

Almost Tight Upper Bounds for the Single Cell and Zone Problems in Three Dimensions*

D. Halperin¹ and M. Sharir²

¹ Robotics Laboratory, Department of Computer Science,
Stanford University, Stanford, CA 94305, USA
danha@flamingo.stanford.edu

² School of Mathematical Sciences, Tel Aviv University,
Ramat Aviv 69 978, Israel
sharir@math.tau.ac.il

and

Courant Institute of Mathematical Sciences, New York University,
New York, NY 10012, USA

Abstract. We consider the problem of bounding the combinatorial complexity of a single cell in an arrangement of n low-degree algebraic surface patches in 3-space. We show that this complexity is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε and on the maximum degree of the given surfaces and of their boundaries. This extends several previous results, almost settles a 9-year-old open problem, and has applications to motion planning of general robot systems with three degrees of freedom. As a corollary of the above result, we show that the overall complexity of all the three-dimensional cells of an arrangement of n low-degree algebraic surface patches, intersected by an additional low-degree algebraic surface patch σ (the so-called *zone* of σ in the arrangement) is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε and on the maximum degree of the given surfaces and of their boundaries.

*Work on this paper by the first author has been supported by a Rothschild Postdoctoral Fellowship, 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 the second author has been supported by NSF Grants CCR-91-22103 and CCR-93-11127, 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.

1. Introduction

Let $\Sigma = \{\sigma_1, \dots, \sigma_n\}$ be a given collection of n low-degree algebraic surface patches in 3-space (see below for a more precise statement of the properties that these surfaces are assumed to satisfy). We denote by $\mathcal{A}(\Sigma)$ the *arrangement* of Σ , i.e., the decomposition of 3-space into (relatively open) cells of various dimensions, each being a maximal connected set contained in the intersection of a fixed subcollection of Σ and not meeting any other surface; we denote j -dimensional cells of $\mathcal{A}(\Sigma)$, for $j = 0, 1, 2$, as *vertices, edges, and faces*, respectively, and the unquantified term *cell* is used to denote three-dimensional cells of $\mathcal{A}(\Sigma)$. The *combinatorial complexity* of a cell C is the number of lower-dimensional cells appearing on its boundary. The problem studied in this paper is to obtain a sharp upper bound on the combinatorial complexity of a single cell in such an arrangement.

One of the main motivations for studying this problem is its applications to *robot motion planning*. Let B be an arbitrary robot system with three degrees of freedom, moving in some environment V filled with obstacles. Any placement of B can be represented by a point in 3-space, whose coordinates are the three parameters controlling the degrees of freedom of B ; this space is called the *configuration space* of B . We want to compute the *free* portion of this space, denoted as FP , and consisting of those placements of B at which it does not meet any obstacle. We note that the boundary of FP consists of placements at which B makes contact with some obstacles, 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, all algebraic of constant maximum degree (and whose relative boundaries are also algebraic of constant maximum degree).

For example, if B is an arbitrary polygonal object with k sides, and V is an open planar polygonal region bounded by m edges, the configuration space of B is a three-dimensional space, each point of which represents a possible placement of B by the parametrization $(x, y, \tan(\theta/2))$, where (x, y) are the coordinates of some fixed reference point on B , and θ is the orientation of B . In this case each “contact surface” is 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 . Each of the resulting $O(km)$ contact surfaces is a two-dimensional algebraic surface patch of degree at most four, 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 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 nonfree (assuming, as is customary, that the boundaries of B and V lie in the closure of their interiors, and that B and V lie in *general position*). 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. (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 is a major problem that has to be tackled, prior to the design of efficient algorithms for computing such a cell.

Here is a brief history of the single-cell problem. In two dimensions it has been shown in [14] that the complexity of a single face in an arrangement of n Jordan arcs, each pair of which intersect in at most some constant number, s , of points, is $O(\lambda_{s+2}(n))$, where $\lambda_q(m)$ is the maximum length of *Davenport–Schinzel sequences* of order q composed of m symbols, and is nearly linear in m for any fixed q (see [1] and [19] for more details). Thus the maximum complexity of a single face is nearly linear in the number of arcs (for any fixed s), as opposed to the complexity of the entire arrangement of the arcs, which can be quadratic in the worst case. Efficient algorithms for computing a single face in a two-dimensional arrangement are given in [14] and [8].

In higher dimensions a prevailing conjecture (see, e.g., [23]) is that the complexity of a single cell in an arrangement $\mathcal{A}(\Sigma)$ as above is at most only slightly larger than $O(n^{d-1})$, which is again roughly “one order of magnitude” smaller than the maximum complexity of the entire arrangement, which can be $\Theta(n^d)$ (see [24]). A stronger version of the conjecture asserts that the maximum complexity of a single cell 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.

These conjectures have been proved only for a few special cases of arrangements. They are largely open in the general case stated above. In fact, no bounds better than $O(n^d)$ are known for the general case, even in three dimensions. The special cases for which better bounds are known include the case of hyperplanes, where the complexity of a single cell, being a convex polytope bounded by at most n hyperplanes, is $O(n^{\lfloor d/2 \rfloor})$ (by the Upper Bound Theorem [22]), the case of spheres, where an $O(n^{\lfloor d/2 \rfloor})$ bound is easy to obtain by lifting the spheres into hyperplanes in $(d+1)$ -space [11], [25], the case of $(d-1)$ -simplices, where an $O(n^{d-1} \log n)$ bound has been recently established in [5], and several special cases in three dimensions that arise in motion planning for various specific robot systems B with three degrees of freedom, including the case of a moving polygon mentioned above, where an $O(n^{2+\varepsilon})$ bound is proved in [17], some restricted cases of that problem, where slightly more improved bounds are obtained [15], [16], and a few other systems (see [15]).

The single-cell problem is a generalization of the related problem of bounding the complexity of the *lower envelope* of Σ , i.e., the portion of the union of the surfaces of Σ , consisting of those points w for which no surface of Σ passes below w . This problem, also rather difficult, is nevertheless easier to analyze, and recent results [18], [27] show that the combinatorial complexity of such an envelope is $O(n^{d-1+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε , d , and the maximum algebraic degree of the given surfaces and of their relative boundaries.

In this paper we derive an improved upper bound for the complexity of a single cell in an arrangement $\mathcal{A}(\Sigma)$ of algebraic surfaces in 3-space, as above. This bound, in three dimensions, is the same as the bound for lower envelopes just mentioned; that is, it is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends, as above, on ε and on the maximum algebraic degree of the given surfaces and of

their relative boundaries. This almost establishes the conjecture mentioned above, in three dimensions.

Our analysis adapts the proof technique of [17], which in turn is based on the analysis technique of [18] and [27] for the case of lower envelopes. The lesson that can be learned from the analysis in [17] is that in the case of a single cell the following two preliminary results are necessary 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 neighborhood of the vertex within the closure of C).
- (b) A sharp bound on the number of vertices bounding “popular” faces of C (faces that are adjacent to C on both “sides”; see [3], [5], [17], and below).

Bounds on these quantities were obtained in [17] using special properties of the surfaces that arise in the case studied there. A main technical contribution of the present paper 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., [21]), but are simpler to derive in three dimensions. The bound (b) is obtained by applying the new probabilistic technique of [17], [18], and [27] to counting only the vertices of popular faces (this idea is in the spirit of the methodology used in [3] and [5]). Once these two bounds are available, the rest of the proof is rather similar to those used in [17], [18], and [27], although certain additional nontrivial adjustments are required.

The paper is organized as follows. In Section 2 we give several preliminary results, including the analysis of the number of x -extreme vertices of a single cell. The main analysis is presented in Section 3. The application of the main result to the *zone* problem, as mentioned in the abstract, is discussed in Section 4, and the paper concludes in Section 5 with further applications of our results and some open problems.

2. Preliminaries

Let $\Sigma = \{\sigma_1, \dots, \sigma_n\}$ be a given collection of n surface patches in 3-space that satisfy the following conditions:

- (i) Each σ_i is monotone in the xy -direction (that is, every vertical line intersects σ_i in at most one point). Moreover, each σ_i is a portion of an algebraic surface of constant maximum degree b .
- (ii) The vertical projection of σ_i onto the xy -plane is a planar region bounded by a constant number of algebraic arcs of constant maximum degree (say, b too).
- (iii) The relative interiors of any triple of the given surfaces intersect in at most s points (by Bezout’s theorem [20] and by property (iv) below, we always have $s \leq b^3$).
- (iv) The surface patches in Σ are in *general position*; one way of defining this is to require that the coefficients of the polynomials defining the surfaces and

their boundaries, are algebraically independent over the rationals (i.e., no multivariate polynomial with rational coefficients vanishes when substituting into it some of the given coefficients), thereby excluding all kinds of “degenerate” configurations; see [17], [18], and [27] for more details.

We remark that the somewhat restrictive condition (iv) and the first part of condition (i) are not essential for the analysis. If the first part of condition (i) does not hold, we can decompose each surface into a constant number of xy -monotone pieces by cutting it along the locus of points with z -vertical tangency. If condition (iv) does not hold, we can argue, by applying, as in [27], a small random perturbation of the polynomials, that the complexity of a single cell in a degenerate arrangement of surfaces is at most proportional to the worst-case complexity of a single cell in arrangements of surfaces in general position.

2.1. The Number of x -Extreme Vertices

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 ; by definition, C is an open set in \mathbb{R}^3 . Recall the following definition (already mentioned above):

Definition 2.1. An x -extreme vertex v of the cell C is a vertex whose x -coordinate is the smallest or largest in the closure of some connected component of $N \cap C$, where N is a sufficiently small ball centered at v .

For each $x_0 \in \mathbb{R}$, let π_{x_0} denote the plane $x = x_0$.

Definition 2.2. A point $w \in \partial C$ is said to be *critical* if there is a neighborhood N of w and a connected component K of $C \cap N$ so that $K \cap \pi_{x_0}$ is disconnected, where x_0 is the x -coordinate of w , but $K \cap \pi_x$ is connected either for all $x < x_0$ or for all $x > x_0$ sufficiently close to x_0 .

Remark. This definition is a special case of the definition of critical points (of the coordinate function x) in Morse theory [21]. Another distinction is that the classical Morse theory applies to smooth manifolds, whereas here ∂C is generally non-smooth.

The main result of this section is:

Theorem 2.3. *The number of x -extreme vertices of C is $O(n^2)$.*

Proof. We first claim that the number of x -extreme vertices of C is proportional to one plus the number of critical points of C . We prove this by an argument borrowed from [6]. We then show that the number of critical points of C is $O(n^2)$ and, by that, complete the proof of the theorem. These two steps are achieved, respectively, in the two following lemmas.

Lemma 2.4. *The number of x -extreme vertices of C is proportional to one plus the number of critical points of C .*

Proof. We sweep C by moving the plane π_x in the direction of increasing x , and keep track of the number I of connected components of $C \cap \pi_x$. This number is initially (at $x = -\infty$) $O(n^2)$, since this is an upper bound on the overall complexity of any planar cross section of the entire arrangement $\mathcal{A}(\Sigma)$. The number I increases by one when π_x sweeps through a local x -minimum of C , or when a connected component of $C \cap \pi_x$ splits into two subcomponents; I decreases by one when π_x sweeps through a local x -maximum of C , or when two components of $C \cap \pi_x$ merge into a single component. (The general position assumption implies that only two components can merge into, or split from, a component of C at any given x .) The number of events at which components can split or merge is equal, by definition, to the number Q of critical points of C .

Consider the following dynamic scheme for assigning weights to components of $C \cap \pi_x$. Initially, at $x = -\infty$, we assign weight -1 to each component of $C \cap \pi_x$. When π_x sweeps through a local x -minimum point of C , a new component of $C \cap \pi_x$ is created, and is also assigned weight -1 . When two components of $C \cap \pi_x$ merge, we assign to the new component, weight equal to 2 plus the sum of the weights of the merged components. When a component shrinks and disappears, its final weight is added to a global count M . When a component is split into two subcomponents, each of them is assigned weight $1 + w/2$, where w is the weight of the split component.

We claim that, at any given time during the sweep, the weight of any component of $C \cap \pi_x$ is always at least -1 , and the weight of a component that was formed by one or more preceding splitting and merging operations is nonnegative. Both claims are easy to prove by induction on the sweep events. If C has no critical points, then it has at most one local x -minimum and at most one local x -maximum, and the claim holds trivially in this case; so suppose C does have critical points. In this case it is easily verified, using induction on the sweep events, that the weight of each component of $C \cap \pi_x$ that shrinks to a point as we reach a local x -maximum of C is nonnegative, so the value of M is always nonnegative. Similarly, all components that survive as x reaches $+\infty$ have nonnegative weight.

Suppose that, at some point during the sweep, there are s_x local x -minima of C to the left of π_x , that $C \cap \pi_{-\infty}$ has t components, that the number of splittings and mergings of cross-sectional components to the left of π_x is $Q_x \leq Q$, and that the current value of the count M is M_x . Then, as is easily verified by the definition of weights and by induction on the sweep events, the total weight of the components of $C \cap \pi_x$ plus M_x is equal to $2Q_x - s_x - t$. Hence, since at $x = +\infty$ the total weight of the components of $C \cap \pi_x$ plus M_x is nonnegative, we have $2Q - s - t \geq 0$, where s is the total number of local x -minima of C . This implies that $s \leq 2Q$. A symmetric argument applies to the number of local x -maxima of C , and thus the claim is established. \square

Lemma 2.5. *The number of critical points of C is $O(n^2)$.*

Proof. For any fixed surface $\sigma \in \Sigma$, the general position assumption is easily seen to imply that there are only $O(1)$ critical points that lie only on σ and on no other surface, so the total number of such points, over all $\sigma \in \Sigma$, is only $O(n)$. The number of critical points that lie on the boundary of one surface of Σ and on a second surface, summed over all pairs of such surfaces, is only $O(n^2)$. This follows from Bezout’s theorem applied to the number of intersections between an algebraic surface and an algebraic arc, in combination with our basic assumptions (i), (ii), and (iv) about the surface patches in Σ .

Consider next critical points that lie in the relative interior of an intersection curve $\gamma_{ij} = \sigma_i \cap \sigma_j$, for some pair of surfaces $\sigma_i, \sigma_j \in \Sigma$ and on no other surface of Σ . If such a point w is a singular point¹ on, say σ_i , then w is an intersection point between σ_j and the curve of singular points on σ_i , and, by the general position assumption, the number of such points, over all pairs $\sigma_i, \sigma_j \in \Sigma$, is clearly $O(n^2)$. We thus may assume that w is nonsingular on both surfaces, and, by the general position assumption, that σ_i and σ_j meet transversally at w . However, then the criticality of w is easily seen to imply that the tangent vector to γ_{ij} at w must be orthogonal to the x -axis, and the number of points on γ_{ij} with this property is $O(1)$ (under the general position assumption). Hence the total number of such points, over all pairs $\sigma_i, \sigma_j \in \Sigma$, is also $O(n^2)$.

Finally, suppose that w is a critical point that is also a point of intersection of three surfaces $\sigma_i, \sigma_j, \sigma_k \in \Sigma$. Arguing as above, we can rule out the case where w is singular on either of these surfaces. Consider now the three intersection curves $\gamma_i = \pi_{x_0} \cap \sigma_i, \gamma_j = \pi_{x_0} \cap \sigma_j, \gamma_k = \pi_{x_0} \cap \sigma_k$, where x_0 is the x -coordinate of w . These curves meet at w , they are all smooth at w , and $K \cap \pi_{x_0}$ lies on a single side of each of the curves. If any two of these curves, say γ_i and γ_j , are tangent to each other at w , then, as is easily verified, the tangent at w to the curve $\sigma_i \cap \sigma_j$ is orthogonal to the x -axis, and, as argued above, the number of such points w is only $O(n^2)$. Otherwise, $K \cap \pi_{x_0}$ must be fully contained, locally near w , in just one of the six regions into which these curves split π_{x_0} locally near w (see Fig. 1). However, this is easily seen to contradict the criticality of w , and thus implies that the total number of critical points of C is $O(n^2)$, as asserted. \square

This also completes the proof of Theorem 2.3. \square




Fig. 1. Transversal nonsingular intersection of three surfaces does not give rise to a critical point.

¹ A point v is *singular* on an algebraic surface in \mathbb{R}^d , defined by $P = 0$, for some polynomial P , if all the partial derivatives $\partial P / \partial x_i$, for $i = 1, \dots, d$, vanish at v ; see [7] and [20].

2.2. Inner Vertices, Sides, and Borders

For the analysis in the following sections, we need to introduce additional terminology relating to cells of various dimensions in the arrangement. We call a vertex v of $\mathcal{A}(\Sigma)$ an *inner vertex*, if v is formed by the intersection of the relative interiors of three distinct surfaces of Σ . Let v be an inner vertex of $C = C_Z(\Sigma)$, which is incident to three surfaces σ_1 , σ_2 , and σ_3 that meet transversally at v , and which is not a singular point on any of these surfaces. For technical reasons, we distinguish between different *sides* of v , adapting the notation of [17] (see also [3], [5], and [13]). Formally, the three tangent planes to the surfaces σ_i at v partition 3-space into eight octants, and a side R of v is any one of these octants. We call the pair (v, R) a *0-border*. We say that (v, R) is a 0-border of C if, when we move from v in any direction that points into R by any sufficiently small distance, we enter C . We will be counting the number of inner 0-borders of C , which means that we count each vertex v of C with multiplicity, once for each side of v that lies in C (in the above sense). We define $\kappa(\Sigma)$ to be the number of inner 0-borders on $\partial C_Z(\Sigma)$. We also denote by $\kappa(n)$ the maximum possible value of $\kappa(\Sigma)$, taken over all collections Σ , as above, with a total of n surfaces, and over all cells of $\mathcal{A}(\Sigma)$.

In the following analysis the notion of a side also needs to be extended to edges and faces bounding C . For an edge e , formed by the intersection of two surfaces σ , σ' , we can assume (by reasons similar to those used in the preceding arguments) that no point on e is singular on either of these surfaces, and that σ and σ' cross each other transversally at each point of e . Then at each point $z \in e$, the plane normal to e at z is split by the two tangent planes to σ , σ' at z into four quadrants. A side of e can be thought of as a *continuous* mapping (in the Hausdorff sense) that maps each point $z \in e$ to one of the quadrants at z (or, rather, to make the Hausdorff continuity well defined, to the intersection of such a quadrant with the unit ball around z). Similarly, a side of a face f can be defined as a continuous mapping from each point $z \in f$ to one of the two unit vectors normal to f at z . If a vertex v is incident to an edge e and lies on another surface σ crossing e transversally, then a side R of v is *consistent* with a side Q of e if the limit of $Q(z)$, as z approaches v , is contained in (the closure of) R , and R is the positive cone spanned by the limit of $Q(z)$ and by the vector tangent to e at z and pointing from z toward e . Consistency between sides of a vertex and an incident face, or between sides of an edge and an incident face, can be defined in a similar manner. If v is a nonsingular vertex of $\mathcal{A}(\Sigma)$, incident to two edges e_1, e_2 , which are contained in the same intersection curve of a pair of surfaces, we say that a side R_1 of e_1 is consistent with a side R_2 of e_2 if the two limits of $R_1(z)$ as z approaches v along e_1 , and of $R_2(z)$ as z approaches v along e_2 , coincide. If R'_1, R'_2 are the two sides of v consistent with R_1, R_2 , respectively, we say that R'_2 is the side of v *opposite* to R'_1 across the third surface defining v ; see Fig. 2. Given an edge e and a side R of e , we say that (e, R) is a *1-border* of C if, when we move from any point $z \in e$ in a direction contained in $R(z)$, we enter C . Similarly, we can define 2-borders (f, R) of C , for a face f and a side R of f .




Fig. 2. R_i is the side of e_i facing the reader, for $i = 1, 2$; the sides R'_1 and R'_2 of v consistent with R_1 and R_2 , respectively, are also facing the reader and are opposite one another across the surface σ_3 .

3. Complexity of a Single Cell

We concentrate on bounding the number of inner vertices of $C_Z(\Sigma)$, and later justify the use of this reduced measure of complexity. Our main result is:

Theorem 3.1. *The number of inner vertices of $C_Z(\Sigma)$ is $O(n^{2+\epsilon})$, for any $\epsilon > 0$, where the constant of proportionality depends on ϵ and on the maximum degree and shape of the surfaces and of their relative boundaries.*

This result is proved in three stages. Sections 3.1 and 3.2 are each dedicated to analyzing the complexity of a different type of inner vertices. This analysis yields recurrences that are then solved in Section 3.3 to give the asserted bound. In Section 3.4 we argue that Theorem 3.1 implies a similar bound on the total complexity of a single cell.

3.1. The Number of Inner Vertices

Let v be an inner vertex of $C = C_Z(\Sigma)$, which is incident to three surfaces σ_1, σ_2 , and σ_3 . We assume that v is not a singular point on any of these surfaces, and that these surfaces meet transversally at v . If, say, v is a singular point of σ_1 , then it lies on the algebraic curve of bounded degree consisting of all singular points on σ_1 . The number of intersection points of this curve with, say σ_2 , is constant (under the assumption of general position), which implies that the number of such vertices is only $O(n^2)$. The assumption concerning transversality is also justified by the general position assumption.

Let (v, R) be an inner 0-border of C (nonsingular, formed by the transversal intersection of its three incident surfaces, say $\sigma_1, \sigma_2, \sigma_3$). The corresponding vertex v is incident to (at least) three edges of C , which we denote by e_{12}, e_{13} , and e_{23} , where each e_{ij} is a portion of the corresponding intersection curve $\gamma_{ij} = \sigma_i \cap \sigma_j$, for $1 \leq i < j \leq 3$; moreover, each edge e_{ij} has a side R_{ij} which is consistent with R , so that (e_{ij}, R_{ij}) is a 1-border of C . If one of these curves, say γ_{12} , contains two edges, e_{12}, e'_{12} , with respective sides R_{12}, R'_{12} , such that e_{12} and e'_{12} have v as a common endpoint, R_{12} and R'_{12} are consistent with each other, and both (e_{12}, R_{12}) and (e'_{12}, R'_{12}) are 1-borders of C , then, as is easily seen, there is a face f on σ_3 which is




Fig. 3. The popular face f borders the cell C on both its sides.

incident to v and which forms with both its sides 2-borders of C . We call such a face a *popular* face of C , borrowing a notation from [3] and [5]; see Fig. 3. We denote by $\pi(\Sigma)$ the maximum number of inner vertices of popular faces bounding a single cell of $\mathcal{A}(\Sigma)$, and let $\pi(n)$ denote the maximum of $\pi(\Sigma)$, over all collections Σ of n surfaces as above (with the same s and b). (Strictly speaking, a vertex v can be incident to more than one popular face, in which case we count it in $\pi(\Sigma)$ with multiplicity, once for every incident popular face.)

A major novel ingredient of the proof is the derivation of a sharp upper bound on the number of vertices of popular faces of C ; in the previous paper [17] such a bound was derived using special properties of the surfaces that arose in the specific motion-planning application that was studied there; here we apply a new technique for obtaining the desired bound in general arrangements. First, in the rest of this subsection, we obtain an upper bound for the complexity $\kappa(n)$ in terms of the function π , and then, in the next subsection, proceed to derive an upper bound for $\pi(n)$.

Thus, up to an additive term of $\pi(n)$, it suffices, for the bound on $\kappa(n)$ that we seek, to consider only inner vertices v (or, rather, inner 0-borders (v, R)) which are not incident to any popular face of C .

Let (v, R) be a 0-border, and let us continue to follow the notations introduced above. For each $1 \leq i < j \leq 3$, the curve γ_{ij} must contain a maximal relatively open x -monotone connected portion β_{ij} having v as an endpoint, such that the 1-border (β_{ij}, R'_{ij}) , where R'_{ij} is the side of β_{ij} consistent with R_{ij} , is disjoint from C . Let z_{ij} denote the other endpoint of β_{ij} .

We define the *index* of v , denoted $j(v)$, to be the number of points of intersections of $\sigma_1, \sigma_2, \sigma_3$ which lie to the right of v (i.e., with $x > x(v)$). Clearly, $0 \leq j(v) \leq s - 1$.

We define $\kappa^{(j)}(\Sigma)$, for $j = 0, \dots, s - 1$, to be the maximum number of 0-borders (v, R) of any fixed cell of $\mathcal{A}(\Sigma)$, whose vertices v are inner vertices of index at most j . We also define $\kappa^{(j)}(n)$ to be the maximum possible value for $\kappa^{(j)}(\Sigma)$, over all collections Σ of n surface patches satisfying conditions (i)–(iv) (with the same b and s). Similarly, we define $\pi^{(j)}(\Sigma)$, for $j = 0, \dots, s - 1$, to be the maximum number of vertices with index at most j of all popular faces bounding any fixed cell of $\mathcal{A}(\Sigma)$, where each such vertex is counted with multiplicity, once for every incident popular face. We also define $\pi^{(j)}(n)$ to be the maximum possible value for $\pi^{(j)}(\Sigma)$, over all collections Σ of n surface patches satisfying conditions (i)–(iv) (with the same b and s).

Our method is to derive a recurrence relationship for $\kappa(n)$, by bounding each of the functions $\kappa^{(j)}$ in terms of $\kappa^{(j-1)}$ (with a special handling of $\kappa^{(0)}$); the solution of the resulting system of recurrences will yield the asserted bounds. Note that we are in fact looking for a bound on the quantity $\kappa^{(s-1)}(n)$, as each vertex in the arrangement is of index at most $s - 1$. The remainder of this subsection is devoted to proving the following:

Lemma 3.2. *For each $j = 0, \dots, s - 1$, and a parameter $\xi \leq n$ we have*

$$\kappa^{(j)}(n) = O\left(\xi^2 \kappa\left(\frac{n}{\xi}\right) + \xi^3 \pi^{(j)}\left(\frac{n}{\xi}\right) + \xi^3 \kappa^{(j-1)}\left(\frac{n}{\xi}\right) + n^2 + \pi^{(j)}(n)\right),$$

where we put $\kappa^{(j-1)} = 0$ when $j = 0$.

Proof. We fix $0 \leq j \leq s - 1$, and assume that the vertex v under consideration has index at most j . First, there are only at most $\pi^{(j)}(n)$ such vertices that are incident to popular faces. The other vertices can be classified into several categories. The first two cases, (a) and (b) below, are easy to charge directly, and the total number of vertices that fall into these categories is shown to be $O(n^2)$. The difficulties arise when the endpoints z_{ij} of the three arcs β_{ij} all lie on the boundary of the cell C . A more involved charging scheme is needed in these cases (resulting in the more involved terms in the recurrence): In case (c) we handle the situation where at least one of the arcs β_{ij} is not intersected by the third surface (say, β_{12} is not intersected by σ_3). The remaining case (d) handles the situation where all arcs β_{ij} are each intersected by the third surface. Each of the cases (c) and (d) is further divided into subcases according to certain parameters that are introduced in the analysis.

In more detail, we assume that the vertex v under consideration is not incident to any popular face, and consider the following cases:

Case (a). All three arcs β_{ij} emerge from v in the direction of increasing x (or all emerge in the direction of decreasing x). In this case v is an x -extreme vertex of C , as is easily checked (using the general position assumption), and Theorem 2.3 implies that the total number of such vertices (and corresponding 0-borders) is $O(n^2)$. We thus assume in what follows that at least one of these arcs emerges from v in the direction of increasing x , and at least one arc emerges in the direction of decreasing x .

Case (b). At least one of the arcs β_{ij} ends at a point z_{ij} which is either an endpoint of the original intersection curve γ_{ij} , or a point of local x -extremum on that curve. We then charge (v, R) to the point z_{ij} , and note that the number of such points is $O(n^2)$, and that each such point is charged only a constant number of times in this manner (e.g., along β_{ij} it can be charged at most once for every side of β_{ij}), thus implying that the number of 0-borders (v, R) of this kind is only $O(n^2)$. Again, in what follows we assume that this situation does not arise, which means, in particular, that each of the three endpoints z_{ij} is a vertex of C ; more precisely, each z_{ij} has a side R_{ij}^* which lies across the third surface defining z_{ij} from a side consistent with R'_{ij} , such that (z_{ij}, R_{ij}^*) is a 0-border of C . See Fig. 4.




Fig. 4. The dotted faces on the surface σ_1 appear on the boundary of C , and near both of them C lies, say, above σ_1 .

Case (c). At least one of the arcs β_{ij} , say for definiteness β_{12} , is such that $\beta_{12} \cup \{z_{12}\}$ is not intersected by the third surface σ_3 . Define the *level* of a point w in 3-space to be the smallest number of surfaces of Σ , whose removal makes w belong to the closure of the cell containing Z , in the resulting subarrangement. If w is a vertex of $\mathcal{A}(\Sigma)$ and R is a side of w , we say that w (resp. (w, R)) lies at *restricted level* $\lambda(w) = k$ (resp. $\lambda((w, R)) = k$) if by removing k surfaces from Σ , none of which is incident to w , we make w a vertex (resp. make (w, R) a 0-border) of the cell containing Z in the resulting subarrangement, and if k is the smallest number with that property.

Let t denote the number of distinct surfaces of Σ that intersect $\beta_{12} \cup \{z_{12}\}$. We fix some threshold parameter $\xi = \xi_j$, to be defined later (we use a different parameter for each j), and consider the following two subcases:

Case (c(i)): $t \geq \xi$. In this case we charge (v, R) to a block of ξ points of intersection between $\beta_{12} \cup \{z_{12}\}$ and the surfaces of Σ , defined as follows. For each surface σ intersecting $\beta_{12} \cup \{z_{12}\}$, choose its point of intersection that lies nearest to v along β_{12} . We obtain at least ξ such designated points, and we charge (v, R) to the block of the first ξ designated points, in their order along β_{12} from v . All those points are inner vertices of $\mathcal{A}(\Sigma)$, and it is clear that none of these vertices can be charged in this manner more than a constant number of times (along β_{12} each such vertex can be charged at most twice for each side of β_{12}). By construction, each of the charged vertices lies at restricted level at most ξ , as is easily verified. Our goal is thus to obtain an upper bound for the number M of inner vertices of $\mathcal{A}(\Sigma)$ that lie at restricted level $\leq \xi$; the number of 0-borders in the present subcase is $O(M/\xi)$.

For this we apply the probabilistic analysis technique of [9] and [26], in the same manner as in [17]. That is, we choose a random sample \mathcal{R} of $r = n/\xi$ surfaces² of Σ , and construct the arrangement $\mathcal{A}(\mathcal{R})$. Let w be an inner vertex of $\mathcal{A}(\Sigma)$ at restricted level $\lambda \leq \xi$, and let \mathcal{C} be a specific collection of λ surfaces, none incident

² Here, and in similar arguments given below, we use the sloppy notation n/ξ , instead of the more accurate value $\lceil n/\xi \rceil$. This is done for clarity of exposition, and does not affect in any significant way the bounds that we derive.

to w , whose removal makes w a vertex of the cell containing Z . The probability that w shows up as a vertex of $C_Z(\mathcal{A})$ is at least $\binom{n-\lambda-3}{r-3} / \binom{n}{r}$: out of the total number $\binom{n}{r}$ of possible samples \mathcal{A} , consider those samples that contain the three surfaces forming w and do not contain any of the λ surfaces of \mathcal{B} ; for each of these samples (and possibly for other samples as well), w is a vertex of $C_Z(\mathcal{A})$. Hence, we have

$$\sum_{\lambda=0}^{\xi} \frac{\binom{n-\lambda-3}{r-3}}{\binom{n}{r}} F_{\lambda} \leq E[\kappa(\mathcal{A})] \leq \kappa(r),$$

where $E[\cdot]$ denotes expectation (with respect to \mathcal{A}), and where F_{λ} is the number of vertices w of $\mathcal{A}(\Sigma)$ at restricted level λ . (Note that $\kappa(\mathcal{A})$ counts the number of 0-borders bounding $C_Z(\mathcal{A})$, which is clearly an upper bound on the number of vertices of that cell.) This can be rewritten as

$$\frac{r(r-1)(r-2)}{n(n-1)(n-2)} \cdot \sum_{\lambda=0}^{\xi} \frac{(n-r)(n-r-1)\cdots(n-r-\lambda+1)}{(n-3)(n-4)\cdots(n-\lambda-2)} F_{\lambda} \leq \kappa(r)$$

or

$$\left(\frac{n-r-\xi+1}{n-\xi-2}\right)^{\xi} \cdot \left(\sum_{\lambda=0}^{\xi} F_{\lambda}\right) \leq \frac{n(n-1)(n-2)}{r(r-1)(r-2)} \kappa(r).$$

As in [9] and [26], it is easily verified that, for $r = n/\xi$, we have

$$\sum_{\lambda=0}^{\xi} F_{\lambda} = O\left(\xi^3 \kappa\left(\frac{n}{\xi}\right)\right);$$

in other words, the number of inner vertices of $\mathcal{A}(\Sigma)$ at restricted level $\leq \xi$ is $O(\xi^3 \kappa(n/\xi))$, which in turn implies that the number of inner 0-borders (v, R) of C in this subcase is $O(\xi^2 \kappa(n/\xi))$.

Case (c(ii)): $t < \xi$. In this case, if we remove these t surfaces from the arrangement, v becomes a vertex of a popular face of C . Indeed, z_{12} has an appropriately consistent side R_{12}^* so that (z_{12}, R_{12}^*) is a 0-border of C . When we remove the t surfaces crossing β_{12} , the cell C expands from R_{12}^* toward v and “reaches” the other side of v consistent with the side R'_{12} of β_{12} , making v a vertex of a popular face (on σ_3) of the cell containing Z in the reduced arrangement. To exploit this observation, we apply the following variant of the preceding random sampling argument. Fix a parameter $r = n/\xi$, and draw a random sample \mathcal{A} of r surfaces of Σ . Let $E'[\mathcal{A}]$ be the expected number of vertices in $\mathcal{A}(\mathcal{A})$ of index $\leq j$ which are incident to popular faces of $C_Z(\mathcal{A})$ (counted with the appropriate multiplicity). By definition, $E'[\mathcal{A}] \leq \pi^{(j)}(r)$. (Note that the index of a vertex does not change when

we pass to a subset \mathcal{R} , as long as the three surfaces defining the vertex belong to \mathcal{R} .) Now, using a similar argument to the one given above, the probability that our vertex v will show up as a vertex of such a popular face of $C_Z(\mathcal{R})$ is $\geq \binom{n-t-3}{r-3} / \binom{n}{r}$: of the $\binom{n}{r}$ possible ways of choosing \mathcal{R} , we consider those samples for which $\sigma_1, \sigma_2, \sigma_3 \in \mathcal{R}$, and none of the other t surfaces crossing β_{12} is chosen in \mathcal{R} ; as already noted, each such choice (and possibly others as well) will make v appear as a vertex of a popular face of the cell under consideration. Hence, we have

$$\sum_{t=0}^{\xi} \frac{\binom{n-t-3}{r-3}}{\binom{n}{r}} G_t \leq E[\mathcal{R}] \leq \pi^{(j)}(r),$$

where G_t is the number of 0-borders (v, R) in the full arrangement that fall into the present subcase, with $j(v) \leq j$ and with exactly t surfaces crossing the corresponding arc β_{12} . Arguing exactly as above, we obtain, for $r = n/\xi$,

$$\sum_{t=0}^{\xi} G_t = O\left(\xi^3 \pi^{(j)}\left(\frac{n}{\xi}\right)\right);$$

in other words, the number of 0-borders in this subcase is $O(\xi^3 \pi^{(j)}(n/\xi))$.

Case (d). In the remaining case (which can occur only if $j(v) > 0$), each of the three arcs $\beta_{ij} \cup \{z_{ij}\}$ intersects the third surface, and we consider one of these arcs, call it β_{12} , which emanates in the positive x -direction. Thus, the third surface σ_3 intersects $\beta_{12} \cup \{z_{12}\}$ in at least one point w ; if there are several such points (no more than j by assumption), we take w to be the point lying furthest from v along β_{12} . Let t denote the number of distinct surfaces of Σ , *excluding* σ_3 , that intersect $\beta_{12} \cup \{z_{12}\}$. We consider the following two subcases:

Case (d(i)): $t \geq \xi$. In this case we charge (v, R) to a block of ξ vertices of the full arrangement $\mathcal{A}(\Sigma)$ which lie along β_{12} , in complete analogy to the construction in case (c(i)) above, except that the surface σ_3 is excluded from the construction. Since each such vertex can be charged in this manner only a constant number of times, and all these vertices lie at restricted level $\leq \xi + 1$, as is easily checked, it follows, exactly as above, that the total number of 0-borders v in this subcase is $O(\xi^2 \kappa(n/\xi))$.

Case (d(ii)): $t < \xi$. In this case, if we remove these t surfaces (without removing σ_3), the point w , together with a side R_w consistent with R'_{12} , must form a 0-border (w, R_w) of the cell containing Z in the reduced arrangement. Indeed, recall that we are assuming that the other endpoint z_{12} of β_{12} forms a 0-border (z_{12}, R^*_{12}) of C , for a side R^*_{12} lying across a surface from a side consistent with R'_{12} . By assumption, the portion of β_{12} between w and z_{12} is not crossed by σ_2 , so, when the other t surfaces crossing β_{12} are removed, the cell C expands from the side R^*_{12} and “reaches” w from z_{12} along the side (β_{12}, R'_{12}) . We charge (v, R) to (w, R_w) . Clearly, each such (w, R_w) is charged in this manner only a constant number of times.

We next estimate the number of 0-borders (w, R_w) of this kind. We apply a random-sampling argument similar to those used above. That is, we fix a parameter $r = n/\xi$, and draw a random sample \mathcal{R} of r surfaces of Σ . Let $E^n[\mathcal{R}]$ be the expected number of 0-borders (w, R_w) of $C_Z(\mathcal{R})$, such that w had index $\leq j - 1$. By definition, $E^n[\mathcal{R}] = E[\kappa^{(j-1)}(\mathcal{R})] \leq \kappa^{(j-1)}(r)$.

The probability that the charged 0-border (w, R_w) will show up as such a 0-border of $C_Z(\mathcal{R})$ is $\geq \binom{n-t-3}{r-3} / \binom{n}{r}$: of the $\binom{n}{r}$ possible ways of choosing \mathcal{R} , we consider those samples for which $\sigma_1, \sigma_2, \sigma_3 \in \mathcal{R}$, and none of the other t surfaces crossing β_{12} is chosen in \mathcal{R} ; each such choice (and possibly other choices too) will make (w, R_w) appear as a 0-border of $C_Z(\mathcal{R})$, as argued above. Hence, we have

$$\sum_{t=0}^{\xi} \frac{\binom{n-t-3}{r-3}}{\binom{n}{r}} H_t = O(E^n[\mathcal{R}]) = O(\kappa^{(j-1)}(r)),$$

where H_t is the number of 0-borders (v, R) in the full arrangement that fall in the present subcase, with $j(v) \leq j$ and with exactly t surfaces crossing the corresponding arc β_{12} (excluding the corresponding surface σ_3). Arguing as above, we obtain, for $r = n/\xi$,

$$\sum_{t=0}^{\xi} H_t = O\left(\xi^3 \kappa^{(j-1)}\left(\frac{n}{\xi}\right)\right);$$

in other words, the number of 0-borders (v, R) in this subcase is $O(\xi^3 \kappa^{(j-1)}(n/\xi))$.

Hence, summing over all cases, we obtain the following recurrence for $\kappa^{(j)}$ (where, for $j = 0$, we put $\kappa^{(j-1)} = 0$ in the right-hand side):

$$\kappa^{(j)}(n) = O\left(\xi^2 \kappa\left(\frac{n}{\xi}\right) + \xi^3 \pi^{(j)}\left(\frac{n}{\xi}\right) + \xi^3 \kappa^{(j-1)}\left(\frac{n}{\xi}\right) + n^2 + \pi^{(j)}(n)\right), \quad (1)$$

as asserted. □

In order to solve the recurrence (1) we first have to bound the functions $\pi^{(j)}(n)$, which is done in the next section.

3.2. The Number of Vertices on Popular Faces

To bound the number of such vertices, we adapt the analysis given above to bound the functions $\pi^{(j)}(n)$, rather than the functions $\kappa^{(j)}(n)$.

Let v be an inner vertex of a popular face f of the cell C ; assume that f lies on a surface $\sigma \in \Sigma$, and that the two other surfaces incident to v are $\sigma_1, \sigma_2 \in \Sigma$. As above, we may assume that v is nonsingular on any of these surfaces, and that these surfaces meet transversally at v . Denote the two sides of f by R^+ and R^- . Let




Fig. 5. The setting of Section 3.2; f is a popular face.

$\gamma_i = \sigma_i \cap \sigma$, for $i = 1, 2$. Each γ_i contains an edge e_i having v as an endpoint and bounding f . See Fig. 5 for an illustration. Let R_i^+, R_i^- , for $i = 1, 2$, denote the two sides of e_i that are consistent with R^+, R^- , respectively.

If, say γ_1 has another edge e'_1 incident to v such that (a) e'_1 bounds another popular face f' on σ , and (b) f and f' share the edge e_2 , then e_2 is a *popular edge* of C , meaning that all four sides of e_2 lie in C locally near e_2 (see Fig. 6). We claim that the number of popular edges of C is $O(n^2)$. This follows from the observation that, in this case, v must be a locally x -extreme vertex of one of these four sides³ (assuming general position), and, by Theorem 2.3, the number of such vertices, and hence also the number of popular edges, is $O(n^2)$. Moreover, if both e_1 and e_2 emanate from v in the positive x -direction, or if both emanate in the negative x -direction, then, as is easily checked, v must be a locally x -extreme vertex of one of the two sides of f , so the number of such vertices is also $O(n^2)$. Hence, in what follows we may assume that neither of the edges e_1, e_2 is adjacent along σ to another popular face of C , and that one of these edges emanates from v in the positive x -direction and one emanates in the negative x -direction.




Fig. 6. Two adjacent popular faces f, f' , giving rise to a popular edge e_2 .

³ Strictly speaking, v is locally x -extreme in one of the four portions of C into which these sides “point”; we allow ourselves here and below this slight abuse of notation.

As above, our method is to derive a recurrence relationship for $\pi(n)$, by bounding each of the functions $\pi^{(j)}$ in terms of $\pi^{(j-1)}$ (with a special handling of $\pi^{(0)}$); the solution of the resulting system of recurrences will yield the asserted bounds. In the remainder of this subsection we prove the following:

Lemma 3.3. *For each $j = 0, \dots, s - 1$, and a parameter $\xi \leq n$ we have*

$$\pi^{(j)}(n) = O\left(\xi^2 \pi\left(\frac{n}{\xi}\right) + \xi n^2 + \xi^3 \pi^{(j-1)}\left(\frac{n}{\xi}\right)\right),$$

we put $\pi^{(j-1)} = 0$ when $j = 0$.

Proof. We fix $0 \leq j \leq s - 1$, and assume that the vertex v under consideration has index $\leq j$. By our assumption, for $i = 1, 2$, the curve γ_i must contain a maximal relatively open x -monotone connected portion β_i having v as an endpoint and satisfying the following property: let R_i^+, R_i^- be the two sides of β_i which are consistent with the two respective sides R_i^+, R_i^- of e_i ; then β_i does not contain any point at which both sides R_i^+, R_i^- lie locally in C . Let z_i denote the other endpoint of β_i . By assumption, one of these arcs, say β_1 , emanates from v in the positive x -direction and the other emanates in the negative x -direction.

Several cases can then arise (similar to the analysis in the previous subsection). Case (a) below deals with vertices that are easy to charge directly, and their total number is shown to be $O(n^2)$. Case (b) handles situations where β_1 is not intersected by the third surface σ_2 ; and in the remaining case, (c), the surface σ_2 intersects β_1 in at least one point. Here too the more involved cases (b) and (c) are further divided into subcases according to certain parameters that are introduced in the analysis.

Case (a). At least one of the arcs β_i ends at a point z_i which is either an endpoint of the original intersection curve γ_i , or a point of local x -extremum on that curve. We then charge v to the point z_i , and note that the number of such points is $O(n^2)$, and that each such point is charged only a constant number of times in this manner, thus implying that the number of vertices v of this kind is only $O(n^2)$. In what follows we assume that this situation does not arise, which means, in particular, that z_1 is a vertex of another popular face of C , whose two sides are consistent with R_1^+, R_1^- , respectively.

Case (b). $\beta_1 \cup \{z_1\}$ is not intersected by the third surface σ_2 . Define the *popularity level* of a point w lying on some surface $\sigma \in \Sigma$ to be the smallest number of other surfaces of Σ whose removal makes w lie in a popular face on σ bounding the cell containing Z in the resulting subarrangement. If w is a vertex of $\mathcal{A}(\Sigma)$, incident to some face $f \subseteq \sigma$, we say that (w, f) lies at *restricted popularity level* $p((w, f)) = k$ if by removing k surfaces from Σ , none of which is incident to w , the face f becomes (after a possible expansion) a popular face (incident to w) of the cell containing Z in the resulting subarrangement, and if k is the smallest number with that property.

Let t denote the number of distinct surfaces of Σ that intersect $\beta_1 \cup \{z_1\}$. We fix some threshold parameter $\xi = \xi_j$, and consider the following two subcases:

Case (b(i)): $t \geq \xi$. In this case we charge the pair (v, f) to a block of ξ points of intersection between $\beta_1 \cup \{z_1\}$ and the surfaces of Σ , defined as follows. For each surface σ' intersecting $\beta_1 \cup \{z_1\}$, choose its point of intersection that lies nearest to v along β_1 . We obtain at least ξ such designated points, and we charge (v, f) to the block of the first ξ designated points, in their order along β_1 from v . All those points are inner vertices of $\mathcal{A}(\Sigma)$, and it is clear that none of these vertices can be charged in this manner more than a constant number of times. By construction, each of the charged vertices w , together with some incident face along σ , lies at restricted popularity level at most ξ : the removal of the at most ξ surfaces intersecting β_1 between v and w (including σ_2 but excluding the surface incident to w) makes the popular face f expand into a bigger, still popular face of the cell containing Z , which has w as a vertex (see Fig. 7). Our goal is thus to obtain an upper bound for the number M of pairs (w, f') of inner vertices w of $\mathcal{A}(\Sigma)$ and incident faces f' that lie at restricted popularity level $\leq \xi$; the number of pairs (v, f) in the present subcase is $O(M/\xi)$.

For this we apply an appropriately modified version of the probabilistic analysis technique used in the previous subsection. That is, we choose a random sample \mathcal{R} of $r = n/\xi$ surfaces of Σ , and construct the arrangement $\mathcal{A}(\mathcal{R})$. Let (w, f') be a pair of an inner vertex w of $\mathcal{A}(\Sigma)$ and an incident face f' , lying on a surface σ , at restricted popularity level $p \leq \xi$, and let \mathcal{Q} be a specific collection of p surfaces, none incident to w , whose removal makes w a vertex of a popular face (containing f') of the cell containing Z . The probability that w shows up as such a vertex in $\mathcal{A}(\mathcal{R})$ is, in complete analogy to the preceding analysis, at least $\binom{n-p-3}{r-3} / \binom{n}{r}$. Hence, we have

$$\sum_{p=0}^{\xi} \frac{\binom{n-p-3}{r-3}}{\binom{n}{r}} F_p \leq E[\pi(\mathcal{R})] \leq \pi(r),$$

where $E[\cdot]$ denotes expectation, and where F_p is the number of vertex-face pairs (w, f') of $\mathcal{A}(\Sigma)$ at restricted popularity level p . As in the preceding analysis, this




Fig. 7. Cases (b) and (c) of the analysis of vertices of popular faces; f' is also a popular face.

implies, for $r = n/\xi$,

$$\sum_{p=0}^{\xi} F_p = O\left(\xi^3\pi\left(\frac{n}{\xi}\right)\right);$$

in other words, the number of inner vertex-face pairs of $\mathcal{A}(\Sigma)$ at restricted popularity level $\leq \xi$ is $O(\xi^3\pi(n/\xi))$, which in turn implies that the number of vertices v of popular faces of C in this subcase is $O(\xi^2\pi(n/\xi))$.

Case (b(ii)): $t < \xi$. In this case, if we remove these t surfaces from the arrangement, v becomes a vertex of a popular edge of C (namely, e_2 or an appropriate extension of it). Indeed (see Fig. 7), z_1 has two sides, R'^+, R'^- , which are consistent with the two sides of a popular face f' of C having z_1 as a vertex, and are also consistent with the two respective sides R_1^+, R_1^- of β_1 . This easily implies that when we remove the t surfaces crossing $\beta_1 \cup \{z_1\}$ (not removing σ_2), the cell C expands from R'^+ and R'^- toward v and the face f' expands into a bigger, still popular face which is bounded by v and by (a possible extension of) its incident edge e_2 ; by definition, this edge is thus a popular edge of the resulting cell. To exploit this observation, we apply the following variant of the preceding random sampling arguments. Fix a parameter $r = n/\xi$, and draw a random sample \mathcal{R} of r surfaces of Σ . Let $E'[\mathcal{R}]$ be the expected number of vertices in $\mathcal{A}(\mathcal{R})$ that are incident to popular edges of the cell containing Z in this arrangement. By the preceding analysis, $E'[\mathcal{R}] = O(r^2)$. Now, using a similar argument to the one given above, the probability that our vertex v will show up as a vertex of such a popular edge in $\mathcal{A}(\mathcal{R})$ is $\geq \binom{n-t-3}{r-3} / \binom{n}{r}$. Hence, we have

$$\sum_{t=0}^{\xi} \frac{\binom{n-t-3}{r-3}}{\binom{n}{r}} G_t \leq E'[\mathcal{R}] = O(r^2),$$

where G_t is the number of vertices v in the full arrangement that fall into the present subcase, with exactly t surfaces crossing the corresponding arc β_1 . Arguing as above, it follows that, for $r = n/\xi$,

$$\sum_{t=0}^{\xi} G_t = O(\xi^3 r^2) = O(\xi n^2);$$

in other words, the number of vertices v in this subcase is $O(\xi n^2)$.

Case (c). In the remaining case (which can occur only when $j(v) > 0$), the surface σ_2 intersects $\beta_1 \cup \{z_1\}$ in at least one point w ; if there are several such points (no more than j by assumption), we take w to be the point lying furthest from v along β_1 . Let t denote the number of distinct surfaces of Σ , excluding σ_2 , which intersect β_1 . We consider the following two subcases:

Case (c(i)): $t \geq \xi$. In this case we charge the pair (v, f) to a block of ξ vertices of the full arrangement $\mathcal{A}(\Sigma)$ which lie along β_1 , in complete analogy to the construction in case (b(i)) above, except that the surface σ_2 is excluded from the construction. Since each such vertex can be charged in this manner only a constant number of times, and since all these vertices lie at restricted popularity level $\leq \xi + 1$, as is easily checked, it follows, exactly as above, that the total number of vertices v that are charged in this way is $O(\xi^2 \pi(n/\xi))$.

Case (c(ii)): $t < \xi$. In this case, if we remove these t surfaces (without removing σ_2), the point w , together with two appropriate sides R_w^+, R_w^- , consistent with R_1^+, R_1^- , respectively, must be a vertex of a popular face f' of the cell containing Z in the resulting subarrangement, so that R_w^+ and R_w^- are consistent with the two sides of f' . Indeed, arguing as above, the other endpoint z_1 of β_1 is a vertex of an appropriate popular face f'' of C , and removal of the t other surfaces crossing β_1 makes f'' expand and "reach" w along the appropriate sides; see Fig. 8 for an illustration. We charge (v, f) to (w, f') . Clearly, each such pair (w, f') is charged in this manner only a constant number of times.

We next estimate the number of pairs (w, f') of this kind. We apply a random-sampling argument similar to those used above. That is, we fix a parameter $r = n/\xi$, and draw a random sample \mathcal{R} of r surfaces of Σ . Let $E^n[\mathcal{R}]$ be the expected number of vertices w of popular faces f' of $C_Z(\mathcal{R})$ which have index $\leq j - 1$ (counted with the appropriate multiplicity). By definition, $E^n[\mathcal{R}] = E[\pi^{(j-1)}(\mathcal{R})] \leq \pi^{(j-1)}(r)$.

Arguing as above, the probability that the charged vertex-face pair (w, f') will show up as such a vertex in $\mathcal{A}(\mathcal{R})$ is $\geq \binom{n-t-3}{r-3} / \binom{n}{r}$. Hence, we have

$$\sum_{t=0}^{\xi} \frac{\binom{n-t-3}{r-3}}{\binom{n}{r}} H_t = O(E^n[\mathcal{R}]) = O(\pi^{(j-1)}(r)),$$

where H_t is the number of vertex-face pairs (v, f) of $C_Z(\Sigma)$ that fall in the present subcase, with exactly t surfaces crossing the corresponding arc β_1 (excluding the




Fig. 8. Case (c(ii)): the vertices v and w are both intersections of the same three surfaces.

corresponding surface σ_2). Again, for $r = n/\xi$ this becomes

$$\sum_{t=0}^{\xi} H_t = O\left(\xi^3 \pi^{(j-1)}\left(\frac{n}{\xi}\right)\right);$$

hence, the number of vertex–face pairs (v, f) in this subcase is $O(\xi^3 \pi^{(j-1)}(n/\xi))$.

Summing over all cases, we obtain the following recurrence for $\pi^{(j)}$ (where, for $j = 0$, we put $\pi^{(j-1)} = 0$ in the right-hand side):

$$\pi^{(j)}(n) = O\left(\xi^2 \pi\left(\frac{n}{\xi}\right) + \xi n^2 + \xi^3 \pi^{(j-1)}\left(\frac{n}{\xi}\right)\right), \tag{2}$$

as asserted. □

3.3. Solving the Recurrences

We next proceed to solve recurrences (1) and (2).

3.3.1. The Recurrence for $\pi^{(j)}$. We start with (2), fix some $\varepsilon > 0$, and claim that its solution is $\pi^{(j)}(n) \leq B_j n^{2+\varepsilon}$, for $j = 0, \dots, s - 1$, where the constants B_j depend on ε, j , and on the maximum degree b . By definition, this implies that $\pi(n) \leq B_{s-1} n^{2+\varepsilon}$.

We prove this claim by induction on n . We first rewrite (2), using a different parameter ξ_j for each j , as

$$\pi^{(0)}(n) \leq c \xi_0^2 \pi\left(\frac{n}{\xi_0}\right) + c \xi_0 n^2 \tag{3}$$

and

$$\pi^{(j)}(n) \leq c \xi_j^2 \pi\left(\frac{n}{\xi_j}\right) + c \xi_j n^2 + c \xi_j^3 \pi^{(j-1)}\left(\frac{n}{\xi_j}\right), \quad j = 1, \dots, s - 1,$$

for appropriate positive constants $c, \xi_0, \dots, \xi_{s-1}$; without loss of generality, we assume $c > 1$. We take ξ_0 to be sufficiently large, and put $\xi_j = \xi_0^{\varepsilon^j}$, for $j = 0, \dots, s - 1$; note that $\xi_j = \xi_{j-1}^{\varepsilon}$, for $j = 1, \dots, s - 1$. We note that, by choosing the B_j 's to be sufficiently large, we can assume that the claimed bounds hold for all $n \leq \xi_0$. (For example, choose the B_j 's to be larger than some appropriate multiple of ξ_0 , and use the fact that all the quantities we want to estimate are bounded by $O(n^3)$, where the constant of proportionality depends only on the maximum degree b of the surfaces and of their boundaries.)

For $n > \xi_0$, we apply the induction hypothesis in the right-hand side of (3), and conclude that the asserted bounds continue to hold for n too, provided that the

following inequalities are satisfied:

$$\frac{cB_{s-1}}{\xi_0^\varepsilon} + c\xi_0^{1-\varepsilon} \leq B_0, \tag{4}$$

$$\frac{cB_{s-1}}{\xi_j^\varepsilon} + \frac{c\xi_j}{\xi_0^\varepsilon} + c\xi_j^{1-\varepsilon}B_{j-1} \leq B_j, \quad j = 1, \dots, s - 1.$$

To achieve this, we choose

$$B_0 > 2c\xi_0^{1-\varepsilon} \quad \text{and} \quad B_j = \frac{(j + 1)c^j}{\xi_j^\varepsilon} \xi_0^\varepsilon B_0,$$

and require that ξ_0 be sufficiently large so that $\xi_0^{\varepsilon^s} > 2sc^s$.

The first inequality in (4) is equivalent to

$$\frac{sc^{s-1} \cdot \xi_0^\varepsilon B_0}{\xi_{s-1}^\varepsilon} \cdot \frac{c}{\xi_0^\varepsilon} + c\xi_0^{1-\varepsilon} \leq B_0$$

or to

$$\frac{sc^s}{\xi_0^{\varepsilon^s}} \cdot B_0 + c\xi_0^{1-\varepsilon} \leq B_0.$$

Since $sc^s/\xi_0^{\varepsilon^s} < \frac{1}{2}$, the choice of B_0 is easily seen to satisfy this inequality. The general inequality in (4) is equivalent to

$$\frac{sc^s \cdot \xi_0^\varepsilon B_0}{\xi_0^{\varepsilon^s}} \cdot \frac{1}{\xi_0^{\varepsilon^{j+1}}} + \frac{c\xi_j}{\xi_0^\varepsilon} + c\xi_0^{\varepsilon^j(1-\varepsilon)} \cdot \frac{jc^{j-1}}{\xi_0^{\varepsilon^j}} \cdot \xi_0^\varepsilon B_0 \leq \frac{(j + 1)c^j}{\xi_0^{\varepsilon^{j+1}}} \cdot \xi_0^\varepsilon B_0.$$

Using again the fact that $sc^s/\xi_0^{\varepsilon^s} < \frac{1}{2}$, the last inequality is easily seen to be implied by

$$\frac{c\xi_j}{\xi_0^\varepsilon} \leq \frac{(c^j - \frac{1}{2})\xi_0^\varepsilon B_0}{\xi_0^{\varepsilon^{j+1}}}$$

or by

$$\frac{c\xi_j^{1+\varepsilon}}{\xi_0^\varepsilon} \leq (c^j - \frac{1}{2})\xi_0^\varepsilon B_0. \tag{5}$$

However, when j increases, the left-hand side of (5) decreases, while the right-hand side increases (using the assumption that $c > 1$). Hence it suffices to verify (5) for $j = 0$, which trivially holds by the choice of B_0 . This inductive step completes the solution of the recurrence for the functions $\pi^{(j)}(n)$.

3.3.2. *The Recurrence for $\kappa^{(j)}$.* We next proceed to solve recurrences (1). If we substitute in these equations the bounds for $\pi^{(j)}(n)$, as just obtained, and use a

different threshold parameter ξ_j for each j , we can rewrite the recurrences as

$$\kappa^{(0)}(n) \leq c\xi_0^2\kappa\left(\frac{n}{\xi_0}\right) + c\xi_0^{1-\varepsilon}B_0n^{2+\varepsilon}$$

and

$$\kappa^{(j)}(n) \leq c\xi_j^2\kappa\left(\frac{n}{\xi_j}\right) + c\xi_j^3\kappa^{(j-1)}\left(\frac{n}{\xi_j}\right) + c\xi_j^{1-\varepsilon}B_jn^{2+\varepsilon}, \quad j = 1, \dots, s - 1,$$

for appropriate positive constants $c, \xi_0, \dots, \xi_{s-1}$ (not necessarily the same as in the recurrences for π); with no loss of generality, we assume $c > 1$.

As above, we fix some $\varepsilon > 0$, and claim that the solution of these equations is $\kappa^{(j)}(n) \leq A_jn^{2+\varepsilon}$, for $j = 0, \dots, s - 1$, where the constants A_j depend on ε, j , and the maximum degree b . By definition, this implies that $\kappa(n) \leq A_{s-1}n^{2+\varepsilon}$.

We prove this claim by induction on n . Again, we take ξ_0 to be sufficiently large, and put $\xi_j = \xi_0^{\varepsilon^j}$, for $j = 0, \dots, s - 1$. By choosing the A_j 's to be sufficiently large, we can assume, as above, that the claimed bounds hold for all $n \leq \xi_0$.

For $n > \xi_0$, we apply the induction hypothesis, and conclude that the asserted bounds continue to hold for n too, provided that the following inequalities are satisfied:

$$\begin{aligned} \frac{cA_{s-1}}{\xi_0^\varepsilon} + c\xi_0^{1-\varepsilon}B_0 &\leq A_0, \\ \frac{cA_{s-1}}{\xi_j^\varepsilon} + c\xi_j^{1-\varepsilon}(A_{j-1} + B_j) &\leq A_j, \quad j = 1, \dots, s - 1. \end{aligned}$$

To achieve this, we choose

$$A_0 > \max_{0 \leq j \leq s-1} \frac{2c\xi_j B_j}{\xi_0^\varepsilon} \quad \text{and} \quad A_j = \frac{(j+1)c^j}{\xi_j^\varepsilon} \xi_0^\varepsilon A_0,$$

and require that ξ_0 be sufficiently large so that $\xi_0^{\varepsilon^s} > 2sc^s$.

We leave it to the reader to verify, in much the same way as above, that this choice of coefficients satisfies the above inequalities. This inductive step completes the solution of the recurrences for the functions $\kappa^{(j)}(n)$, and therefore completes the proof of Theorem 3.1.

3.4. The Overall Result

The overall complexity of a cell C counts the number of vertices, edges, and faces of ∂C , each with an appropriate but constant multiplicity. In order to derive the main result of the paper, it remains to show that the overall combinatorial complexity of C is dominated asymptotically by the bound, given in Theorem 3.1, for the number of its inner vertices.

Note first that noninner vertices of C must arise as intersections of the relative boundary of one surface with another surface, so the number of such vertices is only $O(n^2)$.

Using rather standard arguments (see, e.g., Section 3.1 of [15]), it can easily be shown that asymptotically, and up to an additive factor $O(n^2)$, a bound on the number of vertices in C dominates the number of edges and faces in C . Every bounded edge and face of ∂C can be charged to the point with the smallest x -coordinate lying on its closure. Such a point is either a vertex of C or a locally x -extreme point on one of the surfaces or on one of their intersection curves; the number of such extreme points is clearly only $O(n^2)$. A similar argument counts the number of unbounded edges and faces. Since each vertex or extreme point is charged in this manner only a constant number of times (by our assumption on general position), we easily obtain the above claim. Combined with the observation that the same bound also applies to collections of surfaces not in general position, we thus obtain:

Theorem 3.4. *The combinatorial complexity of a single cell in an arrangement of n algebraic surface patches in 3-space, satisfying conditions (i)–(iii), is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε , s , and b .*

4. Zone Complexity

An interesting application of Theorem 3.4 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.

Theorem 4.1. *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 .*

Proof. We extend the idea used in [12] for the analysis of zones in two-dimensional arrangements. That is, assume first that σ is a connected surface. We cut each $\sigma_i \in \Sigma$ into a constant number of subpatches along its curve of intersection with σ . If we shrink these subpatches away from each other by a small amount, all the cells of the zone become connected to each other, and form a single cell in the arrangement of the new patches. Since these patches are easily seen also to satisfy conditions similar to (i)–(iv), the asserted bound follows. If σ is not connected, it consists of a constant number of connected components, and we apply the above argument to each component separately. \square

5. Conclusion

In this paper we have obtained a near-quadratic bound for the combinatorial complexity of a single cell in an arrangement of n algebraic surface patches of

constant maximum degree, each bounded by a constant number of algebraic arcs of constant maximum degree as well. This almost settles a long-standing conjecture, and provides a fairly satisfactory extension of the two-dimensional Davenport–Schinzel theory developed in [14]. Our proof extends recent techniques developed in [17], [18], and [27], all based on the probabilistic technique of [9] and [26] for deriving bounds on the number of generalized “($\leq k$)-sets” in arrangements.

As noted in the introduction, our result has immediate applications to the general motion-planning problem with three degrees of freedom. That is, for rather general systems with three degrees of freedom, the combinatorial complexity of the connected component of the free configuration space, consisting of all robot placements which are reachable from some given initial free placement, is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where n is the number of contact surfaces, as defined in the Introduction. This still falls short of efficient construction of such a component (in near-quadratic running time, which is one of the major open problems that we pose in this paper; we note that the *entire* arrangement can be efficiently constructed by a simple generalization of the algorithm in [10]). However, at least we know that the complexity of such a component is nearly an order of magnitude smaller than the worst-case complexity of the entire arrangement, and, in most cases, also of the entire free configuration space.

The paper raises several open problems. The first one is to design an efficient algorithm (of near-quadratic complexity) for constructing a single cell in an arrangement of algebraic surfaces in 3-space (in an appropriate “algebraic” model of computation). One method of doing this is to select a random sample \mathcal{R} of r surfaces of Σ , for some large constant r , construct $C_Z(\mathcal{R})$ by brute force, and decompose it into subcells of “constant description complexity.” Then, with high probability, each of these cells is crossed by only $O((n/r) \log r)$ surfaces of Σ , and we can continue the construction recursively within each cell, in a manner similar to that described in [4] and [27]. The problem that arises here, however, is to ensure that the number of resulting subcells be as small as possible (ideally, near-quadratic in r). This is an interesting combinatorial subproblem, which seems to be open, and considerably harder than the problem studied in this paper.

Another open problem that the paper raises is to improve further the bound that we have obtained, to the conjectured bound of $O(n\lambda_s(n))$, or at least to $O(n^2 \cdot \text{polylog}(n))$. Another open problem is to extend our results to arrangements in d dimensions. We believe that this is feasible, and are currently exploring this problem. One main subproblem here is to extend Theorem 2.3 to higher dimensions.

Other open problems are to extend our analysis to obtain sharp bounds on the complexity of many cells in three-dimensional arrangements, and to derive bounds on the sum of squares of cell complexities in an entire arrangement; see [2] and [5] for related work.

References

1. 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.
2. B. Aronov, J. Matoušek, and M. Sharir, On the sum of squares of cell complexities in hyperplane arrangements, *J. Combin. Theory Ser. A* **65** (1994), 311–321.

3. B. Aronov, M. Pellegrini, and M. Sharir, On the zone of a surface in a hyperplane arrangement, *Discrete Comput. Geom.* **9** (1993), 177–186.
4. B. Aronov and M. Sharir, Triangles in space, or building (and analyzing) castles in the air, *Combinatorica* **10** (1990), 137–173.
5. B. Aronov and M. Sharir, Castles in the air revisited, *Discrete Comput. Geom.* **12** (1994), 119–150.
6. B. Aronov and M. Sharir, On translational motion planning in 3-space. *Proc. 10th ACM Symp. on Computational Geometry*, 1994, pp. 21–30.
7. J. Bochnak, M. Coste, and M.-F. Roy, *Géométrie Algébrique Réelle*, Springer-Verlag, Berlin, 1987.
8. B. Chazelle, H. Edelsbrunner, L. J. Guibas, M. Sharir, and J. Snoeyink, Computing a face in an arrangement of line segments and related problems, *SIAM J. Comput.* **22** (1993), 1286–1302.
9. K. Clarkson and P. Shor, Applications of random sampling in computational geometry, II, *Discrete Comput. Geom.* **4** (1989), 387–421.
10. M. de Berg, L. J. Guibas, and D. Halperin, Vertical decompositions for triangles in 3-space, *Proc. 10th ACM Symp. on Computational Geometry*, 1994, pp. 1–10.
11. H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg, 1987.
12. H. Edelsbrunner, L. J. 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.
13. H. Edelsbrunner, R. Seidel, and M. Sharir, On the zone theorem in hyperplane arrangements, *SIAM J. Comput.* **22** (1993), 418–429.
14. 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.
15. D. Halperin, *Algorithmic Motion Planning via Arrangements of Curves and of Surfaces*, Ph.D. Dissertation, Computer Science Department, Tel Aviv University, July 1992.
16. 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.
17. 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.
18. D. Halperin and M. Sharir, New bounds for lower envelopes in three dimensions, with applications to visibility in terrains, *Discrete Comput. Geom.* **12** (1994), 313–326.
19. S. Hart and M. Sharir, Nonlinearity of Davenport–Schinzel sequences and of generalized path compression schemes, *Combinatorica* **6** (1986), 151–177.
20. R. Hartshorne, *Algebraic Geometry*, Springer-Verlag, New York, 1977.
21. M. W. Hirsch, *Differential Topology*, Springer-Verlag, New York, 1976.
22. P. McMullen, The maximum number of faces of a convex polytope, *Mathematika* **17** (1970), 179–184.
23. 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.
24. R. Pollack and M.-F. Roy, On the number of cells defined by a set of polynomials, *C.R. Acad. Sci. Paris Sér. I* **316** (1993), 573–577.
25. J. T. Schwartz and M. Sharir, On the two-dimensional Davenport–Schinzel problem, *J. Symbolic Comput.* **10** (1990), 371–393.
26. M. Sharir, On k -sets in arrangements of curves and surfaces, *Discrete Comput. Geom.* **6** (1991), 593–613.
27. M. Sharir, Almost tight upper bounds for lower envelopes in higher dimensions, *Discrete Comput. Geom.* **12** (1994), 327–345.

Received May 1994, and in revised form December 1994.