Arrangements and their Applications in Robotics: Recent Developments

Dan Halperin, *Stanford University, Stanford, CA, USA*
Micha Sharir, *Tel Aviv University, Tel Aviv, Israel and New York University, NY, USA*

We survey a collection of recent combinatorial and algorithmic results in the study of arrangements of low-degree algebraic surface patches in three or higher dimensions. The new results extend known results involving 2-dimensional arrangements, and they almost settle several conjectures posed eight years ago. Arrangements play a central role in the design and analysis of geometric algorithms, and they arise in a variety of seemingly unrelated applications. In this survey we concentrate on the application of the new results to problems involving collision-free motion planning with three degrees of freedom, and visibility over polyhedral terrains.

1 Introduction

Let $\mathcal{L} = \{\ell_1, \ell_2 \ldots, \ell_n\}$ be a given collection of n lines in the plane. We denote by $\mathcal{A}(\mathcal{L})$ the *arrangement* of $\mathcal{L}$, i.e., the decomposition of the plane into vertices, edges, and faces, induced by the lines in $\mathcal{L}$. A vertex of $\mathcal{A}(\mathcal{L})$ is an intersection point of two lines, an edge is a maximal connected relatively open portion of a line that does not meet any vertex, and a face is a maximal connected open region of the plane not meeting any edge or vertex. See Figure 1. Similarly, an arrangement of hyperplanes in d-dimensional space is the subdivision of that space into cells of dimensions $0, 1, \ldots, d$, each being a maximal connected set contained in the intersection of a fixed subcollection of the hyperplanes and not meeting any other hyperplane.

Arrangements of hyperplanes have been studied extensively in combinatorial and computational geometry [26]. In the preface to his book *Algorithms in Combinatorial Geometry* [26], Edelsbrunner writes: *"These* [geometric] *transforms led me to believe that arrangements of hyperplanes are at the very heart of computational geometry — and this is my belief now more than ever."*

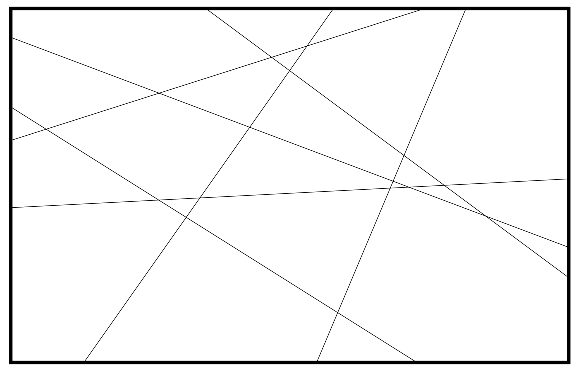

Figure 1: *An arrangement of lines (clipped within a window)*

However, when reasoning about geometric problems that arise in areas such as robotics or computer vision, hyperplanes are often not the most adequate objects to model these problems with. As will be seen below, it is often more appropriate to use objects that are (i) bounded, e.g., segments in the plane or polygons in 3-space; and/or (ii) non-linear, e.g., circular arcs or ellipses in the plane, ruled surfaces in 3-space, etc.

The definition of arrangements can naturally be extended to apply in more general situations, as follows. Let $\Sigma = \{\sigma_1, \ldots, \sigma_n\}$ be a given collection of n low-degree algebraic[1] surface patches in d-space (see Section 3 for a more precise statement of the properties that these surfaces are assumed to satisfy). The *arrangement* $\mathcal{A}(\Sigma)$ of Σ is the decomposition of d-space into (relatively open) cells of various dimensions, each

[1] We restrict ourselves to *algebraic* surface patches, of constant maximum degree, because the patterns of interaction among such surfaces are easier to analyze and to compute. Moreover, in practically all applications, the arrangements that arise do indeed involve low-degree algebraic surfaces.

being a maximal connected set contained in the intersection of a fixed subcollection of Σ and not meeting any other surface of Σ.

Before continuing, let us give a simple example of the use of arrangements of surfaces in robotics. Consider the following simple case of *robot motion planning.* Let B be an arbitrary polygonal object with k sides, moving (translating and rotating) in some open planar polygonal region V bounded by m edges. Any placement of B can be represented by the triple $(x, y, \tan\frac{\theta}{2})$, where (x, y) are the coordinates of some fixed reference point on B, and θ is the orientation of B. We regard these triples as points in 3-space, referred to as the *configuration space* of B. In the motion planning problem, we want to compute the *free* portion of this space, denoted as FP, and consisting of those placements of B at which it is fully contained in V. We note that the boundary of FP consists of placements at which B makes contact with the boundary of V, but does not penetrate into the interior of the complement of V. We can therefore define, within the configuration space of B, a collection of 'contact surfaces', each being either the locus of all placements of B at which some specific corner of B touches some specific edge of V, or the locus of placements at which some side of B touches some vertex of V. Under the above parametrization, it is easily seen that each of the resulting $O(km)$ contact surfaces is a 2-dimensional algebraic surface patch of degree at most 4, and its relative boundary consists of a constant number of algebraic arcs, of constant maximum degree as well.

If B is placed at a free placement $Z \in FP$, and moves continuously from Z, then it remains free as long as the corresponding path traced in configuration space does not hit any contact surface. Moreover, once this path crosses a contact surface, B becomes non-free (assuming, as is customary, *general position* of B and V). It follows that the connected component of FP that contains Z is the cell that contains Z in the arrangement $\mathcal{A}(\Sigma)$ of the contact surfaces, and that the entire FP is the union of a collection of certain cells in this arrangement. Hence, bounding the combinatorial complexity of a single cell in such an arrangement, and designing efficient algorithms for computing such a cell, are natural and major problems in the study of motion planning. As we will see later, these notions can naturally be extended, in more general motion planning problems, to arrangements of other kinds of surfaces and to higher dimensions.

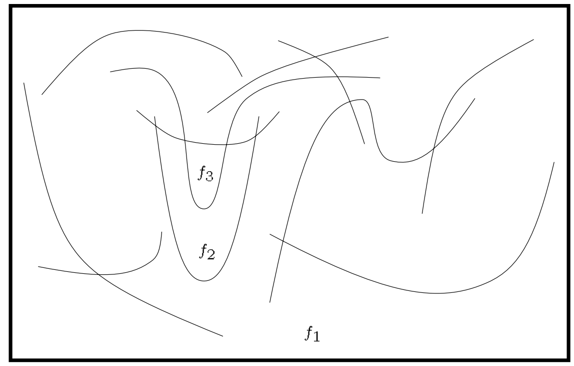

Figure 2: *An arrangement of curves*

To give the feeling of what makes the general setting more difficult to reason about (than arrangements of hyperplanes), let us consider the following problem. Let $\mathcal{A}$ be an arrangement of curves in the plane, and let π be a point in the plane, not lying on any curve. We wish to bound the *combinatorial complexity* of the face f_π of the arrangement that contains the point π, where this complexity is defined as the number of vertices and edges on the boundary of f_π. A face f in an arrangement of lines (see Figure 1) is clearly convex, as it is the intersection of halfplanes bounded by these lines; therefore, no line can contribute more than one edge to the boundary of f, and hence its complexity is at most $2n$. In contrast, a face in an arrangement of arbitrary curves can have a rather convoluted shape (see face f_1 in Figure 2, which makes up for most of the arrangement; the only other 2-dimensional faces are the small faces f_2 and f_3). Even simply-shaped faces in such an arrangement (such as face f_2 in Figure 2) need not be convex, and a curve may contribute more than one edge to the boundary of such a face. We are thus faced with an interesting and rather challenging problem of estimating the maximum complexity of a single face in general planar arrangements. (As just noted, this is a trivial problem for arrangements of lines; however, it is highly non-trivial even in the case of line segments.)

This and related questions are the topic of our survey. The new developments that we survey, however, concern arrangements in three or higher dimensions. The case of planar arrangements has been studied extensively, and rather satisfactorily, during the past decade, and we will mention the main known results in the background section 2.

The paper is organized as follows. In Section 2 we present the problems that this paper addresses and survey previous work on these problems. We state the basic new results in Section 3. We exemplify the usefulness of these results by applying them to problems involving robot motion planning (Section 4) and visibility and aspect graphs (Section 5). Section 6 deals with new results on Minkowski sums of convex polyhedra in three dimensions, which have applications in robot motion planning and in other related areas. The paper concludes in Section 7, with further applications of the new results and with some open problems.

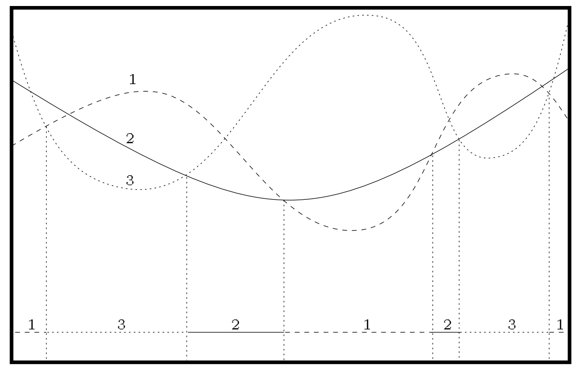

Figure 3: *A lower envelope and its associated sequence "1321231"*

2 Background and Previous Work

The problem posed at the introduction, which calls for estimating the complexity of a single face in a general planar arrangement, and a collection of related problems involving planar arrangements of curves, were settled, in a fairly complete and satisfactory fashion, more than five years ago. The main tool used in obtaining these results is *Davenport-Schinzel sequences*, a combinatorial tool that has found many applications in combinatorial and computational geometry. It is beyond the scope of this survey to give a detailed review of Davenport-Schinzel sequences, but we briefly present the basic relevant definitions and results, and refer the reader to the survey paper [60] or to the forthcoming book [63]; many of the other papers cited in our survey also make use of Davenport-Schinzel sequences.

Let n and s be positive integers. A *Davenport-Schinzel* sequence of order s composed of n distinct symbols (an (n,s)-DS sequence for short) is a sequence $U = (u_1, u_2, \ldots, u_m)$ that satisfies the following two conditions:

(i) $u_i \neq u_{i+1}$, for all $i < m$.

(ii) There do not exist $s+2$ indices, $1 \leq i_1 < i_2 < \cdots < i_{s+2} \leq m$, such that

$$\begin{array}{ccccccc} u_{i_1} & = & u_{i_3} & = & \cdots & = & a \\ u_{i_2} & = & u_{i_4} & = & \cdots & = & b\,, \end{array}$$

and $a \neq b$.

Let $\lambda_s(n)$ denote the maximum length of an (n,s)-DS sequence. It is easy to see that $\lambda_1(n) = n$, and that $\lambda_2(n) = 2n-1$. Hart and Sharir [39] showed that $\lambda_3(n) = \Theta(n\alpha(n))$, where $\alpha(n)$ is the extremely slowly growing functional inverse of Ackermann's function, which, for all practical purposes, can be regarded as a constant. Still, from a theoretical point of view, $\lambda_3(n)$ grows faster than any linear function of n. The best bounds known to date for $\lambda_s(n)$, with $s > 3$, were obtained by Agarwal et al. [4], who showed that, for any fixed s, $\lambda_s(n)$ is an almost linear, slightly superlinear, function of n.

The connection of Davenport-Schinzel sequences to geometric problems lies in the following easy observation: Let $f_1, \ldots, f_n$ be n real-valued continuous functions defined on the real line, with the property that the graphs of any pair of these functions intersect in at most s points. Let $\psi(x) = \min_i f_i(x)$ denote the *lower envelope* of these functions, and let U be the sequence of function indices, in the order in which they appear along the graph of ψ from left to right; see Figure 3. Then U is an (n,s)-DS sequence. Conversely, any (n,s)-DS sequence U can be realized as the 'lower-envelope sequence' of a collection of n functions, with the above properties. Thus, Davenport-Schinzel sequences arise naturally in geometric (and other) problems involving minimization of a collection of univariate functions.

This observation can be extended to the case where the given functions are only partially defined, so that the graph of each function is a connected x-monotone arc. We state the result in the following theorem, and refer the reader to [39] for more details.

Theorem 1 *The complexity of the lower envelope of n x-monotone curves, having the property that no pair of them intersect in more than s points, is $O(\lambda_{s+2}(n))$. The envelope can be calculated, by a simple divide-and-conquer procedure, in time $O(\lambda_{s+2}(n) \log n)$.*

Remark: A slightly faster algorithm, with running time $O(\lambda_{s+1}(n) \log n)$, has been given by Hershberger [42].

However, the application of these sequences goes much further. Guibas et al. [32] obtained the following result:

Theorem 2 [32] *The complexity of a single face in an arrangement of n Jordan arcs, having the property that no pair of arcs intersect in more than s points, is $O(\lambda_{s+2}(n))$. Such a face can be calculated (deterministically) in time $O(\lambda_{s+2}(n) \log^2 n)$.*

This result, as well as most of the algorithmic results reported in this survey, assumes that the shape of the given arcs is sufficiently simple (e.g., that they are all algebraic of constant maximum degree), and assumes a standard model of computation in *computational geometry,* namely, the algorithms use infinite-precision real numbers and each arithmetic operation (or any other simple operation, such as extracting the real roots of any fixed-degree polynomial) has constant cost. See, e.g., [57, Section 1.4].

Remark: There are also several *randomized* algorithms for computing a single face in such an arrangement [16, 18, 23]. The expected running time of these algorithms, over the internal randomizations that they perform, is $O(\lambda_{s+2}(n) \log n)$.

The maximum possible combinatorial complexity of the entire arrangement of such a collection of arcs is $\Theta(n^2)$. Theorems 1 and 2 thus imply that the maximum complexity of the lower envelope and of any single face in an arrangement of 'well-behaved' curves in the plane are asymptotically equal, and they are both smaller by roughly a factor of n than the complexity of the entire arrangement.

In $d \geq 3$ dimensions, a prevailing conjecture (see, e.g., [53]) is that the maximum complexity of a single cell (or of the lower envelope) in an arrangement $\mathcal{A}(\Sigma)$ of n surface patches, as above, is close to $O(n^{d-1})$, which is again smaller by roughly a factor of n than the maximum complexity of the entire arrangement, which can be $\Theta(n^d)$ [56].[2] (The complexity of the envelope, or of a cell, is the number of faces of the arrangement, of all dimensions, that appear along the envelope, or along the cell boundary.) A stronger version of the conjecture asserts that the maximum complexity of a single cell (or of the lower envelope) in such an arrangement is $O(n^{d-2}\lambda_s(n))$, where s is some constant that depends on the maximum degree of the given surfaces and of their boundaries.

For the single cell problem, there are only a few special cases in which the known bounds are better than the naive bound $O(n^d)$ (for the complexity of the entire arrangement). These include the cases of

(i) hyperplanes, where the complexity of a single cell, being a convex polyhedron bounded by at most n hyperplanes, is $O(n^{\lfloor d/2 \rfloor})$ (by the Upper Bound Theorem [50]);

(ii) spheres, where an $O(n^{\lceil d/2 \rceil})$ bound is easy to obtain by lifting the spheres into hyperplanes in $(d+1)$-space [26, 58];

(iii) $(d-1)$-simplices, where an $O(n^{d-1} \log n)$ bound has been established in [8]; and

(iv) several special types of surfaces in three dimensions, that arise in motion planning for various specific robot systems B with three degrees of freedom [34].

Similarly, improved bounds on the complexity of lower envelopes (better than the naive $O(n^d)$ bound) were previously obtained only for families of a few types of surfaces or surface patches, such as hyperplanes, balls, simplices, and, in three dimensions, also for a few other types of surfaces (see [53, 58]). The bounds for hyperplanes and for balls are the same as in the case of a single cell, and in fact are considerably smaller than the conjectured bounds. However, these surfaces are 'too simple'; it is easy to produce examples of a collection of n 'well-behaved' surfaces in d-space whose lower

[2] In these bounds, and in most of the other bounds mentioned in this paper, we are mainly concerned with the dependence of the quantities in question on the number n of curves or surfaces. In other words, we regard n as a large variable, and regard the dimension d and the maximum algebraic degree b of the surfaces as fixed constants. This is justified in most of the applications. The further dependence of the bounds on d and b is hidden in the constants of proportionality.

envelope has complexity slightly larger than $\Omega(n^{d-1})$. For example, the maximum complexity of the lower envelope of n $(d-1)$-simplices in d-space was shown in [27, 53] to be $\Theta(n^{d-1}\lambda_3(n)) = \Theta(n^{d-1}\alpha(n))$, so the conjecture is fully established for this special case.

In summary, these conjectures have been proved only for a few special cases of arrangements, and they were largely open in the general case stated above. In fact, no bounds better than $O(n^d)$ were known for the general case, even in three dimensions.

The situation has changed considerably over the past couple of years, when, in a series of papers [35, 36, 37, 62], the authors have obtained improved bounds for the complexity of lower envelopes and single cells in three and higher dimensions. These results have almost settled the above conjectures in the case of lower envelopes in any dimension, and in the case of single cells in three dimensions. In this survey, we summarize the recent developments, as reported in these papers, with an emphasis on their applications to robot motion planning and computer vision.

We also survey here another, different though related, problem, where sharp bounds were known in the two-dimensional case, and no non-trivial bounds were known in three (or higher) dimensions. The problem is that of bounding the complexity of the *Minkowski sum* of certain polygonal sets in the plane, or of certain polyhedral sets in 3-space.

Let A and B be two sets in $\mathbb{R}^2$ or in $\mathbb{R}^3$. The *Minkowski sum* (or vector sum) of A and B, denoted $A \oplus B$, is the set $\{a + b \mid a \in A,\, b \in B\}$. We will also use the notation $A \ominus B = A \oplus (-B) = \{a - b \mid a \in A,\, b \in B\}$.

Choose a reference point r rigidly attached to B, and suppose that the given placement of B is such that the reference point coincides with the origin. Suppose that B is allowed to translate (without rotating) in the plane or in 3-space, and that we represent each placement of B by the placement of the reference point r. In this case, $A \ominus B$ is the locus of all placements of the reference point for which the corresponding placement of B intersects A. Therefore, the Minkowski sum $A \ominus B$ is a useful construct in (translational) robot motion planning and in related problems [46, 49], and is often referred to as the *C-obstacle* (or *expanded obstacle*) induced by A.

We confine ourselves to Minkowski sums of polygonal sets in the plane, and of polyhedral sets in 3-space. As is well known, the Minkowski sum of polygonal sets is a polygonal set (see [31]), and, similarly, the Minkowski sum of polyhedral sets is a polyhedral set.

One of the main interesting special cases involving Minkowski sums, as above, is the case where B is a convex polygon in the plane, and the set A is the union of a collection of pairwise-disjoint convex polygons, $A_1, A_2, \ldots, A_m$ (it suffices to require that these polygons have pairwise disjoint interiors). We assume that the polygon B has a fixed number of sides, and that all the polygons A_i have a total of n sides. Kedem *et al.* proved the following result:

Theorem 3 [45] *The maximum combinatorial complexity of the boundary of the Minkowski sum $A \ominus B$, for A and B as defined above, is $O(n)$; moreover, this boundary contains at most $6m - 12$ points of intersection between the boundaries of pairs of different sums of the form $A_i \ominus B$. The sum $A \ominus B$ can be computed (deterministically) in $O(n \log^2 n)$ time.*

The three-dimensional variant of this problem has been open since 1986, when the paper [45] appeared. The goal is to bound the combinatorial complexity of the Minkowski sum $A \ominus B$, where B is a convex polyhedron (with a fixed number of faces), and A is the union of k pairwise-disjoint convex polyhedra with a total of n faces. It was conjectured that the complexity of this sum is roughly quadratic in n. A trivial cubic upper bound for this complexity can easily be deduced; however, no subcubic bounds were known, except for the special case where B is a line segment [44, 59]. Recently, a sharp bound for the case where B is a box was obtained by Halperin and Yap [38], and an almost tight bound for the general case was obtained by Aronov and Sharir [10]. The new results and their applications are described in Section 6.

3 Envelopes and Single Cells in Higher Dimensions: New Results

Let $\Sigma = \{\sigma_1, \ldots, \sigma_n\}$ be a given collection of n $(d-1)$-dimensional surfaces or surface patches in d-space. We assume that these surfaces satisfy the following conditions:

(i) Each σ_i is monotone in the $x_1 x_2 \cdots x_{d-1}$-direction (that is, every line parallel to the x_d-axis intersects σ_i in at most one point). Moreover, each σ_i is a portion of a (($d-1$)-dimensional) algebraic surface of constant maximum degree b (i.e., a set of the form $\{(x_1, \ldots, x_d) \mid P(x_1, \ldots, x_d) = 0\}$, for some polynomial P of degree $\leq b$).

(ii) The projection in the x_d-direction of σ_i onto the hyperplane $x_d = 0$ is a *semi-algebraic set* (see, e.g., [12, 40]) defined in terms of a constant number of $(d-1)$-variate polynomials of constant maximum degree, say, b too.

(iii) The relative interiors of any d of the given surfaces intersect in at most s points, for some constant parameter s (by Bezout's theorem [40] and by Property (iv) below, we always have $s \leq b^d$).

(iv) The surface patches in Σ are in *general position*, meaning that the coefficients of the polynomials defining the surfaces and their boundaries are algebraically independent over the rationals. (This assumption excludes degenerate interactions among the surfaces; for example, in three dimensions, it excludes configurations where four surfaces meet at a point, the boundary of one surface meets a curve of intersection of two other surfaces, two boundary curves of distinct surfaces meet at a point, etc.)

We remark that the somewhat restrictive condition (iv) and the first part of condition (i) are not essential for the analysis. If condition (i) does not hold, we can decompose each surface into a constant number of pieces that are monotone in the $x_1 x_2 \cdots x_{d-1}$-direction by cutting it along the locus of points with x_d-vertical tangency. If condition (iv) does not hold, we can argue, by applying a small random perturbation of the given polynomials, that the complexity of the lower envelope (or of a single cell) in a degenerate arrangement of surfaces is at most proportional to the worst-case complexity of the lower envelope (or of a single cell) in arrangements of surfaces in general position (see [62]). Condition (iii) is stated explicitly, because the bounds given below depend more significantly on the parameter s than on any of the other parameters d and b.

3.1 Lower Envelopes in Higher Dimensions

We extend the definition of a lower envelope given in Section 2 to higher dimensions. The lower envelope E_Σ of Σ is the graph of the (partial) function $x_d = E_\Sigma(x_1, \ldots, x_{d-1})$ that maps each point $(x_1, \ldots, x_{d-1})$ to the smallest x_d-coordinate among those of the points of intersection between the x_d-parallel line through $(x_1, \ldots, x_{d-1})$ and the surfaces in Σ (if no such surface exists, E_Σ is undefined at $(x_1, \ldots, x_{d-1})$). If that point lies on the boundary of one or more surfaces, we take the maximal closed segment contained in the x_d-parallel line through $(x_1, \ldots, x_{d-1})$, whose bottom endpoint lies on the envelope and which does not cross the relative interior of any surface in Σ, and say that the envelope is *attained* at $(x_1, \ldots, x_{d-1})$ by all surfaces that touch that segment; if the envelope point does not lie on any surface boundary, we say that the envelope is attained by (the relative interior of) all surfaces incident to that point. If we project E_Σ onto the hyperplane $H : \; x_d = 0$, we obtain a decomposition $\mathcal{M} = \mathcal{M}_\Sigma$ of H into connected relatively-open semi-algebraic sets (which we call *cells*), such that each cell c of $\mathcal{M}$ is a maximal connected portion of H over which E_Σ is attained by a fixed combination of the relative interiors of some surfaces and/or the boundaries of other surfaces (by the general position assumption, the number of such surfaces is at most d). The *combinatorial complexity* of E_Σ is defined to be the number of cells (of all dimensions) of $\mathcal{M}$.

The main result for lower envelopes is:

Theorem 4 [35, 62] *The combinatorial complexity of the lower envelope of a collection of n $(d-1)$-dimensional surface patches in d-space that satisfy conditions (i)–(iv), is $O(n^{d-1+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε, d, s, and on the maximum degree b and the shape of the given surfaces and of their boundaries.*

The main ingredients of the proof of this theorem appear in the paper [35], which addresses a special three-dimensional case, and exploits the probabilistic analysis technique of Clarkson and Shor (see [20, 61]) for obtaining generalized '$(\leq k)$-level' bounds in arrangements. The proof of the general case is given in [62]; it uses induction on d, and extends the random-sampling technique of [35] in several ways, so as to make it apply to the more general situation.

The main idea of the proof is rather simple. We fix some threshold constant parameter k. For each vertex v of the envelope, we choose a curve γ of $\mathcal{A}(\Sigma)$ incident to v, and follow γ from v *away* from the envelope. If we encounter at least k vertices of $\mathcal{A}(\Sigma)$ along γ before reaching the envelope again (or before γ ends), we charge v to the block of the first k such vertices, and observe that the level in $\mathcal{A}(\Sigma)$ of each charged vertex is at most k. Hence the number of envelope vertices v of this kind is proportional to $1/k$ times the number of vertices of $\mathcal{A}(\Sigma)$ at level $\leq k$. The latter number, by the results of Clarkson and Shor [20, 61], is $O(k^d \phi(n/k))$, where $\phi(m)$ denotes the maximum number of vertices of the lower envelope of m surfaces, with the above properties. After bounding the number of envelope vertices for which the above charging scheme does not apply, we obtain a recurrence for $\phi(n)$, which has, roughly, the form $\phi(n) = O(k^{d-1}\phi(n/k))$ plus other 'overhead' terms. This recurrence solves to the asserted bound. See [35, 62] for more details.

In [62], Sharir also presents a randomized algorithm for computing lower envelopes in three dimensions. Alternative randomized algorithms have been given by Boissonnat and Dobrindt [13], and by de Berg, Dobrindt and Schwarzkopf [23]. Perhaps the simplest algorithm for computing lower envelopes in three dimensions is due to Agarwal, Schwarzkopf and Sharir [2]; this algorithm is deterministic and uses a straightforward divide-and-conquer approach. We summarize all these results in the following theorem.

Theorem 5 *The lower envelope of n surface patches in 3-space, satisfying conditions (i)–(iv) above, can be computed in deterministic or in randomized expected time $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, in an appropriate 'algebraic' model of computation.*[3] *The constant of proportionality depends on ε and on the maximum degree of the given surfaces.*

In $d > 3$ dimensions, efficient construction of lower envelopes, with time complexity close to $O(n^{d-1})$, is more problematic to design, for several technical reasons that we will not elaborate here. There have been two recent developments in this direction, both due to Agarwal, Aronov and Sharir [1]:

Theorem 6 *Let Σ be a collection of n surfaces or surface patches in d-space, satisfying conditions (i)–(iv).*

(a) All the vertices, edges, and 2-faces of the lower envelope of Σ can be computed in randomized expected time $O(n^{d-1+\varepsilon})$, for any $\varepsilon > 0$.

(b) In four dimensions, the lower envelope of Σ can be computed in randomized expected time $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$, in the following strong sense: The algorithm produces a data structure, of size $O(n^{3+\varepsilon})$, so that, given any query point (x, y, z) in 3-space, the value of the envelope at (x, y, z), and the function(s) attaining the envelope at this point, can be computed in time $O(\log^2 n)$.

3.2 Single Cells and Zones in Three Dimensions

In this subsection, we consider the arrangement of a collection Σ of n low-degree algebraic surface patches, as above, in three-dimensional space. We are given a point Z not lying on any surface, and define $C = C_Z(\Sigma)$ to be the cell of the arrangement $\mathcal{A}(\Sigma)$ that contains Z.

The *combinatorial complexity* of a cell C is the number of lower-dimensional cells of $\mathcal{A}(\Sigma)$ appearing on the boundary of C. The authors have recently obtained in [36, 37] an almost-tight upper bound for the complexity of a single cell in such an arrangement:

Theorem 7 *The combinatorial complexity of a single cell in an arrangement of n algebraic surface patches in 3-space, satisfying the conditions (i)–(iv), is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε, s and b.*

This bound is asymptotically the same as the bound for the complexity of lower envelopes in three dimensions, as just mentioned. This result almost establishes the conjecture mentioned above, in three dimensions.

The proof of Theorem 7 adapts the analysis technique of [36], used by the authors to tackle a special case that arises in motion planning (see Section 4 below), which in turn is based on the analysis technique of

[3] We assume here a model of computation in which primitive operations involving a constant number of the given surfaces can be performed in constant time. Since we have assumed that all the surfaces and their boundaries are algebraic of constant maximum degree, such a model is plausible, using standard machinery from computational real algebraic geometry for performing each operation of this kind in an exact manner, using rational arithmetic, and in constant time (see, e.g., [41]). Alternatively, we can always fall back to the more powerful model of infinite-precision real arithmetic, described earlier.

[35, 62] for the case of lower envelopes. The lesson one can learn from the analysis in [36] is that, in studying the complexity of a single cell, one needs the following two preliminary results to 'bootstrap' the recurrences appearing in the analysis:

(a) a sharp bound on the number of 'x-extreme' vertices of the cell C (vertices whose x coordinate is smallest or largest in a small connected neighborhood of the vertex within C), and

(b) a sharp bound on the number of vertices bounding 'popular' faces of C (2-faces that are adjacent to C on both 'sides'; see [6, 8, 36], and below).

Bounds on these quantities were obtained in [36], using special properties of the surfaces that arise in the case studied there. The main technical contributions of the study of the general case [37] is a derivation of such bounds in the general algebraic setting assumed above. The bound (a) is obtained using considerations which are related to Morse theory (see e.g. [43]), but are simpler to derive in three dimensions. The bound (b) is obtained by applying the new probabilistic technique of [35, 36, 62] to counting only the vertices of popular faces (this idea is in the spirit of the methodology used in [6, 8]). Once these two bounds are available, the rest of the proof is rather similar to those used in [35, 36, 62], although certain non-trivial adjustments are required.

An interesting application of Theorem 7 is to bound the combinatorial complexity of the *zone* of a surface in an arrangement of other surfaces in 3-space. Specifically, let Σ be a collection of n algebraic surface patches in 3-space, and let σ be another such surface, so that the surfaces in $\Sigma \cup \{\sigma\}$ satisfy conditions (i)–(iv). The zone of σ in $\mathcal{A}(\Sigma)$ is the collection of all cells of $\mathcal{A}(\Sigma)$ that are crossed by σ. The complexity of the zone is the sum of the complexities of all its cells. The following theorem is a simple consequence of Theorem 7, based on an extension of a simple observation used in [28] for the analysis of zones in 2-dimensional arrangements:

Theorem 8 *The combinatorial complexity of the zone of σ in $\mathcal{A}(\Sigma)$ is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε, s and b.*

4 Motion Planning with Three Degrees of Freedom

One of the main motivations for studying the 'single cell' problem is its applications to *robot motion planning*. Extending the discussion given in the introduction, let B be an arbitrary robot system with k degrees of freedom, moving in some environment V filled with obstacles. Any placement of B can be represented by a point in $\mathbb{R}^k$, whose coordinates are the k parameters controlling the degrees of freedom of B; as above, this space is called the *configuration space* of B. We want to compute the *free* portion of this space, denoted as FP, which consists of those placements of B at which it does not meet any obstacle. As above, the boundary of FP consists of placements at which B makes contact with some obstacle(s), but does not penetrate into any of them. Under reasonable assumptions concerning B and V, we can represent the subset of 'contact placements' of B (including those placements at which B makes contact with an obstacle but may also penetrate into other obstacles) as the union of a collection of a finite number of surface patches, referred to as *contact surfaces*, all algebraic of constant maximum degree, and whose relative boundaries are also algebraic of constant maximum degree.

As already noted, if B is placed at a free placement Z, and moves continuously from Z, then it remains free as long as the corresponding path traced in configuration space does not hit any contact surface. Moreover, once this path crosses a contact surface, B becomes non-free (under appropriate general position assumptions). It follows, as above, that the connected component of FP that contains Z is the cell that contains Z in the arrangement $\mathcal{A}(\Sigma)$ of the contact surfaces. (The entire FP is the union of a collection of certain cells in this arrangement.) Hence, the immediate problems that we face are to bound the combinatorial complexity of a single cell in such an arrangement, and to design an efficient algorithm for its construction.

In view of this discussion, Theorem 7 immediately implies the following corollary:

Theorem 9 *Let B be a given robot with three degrees of freedom, and with a fixed number of boundary features, and suppose that B is free to move in an environment filled with obstacles, whose boundaries have a total of n features. We assume that each contact*

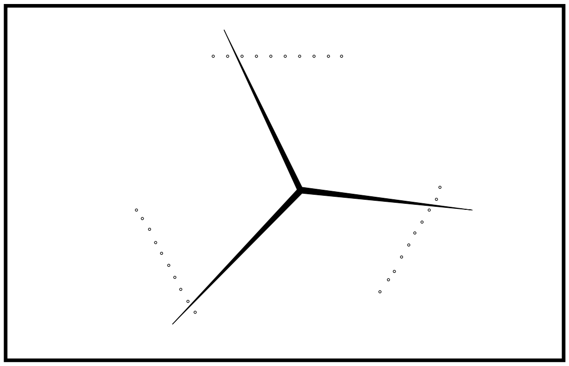

Figure 4: *A rigid polygonal robot moving among n point obstacles in the plane; the free configuration space has $\Theta(n^3)$ connected components.*

surface (representing contacts between some fixed robot feature and some fixed obstacle feature) is an algebraic surface patch of constant degree. Then the combinatorial complexity of a single connected component of the free configuration space of the moving robot is $O(n^{2+\varepsilon})$, where the constant of proportionality depends on ε and on the maximum degree of the contact surfaces and of their boundaries.

This combinatorial result implies that, in the set-up assumed above, the combinatorial complexity of the portion of the configuration space that the robot can reach from a fixed initial placement, is at most near-quadratic (in the complexity of the obstacles). To appreciate this result, we should note that the complexity of the entire free configuration space of such a robot can be $\Theta(n^3)$ (as is illustrated in Figure 4).

However, in order to solve the above motion planning problem, we still need to compute a reasonable representation of such a single cell. So far, we do not have an efficient algorithm (of near-quadratic complexity) for constructing a single cell in general 3-dimensional arrangements. Such efficient algorithms have been obtained only in a few special cases, one of them is that of the arrangement induced by moving (translating and rotating) an arbitrary polygon in a polygonal environment:

Theorem 10 [36] *A single connected component of the free configuration space of a polygon with a constant number of sides, moving in a polygonal region bounded by n edges, can be computed by a randomized algorithm with expected running time $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, with the constant of proportionality depending on ε.*

This algorithm is based on random sampling, and is adapted from an algorithm for computing a single cell in an arrangement of triangles, due to Aronov and Sharir [7]. In fact, the same paradigm can be used to compute the single cell in the general setting of Theorem 7. As it turns out, the running time of the resulting algorithm depends on the number of subcells of 'constant description complexity', into which one can decompose the given cell. To obtain algorithms with good running time, using the above paradigm, one needs such a decomposition scheme that does not create too many subcells (in the above set-up, we need decompositions with only a near-quadratic number of subcells). Such decompositions were obtained for arrangements of triangles in 3-space [7, 24] (see also Theorem 18 below), and for the above case of a polygon moving in a polygonal environment [36]. Thus, the algorithmic open problem here—to extend the algorithm to the general case—can be reduced to the combinatorial problem of obtaining sharp upper bounds on the number of subcells of constant description complexity needed in a decomposition of any single cell in an arrangement of surfaces, as above. (Of course, there is always the possibility of obtaining an alternative efficient algorithm that does not depend on the availability of such an economical decomposition.)

5 Visibility Over a Polyhedral Terrain

In this section we study 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 [14].

Here we will show how the new complexity bounds for lower envelopes (Theorem 4), with some additional machinery, can be used to derive near-tight bounds on the number of views of polyhedral terrains.

5.1 The Envelope of Rays over a Terrain

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. Our goal is to analyze the combinatorial complexity of the lower envelope, but to simplify matters (and with no loss of generality), we only count the number of its *vertices*, namely those lines in $\mathcal{L}_K$ that touch four distinct edges of K. We show:

Theorem 11 [35] *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 (see [35] for a detailed 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 of Section 3.

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 satisfy conditions (i)–(iv) of Section 3 for $d = 3$; we omit here the rather straightforward details. Actually, it is easily seen that the number s of intersections of any triple of these surfaces is at most 2. The paper [35] studies the special case of lower envelopes of collections of surface patches in 3-space, which satisfy conditions (i)–(iv), with the extra assumption that $s = 2$. It is shown in [35] 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 in Theorem 4. These arguments immediately imply Theorem 11.

Remarks: (1) The bound of Theorem 11 has been independently obtained by Pellegrini [54], using a different proof technique.

(2) Recently, de Berg [22] 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 in Theorem 11 is almost tight in the worst case.

We can extend the result of Theorem 11, 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 of Theorem 11, one easily verifies that it applies equally well to rays instead of lines. This is because, after fixing an edge e_0, each 'ray-vertex' of $\mathcal{R}_K$ under consideration, when extended into a full line, becomes a 'line-vertex' of $\mathcal{L}_{K^*}$, where K^* is the portion of K cut off by a halfspace bounded by the vertical plane through e_0. Hence we obtain:

Corollary 1 *The number of vertices of $\mathcal{R}_K$, as defined above, is also $O(n^3 \cdot 2^{c\sqrt{\log n}})$.*

This corollary will be needed in the following subsection.

5.2 The Number of Orthographic Views of a Polyhedral Terrain

We next apply Theorem 11 to obtain a bound on the number of topologically-different orthographic views (i.e., views from infinity) of a polyhedral terrain K with n edges.

Following the analysis of [25], each orthographic view of K can be represented as a point on the sphere at infinity $\mathcal{S}^2$. For each triple (e_1, e_2, e_3) of edges of K we consider the locus $\gamma_{(e_1,e_2,e_3)}$ of views for which these three edges appear to be concurrent (that is, there exists a line parallel to the viewing direction which touches these three edges); each such locus is a curve along $\mathcal{S}^2$.

We next replace each curve $\gamma = \gamma_{(e_1,e_2,e_3)}$, as defined above, by its maximal *visible* portions; a point on γ is said to be visible if the corresponding line that touches the three edges e_1, e_2, e_3 either lies over K or else penetrates below K only at points that lie further away from its contacts with the edges e_1, e_2, e_3; in other words, we require the existence of a ray in the direction opposite to the viewing direction that touches e_1, e_2, and e_3, but otherwise lies fully above K. As is easily verified, each visible portion of γ is delimited either at an original endpoint of γ or at a point whose corresponding ray is a vertex of $\mathcal{R}_K$. Hence, by Corollary 1, the total number of the visible portions of all the loci γ is $O(n^3 \cdot 2^{c\sqrt{\log n}})$. We refer to these visible portions as *arcs of visible triple-contact views* (along $\mathcal{S}^2$).

Now, consider the arrangement, along $\mathcal{S}^2$, of the arcs of visible triple-contact views, and observe that the number of views that we seek to bound is proportional to the complexity of the arrangement of these arcs within $\mathcal{S}^2$. This is a consequence of the easy observation that a combinatorial change in the structure of an orthographic view can occur only in directions where either three edges of K appear to be concurrent (and this concurrency is visible), or a vertex of K and an edge of K appear to be coincident; both types of directions appear along arcs of the above arrangement. The view inside each face is combinatorially fixed. Thus we wish to count the number of faces in the arrangement, which is bounded by the complexity of the entire arrangement. We next apply a result of Cole and Sharir [21], which, rephrased in the context under discussion, states that each meridian of $\mathcal{S}^2$ crosses at most $k = O(n\lambda_4(n))$ arcs of visible triple-contact views. As shown in [25], this implies that the complexity of the arrangement of these arcs is $O(Nk)$, where N is the number of arcs. Hence we obtain:

Theorem 12 [25, 35] *The number of topologically-different orthographic views of a polyhedral terrain with n edges is*

$$O(n^4\lambda_4(n) \cdot 2^{c\sqrt{\log n}}) = O(n^5 \cdot 2^{c'\sqrt{\log n}}) ,$$

for any constant c' slightly larger than c.

de Berg et al. [25] give lower bound constructions for the number of topologically-different views of a polyhedral terrain:

Theorem 13 [25] *The maximum number of topologically-different orthographic views of a polyhedral terrain, with a total of n edges, can be as large as $\Omega(n^5\alpha(n))$. The number of topologically-different* perspective *views (views from any point in space) of such a terrain can be $\Omega(n^8\alpha(n))$.*

Thus the upper bound in Theorem 12 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 $\Theta(n^6)$ [55].

5.3 The Number of Perspective Views of a Terrain

As it turns out, if we wish to extend the above analysis to the case of perspective views, where the viewpoint can be anywhere in 3-space, we require a bound on the number of vertices of the lower envelope of the space $\mathcal{E}_K$ of all line segments that lie over K, defined in analogy with the definition of the spaces $\mathcal{L}_K$ and $\mathcal{R}_K$. Unfortunately, we do not know how to extend our analysis to obtain non-trivial bounds for the complexity of $\mathcal{E}_K$. Nevertheless, Agarwal and Sharir [3] have recently obtained an almost tight bound on the number of topologically-different perspective views of a polyhedral terrain, using a different approach, where they apply Theorem 4 to the lower envelope of certain collections of surfaces in six-dimensional space. By Theorem 13, the following bound is almost tight in the worst case:

Theorem 14 [3] *The number of topologically-different perspective views of a polyhedral terrain with n edges is $O(n^{8+\varepsilon})$, for any $\varepsilon > 0$.*

Remarks: (1) Agarwal and Sharir [3] also analyze the number of topologically-different orthographic views of a polyhedral terrain, applying Theorem 4 to different collections of functions; they obtain the slightly inferior bound $O(n^{5+\varepsilon})$, for any $\varepsilon > 0$.

(2) If K is an arbitrary polyhedral set with n edges, the maximum possible number of topologically-different perspective views of K is $\Theta(n^9)$ [55].

6 Minkowski Sums of Polyhedral Sets

Minkowski sums of geometric figures, as defined in Section 2, play a central role in various applications of geometric computing, including robot motion planning, layout design, part machining, and assembly planning [48]. We have already explained in Section 2 the connection between Minkowski sums and motion planning applications in robotics. In this section we consider a special instance of this application, involving a convex polyhedral body translating among a collection of convex polyhedral obstacles in 3-space.

As mentioned in Section 2, the planar version of this problem, involving a convex polygon translating among convex polygonal obstacles, has been solved by Kedem et al. [45]. Their results, summarized in Theorem 3, imply that the complexity of the entire free configuration space of the translating polygon is only linear. (If the translating polygon is non-convex, the complexity of the free configuration space can be quadratic in the worst case [64].) This result is clearly stronger than the single-face bound of Theorem 2, since it applies to the complexity of the entire free space. See [59] for more background information on translational motion planning for convex objects.

In three dimensions, the complexity of the free configuration space of a non-convex polyhedron with a fixed number of vertices, translating among polyhedral obstacles with a total of n vertices, can be $\Omega(n^3)$ (see, e.g., [29]). It has long been conjectured that if the moving polyhedron is convex then the complexity of the free space is at most nearly quadratic. In [59], Sharir conjectures that the actual complexity is $O(n^2\alpha(n))$, and mentions that this conjectured bound is the best possible: using Davenport-Schinzel sequences, one can construct a polyhedral setting as above, where the free space has complexity $\Omega(n^2\alpha(n))$.

However, only the trivial upper bound of $O(n^3)$ was known for the complexity of the free space in this case. The only non-trivial result in support of the conjecture involves the case where B is a ladder (line segment). This result is described in [59], and was also independently observed by Ke and O'Rourke [44]. Their bound, $O(n^2)$, is slightly better than the bound in the general conjecture.

Six years after the conjecture had been proposed, Halperin and Yap [38] proved it for the case of a translating box:

Theorem 15 [38] *The combinatorial complexity of the entire free configuration space for a box translating among polyhedral obstacles, with a total of n vertices, is $O(n^2\alpha(n))$.*

The proof of the theorem is based on several rather simple observations that enable to reduce the original problem to that of translating a *triangle* in the same environment. To solve the latter problem, the authors adapt a technique devised by Leven and Sharir [47] for the case of a convex polygon translating and rotating among polygonal obstacles in the plane. This technique, in a tricky way, reduces the problem further to a problem involving lower envelopes of univariate functions.

It is not clear whether this bound for the box is tight. Indeed, there is a lower bound of $\Omega(n^2\alpha(n))$ for the complexity of the entire free configuration space in the case of a general translating convex polyhedron, but the best lower bound known for the case of a box is only $\Omega(n^2)$ [38]. Other works related to the translational motion planning problem in three dimensions are [7, 8].

Recently, Aronov and Sharir have (almost) settled the general case of a convex polyhedron translating among polyhedral obstacles. Let $A_1, \ldots, A_k$ be k convex polyhedra in 3 dimensions with pairwise disjoint interiors, and let B be another convex polyhedron, which, with no loss of generality, is assumed to contain the origin O. Suppose A_i has q_i faces and B has p faces, and put $q = \sum_{i=1}^{k} q_i$. Let $P_i = A_i \ominus B$ be the *Minkowski sum* of A_i and $-B$, for $i = 1, \ldots, k$, and let $U = \bigcup_{i=1}^{k} P_i$ be the union of these so-called *expanded obstacles*. As is well known (and noted in Section 2), the complement C of U represents the *free configuration space* FP of B, under the purely-translational motion allowed for B, in the sense that, for each point

$z \in C$, the placement of B, for which the reference point O lies at z, does not intersect any of the obstacles A_i, and all such free placements are represented in this manner.

As is well known, the complexity of each P_i is at most $O(pq_i)$, so the sum of the complexities of the expanded obstacles is $n = O(pq)$. In practice, though, n can be expected to be much smaller than pq, usually linear in $pk + q$.

The main result of this section is:

Theorem 16 [10] *The combinatorial complexity (i.e., the number of vertices, edges, and faces on the boundary) of the union U, and thus also of FP, is $O(nk \log^2 k) = O(pqk \log^2 k)$. This bound is almost tight, as there are constructions where the union complexity is $\Omega(nk\alpha(k))$. The union can be computed by a randomized algorithm whose expected running time is $O(nk \log^3 k)$.*

This should be compared to the recent bound by Aronov and Sharir [9] on the combinatorial complexity of the union of any k convex polyhedra in 3-space with a total of n faces. It is shown there that the maximum complexity of such a union is $O(k^3 + nk \log^2 k)$ and $\Omega(k^3 + nk\alpha(k))$ in the worst case. Thus Theorem 16 shows that the convex polyhedra P_i arising in translational motion planning have special properties that yield the above improved bound, without the cubic term of the general bound.

The proof of Theorem 16 is fairly involved. It is based on the recent inductive analysis technique of [30], but it also involves a careful analysis of the topological structure of the union.

The motion planning application of Theorem 16 is obvious:

Theorem 17 [10] *The combinatorial complexity of the* entire *free configuration space for a convex polyhedron with a fixed number of faces, translating among k convex polyhedral obstacles with a total of n faces, is $O(nk \log^2 k)$ (which is almost tight in the worst case), and it can be computed in randomized expected time $O(nk \log^3 k)$.*

We conclude the presentation of new results with the case where the moving polyhedron is non-convex. As mentioned earlier, the complexity of the entire free space in that case can be $\Omega(n^3)$ in the worst case. However, it might be the case that specific instances of the problem do not result in a free configuration space with such a high complexity. For example, if the environment is not too cluttered with obstacles, the complexity of the resulting *entire* arrangement of contact surfaces can be expected to be much smaller than cubic. To take advantage of such settings, we wish to design an *output sensitive* algorithm for computing arrangements of triangles in 3-space.

In the two-dimensional case, one can solve the corresponding problem of computing arrangements of segments (or of more general arcs) in an output-sensitive manner, by using a variant of the Bentley-Ottmann algorithm for detecting intersections between line segments (or arcs) in the plane [11]. Using this procedure, it is possible to compute the *vertical decomposition* of the collection of constraint segments or arcs, induced by the motion planning instance (see, e.g., [52] for this easy extension). Once the vertical (or trapezoidal) decomposition is computed, determining whether any two given placements of the robot can be reached from each other (by a collision-free translational motion) can be easily done in logarithmic time.

Recently, de Berg *et al.* [24] have extended this result to three-dimensions (see [19, 24] for the definition of vertical decomposition in 3-dimensional arrangements):

Theorem 18 *Given a collection $\mathcal{T}$ of n triangles in general position in three-dimensional space, one can compute the vertical decomposition of the arrangement $\mathcal{A}(\mathcal{T})$ in time $O(n^2 \log n + V \log n)$, where V is the combinatorial complexity of the vertical decomposition. Moreover, for any $\varepsilon > 0$, we have $V = O(n^{2+\varepsilon} + K)$, where K is the complexity of the arrangement $\mathcal{A}(\mathcal{T})$.*

It should be noted, however, that while the algorithm is sensitive to the size of the arrangement, it does not distinguish between free cells and "forbidden" cells. In other words, the complexity parameters K and V mentioned above might be much larger than the complexity of just the 'free' portions of $\mathcal{A}(\mathcal{T})$ (the same disclaimer applies to the planar version of the problem, mentioned above).

7 Conclusion

In this paper we have surveyed a collection of recent combinatorial and algorithmic results involving arrangements of surfaces in three and higher dimensions. These results extend previous results obtained more than 5 years ago for planar arrangements, and have numerous applications in robot motion planning, visibility problems in 3-space, and many other areas.

Additional applications of these new results are given in [1, 2, 51, 62] (see also the forthcoming book [63]). They include applications to Voronoi diagrams in higher dimensions and to dynamic Voronoi diagrams, which are useful tools in many application areas, and also to various problems involving lines in 3-space, including ray shooting among spheres, and computing the width of a point set in 3-space. The most significant of these developments, which is strongly related to the results reported in Section 6, due to Chew et al. [17], shows that the complexity of the Voronoi diagram of a set of n lines in 3-space, induced by a polyhedral convex distance function, is only $O(n^2\alpha(n)\log n)$.

While the results reported in this survey constitute a considerable advancement over previously known results, and they almost settle several long-standing open problems, a lot still remains to be studied. In particular, most of the new results yield improved combinatorial bounds, but fail to produce equally efficient algorithms. For example, it is still an open problem to design an efficient algorithm (of near-quadratic complexity) for constructing a single cell in an arrangement of algebraic surfaces in 3-space. As mentioned in section 4, this problem will be solved if we can show the existence of a decomposition of the single cell into a near-quadratic number of subcells of constant description complexity. This interesting combinatorial subproblem is also open. A few initial results on the algorithmic problem for a single cell are known in three dimensions [7, 24, 36], but the authors are not aware of any result of this kind in four or higher dimensions. Equally challenging is the problem of designing efficient algorithms for computing lower envelopes in five and higher dimensions.

Another open problem that the paper raises is to further tighten the complexity bounds reported here. For example, the bound for the complexity of lower envelopes in d-dimensional arrangements is conjectured to be $O(n^{d-2}\lambda_s(n))$. Except for the bounds for hyperplanes and balls, such sharp bounds are known for lower envelopes of simplices in d-space (their complexity is $O(n^{d-1}\alpha(n)))$ [27, 53]. For a single cell in 3-dimensional arrangements, a bound of $O(n^2\alpha(n))$ is known for arrangements arising in a certain special motion planning problem [33]. It would be interesting to improve the complexity bounds for lower envelopes, at least to $O(n^{d-1}\text{polylog}(n))$, and to improve the bound for a single cell in three dimensions, at least to $O(n^2\text{polylog}(n))$, as has been done in [8] for the case of simplices. Another open problem is to extend the combinatorial bound for the complexity of a single cell to arrangements in $d > 3$ dimensions. We believe that this is doable, and are currently exploring this problem.

Acknowledgements

Work on this paper by Dan Halperin has been supported by a grant from the Stanford Integrated Manufacturing Association (SIMA), by NSF/ARPA Grant IRI-9306544, and by NSF Grant CCR-9215219. Work on this paper by Micha Sharir has been supported by NSF Grant CCR-91-22103, and by grants from the U.S.-Israeli Binational Science Foundation, the G.I.F., the German-Israeli Foundation for Scientific Research and Development, and the Israel Science Fund administered by the Israeli Academy of Sciences.

References

[1] P.K. Agarwal, B. Aronov, and M. Sharir, Computing envelopes in four dimensions with applications, *Proc. 10th ACM Symposium on Computational Geometry*, 1994, pp. 348–358.

[2] P.K. Agarwal, O. Schwarzkopf and M. Sharir, The overlay of lower envelopes and its applications, Manuscript, 1993.

[3] P.K. Agarwal and M. Sharir, On the number of views of polyhedral terrains, *Proc. 5th Canadian Conf. on Computational Geometry*, 1993, pp. 55–60. (To appear in *Discrete Comput. Geom.*)

[4] P.K. 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.

[5] B. Aronov, J. Matoušek and M. Sharir, On the sum of squares of cell complexities in hyperplane arrangements, *Proc. 7th Symp. on Computational Geometry*, 1991, pp. 307–313. (To appear in *J. Comb. Theory Ser. A.*)

[6] B. Aronov, M. Pellegrini and M. Sharir, On the zone of a surface in a hyperplane arrangement, *Discrete Comput. Geom.* 9 (1993), 177–186.

[7] B. Aronov and M. Sharir, Triangles in space, or building (and analyzing) castles in the air, *Combinatorica* 10 (1990), 137–173.

[8] B. Aronov and M. Sharir, Castles in the air revisited, *Proc. 8th ACM Symp. on Computational Geometry*, 1992, pp. 146–156. (To appear in *Discrete Comput. Geom.*)

[9] B. Aronov and M. Sharir, The union of convex polyhedra in three dimensions, *Proc. 34th IEEE Symp. on Foundations of Computer Science*, 1993, pp. 518–527.

[10] B. Aronov and M. Sharir, On translational motion planning in 3-space, *Proc. 10th ACM Symposium on Computational Geometry*, 1994, pp. 21–30.

[11] J.L. Bentley and T.A. Ottmann, Algorithms for reporting and counting geometric intersections, *IEEE Transactions on Computers* 28 (1979), 643–647.

[12] J. Bochnak, M. Coste and M-F. Roy, *Géométrie Algébrique Réelle*, Springer-Verlag, Berlin 1987.

[13] J.-D. Boissonnat and K. Dobrindt, On-line randomized construction of the upper envelope of triangles and surface patches in $\mathbb{R}^3$, Manuscript, 1993.

[14] 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.

[15] B. Chazelle, H. Edelsbrunner, L. Guibas and M. Sharir, A singly-exponential stratification scheme for real semi-algebraic varieties and its applications, *Proc. 16th International Colloquium on Automata, Languages and Programming*, 1989, pp. 179–192. Also in *Theoretical Computer Science* 84 (1991), 77–105.

[16] B. Chazelle, H. Edelsbrunner, L. Guibas, M. Sharir and J. Snoeyink, Computing a face in an arrangement of line segments and related problems, *SIAM J. Computing* 22 (1993), 1286–1302.

[17] L.P. Chew, K. Kedem, M. Sharir and E. Welzl, Voronoi diagrams of lines in 3-space, induced by a polyhedral convex distance function, Manuscript, 1994.

[18] K. Clarkson, Computing a single face in an arrangement of segments, Manuscript, 1990.

[19] K. Clarkson, H. Edelsbrunner, L. Guibas, M. Sharir and E. Welzl, Combinatorial complexity bounds for arrangements of curves and spheres, *Discrete Comput. Geom.* 5 (1990), 99–160.

[20] K. Clarkson and P. Shor, Applications of random sampling in computational geometry, II, *Discrete Comput. Geom.* 4 (1989), 387–421.

[21] R. Cole and M. Sharir, Visibility problems for polyhedral terrains, *J. Symbolic Computation* 7 (1989), 11–30.

[22] M. de Berg, Private communication, 1993.

[23] M. de Berg, K. Dobrindt, and O. Schwarzkopf, On lazy randomized incremental construction, *Proc. 26th ACM Symposium Theory Comput.*, 1994, pp. 105–114.

[24] M. de Berg, L.J. Guibas and D. Halperin, Vertical decompositions for triangles in 3-space, *Proc. 10th ACM Symposium on Computational Geometry*, 1994, pp. 1–10.

[25] M. de Berg, D. Halperin, M. Overmars and M. van Kreveld, Sparse arrangements and the number of views of polyhedral scenes, Manuscript, 1991.

[26] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg, 1987.

[27] H. Edelsbrunner, The upper envelope of piecewise linear functions: Tight complexity bounds in higher dimensions, *Discrete Comput. Geom.* 4 (1989), 337–343.

[28] H. Edelsbrunner, L.J. Guibas, J. Pach, R. Pollack, R. Seidel and M. Sharir, Arrangements of curves in the plane – Topology, combinatorics, and algorithms, *Theoretical Computer Science* 92 (1992), 319–336.

[29] H. Edelsbrunner, L.J. Guibas and M. Sharir, The upper envelope of piecewise linear functions: algorithms and applications, *Discrete Comput. Geom.* 4 (1989), 311–336.

[30] H. Edelsbrunner, R. Seidel and M. Sharir, On the zone theorem in hyperplane arrangements, *SIAM J. Computing* 22 (1993), 418–429.

[31] L. J. Guibas, L. Ramshaw, and J. Stolfi, A kinetic framework for computational geometry, In *Proc. 24th Annu. IEEE Sympos. Found. Comput. Sci.*, 1983, pp. 100–111.

[32] L.J. Guibas, M. Sharir and S. Sifrony, On the general motion planning problem with two degrees of freedom, *Discrete Comput. Geom.* 4 (1989), 491–521.

[33] D. Halperin, On the complexity of a single cell in certain arrangements of surfaces related to motion planning, *Discrete Comput. Geom.* 11 (1994), 1–33.

[34] D. Halperin, *Algorithmic Motion Planning via Arrangements of Curves and of Surfaces,* Ph.D. Dissertation, Computer Science Department, Tel Aviv University, July 1992.

[35] D. Halperin and M. Sharir, New bounds for lower envelopes in three dimensions, with applications to visibility in terrains, *Proc. 9th ACM Symp. on Computational Geometry*, 1993, pp. 11–18. (To appear in *Discrete Comput. Geom.*)

[36] 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, pp. 382–391.

[37] D. Halperin and M. Sharir, Almost tight upper bounds for the single cell and zone problems in three dimensions, *Proc. 10th ACM Symposium on Computational Geometry*, 1994, pp. 11–20.

[38] D. Halperin and C.-K. Yap, Combinatorial complexity of translating a box in polyhedral 3-space, *Proc. 9th ACM Symposium on Computational Geometry*, 1993, pp. 29–37.

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

[40] R. Hartshorne, *Algebraic Geometry*, Springer-Verlag, New York 1977.

[41] 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 the DIMACS Special Year* (J.E. Goodman, R. Pollack and W. Steiger, Eds.), AMS Press, Providence, RI 1991, pp. 137–163.

[42] J. Hershberger, Finding the upper envelope of n line segments in $O(n \log n)$ time, *Information Processing Letters* 33 (1989), 169–174.

[43] M.W. Hirsch, *Differential Topology*, Springer-Verlag, New York 1976.

[44] Y. Ke and J. O'Rourke, An algorithm for moving a ladder in three dimensions, *Technical Report* JHU-87/17, Department of Computer Science, Johns Hopkins University, 1987.

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

[46] J.-C. Latombe, *Robot Motion Planning*, Kluwer Academic Publishers, Boston, 1991.

[47] D. Leven and M. Sharir, On the number of critical free contacts of a convex polygonal object moving in 2-D polygonal space, *Discrete Comput. Geom.* 2 (1987), 255–270.

[48] T. Lozano-Pérez, Spatial planning: A configuration space approach, *IEEE Trans. Comput.* C-32 (1983), pp. 108–120.

[49] T. Lozano-Perez and M. Wesley, An algorithm for planning collision-free paths among polyhedral obstacles, *Comm. ACM* 22 (1979), 560–570.

[50] P. McMullen, The maximum number of faces of a convex polytope, *Mathematika* 17 (1970), 179–184.

[51] S. Mohaban and M. Sharir, Ray shooting amidst spheres in 3 dimensions and related problems, Manuscript, 1993.

[52] K. Mulmuley, *Computational Geometry: An Introduction Through Randomized Algorithms*, Prentice Hall, New York, 1993.

[53] 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.

[54] M. Pellegrini, On lines missing polyhedral sets in 3-space, *Proc. 9th ACM Symp. on Computational Geometry*, 1993, pp. 19–28.

[55] H. Plantinga and C. Dyer, Visibility, occlusion, and the aspect graph, *International J. Computer Vision*, 5 (1990), 137–160.

[56] R. Pollack and M.F. Roy, On the number of cells defined by a set of polynomials, *C.R. Acad. Sci. Paris*, t. 316, Série I (1993), 573–577.

[57] F.P. Preparata and M.I. Shamos, *Computational Geometry—An Introduction*, Springer Verlag, New York, 1985.

[58] J.T. Schwartz and M. Sharir, On the two-dimensional Davenport Schinzel problem, *J. Symbolic Computation* 10 (1990), 371–393.

[59] M. Sharir, Efficient algorithms for planning purely translational collision-free motion in two and three dimensions, *Proceedings of the IEEE International Conference on Robotics and Automation*, Raleigh, North Carolina, 1987, pp. 1326–1330.

[60] M. Sharir, Davenport-Schinzel sequences and their geometric applications, *Theoretical Foundations of Computer Graphics and CAD*, (R.A. Earnshaw, Ed.), NATO ASI Series, F40, Springer Verlag, Berlin, 1988, pp. 253–278.

[61] M. Sharir, On k-sets in arrangements of curves and surfaces, *Discrete Comput. Geom.* 6 (1991), 593–613.

[62] M. Sharir, Almost tight upper bounds for lower envelopes in higher dimensions, *Proc. 34th IEEE Symp. on Foundations of Computer Science*, 1993, pp. 498–507. (To appear in *Discrete Comput. Geom.*)

[63] M. Sharir and P.K. Agarwal, *Davenport Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, to appear.

[64] S. Sifrony and M. Sharir, A new efficient motion planning algorithm for a rod moving in two-dimensional polygonal space, *Algorithmica* 2 (1987), 367–402.