On the Lower Envelope of Bivariate Functions and its Applications

Herbert Edelsbrunner[1], *János Pach*[2,3], *Jacob T. Schwartz*[3] *and Micha Sharir*[3,4]

[1] Computer Science Department, University of Illinois at Urbana-Champaign
[2] Mathematical Institute of the Hungarian Academy of Sciences
[3] Courant Institute of Mathematical Sciences, New York University
[4] School of Mathematical Sciences, Tel Aviv University

Abstract: *We consider the problem of obtaining sharp (nearly quadratic) bounds for the combinatorial complexity of the lower envelope (i.e. pointwise minimum) of a collection of n bivariate (or generally multi-variate) continuous and "simple" functions, and of designing efficient algorithms for the calculation of this envelope. This problem generalizes the well-studied univariate case (whose analysis is based on the theory of Davenport-Schinzel sequences), but appears to be much more difficult and still largely unsolved. It is a central problem that arises in many areas in computational and combinatorial geometry, and has numerous applications including generalized planar Voronoi diagrams, hidden surface elimination for intersecting surfaces, purely translational motion planning, finding common transversals of polyhedra, and more. In this abstract we provide several partial solutions and generalizations of this problem, and apply them to the problems mentioned above. The most significant of our results is that the lower envelope of n triangles in three dimensions has combinatorial complexity at most* $O(n^2\alpha(n))$ *(where* $\alpha(n)$ *is the extremely slowly growing inverse of Ackermann's function), that this bound is tight in the worst case, and that this envelope can be calculated in time* $O(n^2\alpha(n))$.

1. INTRODUCTION; OVERVIEW OF RESULTS

In this paper we study the problems described in the abstract, and derive upper bounds (and some matching lower bounds) on the complexity of the lower envelope of certain collections of bivariate functions. These results solve various special but significant cases of the generalization of the following problem, initially proposed by Davenport and Schinzel [DS], to the case of bivariate functions: Let $f_1(x), \ldots, f_n(x)$ be n continuous univariate functions, each pair of which intersect in at most s points. Let $\lambda_s(n)$ denote the maximum number of maximal connected portions of the graphs of the f_i's in such a collection which compose the graph of their lower envelope. It is known (cf. [DS], [At]) that $\lambda_s(n)$ is also equal to the maximum length of a sequence $U = (u_1, \ldots, u_m)$ of integers (called an *(n,s) Davenport Schinzel sequence*) which satisfies the following conditions:

(i) $1 \le u_i \le n$ for each i.

(ii) For each $i < m$ we have $u_i \ne u_{i+1}$.

(iii) There do not exist $s+2$ indices $1 \le i_1 < i_2 < \cdots < i_{s+2} \le m$ such that $u_{i_1} = u_{i_3} = u_{i_5} = \cdots = a$, $u_{i_2} = u_{i_4} = u_{i_6} = \cdots = b$, and $a \ne b$.

The problem of estimating $\lambda_s(n)$ has been studied repeatedly; see [DS], [Da], [Sz], [At], [HS], [Sh1], [Sh2], [ASS]. It is known that

$\lambda_1(n) = n$; $\lambda_2(n) = 2n-1$ (trivial).

$\lambda_3(n) = \Theta(n\alpha(n))$, where $\alpha(n)$ is the functional inverse of Ackermann's function, and thus grows extremely slowly [HS].

$\lambda_4(n) = \Theta(n \cdot 2^{\alpha(n)})$ [ASS].

$\lambda_{2s}(n) = O(n \cdot 2^{O(\alpha(n)^{s-1})})$, for $s>2$ [ASS].

$\lambda_{2s+1}(n) = O(n \cdot \alpha(n)^{O(\alpha(n)^{s-1})})$, for $s \ge 2$ [ASS].

$\lambda_{2s}(n) = \Omega(n \cdot 2^{\Omega(\alpha(n)^{s-1})})$, for $s>2$ [ASS].

These results have found many applications to diverse problems in computational geometry (see [BS], [CS], [At], [HS], [OSY], [PSS]).

All this gives relatively satisfactory information concerning the "one-dimensional" Davenport-Schinzel problem. However, the two-dimensional generalization of this problem is still largely uninvestigated (with the exception of the simple case where each f_i is a plane [PrM], and a few recent initial studies of more complex cases [KLPS], [SL]; see also [Au] for the case where each f_i is a sphere), and appears to be much harder. In this generalization one considers a collection $F = \{f_1(x,y), \ldots, f_n(x,y)\}$ of n continuous bivariate functions such that any three functions intersect in at most some fixed number s of points, and such that each pair of functions intersect in a curve that has at most some fixed number t of singularities. The goal is to obtain sharp upper and lower bounds on the maximum complexity $\kappa(F)$ (i.e. number of faces, edges, and vertices) of the planar map, which can be called the *minimization diagram* of F, obtained by projecting the pointwise minimum of these functions onto the $x-y$ plane. Each region of this map consists of a maximal connected set of points at which the minimum is attained by a particular function f_i, and the edges

Work on this paper by the first author was supported by Amoco Fnd. Fac. Dev. Comput. Sci. 1-6-44862. Work on this paper by the last three authors was supported by Office of Naval Research Grant N00014-82-K-0381, National Science Foundation Grant No. NSF-DCR-83-20085, and by grants from the Digital Equipment Corporation, and the IBM Corporation.

(resp. vertices) of this map consist of points at which the minimum is attained simultaneously by two (resp. three) functions. We assume here that the functions in F are in "general position", thereby excluding degeneracies at which two functions coincide on a two-dimensional region, or three functions coincide on a one-dimensional set, or four functions coincide at all, etc.

The two-dimensional Davenport-Schinzel problem arises as a central subproblem in many problems in computational and combinatorial geometry. As an illustration, note that almost any conceivable generalization of the (nearest neighbor) Voronoi diagram in the plane can be regarded as the minimization diagram of a certain collection of functions, each measuring the distance from a test point to one of the objects defining the diagram (see [ESe]). Similarly, the boundary of the configuration space of free positions of a moving system with three degrees of freedom can often be defined as the envelope of a certain collection of 2-D surfaces, each representing contact positions with some obstacle. In visibility problems, one often considers some space of rays, and for each ray the object it "sees" is the one whose intersection with the ray is nearest to the view-point. Thus if the ray space is two dimensional (as in the case of rays emanating in the plane from a fixed segment, or rays emanating in 3-space from a fixed point), analysis of the visibility along these rays reduces to the calculation of the lower envelope of an appropriate collection of "distance functions" along the rays. These few examples should suffice to illustrate the significance of the 2-D Davenport Schinzel problem, but more applications will be given below.

This abstract presents several partial solutions to, and some generalizations of this two-dimensional Davenport-Schinzel problem. Clearly, a trivial upper bound for $\kappa(F)$ is $O(n^3)$ (with a constant depending on s and t), and our main goal is to improve that bound. The results presented in the abstract are a combination of results obtained in three different papers [SS], [PS], [ES]. We first give an overview of our results. Some details of the proofs are given in Section 3. More details can be found in the full versions [SS], [PS], and [ES].

1.1. The single intersection and the double intersection cases [SS]

Our first result deals with a particularly favorable case, in which we assume each pair of functions in F to intersect in a connected simple curve which has the additional property that each plane cross-section of the form $x = const$ intersects it in exactly one point. Furthermore, we assume that any three functions in F intersect in at most one point. In this case we prove that $\kappa(F) = O(n)$ (a bound which is about one order of magnitude better than the "general" case - see below). We also present an $O(n \log n)$ algorithm for the calculation of the lower envelope of F; this algorithm is closely related to Shamos' algorithm for the calculation of Voronoi diagrams of points in the plane (see [PrS]). (This result, even without the additional condition about monotonicity of the intersection curves, is also a consequence of the upper bound theorem for oriented matroids [Ma], and can be extended to higher dimensions, but the proof we present is much simpler, and has the advantage that it also yields an algorithm for the calculation of the envelope.) We also discuss the case in which any three functions intersect in at most two points, and any two intersect in a connected simple curve. In this case, extending recent results on properties of the intersection of planar Jordan curves, given in [KLPS], we show that $\kappa(F) = O(n^2)$, and that this bound is tight in the worst case.

1.2. The envelope of functions with favorable cross sections [SS]

We conjecture that for $s > 2$ the complexity $\kappa(F)$ is at most $O(n\lambda_{s'}(n))$, for some s' depending on s (and possibly also on t). That $\kappa(F)$ can indeed be superquadratic is shown in an example given in [SL]; a similar lower bound, but involving only piecewise linear functions, is given in this abstract.

Our first main result proves this conjecture for the special case in which the intersection curve of any pair of the functions in F intersect every plane of the form $x = const$ in at most 2 points. We show that in this case $\kappa(F) = O(n\lambda_{s+2}(n))$, where the constant of proportionality depends on s and t, and present an $O(n\lambda_{s+2}(n) \log n)$ time algorithm for the calculation of the minimization diagram of F. Our analysis proceeds by reducing this restricted case of the 2-D problem to a collection of 1-D problems involving lower or upper envelopes of certain subsets of the intersection curves of pairs of the functions f_i.

The extra condition assumed above is somewhat artificial, but nevertheless covers certain applications in which the functions $f_i(x_0, y)$ have relatively simple form (as functions of y) for each fixed x_0, e.g. are linear or quadratic in y. For example, such a situation arises in analysis of the pattern of changes in the convex hull of n moving points, as is noted in one of the applications that we present. Our result then implies that the number of combinatorial changes in the convex hull is $O(n\lambda_s(n))$, for some constant s depending on the kind of motion of the points, and that these changes can all be found in time $O(n\lambda_s(n) \log n)$. (The particular convex hull problem we discuss as an example was also studied by Atallah [At] using a different technique.)

1.3. The envelope of n triangles in three dimensions [PS], [ES]

Our second major result, and perhaps the most significant of our results, proves a strong form of the above conjecture for the case in which each function in F is piecewise linear, so that all the graphs of these functions have altogether n faces. In this case we can replace F by a collection of $O(n)$ "triangle functions", each of which has a graph that consists of a triangle (in arbitrary position in 3-space) with three adjacent steeply rising unbounded faces. See Fig. 1 for an illustration of the envelope of triangles.

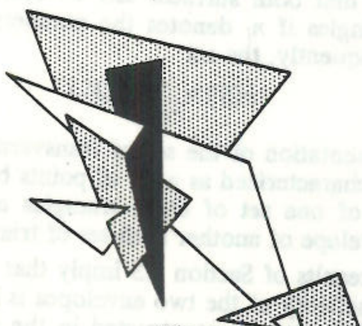

Fig. 1

Unfortunately, the intersection curve of a pair of such triangle functions can have cross sections of the form $x = const$ consisting of three points, so that the preceding result does not apply here. Nevertheless we show that the complexity of the lower envelope of n such functions is at most $O(n\lambda_3(n)) = O(n^2\alpha(n))$. Moreover, using a recent result of [WS], we show that this bound is tight in the worst case. Our analysis, which uses a divide and conquer approach, also yields an algorithm, based on the same technique, for calculating this envelope in time $O(n^2\alpha(n))$. Because of the particular simple form of piecewise linear functions, the problem of estimating the complexity of their envelope has been one of the major open problems in the 2-D Davenport Schinzel theory, and, as shown below, has many applications. We also provide generalization of this tight bound to the envelope of multi-variate piecewise linear functions.

1.4. The complexity of a region bounded by convex plates [PS]

Our last major result considers generalization of the notion of lower envelope to that of the boundary of a single connected component of the complement of the union of n "plates" (compact convex 2-D sets) in 3-space. This useful generalization is much harder to analyze, and so far the only known upper bound on the complexity of such a component was the trivial $O(n^3)$ bound. Using a combinatorial result of Erdős, we show that the complexity of any such component is at most $O(n^{3-\delta})$, for some absolute positive constant $\delta \geq 1/49$. Although this result is weaker than the bounds mentioned above, it applies in more general contexts. We also present generalizations of this result to higher dimensions.

2. APPLICATIONS

The results on the combinatorial complexity of envelopes have many applications in computational geometry. In this abstract we present a few of these applications.

2.1. Hidden Line and Surface Removal

Imagine that we take a picture of a three-dimensional scene. To compute what this picture looks like, assuming opaque objects, is commonly known as the hidden line/surface removal problem. If the surfaces that bound the objects are allowed to intersect then this application is essentially just a reformulation of the envelope problem. The portions of the surfaces visible from the point can be easily interpreted as the lower envelope of the distance functions from the viewpoint to the given surfaces. Thus if the functions in question belong to one of the families for which we have obtained tight bounds on the envelope complexity, then we obtain good bounds on the hidden line/surface removal problem.

For example, consider the case where the scene is modeled by n triangles in three dimensions. If the viewpoint is at infinity, then the triangles themselves can be thought of as the graphs of their own distance functions from the point. Otherwise, we can apply a projective transformation to map the viewpoint to infinity and the triangles to other triangles. (To get a valid picture though we need to map the plane at infinity to a plane, using the same transformation, which acts as a background screen.) With intersections allowed we thus have an $O(n^2\alpha(n))$ time algorithm to compute the view from a fixed point using the results of Section 1.3. This is asymptotically optimal since there are scenes whose images actually have combinatorial complexity $\Omega(n^2\alpha(n))$.

In modeling a three-dimensional scene by triangles it is more common, however, that the triangles do not intersect except possibly at their relative boundaries. In this case it is not difficult to show that the combinatorial complexity of the envelope is $O(n^2)$ in the worst case. In this case our algorithm computes the envelope, resp. the viewed image of the triangles, in $O(n^2)$ time and thus is optimal in the worst case (the same complexity has been achieved also in [De], [MK]). In fact, the non-intersection property allows us to simplify the algorithm considerably and to generalize it to $d \geq 3$ dimensions. For n $(d-1)$-simplices in d dimensions it computes the envelope in $O(n^{d-1})$ time which again is worst-case optimal.

2.2. Translating a Polyhedron in Three Dimensions

This is an instance of the motion planning problem, in which we wish to plan a collision-free path for an object in three dimensions in the presence of a collection of obstacles which the object must avoid. In this section we consider a scenario where the object, B, as well as the obstacles, $A_1, A_2, \ldots, A_m$, are convex polyhedra. The A_i are pairwise disjoint (assuming they are open sets) and stationary. B can be translated (but not rotated) as long as it avoids the obstacles. Thus, a *free placement* of B is defined by a vector (or point) b such that $B+b$ avoids all A_i. By taking the Minkowski difference

$$K_i = A_i - B = \{p = x - y : x \in A_i, y \in B\}$$

we obtain a new "expanded" obstacle with the property that $B+b$ avoids A_i if and only if $b \notin K_i$. It is not difficult to see that the combinatorial complexity of K_i is $O(k \cdot n_i)$ if B and A_i are bounded by k and n_i faces respectively. The set of free placements can thus be represented by the complement K^c of $K = \bigcup_{i=1}^{m} K_i$.

Notice that the connected component of K^c which contains the origin represents the set of free placements of B that can be reached from the initial given placement of B by a continuous translational motion. This set is a connected component of the three-dimensional space minus $O(kn)$ triangles, where $n = \sum_{i=1}^{m} n_i$. By the result in Section 1.4 its combinatorial complexity is subcubic in kn (that is $O((kn)^{3-1/49})$). Little is known beyond this result unless we make some assumptions about the obstacles.

An interesting case is obtained when the union of the (closures of the) A_i is bounded by a polyhedral terrain. This is a continuous (piecewise linear) surface that intersects any vertical line in exactly one point. In this case, the boundary of K^c, the set of free placements of B, is the upper envelope of the K_i. The result in Section 1.3 now implies that its combinatorial complexity is $O(k^2n^2\alpha(kn))$ and that it can be constructed in the same amount of time. Since the boundary of K^c is again a terrain we can compute its intersection with any vertical line in logarithmic time (see e.g. [E]) using two-dimensional point location techniques. This implies that we can decide in logarithmic time whether a given placement of B is free. Furthermore, this can be used to compute in the same time the contact point(s) when B is lowered from any "query" free position until it touches the terrain.

2.3. Stabbing Line Segments and Polytopes

A *transversal* or *stabbing plane* of a set of (connected) objects in three dimensions is a plane that intersects each object in the set. Since a plane meets a connected object if and only if it meets its convex hull, we may as well assume that all objects in the set, S, are convex. This subsection considers the complexity of finding all transversals of S if S is a set of m polytopes bounded by a total number of n faces. We do not assume that the polytopes are disjoint.

Following the approach of [EMPRWW] for the corresponding two-dimensional problem we make use of the dual transform D that maps a point $p=(\pi_1,\pi_2,\pi_3)$ to the non-vertical plane $x_3=-\pi_1x_1-\pi_2x_2+\pi_3$ and a plane $z=ax+by+c$ into the point (a,b,c). This transform preserves incidence and order, in the sense that a point p lies on/above a plane h if and only if the plane $D(p)$ passes through/above the point $D(h)$. For each polytope $P_i \in S$ we can now define $D(P_i)$ as the union of all planes $D(p)$, $p \in P_i$. Because D is incidence preserving, a plane h intersects P_i if and only if $D(h) \in D(P_i)$. The set $D(P_i)$ is determined by the set of planes that are dual to the vertices of P_i. For each vertical line, $D(P_i)$ contains all points that lie between the topmost and the bottommost intersection of the line with any such plane. It follows that $D(P_i)$ can be characterized as the set of points between two piecewise linear surfaces and that both surfaces can be decomposed into $O(n_i)$ triangles if n_i denotes the numbers of faces of P_i. Consequently, the set

$$\sigma(S) = \bigcap_{P_i \in S} D(P_i)$$

is a representation of the set of transversals of S, and it can be characterized as a set of points below a lower envelope of one set of $O(n)$ triangles and above an upper envelope of another such set of triangles.

The results of Section 1.3 imply that the combinatorial complexity of the two envelopes is in $O(n^2\alpha(n))$ and that they can be constructed in the same amount of time. The same is true for the set $\sigma(S)$ itself, as can be checked by straightforward adaptation of the proof of Theorem 8 below. It is interesting to note that, if each P_i in S is a line segment, then the combinatorial and computational complexity reduces to $O(n^2)$. This is because the set $\sigma(S)$ has combinatorial complexity $O(n)$ if S is a set of n line segments in the plane (see [EMPRWW]). As a consequence, we get the recurrence relation $\psi(n) \le 2\psi(\frac{n}{2}) + O(n^2)$ which solves to $O(n^2)$ in the proof of Theorem 8.

2.4. Voronoi Diagrams of Point Clusters

In this section we consider certain problems in Euclidean location theory, where the goal is to analyze the min-max or the max-min of the distance from a test point to groups of resources in the plane. By squaring and removing the quadratic terms from the distance functions, these problems can be reduced to that of calculating the lower or upper envelope of a collection of bivariate piecewise linear functions. Using the results of Section 1.3, these problems can be solved in nearly quadratic time by envelope constructions.

An interesting special case of this kind of problem is the construction of the Voronoi diagram for a set of point clusters in the plane. Here we measure the distance from a point to a cluster by the maximum distance to a point in the cluster. As mentioned above, the complexity of this problem is $O(n^2\alpha(n))$ both in terms of the number of faces and the time required to build it, where n is the total number of points in the clusters. If each cluster contains only one or two points then the complexity of the diagram goes down to $O(n^2)$ and it can be shown that this bound is tight. A dramatic decrease in complexity happens if we require that the convex hulls of any two clusters are disjoint. In this case we can show that the region of each cluster is connected which implies that the overall complexity of the diagram is $O(n)$. No subquadratic algorithm for constructing such a diagram is currently known.

2.5. Decision procedure for linear inequalities

Here we consider the problem of deciding quantified Boolean formulae involving inequalities between linear functions of two variables. Assume that the unquantified portion of such a formula f is given in conjunctive normal form. That is, it has the form $p_1 \wedge p_2 \wedge \cdots \wedge p_m$, where each p_i has the form $L_{i1} \geq 0 \vee L_{i2} \geq 0 \vee \cdots \vee L_{in_i} \geq 0$, where each L_{ij} is a linear form in the two variables x,y. Let $n = \sum_{i=1}^{m} n_i$ be the total number of inequalities. Let h_i denote the pointwise maximum of $L_{i1}, \ldots, L_{in_i}$, and let M be the pointwise minimum of $h_1, \ldots, h_m$. Then it is easily checked that the unquantified part of f is true at some (x,y) if and only if $M(x,y) \geq 0$. Thus to decide, say, the validity of f, we can calculate (the minimization diagram of) M, and then search over it in an appropriate manner. Note that each function h_i is convex and its graph consists of $O(n_i)$ planar faces (and can be calculated in time $O(n_i \log n_i)$). Thus the overall combinatorial complexity of M is $O(n^2\alpha(n))$, and it can be computed in the same amount of time. This yields a nearly quadratic decision procedure for such formulae.

2.6. Convex hulls of general objects

Our last application deals with calculation of convex hulls of objects other than points, or of objects that vary in time. We show that the complexity of such hulls is equivalent to the complexity of the lower envelope of certain "support functions", and the above results yield sharp bounds for this complexity in several cases. One such application have already been mentioned; as another application, we show that the convex hull of n balls in 3-space has combinatorial complexity $O(n^2)$, and this is tight in the worst case.

To see the case of balls in more detail, let $B_1, \ldots, B_n$ be n given balls in three dimensions. For each B_i define a function f_i on the unit sphere S^2 so that for each $u \in S^2$, $f_i(u)$ is the distance from the origin to the plane supporting B_i and having u as its outward-directed normal. It follows easily that $M(u) = \max_i f_i(u)$ gives the distance from the origin to the plane supporting the convex hull C of $B_1, \ldots, B_n$ and having u as its outward-directed normal. In this set-up, calculating C becomes equivalent to the calculation of the minimization diagram of these support functions. For balls it is easily checked that any two functions f_i and f_j intersect in a circle along S^2, and that any three functions intersect in at most two points. Using Theorem 5, we obtain the asserted bound. The lower bound example is given in [SS].

3. TECHNICAL DETAILS

In this section we provide more details about the results reported in this abstract.

3.1. The single intersection case

Here we assume that the collection $\mathbf{F}$ of bivariate functions satisfies the following additional properties:

(1a) For each pair f_i, f_j of distinct functions in $\mathbf{F}$ the curve γ_{ij} defined by $f_i(x,y) = f_j(x,y)$ is either empty or is connected and each of its cross-sections of the form $x=x_0$ consists of exactly one point; moreover, the cross sections $f_i(x_0,y)$, $f_j(x_0,y)$ intersect transversally at that point.

(1b) Each triple f_i, f_j, f_k of distinct functions in $\mathbf{F}$ intersect in at most one point.

Although these conditions are very restrictive, they do arise in certain applications, as will be noted below. In particular, they apply to any collection $\mathbf{F}$ of linear functions (for which the foregoing results are of course well known).

Theorem 1: If $\mathbf{F}$ is a collection of n bivariate functions satisfying the above assumptions, then $\kappa(\mathbf{F}) = O(n)$.

Sketch of proof: For any fixed x_0, each pair of restricted functions $f_i(x_0,y)$ has at most one intersection, and therefore the minimum $M(x_0,y)$ is attained by each function along a single interval (which may be empty). As x_0 increases, the sequence of functions, arranged in the order they attain M along the line $x=x_0$, can change only at critical values of x_0 at which a new function appears along $M(x_0,y)$ between two others, or the interval along which some function attains M shrinks to a point and disappears. It can be shown that a function can disappear from M at most once, so that there are at most $2n$ critical values of x_0. The theorem follows immediately from this. □

Next we describe an algorithm for calculating the minimization diagram M^* of a collection $\mathbf{F}$ of functions satisfying the hypotheses of the preceding theorem. The algorithm runs in $O(n \log n)$ time (assuming certain primitive operations on the f_i's, e.g. calculating the intersection point of a triple of them, to take constant time). It follows from the observations in [ESe] that our algorithm also yields another efficient technique for calculating planar Voronoi diagrams.

We begin by sorting the functions in $\mathbf{F}$ according to their values at $y = +\infty$, and by eliminating functions which are larger than their predecessor in this order throughout the entire plane; call the resulting sequence $f_1, \ldots, f_n$. Partition it into a *lower group* $\mathbf{F}_l = \{f_1, \ldots, f_{n/2}\}$ and an *upper group* $\mathbf{F}_u = \{f_{n/2+1}, \ldots, f_n\}$, and recursively calculate the minimization diagrams M_l^*, M_u^* representing the lower envelopes M_l, M_u of $\mathbf{F}_l$, $\mathbf{F}_u$ respectively. We represent each such diagram by a list of triple intersections of the functions in the corresponding group, arranged in ascending order of their abscissae, and by a list of the functions constituting the minimum, in order of increasing y, for x_0 below the least abscissa of any triple intersection. As shown e.g. in [Co], [CG], [DSST] this representation can be transformed in $O(n \log n)$ time into a linear-size data structure that supports e.g. efficient point location in M^*

To "merge" the two diagrams M_l^*, M_u^* into the diagram M^* for the entire $\mathbf{F}$, we first show

Lemma 2: For each x_0 there exists exactly one y_0 at which $M_l(x_0,y_0) = M_u(x_0,y_0)$. For every $y > y_0$ we have $M(x_0,y) = M_l(x_0,y) < M_u(x_0,y)$ and for every $y < y_0$ we have $M(x_0,y) = M_u(x_0,y) < M_l(x_0,y)$.

Proof: Omitted in this version. □

The preceding lemma shows that the *separating contour*

$$C = \{(x,y) : M_l(x,y) = M_u(x,y)\}$$

is an unbounded x-monotone and connected curve, and that on the upper side (resp. lower side) of C the minimization diagram M^* coincides with M_l^* (resp. with M_u^*).

Take some x_0 below the smallest abscissa of any triple intersection for either of the two groups, merge there the two lists of the functions f_i achieving the minimum in each group separately, to obtain that list for the entire collection, and find a "starting point" on C, in overall linear time.

Suppose at this point we have $f_i = f_j$, for some $f_i \in \mathbf{F}_l$, $f_j \in \mathbf{F}_u$. We then perform a vertical line sweeping, in which we follow C along the curve $\gamma_{ij} : f_i = f_j$, and use a *local advance* procedure for finding the next triple intersection along C, by examining the intersection points of γ_{ij} with the four "neighboring" curves $\gamma_{i-1,i}$, $\gamma_{i,i+1}$, $\gamma_{j-1,j}$, and $\gamma_{j,j+1}$, and choosing the nearest one. As the sweep progresses, we also collect all triple intersections within $\mathbf{F}_l$ and within $\mathbf{F}_u$, which still lie on the appropriate side of C. At each triple intersection along C, we also insert a new region of M^* starting at that point or remove a region just ending there. All this can be done in linear time. More details can be found in [SS]. We thus have

Theorem 3: The minimization diagram of a collection $\mathbf{F}$ of n bivariate functions satisfying the conditions stated at the beginning of this section can be calculated in time $O(n \log n)$.

The following result generalizes Theorem 1 and indicates the essentially topological nature of that result.

Theorem 4: Let f_i, $i=1, \ldots, n$ be a set of n real-valued continuous functions defined in the plane E^2. Suppose that

(a) For each i, j between 1 and n, the set γ_{ij} of points $p \in E^2$ satisfying the condition $f_i(p) = f_j(p)$ is either a closed Jordan curve or an open Jordan curve both of whose ends approach infinity.

(b) For each fixed i, any two of the curves γ_{ij}, γ_{ik} intersect in at most one point, and each such intersection is transversal.

(c) No quadruple intersections satisfying $f_i = f_j = f_k = f_l$ for distinct i, j, k, l exist.

Then the complexity of the minimization diagram of the functions f_i is $O(n)$.

The proof of Theorem 4 is quite involved and is omitted in this version. We note that this theorem is a special case of the upper bound theorem for oriented matroids (see [Ma]), and that it can be extended to arbitrary dimensions. Details can be found in the full version [SS].

3.2. The double intersection case

Here we assume that our collection $\mathbf{F}$ of bivariate functions satisfies the following additional properties (instead of properties (1a) and (1b) in the preceding subsection):

(2a) For each $i \neq j$ the plane curve $f_i(x,y) = f_j(x,y)$ is simple and connected, and is either a closed curve, or is unbounded in both directions, so that it always partitions the x,y plane into two disjoint regions.

(2b) For each triple i, j, k of distinct indices, the equation $f_i(x,y) = f_j(x,y) = f_k(x,y)$ has at most two roots.

Theorem 5: Under the above assumptions, the complexity of the minimization diagram of $\mathbf{F}$ is $\kappa(\mathbf{F}) = O(n^2)$. Moreover there exist collections $\mathbf{F}$ satisfying conditions (2a) and (2b), for which $\kappa(\mathbf{F}) = \Omega(n^2)$.

Sketch of Proof: For each $i=1, \ldots, n$ let σ_i denote the graph of $z = f_i(x,y)$. Fix such a σ_i, and for each $j \neq i$ let γ_{ij} be the vertical projection of $\sigma_i \cap \sigma_j$ onto the xy-plane. By our assumptions each γ_{ij} is either a simple closed Jordan curve, or is a simple open curve both of whose ends approach infinity, and furthermore each pair γ_{ij}, γ_{ik} of these curves intersect in at most two points. For each $j \neq i$ define $K(\gamma_{ij})$ to be the portion of the plane over which $f_j < f_i$. We now adapt and apply the results in [KLPS]. It was shown there that for a collection of $n \geq 3$ Jordan curves in the plane, so that any two of them intersect in at most two points, the number of such intersections along the boundary of the union of their interiors is at most $6n-12$ (a bound that is tight in the worst case). By some rather straightforward technical modifications, we can obtain the following generalization of the result in [KLPS] (see [SS] for a proof).

Lemma 6: Let $\gamma_1, \ldots, \gamma_n$, be a collection of $n \geq 3$ curves in the plane, each of which is either a closed Jordan curve or a simple open curve both of whose ends approach infinity. Suppose that any two of the γ_i have at most two intersections, all of which are distinct, and that all such intersections are transversal. For each γ_i, let $K(\gamma_i)$ be one of the two regions into which γ_i divides the plane. Then the total number of intersections between the curves γ_i which lie on the boundary of $\bigcup_i K(\gamma_i)$ is at most $6n-6$.

This lemma thus implies that the number of intersections of the curves γ_{ij} which lie on the boundary of $\bigcup_{j \neq i} K(\gamma_{ij})$ is at most $6(n-1)-6 = 6n-12$. But these intersection points stand in a 1-1 correspondence with the points of triple intersection of the functions f_k which lie on the intersection of σ_i with the graph of the minimum M.

Repeating the above analysis for each σ_i and observing that each triple intersection point on M will be counted by this process three times, it follows that the number of such corners is at most $\frac{1}{3}n(6n-12) \le 2n^2$, which completes the proof of the first part of Theorem 5.

It is also easy to give examples of collections $\mathbf{F}$ of n functions for which $\kappa(\mathbf{F})=\Omega(n^2)$. For example, one can take $\mathbf{F} = \{f_1, \ldots, f_n, g_1, \ldots, g_n\}$ where

$$f_i(x,y) = (x-i)^2$$

$$g_i(x,y) = a_i y + b_i$$

where the a_i, b_i are chosen so that each g_i appears along the lower envelope of the functions g_k, and so that each intersection of two functions g_i, g_j that lies on the lower envelope has z-coordinate between 0 and 1/4. It is easy to see that $\mathbf{F}$ satisfies the conditions assumed in this subsection, and that $\kappa(\mathbf{F})=\Omega(n^2)$. This completes the proof of Theorem 5. □

3.3. Functions with favorable cross sections.

Again, let $\mathbf{F} = \{f_1(x,y), \ldots, f_n(x,y)\}$ be a collection of n bivariate functions satisfying our basic assumptions, plus the following special properties:

(3a) For each $i \ne j$ and each x_0, the equation $f_i(x_0,y) = f_j(x_0,y)$ has at most two roots $r_{ij}^-(x_0) \le r_{ij}^+(x_0)$.

(3b) Call any point at which any r_{ij}^+ or r_{ij}^- ceases to be defined or has a discontinuity a *singular* point of r_{ij}^+ (resp. of r_{ij}^-). We assume that for each i,j the functions r_{ij}^+, r_{ij}^- have at most t singular points.

(3c) We assume that no four distinct functions f_i, f_j, f_k, f_l become identical at any point, and that for each triple i,j,k of distinct indices, the equations $f_i(x,y) = f_j(x,y) = f_k(x,y)$ have at most s roots.

In what follows we take both t and s to be fixed and independent of n.

Theorem 7: Under the above assumptions, $\kappa(\mathbf{F})$ has the bound $O(n\lambda_{s+2}(n))$, where the constant of proportionality depends on s and t.

Proof: For each given x_0, a pair of functions $f_i(x_0,y)$, $f_j(x_0,y)$ are said to stand in a *definite relation* over a finite or infinite interval of y if one of the relationships $f_i > f_j$ or $f_j > f_i$ holds throughout this interval. Plainly f_i and f_j stand in definite relation below $r_{ij}^-(x_0)$, above $r_{ij}^+(x_0)$, and also between these two roots, if both exist, for every x_0.

We will write $r_{ij}^-(x_0)$ as $\phi_{ij}^-(x_0)$ if $f_i > f_j$ below $r_{ij}^-(x_0)$, otherwise we will write $r_{ij}^-(x_0)$ as $\phi_{ji}^-(x_0)$; similarly, we will write $r_{ij}^+(x_0)$ as $\phi_{ij}^+(x_0)$ if $f_i > f_j$ above $r_{ij}^+(x_0)$, and otherwise we will write $r_{ij}^+(x_0)$ as $\phi_{ji}^+(x_0)$.

Write γ_{ij} for the solution set of $f_i(x,y) = f_j(x,y)$. Let $p = [x_0, y_0]$ be a triple intersection point at which $f_i(p) = f_j(p) = f_k(p) = \min_l f_l(p)$. Then p lies either on the graph of r_{ij}^+ or that of r_{ij}^-, so $f_i(x_0,y)$ and $f_j(x_0,y)$ stand in definite relation either for $y > y_0$ or for $y < y_0$. The same remark applies to the pairs f_i, f_k and f_j, f_k, giving us three definite relations, of which at least two must apply on the same side (above or below) of y_0. Hence two cases arise:

Case A: All three of the ordering relations just alluded to apply on the same side of y_0. In this case, say we have $f_i > f_j$, $f_i > f_k$ on, say $y > y_0$. Then plainly $\phi_{ij}^+(x_0) = \phi_{ik}^+(x_0)$, and no other function ϕ_{il}^+ can pass below p, since then we would have $f_i(p) > \min_m f_m(p)$. Thus

$$\phi_{ij}^+(x_0) = \phi_{ik}^+(x_0) = \min_l \phi_{il}^+(x_0).$$

If $f_i > f_j$ and $f_i > f_k$ for $y < y_0$, we can prove in exactly the same way that

$$\phi_{ij}^-(x_0) = \phi_{ik}^-(x_0) = \max_l \phi_{il}^-(x_0).$$

It now follows from [At] that, for any fixed i, the number of such triple intersections is $O(\lambda_{s+2}(n))$, because these intersections correspond to "breakpoints" along either the "top i-envelope" $\psi_i^+(x) = \min_{j \ne i} \phi_{ij}^+(x)$, or the "bottom i-envelope" $\psi_i^-(x) = \max_{j \ne i} \phi_{ij}^-(x)$, and assumption (3c) is easily seen to imply that any pair of functions ϕ_{ij}^+, ϕ_{ik}^+, or ϕ_{ij}^-, ϕ_{ik}^-, intersect in at most s points. Summing over all n possible values of i, it follows immediately that the total number of this type of triple intersections is $O(n\lambda_{s+2}(n))$.

Case B: Two definite relationships between f_i, f_j and f_k are available on one side of y_0, and one other relationship is available on the other side.

It is easy to check, by case analysis, that the only combination of relationships which cannot be reduced to case A above, is that in which $f_i < f_j$, $f_i < f_k$ above y_0, and (say) $f_j < f_k$ below y_0 (or vice versa).

Since $f_k > f_i$ above y_0 and $f_k(p) = \min_l f_l(p)$, it follows as above that no other function ϕ_{kl}^+ can have a graph which passes below p, so that $\phi_{ki}^+(x_0) = \min_l \phi_{kl}^+(x_0)$, i.e. p lies on the graph of the top k-envelope $\psi_k^+(x) = \min_l \phi_{kl}^+(x)$. Similarly we find that p lies on the graph of $\psi_k^-(x) = \max_l \phi_{kl}^-(x)$.

We can then bound the number of these intersections by dividing the real axis into $O(\lambda_{s+2}(n))$ subintervals by all breakpoints along ψ_k^+, ψ_k^- and all singular points of all the functions ϕ_{kl}^+, ϕ_{kl}^-. It is easily checked that in any such subinterval I, ψ_k^+ and ψ_k^- intersect in at most s points. Thus, for each k, there are at most $O(s\lambda_{s+2}(n))$ new triple intersections, and hence $O(n\lambda_{s+2}(n))$ such intersections in total.

Theorem 7 is immediate from this. □

The above analysis can easily be converted into an algorithm for calculating M, which runs in time $O(n\lambda_{s+2}(n)\log n)$, assuming that each operation involving a specific pair or triple of the functions f_i

only requires constant time. For lack of space, we omit details of this algorithm in this abstract, and refer the reader to [SS].

3.4. The lower envelope of piecewise linear functions

As noted earlier, it suffices to consider the case in which each f_i is a partially defined function whose graph is an arbitrary (non-vertical) triangle in 3-space. Let $\mathbf{F} = \{f_1, \ldots, f_n\}$ be a collection of n such triangles in general position in 3-space, and let $f_1^*, \ldots, f_n^*$ denote their vertical projections onto the $x-y$ plane. The assumption that the triangles are in general position is made for exposition sake, and can be removed. Let $\psi(n)$ denote the maximum number of facets (two-dimensional faces) in the minimization diagram M^* associated with the lower envelope M of any collection $\mathbf{F}$ of n triangles in 3-space.

Theorem 8: (a) The number of vertices, edges, and facets in the lower envelope of any collection of n (non-vertical) triangles in 3-space is at most $O(n\lambda_3(n)) = O(n^2\alpha(n))$, and this bound is tight in the worst case.

(b) The total number of faces (of all dimensions) in the lower envelope of any collection of n d-simplices in $\mathbf{R}^d$ is at most $O(n^{d-1}\alpha(n))$, and this bound is tight in the worst case.

Sketch of Proof: We will concentrate on proving part (a) of the theorem, and then will add a brief comment on the proof of (b). Let $\mathbf{F} = \{f_1, \ldots, f_n\}$ be a collection of triangles in 3-space, and consider their $x-y$ projections $f_1^*, \ldots, f_n^*$. Partition this collection into two disjoint subcollections $\mathbf{F}_1$, $\mathbf{F}_2$, each consisting of $n/2$ triangles. Let M_1 and M_2 denote the lower envelopes of the triangles in $\mathbf{F}_1$ and $\mathbf{F}_2$, respectively. The number of facets of M_1 and of M_2 are both at most $\psi(n/2)$ by definition.

However, the complexity of the overall lower envelope M can in general be larger than the sum of the complexities of the "subenvelopes" M_1 and M_2. The reason is that a facet R of one of the subenvelopes M_i can be split into several facets in M due to the addition of the other envelope. To overcome this difficulty, consider one of these subenvelopes, say M_1, and superimpose the $3n$ lines containing the edges of all the projections $f_1^*, \ldots, f_n^*$ on the map M_1^* to produce a refined planar map $\overline{M}_1$ (see Fig. 2).

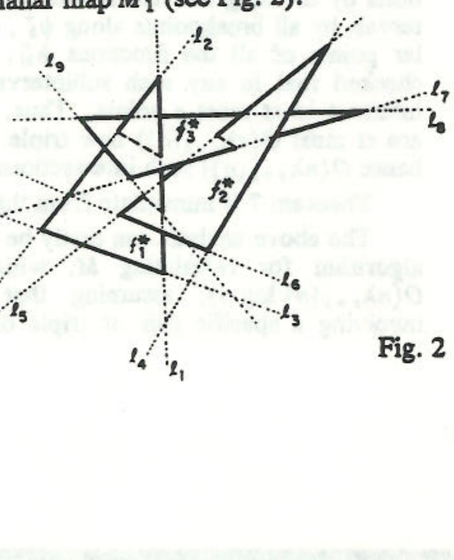

Fig. 2

Lemma 9: Let R be a region of $\overline{M}_1$ which is contained in the projection of some facet of M_1 which is part of some triangle f. Then the portion of R over which f is part of the overall envelope M is connected.

Proof: Let R' be that portion of R. The $3n$ added lines partition the $x-y$ plane into a collection of openly disjoint convex polygonal "base cells", so that no edge of any of the triangles in $\mathbf{F}$ projects to the interior of any of these cells. Let Q be the base cell containing R, and let $\mathbf{F}_Q$ denote the subcollection of all triangles f_j whose projections f_j^* contain Q. Note that, when restricted over Q, the upper envelope M is the same as the upper envelope of the planes containing the triangles in $\mathbf{F}_Q$. In particular, the portion R' over which f_i attains M is convex and thus connected. □

Lemma 9 implies that the sum of the number of facets of $\overline{M}_1$ and of the corresponding refined map $\overline{M}_2$ is an upper bound for the number of facets of M. Our key observation is

Lemma 10: The number $\overline{t}$ of regions in $\overline{M}_1$ is at most the number t^* of regions of M_1^* plus $O(n^2\alpha(n))$.

Proof: Let R be a region of M_1^* which is split into k_R subregions by the addition of the lines $l_1, \ldots, l_{3n}$ containing the edges of the projections $f_1^*, \ldots, f_n^*$ of all triangles in $\mathbf{F}$. Suppose that R is the projection of a connected portion of some triangle $f \in \mathbf{F}_1$ which appears on the lower envelope M_1. For each $1 \le i \le 3n$ let $p_i(R)$ denote the number of connected portions of $R \cap l_i$, and let $q(R)$ denote the number of intersection points of the lines l_i inside R. It is then easily checked that $k_R \le 1 + q(R) + \sum_i p_i(R)$. (This is best seen by adding the lines l_i one at a time.) Hence, if we sum these inequalities over all facets R of M_1^*, we obtain

$$\overline{t} \le t^* + \sum_R q(R) + \sum_{i,R} p_i(R).$$

But clearly $\sum_R q(R) = O(n^2)$. As to the other sum, note that for each i the sum $\sum_R p_i(R)$ is just the complexity of the lower envelope M_1 restricted over the line l_i. But since each of the $n/2$ triangles in $\mathbf{F}_1$, when restricted over l_i, becomes a straight line segment, it follows from the standard one-dimensional Davenport-Schinzel theory (cf. [HS]) that

$$\sum_R p_i(R) \le \lambda_3(n/2) \le \lambda_3(n) = O(n\alpha(n)).$$

Thus, summing over all lines l_i, we obtain

$$\overline{t} \le t^* + O(n^2) + O(n^2\alpha(n)) = t^* + O(n^2\alpha(n)).$$

□

Since a similar inequality applies to the map $\overline{M}_2$, we can now obtain the desired recurrence formula for ψ:

$$\psi(n) = 2\psi(n/2) + O(n^2\alpha(n)).$$

The solution of this formula is readily seen to be $\psi(n) = O(n^2\alpha(n))$.

We also derive similar bounds for the number of vertices and edges of M, and show that similar bounds hold also for collections $\mathbf{F}$ not in general position. This completes the first part of the proof of our theorem.

The lower bound assertion in the theorem follows from the recent result of [WS] that constructs a collection of n line segments in the, say $x-z$ plane, whose lower envelope consists of $\Omega(n\alpha(n))$ subsegments. By taking the Cartesian product of each of these segments with a large interval on the y axis, we obtain a collection of n rectangles, to which we add n descending sharp and narrow wedges whose lower edges are all parallel to the $x-z$ plane, and are all at the same depth, so that they cut through the entire lower envelope of the first n rectangles from above. This yields a collection of $O(n)$ triangles whose lower envelope has complexity $\Omega(n^2\alpha(n))$.

Let us next consider part (b) of Theorem 8. Notice that in order to prove the upper bound of part (a) of the theorem, we used a divide-and-conquer argument (as opposed to a divide-and-conquer algorithm) to get an upper bound on the number of facets. The upper bound on the number of vertices and edges follows by Euler's relation. In $d \geq 4$ dimensions, Euler's relation implies the desired upper bound of $O(n^{d-1}\alpha(n))$ on the number of vertices and edges, provided the same bound holds for all higher-dimensional faces of the envelope. A straightforward generalization of the above argument can establish this bound only for facets, that is, $(d-1)$-faces. To fill the gap we take advantage of the fact that the homogeneous solution of the recurrence relation for $\psi(n)$ is much smaller than the additive term that dominates the solution. This allows us to obtain the envelope by merging several subenvelopes of non-disjoint sets of simplices. Every k-face of M appears also in one of the subenvelopes, if every $(d-k)$-tuple of simplices is fully contained in at least one set. This idea eventually leads to the desired bounds, thus completing the proof of Theorem 8. □

Calculating the Lower Envelope

The preceding analysis easily yields an algorithm for calculating the lower envelope of a collection $\mathbf{F}$ of n triangles in 3-space. Specifically, we partition $\mathbf{F}$ into two subcollections $\mathbf{F}_1$, $\mathbf{F}_2$ of equal size, and calculate recursively the lower envelope of each subcollection. Then we take the $3n$ lines $l_1, \ldots, l_{3n}$ in the $x-y$ plane containing the projections of the sides of the triangles in $\mathbf{F}$, and calculate the arrangement they form in that plane (in time $O(n^2)$ [EOS]). For each cell Q of the arrangement we find all triangles which appear above Q in the lower envelope of either $\mathbf{F}_1$ or of $\mathbf{F}_2$, and then calculate the lower envelope of the planes containing these triangles, in time $O(k_Q \log k_Q)$, where k_Q is the total number of such triangles (see [PrM], [PrS], or [E]). The lower envelope M of $\mathbf{F}$ is then equal to the union of all these subenvelopes, except that some faces of M may have been split into several subfaces by the lines l_i. For simplicity of the algorithm we leave these faces split, and do not join them together, except for the final envelope of the entire collection, if so desired. This last step is accomplished simply by scanning each line l_i and by joining matching faces into a single facet. Here, two facets *match* if they are separated by a common edge on l_i and if they come from the same triangle. Two edges *match* if they are separated by a common vertex on l_i and if they come from the same triangle edge or from the intersection of the same pair of triangles. Since the preceding analysis implies that $\sum_Q k_Q = O(n^2\alpha(n))$, it follows that the running time of the merge step of this algorithm, and also of the entire algorithm, is at most $O(n^2\alpha(n) \log n)$.

We can improve the time complexity of the algorithm to $O(n^2\alpha(n))$, if we partition $\mathbf{F}$ planfully into $\mathbf{F}_1$ and $\mathbf{F}_2$, rather than arbitrarily. For each triangle define its *slope* as the slope of the line of intersection between the $x-z$ plane and the plane that contains the triangle. When we partition $\mathbf{F}$ we do it such that $\mathbf{F}_1$ contains the triangles with the $n/2$ smallest slopes, and $\mathbf{F}_2$ contains the triangles with the $n/2$ largest slopes. The envelope restricted to the region above a cell can now be computed in linear time from the corresponding pieces of the two subenvelopes. This method is explained in dual space in [PrS] and [E].

3.5. The boundary of a region enclosed by convex plates

Let $\mathbf{S} = \{S_1, \ldots, S_n\}$ be a collection of n $(d-1)$-dimensional compact convex sets (plates) in $\mathbb{R}^d$. If we delete from $\mathbb{R}^d$ all points belonging to at least one of these plates, then $\mathbb{R}^d$ may split up into a number of connected components. Let C denote such a component.

Theorem 11: For every $d \geq 3$, there exists a constant $\epsilon(d)>0$ such that, given any collection $\mathbf{S} = \{S_1, \ldots, S_n\}$ of $(d-1)$-dimensional convex plates arranged in $\mathbb{R}^d$, the combinatorial complexity of the boundary of any single connected component C of $\mathbb{R}^d - \bigcup_{i=1}^{n} S_i$ is at most $O(n^{d-\epsilon(d)})$.

Assume without loss of generality that the plates are in general position. Then any vertex of the given component C belongs to exactly d plates, and using the fact that any d plates have at most one point in common, we obtain that C has at most $\binom{n}{d} = O(n^d)$ vertices.

On the other hand, it is easy to see that the total combinatorial complexity of C (i.e. the number of all i-dimensional faces over all $0 \leq i \leq d$) is proportional to the number of its vertices.

Hence, it is sufficient to prove

Theorem 11': For every $d \geq 3$, there exists a constant $\epsilon(d)>0$ such that, given any collection $\mathbf{S} = \{S_1, \ldots, S_n\}$ of $(d-1)$-dimensional convex plates arranged in $\mathbb{R}^d$ in general positions, the number of points belonging to d members of S and lying on the boundary of a given component C of $\mathbb{R}^d - \cup S$ is at most $O(n^{d-\epsilon(d)})$. ($\cup S$ is the shorthand for $\bigcup_{i=1}^{n} S_i$.)

We need some preparation. A *d-uniform hypergraph* is a set system whose members (the so-called *hyperedges)* are d-element sets.

Definition: Let $H = H(S)$ be a d-uniform hypergraph whose vertex set is S and whose hyperedges are those d-tuples $\{S^{(1)}, \ldots, S^{(d)}\} \subset S$ for which $\bigcap_{i=1}^{d} S^{(i)}$ lies on the boundary of the given component C of $\mathbb{R}^d - \cup S$.

Let $K^{(r)}(m_1,m_2, \ldots ,m_r)$ denote an r-uniform hypergraph with $m_1+m_2+\cdots+m_r$ vertices, whose vertex set is $V_1 \cup V_2 \cup \cdots \cup V_r$, $|V_i| = m_i$ $(1 \le i \le r)$ and which consists of all r-tuples containing exactly one element from each V_i. If $m_1 = m_2 = \cdots m_r = m$, then we will write $K^{(r)}(m)$ for $K^{(r)}(m,m, \ldots ,m)$.

Our proof is based on showing that for some sufficiently large constant $m = m(d)$, the hypergraph $H(S)$ does not contain a subhypergraph isomorphic to $K^{(r)}(m)$. Then we can apply the following well-known combinatorial result of Erdös [Er].

Theorem 12 (Erdös): Let H be an r-uniform hypergraph on n vertices containing no subhypergraph isomorphic to $K^{(r)}(m)$. Then $|H| \le n^{r-(1/m)^{r-1}}$.

This theorem will establish our claim. For lack of space we will sketch our approach only for $d=3$. In this case we will need the following slightly stronger version of Erdös's theorem:

Theorem 13: Given any natural numbers $r,m \ge 2$, $M \ge m$, there exists a constant $C(r,m,M)=C$ such that the number of hyperedges of any r-uniform hypergraph on n vertices, which does not contain a subhypergraph isomorphic to $K^{(r)}(m, \ldots ,m,M)$, is at most $Cn^{r-(1/m)^{r-1}}$. □

The proof of our bound is somewhat technical, and some details are omitted in this version.

Lemma 14. Let $\Sigma = \{\sigma_1,\sigma_2,\sigma_3\}$ and $\Sigma' = \{\sigma_1',\sigma_2', \ldots ,\sigma_t'\}$ be two systems of straight line segments in $\mathbb{R}^2$ such that

(i) $\sigma_i \cap \sigma_j' \ne \emptyset$ for every $1 \le i \le 3$, $1 \le j \le t$; and

(ii) all intersection points $p_{ij} = \sigma_i \cap \sigma_j'$ are on the boundary of the same connected component of $\mathbb{R}^2 - \cup\Sigma - \cup\Sigma'$.

Then $t \le 6$ and this bound cannot be improved.

Proof: Elementary and omitted here. □

Definition: Given any system Π of (2-dimensional) planes in $\mathbb{R}^3$, and two planes P_1 and P_2 in general positions, we say that P_1 and P_2 are *equivalent with respect to* Π if there is a single rotation or translation which takes P_1 to P_2 so that during the motion the plane

(i) never passes through any point belonging to three members of Π;

(ii) is never parallel to the intersection line of any two members of Π.

Lemma 15: The preceding definition yields an equivalence relation on the family of all planes which are in general position with respect to Π, and the number of equivalence classes is at most $|\Pi|^9$.

Proof: Omitted in this version. □

Let $H(S)$ be our 3-uniform hypergraph as defined above for some given (say the unbounded) component of $\mathbb{R}^3 - \cup S$. We will show that $H(S)$ cannot contain a subhypergraph isomorphic to $K^{(3)}(7,7,M)$ for some integer M independent of n. In fact, we can prove the following somewhat stronger result.

Lemma 16: $H(S)$ does not contain a subhypergraph isomorphic to $K^{(3)}(3,7,2\cdot 10^9+1)$.

Proof: Assume, in order to obtain a contradiction, that there are three subsystems $T, T', T'' \subset S$ such that

(i) $|T|=3$, $|T'|=7$, $|T''| = 2\cdot 10^9+1$, and

(ii) $S \cap S' \cap S'' \ne \emptyset$ and lies on the boundary of the unbounded component of $\mathbb{R}^3 - \cup S$ for every $S \in T$, $S' \in T'$, $S'' \in T''$.

Let P, P' and P'' denote the systems of planes containing the plates belonging to T, T' and T'', respectively. Applying Lemma 15 with $\Pi = P \cup P'$, we obtain that there exist 3 plates $S_1'', S_2'', S_3'' \in T''$ such that the corresponding planes P_1'', P_2'', P_3'' are pairwise equivalent with respect to $P \cup P'$.

Along each plane P_i'', its intersections with the plates in T, T' form a pattern of segments which, by Lemma 14, must contain an intersection point completely enclosed by a simple closed polygon p, all of whose sides are portions of some intersection segments with the plates in T, T'. This, and the equivalence of the planes P_i'', are easily seen to imply the existence of a triple intersection between plates in T, T', T'' respectively, which is enclosed in a bounded component of $\mathbb{R}^3 - \cup S$, a contradiction (see [PS] for more details). □

Applying Theorem 13, we thus conclude

Theorem 17: Given any collection $S = \{S_1, \ldots ,S_n\}$ of 2-dimensional convex plates scattered in $\mathbb{R}^3$, the combinatorial complexity of the boundary of any given connected component of $\mathbb{R}^3 - \cup S$ is at most $O(n^{3-\frac{1}{49}})$.

Acknowledgement

The authors would like to thank Leonidas Guibas for valuable and stimulating discussions of some of the problems presented here. In particular, the algorithm for calculating the lower envelope of piecewise linear functions has benefited from certain improvements suggested by Guibas. The first author would also like to thank Emo Welzl for discussions on stabbing line segments in three dimensions and on Voronoi diagrams of point clusters.

M. Sharir and P. Shor, Improved upper and ds for the length of general Davenport- uences, in preparation.

, Dynamic computational geometry, *Proc. 24th* *'oundations of Computer Science*, 1983, pp. 92-

in *Comp. Maths. with Appls.* 11 (1985) pp.)

hammer, Improved algorithms for discs and ng power diagrams, Tech. Rept., Technical y of Graz.

an and M. Sharir, On shortest paths between two polyhedra, Tech. Rept. 180, Comp. Science Courant Institute, New York University, Sept. to appear in *J. ACM*.)

e and M. Sharir, Visibility of a polyhedral surface poi Tech. Rept. 266, Computer Science Dept., nt Ins te, New York University, December 1986. pear in *J. Symbolic Computation*.)

avenport, A combinatorial problem connected with rential equations, II, *Acta Arithmetica* 17 (1971), pp. 372.

Davenport and A. Schinzel, A combinatorial problem nected with differential equations, *Amer. J. Math.* 87 65), pp. 684-694.

Devai, Quadratic bounds for hidden line elimination, *oc. 2nd ACM Symp. on Computational Geometry*, 1986, . 269-275.

. Driscoll, N. Sarnak, D. Sleator and R.E. Tarjan, Making data structures persistent, *Proc. 18th Symp. on Theory of Computing*, 1986, pp. 109-121.

H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg, Germany, 1987.

WW] H. Edelsbrunner, H.A. Maurer, F.P. Preparata, A.L. Rosenberg, E. Welzl and D. Wood, Stabbing line segments, *BIT* 22 (1982) pp. 274-281.

H. Edelsbrunner, J. O'Rourke and R. Seidel, Constructin arrangements of lines and hyperplanes with applicati *Proc. 24th IEEE Symp. on Foundations of Computer Science*, 1983, pp. 83-91.

] H. Edelsbrunner and R. Seidel, Voronoi diagrams and arrangements, *Discrete Comput. Geom.* 1 (1986) pp. 25-44.

] H. Edelsbrunner and M. Sharir, The upper envelope of piecewise linear functions: Algorithms and applications, in preparation.

r] P. Erdős, On extremal problems of graphs and generalized graphs, *Israel J. Math* 2 (1964), 183-190.

HS] S. Hart and M. Sharir, Nonlinearity of Davenport-Schinzel sequences and of generalized path compression schemes, *Combinatorica* 6 (1986), pp. 151-177.

[KLPS] K. Kedem, R. Livne, J. Pach and M. Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, *Discrete and Computational Geometry* 1 (1986), pp. 59-71.

[Ma] A. Mandel, Topology of oriented matroids, Ph.D. Dissertation, Univ. of Waterloo, 1981.

[MK] M. McKenna, Worst-case optimal hidden surface removal, Tech. Rept. JHU/EECS-86/05, Johns Hopkins University, 1986.

[OSY] O'Dunlaing, C., Sharir, M. and Yap, C., Generalized Voronoi diagrams for a ladder: II. Efficient construction of the diagram, *Algorithmica* 2 (1987) pp. 27-59.

[PrM] F.P. Preparata and D.E. Muller, Finding the intersection of n half spaces in time $O(n \log n)$, *Theoretical Computer Science* 8 (1979) pp. 44-55.

[PS] J. Pach and M. Sharir, The upper envelope of piecewise linear functions and the boundary of a region enclosed by convex plates: Combinatorial analysis, Tech. Rept. 279, Comp. Science Dept., New York University, March 1987.

[PrS] F.P. Preparata and M.I. Shamos, *Computational Geometry*, Springer Verlag, New York, 1985.

[PSS] R. Pollack, M. Sharir and S. Sifrony, Separating two simple polygons by a sequence of translations, Tech. Rept. 215, Computer Science Dept., Courant Institute, New York University, April 1986. (to appear in *Discrete Comput. Geom.*)

[SS] J.T. Schwartz and M. Sharir, On the two-dimensional Davenport-Schinzel problem, in preparation.

[Sh1] M. Sharir, Almost linear upper bounds on the length of general Davenport-Schinzel sequences, Tech. Report 29/85, The Eskenasy Institute of Computer Sciences, Tel Aviv University, February 1985. (to appear in *Combinatorica*.)

[Sh2] M. Sharir, Improved lower bounds on the length of Davenport Schinzel sequences, Tech. Rept. 204, Comp. Sci. Dept., Courant Institute, New York University, February 1986. (to appear in *Combinatorica*.)

[SL] M. Sharir and R. Livne, On minima of functions, intersection patterns of curves, and Davenport-Schinzel Sequences, *Proc. 26th IEEE Symp. on Foundations of Computer Science* 1985, pp. 312-320.

[Sz] E. Szemeredi, On a problem by Davenport and Schinzel, *Acta Arithmetica* 25 (1974), pp. 213-224.

[WS] A. Wiernik and M. Sharir, Planar realization of nonlinear Davenport Schinzel sequences by segments, Tech. Rept. 224, Comp. Science Dept., Courant Institute, New York University, June 1986. (to appear in *Discrete Comput. Geom.*)