Arrangements in Higher Dimensions: Voronoi Diagrams, Motion Planning, and Other Applications*

Micha Sharir**

Abstract. We review recent progress in the study of arrangements of surfaces in higher dimensions. This progress involves new and nearly tight bounds on the complexity of lower envelopes, single cells, zones, and other substructures in such arrangements, and the design of efficient algorithms (near optimal in the worst case) for constructing and manipulating these structures. We then present applications of the new results to motion planning, Voronoi diagrams, visibility, and geometric optimization.

The combinatorial, algebraic, and topological analysis of arrangements of surfaces in higher dimensions has become one of the most active areas of research in computational geometry during the past 5 years. This is partly due to the fact that many geometric problems in diverse areas can be reduced to questions involving such arrangements. A typical example is the following general *motion planning problem*. Assume that we have a robot system B with d degrees of freedom, i.e., we can represent each placement of B as a point in d-space. Suppose that the workspace of B is cluttered with obstacles, whose shapes and locations are known. For each combination of a geometric feature (vertex, edge, face) of an obstacle and a similar feature of B, define their *contact surface* as the set of all points in d-space that represent a placement of B in which contact is made between these specific features. Let Z be a point corresponding to a given initial *free* placement of B, in which it does not intersect any obstacle. Then the set of all free placements of B that can be reached from Z via a collision-free continuous motion will obviously correspond to the cell containing Z in the arrangement of the contact surfaces. Thus, the robot motion planning problem leads more or less directly to the problem of computing a single cell in an arrangement of surfaces in higher dimensions. The *combinatorial complexity* of this cell, i.e., the total number of lower-dimensional faces appearing on its boundary serves as a trivial lower bound for the running time of the motion planning problem (assuming the

* Work on this paper has been supported by NSF Grants CCR-91-22103 and CCR-93-11127, by a Max-Planck Research Award, and by grants from the U.S.-Israeli Binational Science Foundation, the Israel Science Fund administered by the Israeli Academy of Sciences, and the G.I.F., the German-Israeli Foundation for Scientific Research and Development.

** School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel, and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA

entire cell has to be output). It turns out that in most instances this bound can be almost matched by suitable algorithms.

Other applications that call for combinatorial analysis of arrangements involve geometric algorithms for constructing arrangements, which are based on randomized or ε-net techniques, and whose running times are usually directly influenced by the combinatorial complexity of the relevant geometric substructures of the arrangements that they manipulate; this will be explained in more detail below. We will also describe below more applications of higher-dimensional arrangements to problems in visibility, in geometric optimization, and involving generalized Voronoi diagrams. For some basic terminology related to arrangements, the reader is referred to [30, 39, 45, 65].

This survey describes many recent advances in the study of combinatorial, topological, and algorithmic problems involving arrangements of algebraic surfaces in higher dimensions. In these studies there are three main relevant parameters: the number n of surfaces, their maximum algebraic degree b, and the dimension d. In the approach taken here, we are mainly interested in a 'combinatorial' approach, in which we want to calibrate the dependence of the complexity of the various structures and algorithms on the number n of surfaces, assuming that the maximum degree, as well as any other factor that does not depend on n, is constant. In this way, all issues related to the algebraic complexity of the problem are 'swept under the rug'. These issues should be (and indeed have been) picked up in the complementary 'algebraic' mode of research, where the dependence on the maximum degree b is more relevant; see [48, 62] for studies of this kind.

We should emphasize that, although quite a few of the problems reviewed here are combinatorial in nature, most of them are motivated by algorithmic applications. As mentioned above, it is an interesting feature of the area that the complexity of efficient algorithms for constructing various substructures in arrangements depends mainly on the combinatorial complexity of these structures. This is, of course, always true in terms of lower bounds (as the algorithm must at least output the desired structure), but there is also a strong influence of the combinatorial complexity on the running time of the algorithms. A typical example involves algorithms that are based on *vertical decompositions* in arrangements—see below for details.

During the past three years, significant progress has been made on the problem of bounding the complexity of the *lower envelope* (pointwise minimum) of a collection of multivariate functions. This problem has been open since 1986, when it was shown in [46] that the combinatorial complexity of the lower envelope of n univariate continuous functions, each pair of which intersect in at most s points, is at most $\lambda_s(n)$, the maximum length of an (n, s)-*Davenport–Schinzel sequence*. This bound is slightly super-linear in n, for any fixed s (for example, it is $\Theta(n\alpha(n))$ for $s = 3$, where $\alpha(n)$ is the extremely slowly growing inverse of Ackermann's function [46]; see also [8, 65]). Since the complexity of the arrangement of such a collection of functions can be $\Theta(n^2)$ in the worst case, this result shows that the complexity of the lower envelope is smaller than the

overall complexity of the arrangement by nearly a factor of n.

It was then conjectured that a similar phenomenon occurs in higher dimensions. That is, the combinatorial complexity of the lower envelope of a collection $\mathcal{F}$ of n 'well-behaved' d-variate functions should be close to $O(n^d)$ (as opposed to $\Theta(n^{d+1})$, which can be the complexity of the entire arrangement of the function graphs). More precisely, according to a stronger version of this conjecture, this quantity should be at most $O(n^{d-1}\lambda_s(n))$, for some constant s depending on the shape of the functions in $\mathcal{F}$. These conjectures have been confirmed only in some special cases, including the case in which the graphs of the functions are d-simplices in $\mathbb{R}^{d+1}$, where a tight worst-case bound, $\Theta(n^d\alpha(n))$, was established in [31, 59]. (The case $d = 1$, involving n segments in the plane, where the bound is $\Theta(n\alpha(n))$, had been analyzed earlier, in [46, 69].) There are also some even more special cases, like the case of hyperplanes, where the maximum complexity of their lower envelope is known to be $\Theta(n^{\lfloor (d+1)/2 \rfloor})$, by the so-called *Upper Bound Theorem* for convex polytopes [57]. The case of balls also admits a much better bound, using a standard lifting transformation (see [30]). However, the general problem remained open.

Last year this problem was almost completely settled in [43] and [64]: Let $\mathcal{F}$ be a collection of (possibly partially-defined) d-variate functions, such that all functions in $\mathcal{F}$ are algebraic of constant maximum degree and, in case of partial functions, the domain of definition of each function is a semi-algebraic set defined by a constant number of polynomial equalities and inequalities of constant maximum degree. We refer to such a region as having *constant description complexity*. It was shown that, for any $\varepsilon > 0$, the combinatorial complexity of the lower envelope of $\mathcal{F}$ is $O(n^{d+\varepsilon})$, where the constant of proportionality depends on ε, d, and on the maximum degree of the functions and of the polynomials defining their domains. Thus, apart from a small remaining gap, the above conjecture has been settled in the affirmative.

The proof is based on the probabilistic method developed by Clarkson and Shor [25]. Informally (and not very precisely), one charges each vertex p of the envelope to a block of k 'nearby' vertices that lie along an edge leading from p away from the envelope (here k is some sufficiently large constant parameter). Each of the charged vertices lies at *level* at most k in the arrangement of the function graphs (that is, at most k graphs lie below such a vertex). The Clarkson-Shor technique allows us to bound the number of such nearby vertices by a term equal to $O(k^{d+1})$ times the number of vertices of the lower envelope of a random sample of n/k functions of $\mathcal{F}$. This implies that the number of vertices of the envelope can be bounded by roughly $O(k^d)$ times the number of vertices of the envelope of a random sample of size n/k, which leads to a recurrence, whose solution gives the asserted bounds. (We caution that this description glosses over many technical details, given in the papers cited above, and is given only as an intuitive explanation of the main idea of the proof.)

This result was then followed by several further developments. Lower envelopes (just like single cells) naturally arise in motion planning, scene analysis, Voronoi diagrams, and geometric optimization. The new results can therefore be

applied to obtain improved algorithmic and combinatorial bounds for a variety of problems. Some of these applications have already appeared in the literature (and will be mentioned below), but we believe that many more await to be discovered. The above results have also opened up the door to many significant new research problems, and, in our opinion, it will take at least several years to settle most of them.

Algorithms for Lower Envelopes

Once the combinatorial complexity of lower envelopes of multivariate functions has been (more or less) resolved, the next task is to derive efficient algorithms for computing such lower envelopes. One of the strongest forms of such a computation is as follows. We are given a collection $\mathcal{F}$ of d-variate algebraic functions satisfying the above conditions. We want to compute the lower envelope $E_{\mathcal{F}}$ and store it in some data structure, so that, given a query point $p \in \mathbb{R}^d$, we can efficiently compute the value $E_{\mathcal{F}}(p)$, and the function(s) attaining $E_{\mathcal{F}}$ at p. (Of course, we need to assume here an appropriate model of computation, where various primitive operations on a constant number of functions can be each performed in constant time. There are several different models of this kind, such as the exact arithmetic model in real algebraic geometry; see, e.g., [48].)

This task has recently been accomplished for the case of bivariate functions in several papers [5, 16, 17, 27, 64]. Some of these techniques use randomized algorithms, and their expected running time is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, which is comparable with the maximum complexity of such an envelope. The simplest algorithm is probably the one given in [5]. It is deterministic and uses divide-and-conquer. Its analysis is based on an interesting property of the *overlay* of (the xy-projections of) two lower envelopes of bivariate functions, that the complexity of such an overlay is also $O(n^{2+\varepsilon})$, where n is the total number of functions. (To appreciate this result, observe that, in general, the overlay of two planar maps of complexity N each can have $\Theta(N^2)$ complexity.)

In higher dimensions, the only result known so far is that lower envelopes of trivariate functions satisfying the above properties can be computed, in the above strong sense, in randomized expected time $O(n^{3+\varepsilon})$ [1]. For $d > 3$, it is also shown in [1] that all vertices, edges and 2-faces of the lower envelope of n d-variate functions, as above, can be computed in randomized expected time $O(n^{d+\varepsilon})$. It is still an open problem whether such a lower envelope can be computed within similar time bounds in the above stronger sense, and this problem should certainly be investigated further. Another, more difficult problem is to devise *output-sensitive* algorithms, whose complexity depends on the actual combinatorial complexity of the envelope. It would also be interesting to develop algorithms for certain special classes of functions, where better bounds are known for the complexity of the envelope, e.g., for envelopes of piecewise-linear functions (see below for more details).

Some of the applications of the algorithms produced so far are mentioned below. It can be expected that the proposed extensions of these algorithms will also find other interesting applications.

Single Cells

Lower envelopes are closely related to other substructures in arrangements, notably *single cells* and *zones*. In two dimensions, it was shown in [40] that the complexity of a single face in an arrangement of n arcs, each pair of which intersect in at most s points, is $O(\lambda_{s+2}(n))$, and so is of the same asymptotic order of magnitude as the complexity of the lower envelope of such a collection of arcs. Again, the prevailing conjecture is that the same holds in higher dimensions. That is, the complexity of a single cell in an arrangement of n algebraic surfaces in d-space satisfying the above assumptions is close to $O(n^{d-1})$, or, in a stronger form, this complexity should be $O(n^{d-2}\lambda_s(n))$, for some appropriate constant s. The weaker version of this conjecture for the 3-dimensional case has recently been confirmed in [44]: Let $\mathcal{A}$ be an arrangement of n 2-dimensional surface patches in $\mathbb{R}^3$, all of them algebraic of constant description complexity. It was proved in [44] that, for any $\varepsilon > 0$, the complexity of a single cell in $\mathcal{A}$ is $O(n^{2+\varepsilon})$, where the constant of proportionality depends on ε and on the maximum degree of the surfaces and of their boundaries. The proof is based on an extension of the argument developed for lower envelopes, but one has to add several nontrivial ingredients, in order to handle the more complex topology of a single cell. The analysis in [44] seems to extend to higher dimensions, except for some key steps that seem to require the introduction of new algebraic geometry techniques.

The results of [44] mentioned above easily imply that, for fairly general robot systems with 3 degrees of freedom, the complexity of the space of all free placements of the system, reachable from a given initial placement, is $O(n^{2+\varepsilon})$, a significant improvement over the previous, naive bound $O(n^3)$. The corresponding algorithmic problem, of devising an efficient (near-quadratic) algorithm for computing such a cell, has very recently been solved in [63]. We will say more about this result when we discuss vertical decompositions below. Prior to this result, several other near-quadratic algorithms were proposed for some special classes of surfaces [5, 13, 41, 42]. For example, the paper [42] gives a near-quadratic algorithm for the single cell problem in the special case of arrangements that arise in the motion planning problem for a (nonconvex) polygonal robot moving (translating and rotating) in a planar polygonal region. However, this algorithm exploits the special structure of the surfaces that arise in this case, and does not extend to the general case. The algorithm given in [5] also provides a near-quadratic solution for the case that all the surfaces are graphs of totally-defined continuous algebraic bivariate functions (so that the cell in question is xy-monotone).

In higher dimensions, we mention the special case of a single cell in an arrangement of n $(d-1)$-simplices in $\mathbb{R}^d$. It was shown in [13] that the complexity of such a cell is $O(n^{d-1}\log n)$; a simplified proof was recently given in [67]. This bound is much sharper than the general bound stated above; the best lower bound known for this complexity is $\Omega(n^{d-1}\alpha(n))$, so a small gap between the upper and lower bounds still remains.

Zones

Given an arrangement $\mathcal{A}$ of surfaces in $\mathbb{R}^d$, and another surface σ_0, the *zone* of σ_0 is the collection of all cells of the arrangement $\mathcal{A}$ that σ_0 crosses, and the complexity of the zone is the sum of complexities of all these cells. The 'classical' Zone Theorem [30, 34] asserts that the maximum complexity of the zone of a hyperplane in an arrangement of n hyperplanes in $\mathbb{R}^d$ is $\Theta(n^{d-1})$, where the constant of proportionality depends on d. This has been extended in [10] to the zone of an algebraic or convex surface (of any dimension $p < d$) in an arrangement of hyperplanes. The bound on the complexity of such a zone is $O(n^{\lfloor (d+p)/2 \rfloor} \log^c n)$, and $\Omega(n^{\lfloor (d+p)/2 \rfloor})$ in the worst case, where $c = 1$ when $d-p$ is odd and $c = 0$ when $d-p$ is even. It is not clear whether the logarithmic factor is really needed, or that it is just an artifact of the proof technique.

The result of [44] can easily be extended to obtain a bound of $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, on the complexity of the zone of an algebraic surface σ_0 (of constant description complexity) in an arrangement of n algebraic surfaces in $\mathbb{R}^3$, as above. Intuitively, the proof proceeds by cutting each of the given surfaces along its intersection curve with σ_0, and by shrinking the surface away from that curve, thus leaving a 'tiny' gap there. These modifications transform the zone of σ into a single cell in the arrangement of the new surfaces, and the result of [44] can then be applied. (The same technique has been used earlier in [32], to obtain a near-linear bound on the complexity of the zone of an arc in a 2-dimensional arrangement of arcs.) Once the bound on the complexity of a single cell is extended to higher dimensions, it should lead right away to a similar bound for a zone of a surface. A similar technique implies that the complexity of the zone of an algebraic or convex surface in an arrangement of n $(d-1)$-simplices in $\mathbb{R}^d$ is $O(n^{d-1} \log n)$ [13, 67].

Generalized Voronoi Diagrams

One of the interesting applications of the new lower bounds on the complexity of lower envelopes is to generalized Voronoi diagrams in higher dimensions. Let S be a set of n 'simply-shaped' pairwise-disjoint convex objects in d-space (or in higher dimensions), and let ρ be some metric. The *Voronoi diagram* $\mathrm{Vor}(S)$ of S under the metric ρ is defined, as usual, as the decomposition of d-space into *Voronoi cells* $V(s)$, for $s \in S$, where

$$V(s) = \{x \in \mathbb{R}^d \mid \rho(x,s) \leq \rho(x,s') \text{ for all } s' \in S\}.$$

The problem is to study the combinatorial complexity of $\mathrm{Vor}(S)$, and to devise efficient algorithms for its construction. In the classical case, in which ρ is the Euclidean metric and the objects in S are singletons (points), the maximum possible complexity of $\mathrm{Vor}(S)$ is $\Theta(n^{\lceil d/2 \rceil})$ (see, e.g., [30]). In three dimensions, this bound is $\Theta(n^2)$. It has been a long-standing open problem whether a similar quadratic or near-quadratic bound holds in 3-space for more general objects and metrics. As is well known [33], the Voronoi diagram $\mathrm{Vor}(S)$ is the 'minimization diagram' (projection onto the xyz-hyperplane) of the lower envelope

$\min_{s \in S} \rho(x, s)$ (in 4-space). Under reasonable assumptions on the shape of the objects in S and on the metric ρ, the resulting trivariate functions $\rho(x, s)$ can be assumed to be piecewise-algebraic of constant maximum degree, and the recent results concerning lower envelopes, as reported above, give an upper bound of $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$, for the complexity of $\mathrm{Vor}(S)$. Thus the problem stated above calls for improving this bound by roughly another factor of n. It thus appears to be a considerably more difficult problem than that of lower envelopes, and the only hope of making progress there is to exploit the special structure of the functions $\rho(x, s)$.

Fortunately, some progress on this problem was recently done. It was shown in [23] that the complexity of the Voronoi diagram is $O(n^2 \alpha(n) \log n)$, for the case where the objects of S are lines, and the metric ρ is a convex distance function induced by a convex *polytope* with a constant number of facets. (Note that the L_1 and L_∞ metrics are special cases of such distance functions. Note also that such a distance function is not necessarily a metric, because it will fail to be symmetric if the defining polytope is not centrally symmetric.) The best known lower bound for the complexity of the diagram in this special case is $\Omega(n^2 \alpha(n))$. In another recent paper [18], it is shown that the maximum complexity of the L_1-Voronoi diagram of a set of n points in $\mathbb{R}^3$ is $\Theta(n^2)$. However, no near-quadratic bound is known for point sites and more general polyhedral convex distance functions. We hope that these results will open up this research direction, and lead to many subsequent results. The most intriguing unsolved problem is to obtain a similar bound for a set S of n lines in space but under the Euclidean distance. The proof technique of [23] breaks down in this case. Other, more tractable open problems are to extend these results to sets of more general convex objects (e.g., convex polytopes, or just singleton points) under the same polyhedral convex distance functions.

An interesting special case of these problems involves *dynamic Voronoi diagrams* for moving points in the plane. Let S be a set of n points in the plane, each moving along some line at some fixed velocity. The goal is to bound the number of combinatorial changes of $\mathrm{Vor}(S)$ over time. This dynamic Voronoi diagram can easily be transformed into a 3-dimensional Voronoi diagram, by adding the time t as a third coordinate. The points become lines in 3-space, and the metric is a distance function induced by a horizontal disc (that is, the distance from a point $p(x_0, y_0, t_0)$ to a line ℓ is the Euclidean distance from p to the point of intersection of ℓ with the horizontal plane $t = t_0$). Here too the open problem is to derive a near-quadratic bound on the complexity of the diagram. Cubic or near-cubic bounds are known for this problem, even under more general settings [36, 38, 64], but subcubic bounds are known only in some very special cases [22].

Next, consider the problem of bounding the complexity of generalized Voronoi diagrams in higher dimensions. As mentioned above, when the objects in S are n points in $\mathbb{R}^d$ and the metric is Euclidean, the complexity of $\mathrm{Vor}(S)$ is $O(n^{\lceil d/2 \rceil})$. As d increases, this becomes drastically smaller than the naive $O(n^{d+1})$ bound or the improved bound, $O(n^{d+\varepsilon})$, obtained by viewing the Voronoi diagram as a lower envelope in $\mathbb{R}^{d+1}$. The same bound of $O(n^{\lceil d/2 \rceil})$ has recently been

obtained in [18] for the complexity of the L_∞-diagram of n points in d-space (it was also shown that this bound is tight in the worst case). It is thus tempting to conjecture that the maximum complexity of generalized Voronoi diagrams in higher dimensions is close to this bound. Unfortunately, this was recently shown to be false in [11], where a lower bound of $\Omega(n^{d-1})$ is given. The sites used in this construction are convex polytopes, and the distance is either Euclidean or a polyhedral convex distance function. For $d = 3$, this lower bound does not contradict the conjecture made above, that the complexity of generalized Voronoi diagrams should be at most near-quadratic in this case. Also, in higher dimensions, the conjecture mentioned above is still not refuted when the sites are singleton points. Finally, for the general case, the construction of [11] still leaves a gap of roughly a factor of n between the known upper and lower bounds.

Union of Geometric Objects

A subproblem related to generalized Voronoi diagrams is as follows. Let S and ρ be as above (say, for the 3-dimensional case). Let K denote the region consisting of all points $x \in \mathbb{R}^3$ whose smallest distance from a site in S is at most r, for some fixed parameter $r > 0$. Then $K = \bigcup_{s \in S} B(s, r)$, where $B(s, r) = \{x \in \mathbb{R}^3 \mid \rho(x, s) \leq r\}$. We thus face the problem of bounding the combinatorial complexity of the union of n objects in 3-space (of some special type). For example, if S is a set of lines and ρ is the Euclidean distance, the objects are n congruent infinite cylinders in 3-space. In general, if the metric ρ is a distance function induced by some convex body P, the resulting objects are the *Minkowski sums* $s \oplus (-rP)$, for $s \in S$, where $A \oplus B = \{x + y \mid x \in A,\ y \in B\}$. Of course, this problem can also be stated in any higher dimension.

Since it has been conjectured that the complexity of the whole Voronoi diagram should be near-quadratic (in 3-space), the same conjecture should apply to the (simpler) structure K (whose boundary can be thought of as a 'cross-section' of the diagram at 'height' r). Recently, this conjecture has been confirmed in [14], in the special case where both P and the objects of S are convex polyhedra [14]. Let us discuss this result in more detail. An earlier paper [12] has studied the case involving the union of k arbitrary convex polyhedra in 3-space, with a total of n faces. It was shown there that the complexity of the union is $O(k^3 + nk \log^2 k)$, and can be $\Omega(k^3 + nk\alpha(k))$ in the worst case. The upper bound was subsequently improved to $O(k^3 + nk \log k)$ [15], which still leaves a small gap between the upper and lower bounds. In the subsequent paper [14], these bounds were improved in the special case where the polyhedra in question are Minkowski sums of the form $s_i \oplus P$, where the s_i's are k pairwise-disjoint convex polyhedra, P is a convex polyhedron, and the total number of faces of these Minkowski sums is n. The improved bounds are $O(nk \log k)$ and $\Omega(nk\alpha(k))$. They are indeed near-quadratic, as conjectured (in fact, they are much better than quadratic when $k \ll n$).

However, the case where P is a ball (namely, the case of the Euclidean distance) is still open. The simplest unsolved instance of this problem is to establish

a near-quadratic upper bound for the complexity of the union of n congruent infinite cylinders in 3-space.

There are various extensions that are also interesting to consider. First, it would be interesting to study the problem in higher dimensions. This is likely to be much more difficult, so one should look at relatively simple cases, like the union of axis-parallel hypercubes, or of other simply-shaped objects. The case of axis-parallel hypercubes has recently been solved in [18], where it was shown that the maximum complexity of the union of n such hypercubes in d-space is $\Theta(n^{\lceil d/2 \rceil})$, and this improves to $\Theta(n^{\lfloor d/2 \rfloor})$ when all the hypercubes have the same size. Second, we can consider the case of more general objects (not necessarily Minkowski sums) which satisfy some 'fatness' properties, extending results obtained in [56] for 'fat' triangles in the plane. For example, what is the complexity of the union of n arbitrary (non-isothetic) unit cubes in $\mathbb{R}^3$?

Vertical Decomposition

In many algorithmic applications, one needs to be able to decompose a d-dimensional arrangement, or certain portions thereof, into a small number of subcells, each having constant description complexity. In a typical setup where this problem arises, we need to process in a certain manner an arrangement of n surfaces in d-space. We choose a random sample of r of the surfaces, for some sufficiently large constant r, construct the arrangement of these r surfaces, and decompose it into subcells as above. Since no such subcell is crossed by any surface in the random sample, it follows by standard ε-net theory [24, 47, 53] that, with high probability, none of these subcells is crossed by more than $O(\frac{n}{r} \log r)$ of the n given surfaces. (For this result to hold, it is essential that each of these subcells have constant description complexity.) This allows us to break the problem into recursive subproblems, one for each of these subcells, solve each subproblem separately, and then combine their outputs to obtain a solution for the original problem. The efficiency of this method crucially depends on the number of subcells. The smaller this number is, the faster is the resulting algorithm. (We note that the construction of a 'good' sample of r surfaces can also be performed deterministically, e.g., using the techniques of Matoušek [54].)

The only general-purpose known technique for decomposing an arrangement of surfaces into subcells of constant description complexity is the *vertical decomposition* technique. In this method, we erect a vertical 'wall' up and down (in the x_d-direction) from each $(d-2)$-dimensional face of the arrangement, and extend these walls until they hit another surface. This results in a decomposition of the arrangement into subcells so that each subcell has a unique top facet and a unique bottom facet, and each vertical line cuts it in a connected (possibly empty) interval. We next project each resulting subcell on the hyperplane $x_d = 0$, and apply recursively the same technique within each resulting $(d-1)$-dimensional projected cell, and then 'lift' this decomposition back into d-space, by extending each subcell c in the projection into the vertical cylinder $c \times \mathbb{R}$, and by cutting the original cell by these cylinders. We continue the recursion in this manner until we reach $d = 1$, and thereby obtain the vertical decomposition of

the given arrangement. The resulting subcells have the desired properties. Furthermore, if we assume that the originally given surfaces are algebraic of constant maximum degree, then the resulting subcells are semi-algebraic and are defined by a constant number of polynomials of constant maximum degree (although the latter degree can grow quite fast with d). In what follows, we ignore the algebraic complexity of the subcells of the vertical decomposition, and will be mainly interested in bounding their number as a function of n, the number of given surfaces.

It was shown in [20] that the number of cells in such a vertical decomposition is $O(n^{2d-3}\beta(n))$, where $\beta(n)$ is a slowly growing function of n (related to the inverse Ackermann's function). However, the only known lower bound is the trivial $\Omega(n^d)$, so there is a considerable gap here, for $d > 3$; for $d = 3$ the two bounds nearly coincide. Improving the upper bound appears to be a very difficult task. This problem has been open since 1989; it seems difficult enough to preempt, at the present state of knowledge, any specific conjecture on the true maximum complexity of the vertical decomposition in $d > 3$ dimensions.

The bound stated above applies to the vertical decomposition of an entire arrangement of surfaces. In many applications, however, one is interested in the vertical decomposition of only a portion of the arrangement, e.g., a single cell, the region lying below the lower envelope of the given surfaces, the zone of some surface, a specific collection of cells of the arrangement, etc. Since, in general, the complexity of such a portion is known (or conjectured) to be smaller than the complexity of the entire arrangement, one would like to conjecture that a similar phenomenon applies to vertical decompositions. Very recently, it was shown in [63] that the complexity of the vertical decomposition of a single cell in an arrangement of n surface patches in 3-space, as above, is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. As mentioned above, this leads to a near-quadratic algorithm for computing such a single cell, which implies that motion planning for fairly general systems with three degrees of freedom can be performed in near-quadratic time, thus settling a major open problem in the area. A challenging open problem is to obtain improved bounds for the complexity of the vertical decomposition of the region lying below the lower envelope of n d-variate functions, for $d \geq 3$.

Finally, an interesting special case is that of hyperplanes. For such arrangements, the vertical decomposition is a too cumbersome construct, because there are other easy methods for decomposing each cell into simplices, whose total number is $O(n^d)$. Still, it is probably a useful exercise to understand the complexity of the vertical decomposition of an arrangement of n hyperplanes in d-space. A recent result of [37] gives an almost tight bound of $O(n^4 \log n)$ for this problem in 4-space, but nothing significantly better than the general bound is known for $d > 4$. Another interesting special case is that of triangles in 3-space. This has been studied by [28, 67], where almost tight bounds were obtained for the case of a single cell ($O(n^2 \log^2 n)$), and for the entire arrangement ($O(n^2\alpha(n)\log n + K)$, where K is the complexity of the undecomposed arrangement). The first bound is slightly better than the general bound of [63] mentioned above. The paper [67] also derives sharp complexity bounds for the vertical decomposition of many cells

in such an arrangement, including the case of all nonconvex cells.

Other Applications

We conclude this survey by mentioning some additional applications of the new advances in the study of arrangements. We have already discussed in some detail the motion planning application, and have seen how the new results lead to a near-optimal algorithm for the general motion planning problem with three degrees of freedom. Here we discuss two other kinds of applications: to visibility problems in three dimensions, and to geometric optimization.

Visibility in Three Dimensions: Let us consider a special case of the so-called *aspect graph* problem, which has recently attracted much attention, especially in the context of three-dimensional scene analysis and object recognition in computer vision. The aspect graph of a scene represents all topologically-different views of the scene. For background and a survey of recent research on aspect graphs, see [19]. Here we will show how the new complexity bounds for lower envelopes, with some additional machinery, can be used to derive near-tight bounds on the number of views of polyhedral terrains.

Let K be a *polyhedral terrain* in 3-space; that is, K is the graph of a continuous piecewise-linear bivariate function, so it intersects each vertical line in exactly one point. Let n denote the number of edges of K. A line ℓ is said to *lie over* K if every point on ℓ lies on or above K. Let $\mathcal{L}_K$ denote the space of all lines that lie over K. (Since lines in 3-space can be parametrized by four real parameters, we can regard $\mathcal{L}_K$ as a subset of 4-space.) The *lower envelope* of $\mathcal{L}_K$ consists of those lines in $\mathcal{L}_K$ that touch at least one edge of K. Assuming general position of the edges of K, a line in $\mathcal{L}_K$ (or any line, for that matter) can touch at most four edges of K. We estimate the combinatorial complexity of this lower envelope, in terms of the number of its *vertices*, namely those lines in $\mathcal{L}_K$ that touch four distinct edges of K. It was shown in [43] that the number of vertices of $\mathcal{L}_K$, as defined above, is $O(n^3 \cdot 2^{c\sqrt{\log n}})$, for some absolute positive constant c.

We give here a sketch of the proof. We fix an edge e_0 of K, and bound the number of lines of $\mathcal{L}_K$ that touch e_0 and three other edges of K, with the additional proviso that the three other contact points all lie on one fixed side of the vertical plane passing through e_0. We then multiply this bound by the number n of edges, to obtain a bound on the overall number of vertices of $\mathcal{L}_K$. We first rephrase this problem in terms of the lower envelope of a certain collection of surface patches in 3-space, one patch for each edge of K (other than e_0), and then exploit the results on lower envelopes reviewed above.

The space $\mathcal{L}_{e_0}$ of oriented lines that touch e_0 is 3-dimensional: each such line ℓ can be specified by a triple (t, k, ζ), where t is the point of contact with e_0 (or, more precisely, the distance of that point from one designated endpoint of e_0), and $k = \tan\theta$, $\zeta = -\cot\phi$, where (θ, ϕ) are the spherical coordinates of the

direction of ℓ, that is, θ is the orientation of the xy-projection of ℓ, and ϕ is the angle between ℓ and the positive z-axis.

For each edge $e \neq e_0$ of K, let σ_e be the surface patch in $\mathcal{L}_{e_0}$ consisting of all points (t, k, ζ) representing lines that touch e and are oriented from e_0 to e. Note that if $(t, k, \zeta) \in \sigma_e$ then $\zeta' > \zeta$ iff the line (t, k, ζ') passes below e. It thus follows that a line ℓ in $\mathcal{L}_{e_0}$ is a vertex of the lower envelope of $\mathcal{L}_K$ if and only if ℓ is a vertex of the lower envelope of the surfaces σ_e in the $tk\zeta$-space, where the height of a point is its ζ-coordinate. It is easy to show that these surfaces are algebraic of constant description complexity. Actually, it is easily seen that the number s of intersections of any triple of these surfaces is at most 2. The paper [43] studies the special case of lower envelopes of collections of such algebraic surface patches in 3-space, with the extra assumption that $s = 2$. It is shown there that the complexity of the lower envelope of such a collection is $O(n^2 \cdot 2^{c\sqrt{\log n}})$, for some absolute positive constant c, a bound that is slightly better than the general bound stated above. These arguments immediately complete the proof. (This bound has been independently obtained by Pellegrini [60], using a different proof technique.) Recently, de Berg [26] has given a lower bound construction, in which the lower envelope of $\mathcal{L}_K$ has complexity $\Omega(n^3)$, implying that the upper bound stated above is almost tight in the worst case.

We can extend the above result as follows. Let K be a polyhedral terrain, as above. Let $\mathcal{R}_K$ denote the space of all rays in 3-space with the property that each point on such a ray lies on or above K. We define the lower envelope of $\mathcal{R}_K$ and its vertices in complete analogy to the case of $\mathcal{L}_K$. By inspecting the proof sketched above, one easily verifies that it applies equally well to rays instead of lines. Hence we obtain that the number of vertices of $\mathcal{R}_K$, as defined above, is also $O(n^3 \cdot 2^{c\sqrt{\log n}})$.

We can apply this bound to obtain a bound of $O(n^5 \cdot 2^{c'\sqrt{\log n}})$, for any $c' > c$, on the number of topologically-different orthographic views (i.e., views from infinity) of a polyhedral terrain K with n edges. We omit here details of this analysis, which can be found in [43]. The paper [29] gives a lower bound construction that produces $\Omega(n^5\alpha(n))$ topologically-different orthographic views of a polyhedral terrain, so the above bound is almost tight in the worst case. It is also instructive to note that, if K is an arbitrary polyhedral set in 3-space with n edges, then the maximum possible number of topologically-different orthographic views of K is $\Theta(n^6)$ [61].

Consider next the extension of the above analysis to bound the number of perspective views of a terrain. As shown recently in [6], the problem can be reduced to the analysis of $O(n^3)$ lower envelopes of appropriate collections of 5-variate functions. This leads to an overall bound of $O(n^{8+\varepsilon})$, for any $\varepsilon > 0$, for the number of topologically-different perspective views of a polyhedral terrain with n edges. This bound is also known to be almost tight in the worst case, as follows from another lower bound construction given in [29]. Again, in contrast, If K is an arbitrary polyhedral set with n edges, the maximum possible number of topologically-different perspective views of K is $\Theta(n^9)$ [61].

Geometric Optimization: In the past few years, many problems in geometric optimization have been attacked by techniques that reduce the problem to a problem involving arrangements of surfaces in higher dimensions. These reduced problems sometimes call for the construction of, and searching in lower envelopes or other substructures in such arrangements. Hence the area of geometric optimization is a natural extension, and a good application area, of the study of arrangements, as described above.

One of the basic techniques for geometric optimization is the *parametric searching* technique, originally proposed by Megiddo [58]. It has been used to solve a wide variety of geometric optimization problems, including many of those that involve arrangements. Some specific results of this kind include:

- **Selecting distances in the plane:** Given a set S of n points in $\mathbb{R}^2$ and a parameter $k \leq \binom{n}{2}$, find the k-th largest distance among the points of S [2]. Here the problem reduces to the construction and searching in 2-dimensional arrangements of congruent disks.
- **The segment center problem:** Given a set S of n points in $\mathbb{R}^2$, and a line segment e, find a placement of e that minimizes the largest distance from the points of S to e [35]. Using lower envelopes of bivariate functions, the problem can be solved in $O(n^{1+\varepsilon})$ time, for any $\varepsilon > 0$, improving substantially a previous near-quadratic solution given in [4].
- **Extremal polygon placement:** Given a convex polygon P and a closed polygonal environment Q, find the largest similar copy of P that is fully contained in Q [66]. This is just an extension of the corresponding motion planning problem, where the size of P is fixed. The running time of the algorithm is almost the same as that of the motion planning algorithm given in [51, 52].
- **Width in three dimensions:** Compute the width of a set S of n points in $\mathbb{R}^2$; this is the smallest distance between two parallel planes enclosing S between them. This problem has been studied in a series of papers [1, 7, 21], and the current best bound is $O(n^{3/2+\varepsilon})$, for any $\varepsilon > 0$ [7]. The technique used in attacking this and the two following problems reduce them to problems involving lower envelopes in 4 dimensions, where we need to construct and to search in such an envelope.
- **Biggest stick in a simple polygon:** Compute the longest line segment that can fit inside a given simple polygon with n edges. The current best solution is $O(n^{3/2+\varepsilon})$, for any $\varepsilon > 0$ [7] (see also [1, 9]).
- **Smallest-width annulus:** Compute the annulus of smallest width that encloses a given set of n points in the plane. Again, the current best solution is $O(n^{3/2+\varepsilon})$, for any $\varepsilon > 0$ [7] (see also [1, 9]).
- **Geometric matching:** Consider the problem where we are given two sets S_1, S_2 of n points in the plane, and we wish to compute a minimum-weight matching in the complete bipartite graph $S_1 \times S_2$, where the weight of an edge (p, q) is the Euclidean distance between p and q. One can also consider the analogous nonbipartite version of the problem, which involves just one set S of $2n$ points, and the complete graph on S. The goal is to explore the

underlying geometric structure of these graphs, to obtain faster algorithms than those available for general abstract graphs.
It was shown in [68] that both the bipartite and the nonbipartite versions of the problem can be solved in time close to $n^{2.5}$. Recently, a fairly sophisticated application of vertical decomposition in 3-dimensional arrangements, given in [3], has improved the running time for the bipartite case to $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. This technique does not yet extend to the nonbipartite case, which remains an interesting open problem.

This list is by no means exhaustive.

A final comment is that, although the parametric searching technique yields algorithms that are efficient theoretically, they are usually quite difficult to implement, or even just to describe, because an efficient implementation of algorithms based on parametric searching requires the existence of a fast parallel algorithm for some related problem. Moreover, in most of the applications, one needs to compute the roots of high-degree polynomials, which, if done exactly, slows down the algorithm considerably.

Some effort was made recently to replace parametric searching by alternative, simpler, and more geometric-oriented techniques. The alternative methods that have been used for problems involving arrangements include randomization [7, 55], and expander graphs [49, 50]. However, these techniques work so far only in some special cases, and no general technique is known.

References

1. P.K. Agarwal, B. Aronov and M. Sharir, Computing envelopes in four dimensions with applications, *Proc. 10th ACM Symp. on Computational Geometry* (1994), 348–358.
2. P.K. Agarwal, B. Aronov, M. Sharir and S. Suri, Selecting distances in the plane, *Algorithmica* 9 (1993), 495–514.
3. P.K. Agarwal, A. Efrat and M. Sharir, Vertical decompositions of shallow levels in arrangements and their applications, *Proc. 11th ACM Symp. on Computational Geometry* (1995).
4. P.K. Agarwal, A. Efrat, M. Sharir and S. Toledo, Computing a segment-center for a planar point set, *J. Algorithms* 15 (1993), 314–323.
5. P.K. Agarwal, O. Schwarzkopf and M. Sharir, The overlay of lower envelopes in 3-space and its applications, to appear in *Discrete Comput. Geom.*
6. P.K. Agarwal and M. Sharir, On the number of views of polyhedral terrains, *Discrete Comput. Geom.* 12 (1994), 177–182.
7. P.K. Agarwal and M. Sharir, Efficient randomized algorithms for some geometric optimization problems, *Proc. 11th ACM Symp. on Computational Geometry*, 1995.
8. P. Agarwal, M. Sharir and P. Shor, Sharp upper and lower bounds for the length of general Davenport Schinzel sequences, *J. Combin. Theory, Ser. A* 52 (1989), 228–274.
9. P.K. Agarwal, M. Sharir and S. Toledo, New applications of parametric searching in computational geometry. *J. Algorithms* 17 (1994), 292–318.
10. B. Aronov, M. Pellegrini and M. Sharir, On the zone of a surface in a hyperplane arrangement, *Discrete Comput. Geom.* 9 (1993), 177–186.

11. B. Aronov, personal communication, 1995.
12. B. Aronov and M. Sharir, The union of convex polyhedra in three dimensions, *Proc. 34th IEEE Symp. on Foundations of Computer Science* (1993), 518–527.
13. B. Aronov and M. Sharir, Castles in the air revisited, *Discrete Comput. Geom.* 12 (1994), 119–150.
14. B. Aronov and M. Sharir, On translational motion planning in three dimensions, *Proc. 10th ACM Symp. on Computationl Geometry* (1994), 21–30.
15. B. Aronov, M. Sharir and B. Tagansky, The union of convex polyhedra in three dimensions, to appear in *SIAM J. Comput.* (a revised version of [12]).
16. J.D. Boissonnat and K. Dobrindt, Randomized construction of the upper envelope of triangles in $\mathbb{R}^3$, *Proc. 4th Canadian Conf. on Computational Geometry* (1992), 311–315.
17. J.D. Boissonnat and K. Dobrindt, On-line randomized construction of the upper envelope of triangles and surface patches in $\mathbb{R}^3$, to appear in *Comp. Geom. Theory Appls.*
18. J.D. Boissonnat, M. Sharir, B. Tagansky and M. Yvinec, Voronoi diagrams in higher dimensions under certain polyhedral convex distance functions, *Proc. 11th ACM Symp. on Computational Geometry* (1995).
19. K.W. Bowyer and C.R. Dyer, Aspect graphs: An introduction and survey of recent results, *Int. J. of Imaging Systems and Technology* 2 (1990), 315–328.
20. B. Chazelle, H. Edelsbrunner, L. Guibas and M. Sharir, A singly exponential stratification scheme for real semi-algebraic varieties and its applications, *Proc. 16th Int. Colloq. on Automata, Languages and Programming* (1989), 179–193.
21. B. Chazelle, H. Edelsbrunner, L. Guibas and M. Sharir, Diameter, width, closest line pair, and parametric searching, *Discrete Comput. Geom.* 10 (1993), 183–196.
22. L.P. Chew, Near-quadratic bounds for the L_1 Voronoi diagram of moving points, *Proc. 5th Canadian Conf. on Computational Geometry* (1993), 364–369.
23. L.P. Chew, K. Kedem, M. Sharir, B. Tagansky and E. Welzl, Voronoi diagrams of lines in three dimensions under a polyhedral convex distance function, *Proc. 6th ACM-SIAM Symp. on Discrete Algorithms* (1995), 197–204.
24. K.L. Clarkson, New applications of random sampling in computational geometry, *Discrete Comput. Geom.* 2 (1987), 195–222.
25. K.L. Clarkson and P.W. Shor, Applications of random sampling in computational geometry, II, *Discrete Comput. Geom.* 4 (1989), 387–421.
26. M. de Berg, personal communication, 1993.
27. M. de Berg, K. Dobrindt and O. Schwarzkopf, On lazy randomized incremental construction, *Proc. 26th ACM Symp. on Theory of Computing* (1994), 105–114.
28. M. de Berg, L. Guibas and D. Halperin, Vertical decomposition for triangles in 3-space, *Proc. 10th ACM Symp. on Computational Geometry* (1994), 1–10.
29. M. de Berg, D. Halperin, M. Overmars and M. van Kreveld, Sparse arrangements and the number of views of polyhedral scenes, Manuscript, 1991.
30. H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg 1987.
31. H. Edelsbrunner, The upper envelope of piecewise linear functions: Tight complexity bounds in higher dimensions, *Discrete Comput. Geom.* 4 (1989), 337–343.
32. H. Edelsbrunner, L. Guibas, J. Pach, R. Pollack, R. Seidel and M. Sharir, Arrangements of curves in the plane: topology, combinatorics, and algorithms, *Theoret. Comput. Sci.* 92 (1992), 319–336.
33. H. Edelsbrunner and R. Seidel, Voronoi diagrams and arrangements, *Discrete Comput. Geom.* 1 (1986), 25–44.

34. H. Edelsbrunner, R. Seidel and M. Sharir, On the zone theorem for hyperplane arrangements, *SIAM J. Comput.* 22 (1993), 418–429.
35. A. Efrat and M. Sharir, A near-linear algorithm for the planar segment center problem, to appear in *Discrete Comput. Geom.*
36. J.-J. Fu and R.C.T. Lee, Voronoi diagrams of moving points in the plane, *Internat. J. Comput. Geom. Appl.* 1 (1994), 23–32.
37. L. Guibas, D. Halperin, J. Matoušek and M. Sharir, On vertical decomposition of arrangements of hyperplanes in four dimensions, *Discrete Comput. Geom.* 14 (1995) (in press).
38. L. Guibas, J. Mitchell and T. Roos, Voronoi diagrams of moving points in the plane, *Proc. 17th Internat. Workshop Graph-Theoret. Concepts Computer Science*, Lecture Notes in Comp. Sci., vol. 570, Springer-Verlag, pp. 113–125.
39. L. Guibas and M. Sharir, Combinatorics and algorithms of arrangements, in *New Trends in Discrete and Computational Geometry,* (J. Pach, Ed.), Springer-Verlag, 1993, 9–36.
40. L. Guibas, M. Sharir and S. Sifrony, On the general motion planning problem with two degrees of freedom, *Discrete Comput. Geom.* 4 (1989), 491–521.
41. D. Halperin, On the complexity of a single cell in certain arrangements of surfaces in 3-space, *Discrete Comput. Geom.* 11 (1994), 1–33.
42. D. Halperin and M. Sharir, Near-quadratic bounds for the motion planning problem for a polygon in a polygonal environment, *Proc. 34th IEEE Symp. on Foundations of Computer Science* (1993), 382–391.
43. D. Halperin and M. Sharir, New bounds for lower envelopes in three dimensions with applications to visibility of terrains, *Discrete Comput. Geom.* 12 (1994), 313–326.
44. D. Halperin and M. Sharir, Almost tight upper bounds for the single cell and zone problems in three dimensions, *Proc. 10th ACM Symp. on Computational Geometry* (1994), 11–20.
45. D. Halperin and M. Sharir, Arrangements and their applications in robotics: Recent developments, in *The Algorithmic Foundations of Robotics*, K. Goldberg, D. Halperin, J.C. Latombe and R. Wilson, Eds., A.K. Peters, Boston, MA, 1995, 495–511.
46. S. Hart and M. Sharir, Nonlinearity of Davenport–Schinzel sequences and of generalized path compression schemes, *Combinatorica* 6 (1986), 151–177.
47. D. Haussler and E. Welzl, ϵ-nets and simplex range queries, *Discrete Comput. Geom.* 2 (1987), 127–151.
48. J. Heintz, T. Recio and M.F. Roy, Algorithms in real algebraic geometry and applications to computational geometry, in *Discrete and Computational Geometry: Papers from DIMACS Special Year,* (J. Goodman, R. Pollack, and W. Steiger, Eds.), American Mathematical Society, Providence, RI, 137–163.
49. M. Katz and M. Sharir, Optimal slope selection via expanders, *Inform. Process. Lett.* 47 (1993), 115–122.
50. M. Katz and M. Sharir, An expander-based approach to geometric optimization, *Proc. 9th ACM Symp. on Computational Geometry* (1993), 198–207.
51. K. Kedem and M. Sharir, An efficient motion planning algorithm for a convex rigid polygonal object in 2–dimensional polygonal space, *Discrete Comput. Geom.* 5 (1990), 43–75.
52. K. Kedem, M. Sharir and S. Toledo, On critical orientations in the Kedem-Sharir motion planning algorithm for a convex polygon in the plane, *Proc. 5th Canadian Conference on Computational Geometry* (1993), 204–209.

53. J. Komlós, J. Pach and G. Woeginger, Almost tight bound on epsilon–nets, *Discrete and Computational Geometry* 7 (1992), 163–173.
54. J. Matoušek, Approximations and optimal geometric divide-and-conquer, *Proc. 23rd ACM Symp. on Theory of Computing* (1991), 506–511.
55. J. Matoušek, Randomized optimal algorithm for slope selection, *Inform. Process. Lett.* 39 (1991), 183–187.
56. J. Matoušek, J. Pach, M. Sharir, S. Sifrony and E. Welzl, Fat triangles determine linearly many holes, *SIAM J. Comput.* 23 (1994), 154–169.
57. P. McMullen and G. C. Shephard, *Convex Polytopes and the Upper Bound Conjecture*, Lecture Notes Ser. 3, Cambridge University Press, Cambridge, England, 1971.
58. N. Megiddo, Applying parallel computation algorithms in the design of serial algorithms, *J. ACM* 30, 852–865.
59. J. Pach and M. Sharir, The upper envelope of piecewise linear functions and the boundary of a region enclosed by convex plates: Combinatorial analysis, *Discrete Comput. Geom.* 4 (1989), 291–309.
60. M. Pellegrini, On lines missing polyhedral sets in 3-space, *Proc. 9th ACM Symp. on Computational Geometry* (1993), 19–28.
61. H. Plantinga and C. Dyer, Visibility, occlusion, and the aspect graph, *International J. Computer Vision*, 5 (1990), 137–160.
62. J.T. Schwartz and M. Sharir, On the Piano Movers' problem: II. General techniques for computing topological properties of real algebraic manifolds, *Advances in Appl. Math.* 4 (1983), 298–351.
63. O. Schwarzkopf and M. Sharir, Vertical decomposition of a single cell in 3-dimensional arrangements and its applications, in preparation.
64. M. Sharir, Almost tight upper bounds for lower envelopes in higher dimensions, *Discrete Comput. Geom.* 12 (1994), 327–345.
65. M. Sharir and P.K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.
66. M. Sharir and S. Toledo, Extremal polygon containment problems, *Comput. Geom. Theory Appls.* 4 (1994), 99–118.
67. B. Tagansky, A new technique for analyzing substructures in arrangements, *Proc. 11th ACM Symp. on Computational Geometry* (1995).
68. P.M. Vaidya, Geometry helps in matching, *SIAM J. Comput.* 18 (1989), 1201–1225.
69. A. Wiernik and M. Sharir, Planar realization of nonlinear Davenport–Schinzel sequences by segments, *Discrete Comput. Geom.* 3 (1988), 15–47.

This article was processed using the LaTeX macro package with LLNCS style