2. A Building Block. We describe the nonrecursive part of the construction. Let $aa_1a_2a_3$ be an arbitrary pyramid with apex a . For the sake of exposition, we assume that the base is horizontal. Let a_4 be the centroid of the triangle $a_1a_2a_3$ and let b_i be the midpoint of aa_i ($1 \leq i \leq 4$) (Fig. 5.3). It should be clear that $b_1b_2b_3$ is parallel and similar to $a_1a_2a_3$. Now let c be a point on the segment ab_4 such that for some ϵ_1 ($0 < \epsilon_1 < 1$) we have

$$|cb_4| = \epsilon_1 |ab_4|. \tag{1}$$

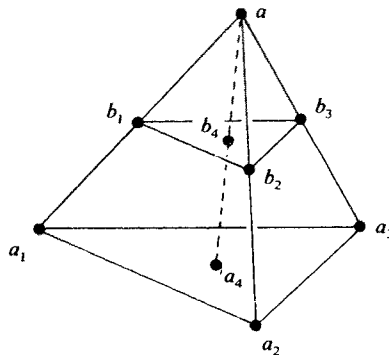


Fig. 5.3

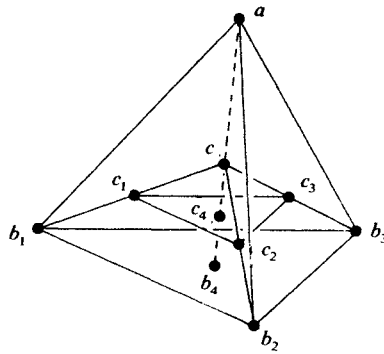


Fig. 5.4

Let c_i be the midpoint of cb_i ($1 \leq i \leq 4$). Once again the triangle $c_1c_2c_3$ is parallel and similar to $b_1b_2b_3$ and c_4 is its centroid (Fig. 5.4). Let C be the horizontal circle centered at c_4 , whose radius is half that of the circle inscribed in the triangle $c_1c_2c_3$ (Fig. 5.5). Let $c_5c_6c_7c_8$ be four points on C forming a rectangle. This rectangle is characterized by a direction and a width. Let θ be the angle between c_5c_6 and a fixed direction X chosen arbitrarily on the plane supporting C :

$$\theta = \angle(c_5c_6, X). \tag{2}$$

The angle θ is measured counterclockwise from c_5c_6 to X ($0 \leq \theta < 2\pi$). For any real ε_2 such that $0 < \varepsilon_2 < 2|c_4c_5|$, let

$$|c_5c_6| = \varepsilon_2. \tag{3}$$

We erect a parallelepiped above $c_5c_6c_7c_8$ of height ε_2 . More specifically, the rectangle $d_5d_6d_7d_8$ is parallel to $c_5c_6c_7c_8$, and its vertical projection coincides with that of $c_5c_6c_7c_8$. Also, we have $|c_5c_6| = |c_5d_5| = \varepsilon_2$. Let d_4 be the centroid of $d_5d_6d_7d_8$ (Fig. 5.6). It is crucial that ε_2 should be chosen small enough so that

$$d_5, d_6, d_7, d_8 \in \text{CH}(b_1, b_2, b_3, c_1, c_2, c_3 | c_1c_2c_3). \tag{4}$$

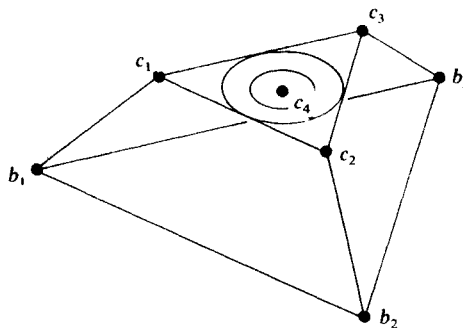


Fig. 5.5

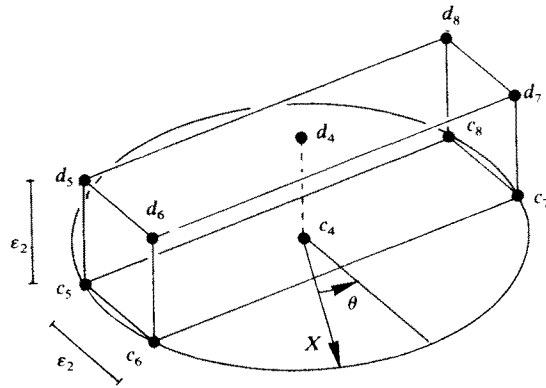


Fig. 5.6

Since $b_1b_2b_3$, $c_1c_2c_3$, and $d_5d_6d_7d_8$ are all parallel, this implies in particular that the ten points thus enumerated are vertices of their own convex hull.

Consider the rectangle $d_5d_6d_7d_8$. Let e_1, \dots, e_λ (resp. f_1, \dots, f_λ) be the points that divide d_5d_8 (resp. d_6d_7) in $\lambda + 1$ equal-size segments and let C_1 (resp. C_2) be the circle passing through d_5, d_8 (resp. d_6, d_7), and d_4 . We define e'_i (resp. f'_i) as the intersection of $e_i f_i$ with C_1 (resp. C_2). Let F_1 (resp. F_2) be the convex hull of $d_8, d_5, e'_1, \dots, e'_\lambda$ (resp. $d_7, d_6, f'_1, \dots, f'_\lambda$) (Fig. 5.7). Look at F_1 and F_2 as two flaps which we can open up. We do so by half the maximum amount we can, where “we can” means ensuring that both flaps remain strictly within the nonempty polytope $\text{CH}(c_1, c_2, c_3, d_5, d_6, d_7, d_8 | d_5d_6d_7d_8)$. (5)

We can easily verify that, indeed, the flaps can always be opened up by a nonzero angle. This is illustrated in Fig. 5.8, which distorts the proportions for clarity. Each edge $e'_i f'_i$ is parallel to $d_5d_6d_7d_8$ and is an edge of the convex hull of $d_5, d_6, d_7, d_8, e'_1, \dots, e'_\lambda, f'_1, \dots, f'_\lambda$. We now truncate little pieces off that convex hull. For each i ($1 \leq i \leq \lambda$) we define some additional points. Let g_1, g_2 , and h_3 be respectively the midpoints of $e'_{i-1}e'_i$, $e'_{i+1}e'_i$, and $f'_i e'_i$. For convenience we

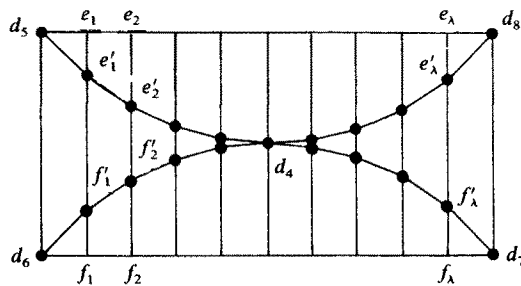


Fig. 5.7

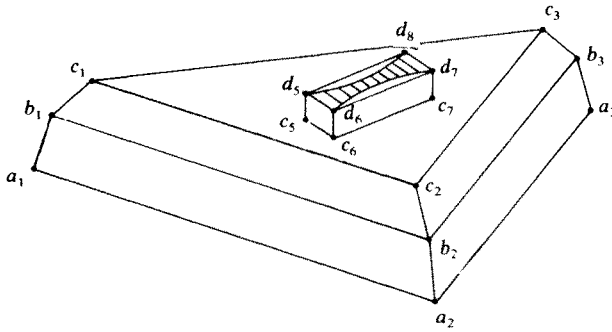


Fig. 5.8

pose $e'_0 = d_5$ and $e'_{\lambda+1} = d_8$. Similarly, for $j = 1, 2$, let h_j be the midpoint of $g_j e'_1$. Obviously the triangle $h_1 h_2 h_3$ is parallel and similar to $g_1 g_2 f'_1$ (Fig. 5.9). Let T and T' be two triangles coplanar with and similar to $h_1 h_2 h_3$, whose centroids coincide with that of $h_1 h_2 h_3$, and whose areas are respectively two-thirds and one-third the area of $h_1 h_2 h_3$. Recall that there is a distinct triangle T' for each i between 1 and λ . Since none of their supporting planes are identical it is always possible to find a point m_i in each T' such that the points $\{m_1, \dots, m_\lambda\}$ are collectively in general position. Clearly, a random pick will give such a set with probability 1. Therefore there exists some $r > 0$ such that the spheres $\{S(m_1, r), \dots, S(m_\lambda, r)\}$ are also in general position. (6)

What we mean here is that any four points on distinct spheres cannot be coplanar. We choose r small enough so that the circle C_i formed by the intersection of $S(m_i, r)$ and the plane supporting T lies entirely inside T (Fig. 5.10). (7)

Now, let $\alpha_1 \alpha_2 \alpha_3$ be an equilateral triangle with vertices on C_i and no edge parallel to the edges of $h_1 h_2 h_3$. The latter condition is added simply for convenience: in this manner, indeed, the convex hull of $\alpha_1, \alpha_2, \alpha_3, g_1, g_2$, and f'_1 always has precisely eight facets. Let α be the apex of a pyramid $\alpha \alpha_1 \alpha_2 \alpha_3$, such

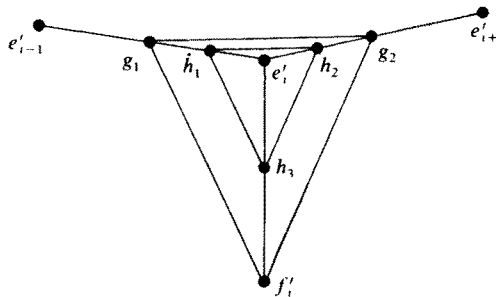


Fig. 5.9

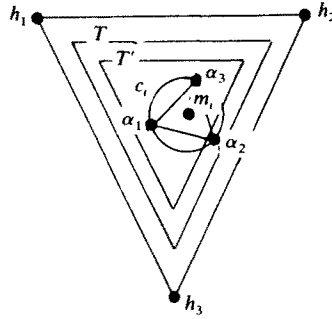


Fig. 5.10

that α lies inside

$$CH(g_1, g_2, f'_i, \alpha_1, \alpha_2, \alpha_3 | \alpha_1 \alpha_2 \alpha_3) \cap S(m_i, r). \tag{8}$$

We put $\varepsilon_3 = |m_i \alpha|$; note that $g_1, g_2, f'_i, \alpha_1, \alpha_2, \alpha_3$, and α are the vertices of their convex hull (Fig. 5.11). Let $J_i = \{\alpha, \alpha_1, \alpha_2, \alpha_3\}$ and $B_i = J_i \cup \{g_1, g_2, f'_i\}$. A *block* is defined as the convex hull of

$$\{a_1, a_2, a_3, b_1, b_2, b_3, c_1, c_2, c_3, d_5, d_6, d_7, d_8\} \cup \left(\bigcup_{1 \leq i \leq \lambda} B_i \right).$$

By construction, the vertices of the block are precisely the points in the above set and none other. The pyramid formed by the convex hull of J_i is called a *joint*. A block has λ joints; it is *based* on a pyramid $(aa_1a_2a_3)$ and parametrized by the triplet $(\varepsilon_1, \varepsilon_2, \varepsilon_3, \theta)$. Figure 5.12 gives an overview of the upper part of a block. The specification of a block is determined solely by its base pyramid and the four parameters $\varepsilon_1, \varepsilon_2, \varepsilon_3$, and θ . It is uniquely defined up to scaling, assuming a deterministic choice of the points m_1, \dots, m_λ . Indeed, the steps which were described in a probabilistic or existential manner can always be made deterministic (conditions (6) and (7)). At any rate conditions (1)–(8) summarize the main specifications of a block.

Lemma 5.2. *For any (well-grounded) pyramid $aa_1a_2a_3$, there exists a function $t: [0, 1] \mapsto [0, 1]$ with the following property: for any ε_1 ($0 < \varepsilon_1 < 1$), any ε_2 ($0 < \varepsilon_2 < t(\varepsilon_1)$), and any θ ($0 \leq \theta < 2\pi$), there exists ν such that, for any ε_3 ($0 < \varepsilon_3 < \nu$), there is a block based on $aa_1a_2a_3$ with parameters $(\varepsilon_1, \varepsilon_2, \varepsilon_3, \theta)$.*

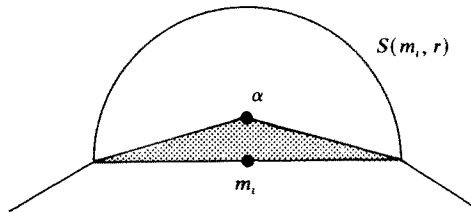


Fig. 5.11

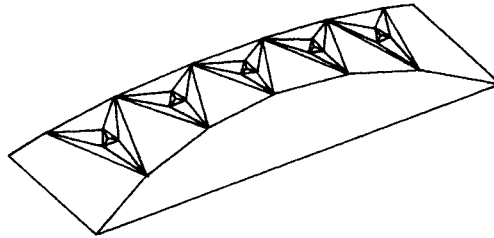


Fig. 5.12

5.2.3. *Tuning a Cap.* Let P be a well-grounded pyramid. We define a *cap of order k* , based on P , as any polytope C such that if

1. $k = 0$, then $C = P$; else if
2. $k > 0$, then C is a block based on P , where each of its joints has been replaced by a cap of order $k - 1$ based on the joint.

The λ caps of order $k - 1$ are called the *child caps* of C . Note that, although we no longer assume that the base of the cap should be horizontal, well-groundedness is always satisfied, since by construction, all facets are visible from point $(0, 0 + \infty)$, except for the base of the initial payment.

Lemma 5.3. *A cap of order k has exactly four facets if $k = 0$, and $10\lambda^k + 23[(\lambda^k - 1)/(\lambda - 1)] - 6$ facets if $k > 0$.*

Proof. By simple examination we find that a block has $10\lambda + 17$ facets, therefore the number N_k of facets in a cap of order k follows the recurrence relation: $N_0 = 4$ and $N_k = 6\lambda + 17 + \lambda N_{k-1}$ for $k \geq 1$. □

Before defining the concept of a tuned cap we must introduce two notions: the *shadow* and the *estate*. Informally, the shadow includes the *bad* planes of E^3 , that is, the planes intersecting many facets of a block. The estate is a region of safety which avoids the shadow. Fortunately, it is possible to describe these notions by mapping them in two dimensions via Δ , which considerably simplifies the discussion.

Using the previous notation, let C be an arbitrary cap of order $k > 0$, and let Ω be the transform of the plane supporting the base of P via Δ . It is clear that the closed region $\{\Delta(\pi) \mid \text{plane } \pi \text{ intersecting both triangles } a_1a_2a_3 \text{ and } b_1b_2b_3\}$ surrounds but does not contain Ω . Let D^* be the largest open disk centered at Ω that does not intersect the region: D^* is called the *dual base* of the cap. Any square parallel and congruent to $c_5c_6d_6d_5$ that lies adjacent to d_5d_8 and d_6d_7 , outside the parallelepiped formed by the c_i 's and d_i 's, is called a *screen*. Let d_9 (resp. d_{10}) be the midpoint of d_5d_6 (resp. d_7d_8) and let \mathcal{L} be the set of lines that intersect two screens distant from each other by

$$\frac{1}{2}|e_1e_2| = \frac{1}{2(\lambda + 1)}|d_5d_8|.$$

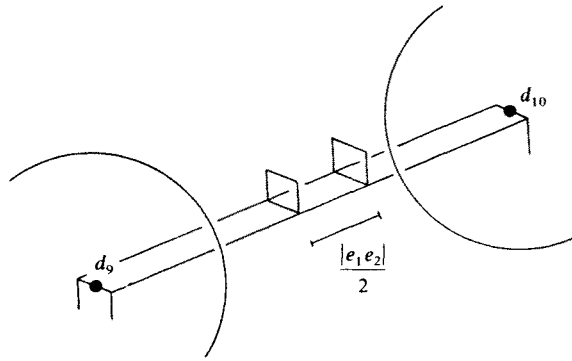


Fig. 5.13

We define S_0 (resp. S_{10}) as the smallest sphere centered at d_0 (resp. d_{10}) that intersects *every* line of \mathcal{L} (Fig. 5.13)—note that \mathcal{L} is defined over *all* pairs of screens. By choosing ε_2 small enough we can force the spheres to be as small as desired. This allows us to define the *shadow* of the cap as the closed region

$$\{\Delta(\pi) \mid \text{plane } \pi \text{ intersecting } S_0 \text{ and } S_{10}\} \setminus D^*.$$

If the cap is of order 0, then its shadow is the empty set.

We can now introduce the notion of an estate. We say that the cap C admits of an *estate* E if:

1. E is a bow whose articulation contains the dual base of C .
2. If $k > 0$ then (i) the shadow \mathcal{S} of C lies inside E , and (ii) each child cap admits of an estate which lies entirely inside $E \setminus \mathcal{S}$.

Intuitively, the estate covers the set of bad planes (or rather its map via Δ). A plane is said to be bad if its map through Δ lies either in the shadow or in the estate of the child caps. Condition 2 carefully prevents the shadow and the estates of the child caps from overlapping, thus avoiding compounding unfavorable situations. The presence of the articulation is somewhat unfortunate because it complicates the argument. It is unavoidable, however, because the map of planes almost parallel to the base cannot be made arbitrarily thin independently of the other conditions.

We say that the cap C is *tuned* to (P, E) if it is based on the well-grounded pyramid P and it admits of the estate E . Furthermore, if we consider the shadows of C , of its children, its grandchildren, etc., it must be the case that *no* point of the plane can belong to more than two shadows: this is called the *sparsity condition* of a tuned cap. Of course, we would prefer to say “one” instead of “two”: but having shadows intersect is once again unavoidable.

We show later that tuned caps cannot be cut through too many facets by any given plane. In the meantime we must concern ourselves with a more fundamental question: do there exist tuned caps of arbitrarily large order? We continue to use the notation of the previous section to describe the cap C whose block has parameters $(\varepsilon_1, \varepsilon_2, \varepsilon_3, \theta)$.

Let T^* be the triangle formed by the transforms via Δ of the planes supporting cc_1c_2 , cc_2c_3 , and cc_3c_1 , and let $D(\varepsilon_1)$ be the smallest disk centered at Ω that contains T^* . Clearly, every plane supporting a facet of the cap C "above" $c_1c_2c_3$ maps into a point inside $T^* \subseteq D(\varepsilon_1)$. An interesting property of this disk is that it converges toward Ω as ε_1 goes to 0. We omit the proof, which is trivial.

Lemma 5.4. *The radius of $D(\varepsilon_1)$ tends to 0 as ε_1 goes to 0.*

Let Λ be the transform via Δ of the set of nonvertical planes containing d_9d_{10} . Since the plane supporting $d_5d_6d_7d_8$ contains d_9d_{10} and is parallel to the base of P , Λ is a line passing through Ω . We now show that regardless of the value of ε_1 , we can always isolate the shadow of the cap inside an arbitrarily small double wedge.

Lemma 5.5. *For any ε_1 ($0 < \varepsilon_1 < 1$) and μ ($0 < \mu < \pi/2$) there exists ν ($0 < \nu < 1$) such that, for any ε_2 ($0 < \varepsilon_2 < \nu$), the shadow of the cap, if defined, lies inside the double wedge $W(\Omega, \Lambda, \mu)$.*

Proof. Let $r(\varepsilon_2)$ be the common radius of the spheres S_9 and S_{10} . We can easily show that $r(\varepsilon_2)$ goes to 0 with ε_2 . Let p be the corner, opposite of d_8 , of the screen at a distance $\frac{1}{2}|e_1e_2|$ from d_8 . Any sphere centered at d_9 passing through the intersection of the lines through d_8p and the plane supporting $c_5c_6d_6d_5$ must contain S_9 . Consequently, we have $r(\varepsilon_2) < 2\sqrt{2}(\lambda + 1)\varepsilon_2$, hence our claim. We must now examine the shape of the shadow of C in some detail (Fig. 5.14). Let (x_1, y_1, z_1) and (x_2, y_2, z_2) be the coordinates of d_9 and d_{10} , respectively. The equation of Λ is

$$\Lambda: (x_1 - x_2)x + (y_1 - y_2)y + z_1 - z_2 = 0. \tag{9}$$

The nonvertical planes that intersect S_9 and S_{10} are characterized by an equation of the form $ax + by + z = \alpha$, with

$$\frac{|ax_i + by_i + z_i - \alpha|}{\sqrt{a^2 + b^2 + 1}} \leq r(\varepsilon_2)$$

for $i = 1, 2$. This implies

$$|a(x_1 - x_2) + b(y_1 - y_2) + z_1 - z_2| \leq 2r(\varepsilon_2)\sqrt{a^2 + b^2 + 1}. \tag{10}$$

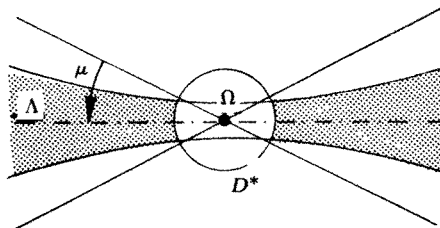


Fig. 5.14

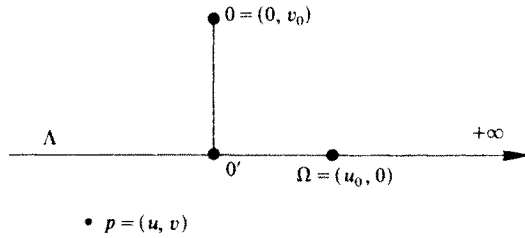


Fig. 5.15

Since the shadow decreases (in the set-inclusion sense) with ε_2 , it suffices to show that there exists some $\varepsilon_2 > 0$ such that all points (a, b) satisfying (10) and lying outside D^* also lie inside $W(\Omega, \Lambda, \mu)$. Let us change the reference system (Fig. 5.15): the new origin is the orthogonal projection of O on Λ , with Λ and $O'O$ providing respectively the axes $(O'u, O'v)$, with $O : (0, v_0), \Omega : (u_0, 0)$. Let (u, v) be the new coordinates of $p = (a, b)$. If the point p satisfies (10) then, from (9), its distance $|v|$ to Λ clearly satisfies

$$|v| = \frac{|a(x_1 - x_2) + b(y_1 - y_2) + z_1 - z_2|}{d} \leq \frac{2}{d} r(\varepsilon_2) \sqrt{u^2 + (v - v_0)^2 + 1}, \quad (11)$$

where $d = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$. Let ρ be the radius of D^* . Then, any point (u, v) satisfying $|v| < |(u - u_0) \tan \mu|$ certainly lies inside $W(\Omega, \Lambda, \mu)$. Since every point of the shadow lies outside D^* , the condition

$$|v| < \max(|u - u_0| \tan \mu, \rho \sin \mu) \quad (12)$$

will force every point (u, v) of the shadow to be in $W(\Omega, \Lambda, \mu)$. The proof will be complete once we have shown that there exists $\varepsilon_2 > 0$ such that (11) implies (12). From our starting observation we know that if $\varepsilon_2 < v_1 < 1$, then $r(\varepsilon_2)$, abbreviated here as r , is sufficiently small to have $d^2/2 - 4r^2 > 0$ and

$$64r^4 v_0^2 - 4\left(\frac{d^2}{2} - 4r^2\right)(1 - 4rv_0^2)r < 0.$$

This implies that for all values of v we have

$$\frac{d^2}{2} v^2 - r < d^2 v^2 - 4r^2 (v - v_0)^2. \quad (13)$$

From (11) and (13) we find that

$$v^2 < \frac{2r}{d^2} (1 + 4(1 + u^2)r), \quad (14)$$

hence $|v| < |u - u_0| \tan \mu$, for r sufficiently small and $|u| > u_1 > 0$. Thus (11) implies (12) for any ε_2 and u such that $\varepsilon_2 < \nu_2 < \nu_1$ and $|u| > u_1$. For any u ($|u| \leq u_1$), we still have (14) if $\varepsilon_2 < \nu_1$, therefore

$$v^2 < \frac{2r}{d^2}(1 + 4(1 + u_1^2)r) < (\rho \sin \mu)^2,$$

with the right-hand side inequality holding for any r small enough, that is, for any $\varepsilon_2 < \nu_3 \leq \nu_1$. From Lemma 5.2, it follows that choosing $\nu = \min(\varepsilon_1, \nu_2, \nu_3)$ will complete the proof. \square

Lemma 5.6. *As θ varies in $[0, 2\pi)$ from 0 to 2π , the line Λ rotates entirely around Ω .*

Proof. Recall that Λ is obtained by mapping via Δ all the nonvertical planes passing through d_9 and d_{10} . If L is the line containing d_9 and d_{10} , then Λ is the intersection of $(z = 1)$ with the plane normal to L passing through O . All the lines Λ are obtained by rotating this plane around $O\Omega$. Ensuring that the base of the pyramid is not vertical, which is true from the well-groundedness of P , guarantees the existence of Ω and establishes the lemma. \square

Lemma 5.7. *Let E be a bow with Λ for axis and a disk D centered at Ω for articulation. Let D_r designate the disk centered at Ω of radius r . Then there exist two positive reals r and μ such that a bow F can be found with articulation D_r , satisfying $F \subseteq E$ and $F \cap (W(\Omega, \Lambda, \mu) \setminus D) = \emptyset$.*

Proof. For convenience, let us place Ω at the origin and Λ collinear with the x -axis. Let γ be the (angular) slope of the edge of E extending into the northeast quadrant, and let $\gamma_1 = i\gamma/5$ and $\mu = \gamma_1$. Also, let $\gamma(r)$ be the slope of the line tangent to the boundary of D_r passing through the intersection of the boundary of D with the line through Ω of slope γ_2 (of the two candidate lines, we pick the one with the larger slope). As r varies from the radius of D to 0, $\gamma(r)$ decreases monotonically from a value greater than $\pi/2$ to γ_2 . Choose r so that $\gamma(r) = \gamma_3$. Taking another line tangent to the boundary of D_r with slope γ_4 completes the construction of the bow F . We leave it as an exercise to check that $F \subseteq E$ and that F does not intersect $W(\Omega, \Lambda, \mu)$ outside of D (Fig. 5.16). \square

The last four lemmas were in preparation for the next result, which is the centerpiece of the overall argument.

Lemma 5.8. *Let P be an arbitrary well-grounded pyramid and let k be any nonnegative integer. Then, for any bow E whose articulation contains the dual base of P , there exists a cap of order k based on P with estate E .*

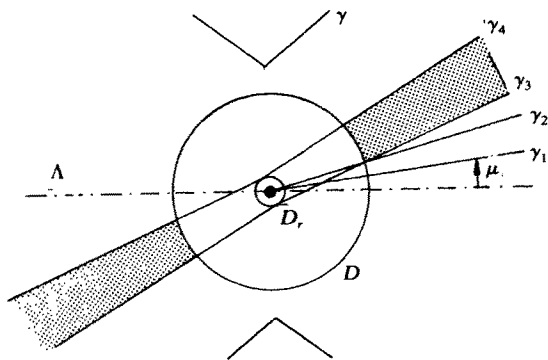


Fig. 5.16

Proof. We prove the lemma by induction on k . The case $k = 0$ is obvious, so let us assume that $k > 0$. First, we shrink E (if necessary) so that the dual base of P precisely coincides with the articulation of E (we omit the proof that this is always possible). Referring now to Lemma 5.7, we identify D with D^* and D_r with $D(\varepsilon_1)$. By virtue of Lemma 5.4, we can choose ε_1 small enough so that the radius of $D(\varepsilon_1)$ is smaller than r , yet strictly positive. Because of Lemma 5.6, we know that we can choose θ such that Λ coincides with the axis of E . Also, from Lemma 5.5, we can choose ε_2 such that the shadow of the cap C lies inside $W(\Omega, \Lambda, \mu)$, with μ specified as in Lemma 5.7. The same lemma then shows the existence of a bow F whose articulation contains $D(\varepsilon_1)$ and which avoids entirely the shadow of the cap C . Note that these specifications are realizable because of Lemma 5.2. Every joint of C has its dual base inside T^* , hence inside $D(\varepsilon_1)$ (recall that T^* is the triangle formed by the transforms via Δ of the planes supporting cc_1c_2 , cc_2c_3 , and cc_3c_1). Consequently, by induction hypothesis, each joint of C can be replaced by a cap of order $k - 1$ based on it, with estate F . This establishes the lemma. \square

We have now arrived at one of the most delicate phases of our discussion: proving that Lemma 5.8 still holds if we add the sparsity condition of a tuned cap. To this end we introduce a few notions. Points of the plane are colored *black* or *white*. Let E be a bow centered at v . We say that E is *hollow* if, for any line L completely in the interior of E and any disk \mathcal{D}^+ centered at v , there exists a white disk centered on L in $E \setminus \mathcal{D}^+$ (a region is white if all of its points are white). We now go back to the proof of Lemma 5.8 and refine some of the steps. We start with the same hypotheses, except for E which we now assume to be hollow. Once we have the bow F , we must proceed with a more elaborate construction in order to avoid having too many overlapping estates.

Let v_1, \dots, v_λ be the centers of the dual bases of the child caps of C . Since no two bases of these caps are parallel the centers are all distinct. Obviously, since E is hollow, so is F . Therefore there exists a white disk \mathcal{D} with the following characteristics: \mathcal{D} is centered outside the articulation of F , but on its axis; furthermore, for each i ($1 \leq i \leq \lambda$) and every point p in \mathcal{D} , there exists a double

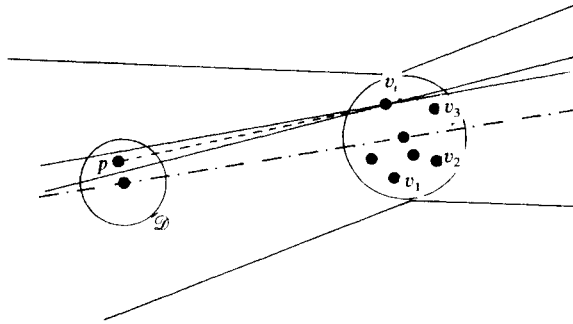


Fig. 5.17

wedge $W(v_i, pv_i, \alpha_i)$ lying entirely inside F , for $\alpha_i > 0$ (Fig. 5.17). This can be achieved by choosing a small disk whose center lies sufficiently far from the points v_i . Let \mathcal{D}' be a disk strictly inside \mathcal{D} , and let L_1, \dots, L_λ be the lines passing through v_1, \dots, v_λ and tangent to the boundary of \mathcal{D}' (pick only one of the two tangents for each v_i). We omit the proof that with a careful choice of \mathcal{D}' it is always possible to ensure that:

1. The tangent points are all distinct.
2. The $\binom{\lambda}{2}$ intersections of the form $L_i \cap L_j$ are all distinct and lie inside \mathcal{D} .

With these conditions, there exist double wedges $W(v_i, L_i, \beta_i)$ (for β_i small enough, but positive) such that the intersection of any two is white and the intersection of any three is empty. By translating slightly the boundaries of the double wedges, it is possible to prove the same result with bows centered at the v_i 's (instead of double wedges). Since making ϵ_3 arbitrarily small will make the dual bases of the child caps tend toward their centers, we can find estates for these child caps such that: (i) the intersection of any two lies in the white disk; (ii) the intersection of any three is empty. At this stage we color black each pairwise intersection of estates, and we pursue this process recursively with respect to the child caps (it is clear that each estate remains hollow). The purpose of this modification is the following.

Lemma 5.9. *For any $k \geq 0$ there exists a tuned cap of order k .*

Proof. The process described above can be modeled by a λ -ary tree, with each node corresponding to the creation and coloring of λ estates. It is important to recall that estates inherit the coloring applied at higher nodes. Let $E(v)$ denote the estate associated with node v . To begin we show that if v_1, v_2 , and v_3 are three nodes of the tree, no two of which are in an ancestor–descendant relationship, then

$$E(v_1) \cap E(v_2) \cap E(v_3) = \emptyset.$$

Assume that this is not true and that there exists a point q in all three estates. We will see that q must then be colored black twice, which is impossible. Let v_4 be the nearest common ancestor of the three nodes. By construction, v_4 cannot be the parent of three distinct ancestors of the nodes (where it is understood that a node includes itself among its ancestors). Therefore there must exist two of them, say v_1 and v_2 , whose nearest common ancestor v_5 is a proper descendant of v_4 . Let v_6 (resp. v_7) be the child of v_4 which is an ancestor of v_5 (resp. v_3), and let v_8 (resp. v_9) be the child of v_5 which is an ancestor of v_1 (resp. v_2). Since the estate of a node lies in the estate of any of its ancestors, $E(v_6) \cap E(v_7) \neq \emptyset$, therefore q is colored black at node v_4 . Similarly, $E(v_8) \cap E(v_9) \neq \emptyset$, so q is colored black a second time at node v_5 , which establishes our claim. One immediate consequence is that if the shadows associated with three nodes v_1, v_2, v_3 have a common intersection, then one node, say v_1 , must be the descendant of another, say v_2 . But this is ruled out by the definition of an estate. The sparsity condition is therefore satisfied and the proof is complete. \square

5.2.4. The Upper Bound. It now remains to show that caps, indeed, have small cross-section spans. Let C be a cap of order k . We define the *construction tree* T of C as follows. If $k=0$ then T is a tree consisting of a single node, which we conceptually associate with the three nonbase facets of a cap of order 0. If $k>0$ then T is a λ -ary tree, whose root is associated with the facets of C that do not belong to any child cap. The λ subtrees below the root of T are defined recursively with respect to each of the child caps of C . In this manner each facet of C is associated with a distinct node of T ; conversely, with each node v is associated a cap $C(v)$. Let π be an arbitrary plane of E^3 : we *mark* each node of T associated with at least one facet which π intersects. Let G be the minimal Steiner tree of the marked nodes of T .

Lemma 5.10. *Every node of G is marked.*

Proof. Let ∂C be the boundary of the cap C and let p be a point of ∂C . Since $\partial C \setminus \{p\}$ is homotopic to E^2 , we can define the notion of a simple closed curve, to which we can then apply the Jordan Curve Theorem. In particular, because of convexity, any plane intersecting C (not tangent to it) intersects ∂C in a simple closed curve. Let B denote the base of the pyramid associated with a node v distinct from the root, and let w be the parent of v . We have the obvious implication: "If v is marked but w is not, then $\pi \cap \partial C$ lies entirely on one side of $B \cap \partial C$." It follows from it that only v and its descendants can be marked. \square

We say that π *splits* a node if it intersects the base of its corresponding pyramid. If π intersects a facet at node v but does not split v , we say that π *scratches* the node.

Lemma 5.11. *At most one node can be scratched: the root of G .*

Proof. Let v be a scratched node distinct from the root of T and let w be its parent. The proof of Lemma 5.10 shows that π cannot intersect any facet at node

w , so this node is not marked, hence not in G . Since G is a tree the proof is complete. \square

Lemma 5.12. *Suppose that π splits node v . If π intersects either more than 100 facets at node v or at least two child caps, then $\Delta(\pi)$ lies in the shadow of $C(v)$.*

Proof. Assume that π intersects 100 facets at node v (there is nothing magic in this figure; we just need a large enough integer). Consider the screens passing through $e_0, \dots, e_{\lambda+1}$. Then π intersects at least two of them, so it intersects both spheres S_9 and S_{10} . If now π intersects two child caps, a similar reasoning will lead to the same conclusion. We should now use the screens passing through the 2λ points of the form h_1 or h_2 . In both cases, by convexity, it is clear that since π intersects the base $a_1a_2a_3$ it must also intersect $b_1b_2b_3$. This shows that $\Delta(\pi)$ lies outside the dual base of $C(v)$. Consequently, $\Delta(\pi)$ lies in the shadow of $C(v)$. \square

Lemma 5.13. *In G no node has more than three children, and at most three nodes have each more than one child.*

Proof. The enclosure of child caps in spheres that are in general position guarantees that no more than three child caps can be intersected by a given plane. Suppose that three nodes of G distinct from the root of G have each at least two children in G . By Lemma 5.11 these nodes are split by π . But, by Lemma 5.12, $\Delta(\pi)$ must then lie in their respective shadows. This contradicts the sparsity condition which stipulates that no three shadows can intersect at the same point. \square

Lemma 5.14. *Any plane intersects $O(k + \lambda)$ facets of a tuned cap of order k .*

Proof. By Lemma 5.3, no plane can intersect more than $10\lambda + 17$ facets at any given node of T . Because of the sparsity condition and Lemmas 5.11 and 5.12, a plane will intersect fewer than 100 facets in all but at most three nodes: one scratched node and two split ones (why?). From Lemma 5.13 we know that G has $O(k)$ nodes, so a total of $O(k + \lambda)$ facets can be intersected by a single plane. \square

Lemma 5.15. *For all $d \geq 3$, we have $c_d(n) = O(\log_d n + \log n / \log \log n)$.*

Proof. It is clear that beyond a small constant d_0 , there exist tuned caps of order k and arbitrary degree $\geq d_0$. For values below d_0 , we can very simply ensure degree three by placing one point close to each facet with too many edges and taking the convex hull of the result. If the points are sufficiently close to, say, the mass center of the vertices of the facet in question, Lemma 5.14 will still be satisfied (Fig. 5.18). What should be done now if the number of desired facets cannot be achieved precisely by any cap? During the top-down construction of the cap of smallest order with at least n facets we will stop as soon as replacing

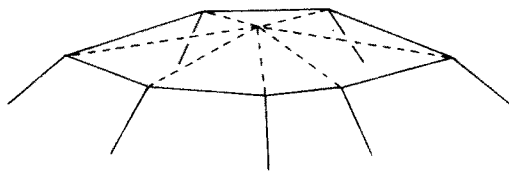


Fig. 5.18

a joint by a cap of order 1 would cause the number of facets to exceed n . At that point, only $O(\lambda)$ facets will be missing. These can be added by either introducing new vertices as before or by truncating vertices: Lemma 5.14 will still hold. Note, however, that not all values of d and n can be achieved simultaneously. For example, simplicial polytopes ($d = 3$) always have an even number of facets.

By Lemma 5.3, a cap of order k has $10\lambda^k + 23[(\lambda^k - 1)/(\lambda - 1)] - 6$ facets, therefore we have $k \leq \log_\lambda n$ for $\lambda > 1$. Since the degree of a cap is given by the size of its two flaps, which is $\Theta(\lambda)$, we have $\lambda = \Theta(d)$ and $k = O(\log_d n)$, therefore any plane intersects $O(d + \log_d n)$ facets. If $3 \leq d < \log n / \log \log n$ then, for n large enough, $d < \log_d n$, therefore $c_d(n) = O(\log_d n)$. Otherwise, we consider caps of degree $\max(3, \lfloor \log n / \log \log n \rfloor)$, from which we find $c_d(n) = O(\log n / \log \log n)$. \square

5.3. Cross-Sections and Silhouettes

In this section we demonstrate that the cross-section problem for a given polytope P is the same as the silhouette problem for its dual polytope, Q . Intuitively, the silhouette of Q from a point q outside the polytope is the contour of Q as seen from q —we give a more formal definition of a silhouette below. The correspondence between the two problems is interesting for several reasons. For one thing it shows that Theorem 3—which summarizes our results on the cross-section problem—can be reinterpreted to solve the silhouette problem. Also, it seems that from an expository point of view silhouettes are preferable to cross-sections because the equivalence classes of points (from where we look at a polytope) are easier to visualize than the equivalence classes of planes (cutting a given polytope). Section 5.4 takes advantage of this fact and describes an algorithm that computes the cross-section span of a polytope in terms of silhouettes in dual space.

Let Q be a convex polytope and q be a point outside the polytope. We assume in the following that the origin O belongs to the interior of the polytope. The *silhouette* of Q with respect to q is defined as the collection of faces f of Q that allow a supporting plane h of Q such that q lies in h and f is the relative interior of $Q \cap h$. Note that the silhouette is a set of edges and vertices of Q , unless q is coplanar with a facet of Q . The *size* of the silhouette is its number of vertices, and we define the *silhouette span* of Q , denoted $s(Q)$, as the size of the largest silhouette of Q . We might observe that whether or not we allow silhouettes to be defined with respect to points at infinity is of no consequence since there are

always small perturbations of the point of observation that do not decrease the size of the silhouette. Let $d \geq 3$ be an integer; we define $s_d(n)$ as the minimum value of $s(Q)$ over all polytopes Q of n vertices and vertex degrees at most d (note that vertex degree dualizes the notion of face degree). When no restriction is placed on the degree we simply write $s(n)$. Note that $s(n) = s_{n-1}(n)$.

To establish the equivalence of the cross-section and silhouette problems we introduce a geometric transform, D , that maps points to planes and vice versa. Let $p = (\pi_1, \pi_2, \pi_3)$ be a point distinct from the origin. We define a plane h normal to Op : if x is the point of coordinates $x = (\xi_1, \xi_2, \xi_3)$, we have

$$h = \{x | \langle x, p \rangle = \xi_1 \pi_1 + \xi_2 \pi_2 + \xi_3 \pi_3 = 1\}.$$

The transform D maps p to h and vice versa: by abuse of notation we write $h = D(p)$ and $p = D(h)$. Notice that D is not defined if either $p = O$ or h contains O ; this is assumed not to be the case in our discussion below. Given a plane $h = \{x | \langle x, v \rangle = 1\}$, we define two half-spaces $h^+ = \{x | \langle x, v \rangle < 1\}$ and $h^- = \{x | \langle x, v \rangle > 1\}$. Clearly, h^+ is the side of h that contains the origin, while h^- denotes the other side of h . It is elementary to prove that D preserves both incidence and order relationships.

Lemma 5.16. *We have the following incidence and order invariants: (i) a point p belongs to a plane h if and only if the point $D(h)$ lies in the plane $D(p)$; (ii) a point p belongs to h^+ if and only if $D(h) \in D(p)^+$, and p belongs to h^- if and only if $D(h) \in D(p)^-$.*

Next, we extend the domain of D to include polytopes. Let P be a polytope in E^3 whose interior contains the origin. We define the *dual polytope* of P as

$$Q = D(P) = \bigcap_{p \in P} \text{closure}(D(p)^+).$$

We prove below that there is a one-to-one correspondence between the facets of P and the vertices of Q [G]: a plane h intersects a facet of P if and only if its corresponding vertex v belongs to the silhouette of Q with respect to $D(h)$. This allows us to conclude with the following result.

Lemma 5.17. *The cross-section span of P is equal to the silhouette span of Q .*

Proof. Let f be a facet of P and let $v = \gamma(f)$ be the vertex of Q that belongs to all planes $D(p)$, $p \in f$. By definition of Q , γ is a bijective function mapping facets of P to vertices of Q . If h is a plane that intersects f , then there is a point $p \in h \cap f$. The plane $D(p)$ intersects Q in a unique vertex, namely v . Furthermore, from Lemma 5.16(i), the plane $D(p)$ contains the point $D(h)$. Thus, v belongs to the silhouette of Q with respect to point $D(h)$. The argument can be reversed to prove that if v belongs to the silhouette of Q with respect to $D(h)$, then $f = \gamma^{-1}(v)$ intersects the plane h . Incidentally, note that the spans of P and Q are unaffected if we discount all planes through the origin as well as parallel views of Q . The reason is that faces are defined as relatively open sets, so we can perturb a plane slightly without decreasing the span of the cross-section. \square

5.4. Computing the Cross-Section Span

This section describes a cubic algorithm for computing the silhouette span of a convex polytope Q . By the results of Section 5.3, this algorithm can be interpreted as computing the cross-section span of Q 's dual polytope, P . In fact, we develop the algorithm using arguments in both spaces.

Let P be an n -facet convex polytope which contains the origin and let Q be its dual polytope. Our method can be viewed as a search through all silhouettes of Q . Note, however, that we need not consider degenerate silhouettes, that is, silhouettes that contain facets of Q . In a first stage we describe the equivalence classes of points that define the same silhouette. Secondly, we indicate in which order these equivalence classes will be searched and we explain what this means in dual space. Finally, we map the problem to several instances of a two-dimensional line segment problem for which the best known method takes quadratic time and linear storage.

Let H be the set of planes that contain a facet of Q each. By Euler's formula, the cardinality of H , denoted m , is at most $2n - 4$. The set H dissects E^3 into a cell complex called the *arrangement* $\mathcal{A}(H)$ of H . We are concerned here with the *cells* of $\mathcal{A}(H)$, defined as the connected components of

$$E^3 \setminus \bigcup_{h \in H} h.$$

For example, the interior of Q is one of the cells of $\mathcal{A}(H)$. We refer to [E] and [Z] for a proof that $\mathcal{A}(H)$ contains at most $\binom{m}{3} + \binom{m}{2} + m + 1$ cells. This upper bound is tight if $\mathcal{A}(H)$ is *simple*, that is, if every three planes intersect in one point and no point belongs to more than three planes. For convenience, we henceforth assume that $\mathcal{A}(H)$ is simple. A procedure known as *simulation of simplicity* (see Chapter 9 of [E]) can be invoked to treat nonsimple arrangements just as though they were simple. The cells of $\mathcal{A}(H)$ constitute the aforementioned equivalence classes of points. More specifically, we have the following result.

Lemma 5.18. *If p and q are points of the same cell then the silhouettes of Q with respect to p and q are the same. Similarly, if p and q are points of different cells then the silhouettes of Q with respect to p and q are different, unless every plane in H separates point p from point q .*

The second part of Lemma 5.18 states that the silhouettes are the same when p and q belong to opposite unbounded cells. This relates to the fact that two parallel views of Q generate the same silhouette if they are taken from opposite sides. To avoid computational difficulties, we treat opposite unbounded cells separately. This will not affect the asymptotic time complexity of the algorithm.

Our algorithm visits the cells of $\mathcal{A}(H)$ in m steps. At each step we pick a plane h in H and visit all cells that have a facet in h and lie on the other side of h than Q . These are exactly the cells intersecting the plane h_ϵ that is obtained by moving h by a distance ϵ away from Q , for ϵ sufficiently small. In dual space,

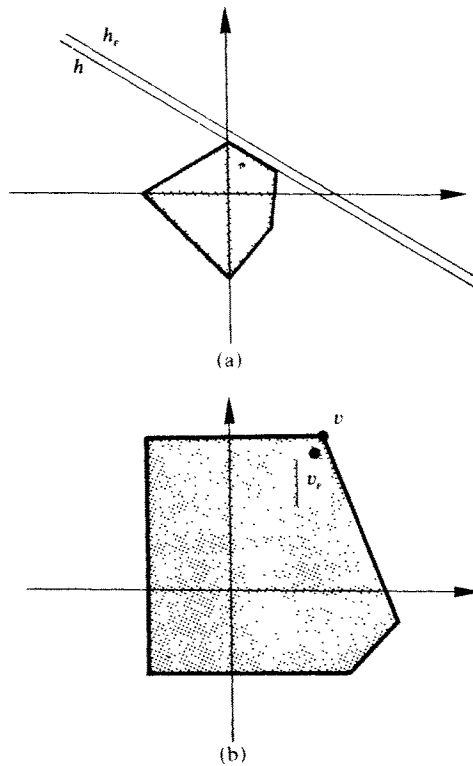


Fig. 5.19

h corresponds to a vertex $v = D(h)$ of P , the dual polytope of Q , and h_ϵ corresponds to a point $v_\epsilon = D(h_\epsilon)$ close to v and inside P . Figure 5.19 illustrates this argument using two dimensions as the embedding space. Figure 5.19(a) shows Q as a convex polygon and the planes h and h_ϵ as lines. Figure 5.19(b) depicts P , the dual polygon of Q , as well as the points v and v_ϵ . A point q in h_ϵ corresponds to a plane $D(q)$ that contains v_ϵ , and the silhouette defined by q corresponds to the cross-section defined by $D(q)$.

We are now ready to reduce the problem to a two-dimensional problem involving line segments and lines. Let g be the plane through the origin normal to Ov_ϵ . We map each edge of P into g by a central projection about v_ϵ . In order to obtain a plane subdivision of g we do not project the pieces of P 's edges separated from g by the parallel plane through v_ϵ . It is easy to keep track of how many such edges a plane intersects and to correct the result for this plane accordingly. The problem is now to find a line in g that intersects the maximum number of edges of this subdivision. Thus, we arrive at the same problem as in Section 3.4. This problem can be solved in $O(n^2)$ time and $O(n)$ storage using the topological sweep method given in [EG]. Since the cross-section span requires running this algorithm $O(n)$ times, once for each vertex of P , this yields a method that takes $O(n^3)$ time and $O(n)$ storage.

It is perhaps worthwhile to mention that the techniques above can be used to find the maximum silhouette in $O(n^2)$ time and $O(n)$ storage, provided that the locations of the sources of light are restricted to a single plane (even if this plane is at infinity). If the locations are further restricted to a single line in E^3 , then a straightforward $O(n \log n)$ algorithm can be given using the same geometric setting which led to the cubic and the quadratic algorithm.

5.5. Discussion

Theorem 1, Lemma 5.1, and Lemma 5.15 together prove the combinatorial part of Theorem 3. The computational results are provided in Section 5.4. The equivalence result of Lemma 5.17 establishes Theorem 4 as a direct corollary of Theorem 3.

Theorem 3 (Cross-Section Span). *The cross-section span of a convex polytope of n facets can be computed in $O(n^3)$ time and $O(n)$ space. The minimum cross-section span over all n -facet polytopes is $\Theta(\log n / \log \log n)$. If the degree of the polytope is d or less then the bound is $\Theta(\log_a n + \log n / \log \log n)$.*

Theorem 4 (Silhouette Span). *The silhouette span of a convex polytope of n facets can be computed in $O(n^3)$ time and $O(n)$ space. The minimum silhouette span over all n -facet polytopes is $\Theta(\log n / \log \log n)$. If no vertex of the polytope is incident to more than d edges then the bound is $\Theta(\log_a n + \log n / \log \log n)$.*

6. Max–Min Problems

In this section we provide a proof of Theorem 5, which gives bounds on the line width and monotone width of convex subdivisions and the cross-section width of convex polytopes. Recall that the *line width* of a subdivision is the minimum number of regions that any line can intersect. The *monotone width* is the minimum value over all directions of the maximum length of any contact chain. Finally, the *cross-section width* of a convex polytope is the minimum number of facets that any plane through the origin intersects. Similarly, the *silhouette width* of a polytope is the minimum number of edges in any parallel view. Recall that a parallel view is a silhouette with the source of light at infinity.

Trivially, the line width and the monotone width of a subdivision of size n are both less than or equal to n . As can be seen from Fig. 6.1(a), there exists a subdivision of line and monotone width proportional to n . Note that the degree of this subdivision is 5 and that the edges assume only two distinct slopes.

We prove the only nontrivial statement of Theorem 5 below: the lower bound on the cross-section width of a convex polytope. We can assume that the polytope P contains the origin O in its interior and the cutting plane passes through the origin (else the width vanishes). By the duality result of Section 5, the cross-section problem for P is equivalent to the silhouette problem for its dual polytope, denoted Q . To prove that the maximum cross-section width is proportional to n ,

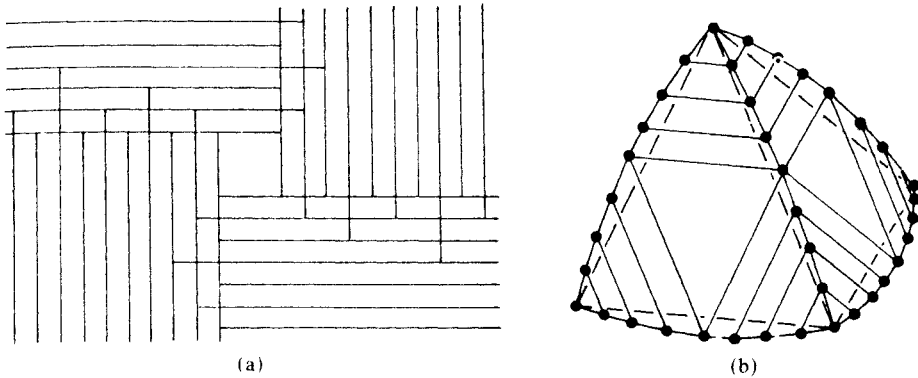


Fig. 6.1

the number of facets, it is thus sufficient to construct a polytope Q of n vertices such that every parallel view yields a linear silhouette.

Lemma 6.1. *There is a convex polytope of n vertices in E^3 (each adjacent to at most three other vertices) such that the orthogonal projection onto any plane is a polygon of $\Omega(n)$ vertices.*

Proof. The polytope Q is shown in Fig. 6.1(b). Take a regular tetrahedron \mathcal{T} of side length equal to 1, and replace each edge by a circular arc of sufficiently large radius, such that the plane that contains the circle cuts the tetrahedron in two equal-size pieces. Each one of the six arcs contains $(n-4)/6$ vertices of Q , not counting the endpoints. We assume that $(n-4)/6$ is integral, and we place the vertices uniformly on the arcs (i.e., at regular intervals). Next, we prove that every parallel view of Q contains all vertices of at least one arc, provided that the radius of the arcs is larger than some threshold. To demonstrate the existence of such a threshold we must argue about the area and the angles of a planar projection of the regular tetrahedron \mathcal{T} .

Claim 1. The length of every edge of such a projection is at most 1.

Claim 2. Every projection of \mathcal{T} has an area greater than some constant c_A .

Claim 1 is obvious. Claim 2 follows from the fact that the area of the projection of Q always exceeds the area of the projection of the inscribed sphere of the initial tetrahedron: this area is equal to $6\pi/144$.

The projection of \mathcal{T} is either a triangle or a quadrilateral. If it is a triangle then its edges belong to the silhouette. The projections of the other edges decompose the triangle into three smaller triangles (Fig. 6.2(a)). Let Δ be the largest of these triangles. Since its area is at least $c_A/3$ and the length of each edge is at most 1 (Claims 1 and 2), we know that there is a positive lower bound c_α on the size of each angle of Δ , as well as a positive lower bound c_l on the length of each edge of Δ . We choose c_α and c_l so as to be also lower bounds on the angles and edges of the largest of the four triangles that we get if the projection of \mathcal{T} is a quadrilateral (Fig. 6.2(b)).

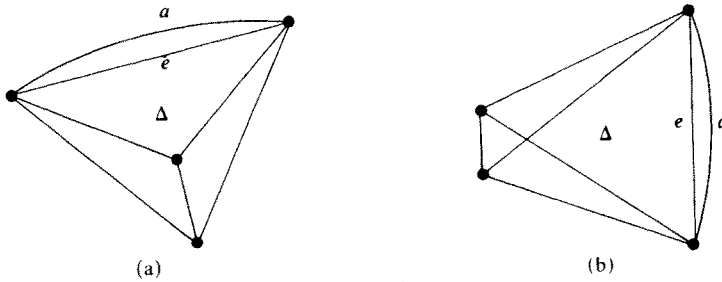


Fig. 6.2

Now let e be an edge of the parallel view of the triangle Δ . The only possibility that not all vertices on the circular arc a of e are vertices of the parallel view of Q is that the projections of other circular arcs intersect the projection of a . Let b be such an arc and let β be the angle between the projection of b and the projection of its edge.

First, consider the case where the edge of b projects onto an edge of Δ . Then arc b can intersect arc a only if $\beta \geq c_\alpha$. Let γ be the angle between arc b and its edge in three dimensions. There is a constant c such that $\beta < c\gamma$: the reason is that the length of the projection of the edge is at least c_1 , which prohibits β from being too large. Thus, we can choose γ small enough to prevent β from exceeding c_α , which ensures that the projections of arcs a and b cannot intersect. Next, let the edge of b project onto an edge that does not bound Δ . The same argument as for the other edges applies, unless this edge does not even share an endpoint with e —in this case, the length of this edge can be arbitrarily small. Fortunately, each of its points is at least a fixed minimum distance away from every point of e . This fixed distance guarantees that there is no interference if the radii of the circles are chosen above some threshold.

If the radius of all six arcs is chosen to be the same, and if the distance between any two consecutive points on any arc is also the same, then Q , the convex hull of these points, has degree 6. Thus, Lemma 6.1 also holds if we bound the vertex degrees from above by 6. Note that we can easily achieve maximum degree 3 if we truncate the polytope at every vertex. \square

7. Conclusions

Possible extensions of our work go mostly in the direction of better algorithms for the problems considered. Beating quadratic time is likely to require a novel idea, and seems elusive at this time. The combinatorial bounds given in the paper are by and large optimal, so working out the constant factors appears to be the next (exciting) task at hand. A few gaps remain in some of the bounds we gave, in particular regarding monotone spans with restrictions on the number of slopes. Another interesting question is to inquire about the expected value of spans in some reasonable probabilistic model. The problem of cross-sections seems the most natural and worthy of interest in this regard.

Extensions to higher dimensions are interesting in their own right. Unfortunately, face-counts cease to be linearly related in four dimensions and above, so the problems lose some of their natural appeal. An interesting question which we leave as an open problem is as follows: using linear preprocessing, it is possible to intersect an n -facet convex polytope with an arbitrary plane in time $O(k + \log n)$, where k is the size of the output?

Acknowledgments

We wish to thank J. Pach and E. Szemerédi for valuable discussions on several of the problems studied in this paper.

References

- [CH] Connelly, R. and Henderson, D. W., A convex 3-complex is not simplicially isomorphic to a strictly convex complex, *Math. Proc. Cambridge Philos. Soc.* **88** (1980), 299–306.
- [E] Edelsbrunner, H., *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg, 1987.
- [EG] Edelsbrunner, H. and Guibas, L. J., Topologically sweeping an arrangement, Tech. Rep. 9, DEC/SRC, April 1986. Also, *Proc. 18th STOC*, 1986, pp. 389–403.
- [G] Grünbaum, B., *Convex Polytopes*, Wiley, New York, 1967.
- [GS] Guibas, L. J. and Stolfi, J., Primitives for the manipulation of general subdivisions and the computation of Voronoi diagrams, *ACM Trans. Graphics* **4** (1985), 75–123.
- [GY] Guibas, L. J. and Yao, F. F., On translating a set of rectangles, in *Advances in Computing Research*, Vol. 1 (F. P. Preparata, ed.), JAI Press, Greenwich, CT, 1983, pp. 61–77.
- [H] Hall, M., *Combinatorial Theory*, Blaisdell, Waltham, MA, 1967.
- [MP] Moser, W. and Pach, J., Research Problems in Discrete Geometry, Manuscript, 1986.
- [Z] Zaslavsky, Th., Facing up to arrangements: face-count formulas for partitions of space by hyperplanes, *Mem. Amer. Math. Soc.* **154** (1975), 1–71.

Received March 16, 1987, and in revised form March 17, 1988.