Vertical Decomposition of a Single Cell in a Three-Dimensional Arrangement of Surfaces and its Applications*

Otfried Schwarzkopf† Micha Sharir‡

Abstract

Let Σ be a collection of n algebraic surface patches of constant maximum degree in $\mathbb{R}^3$. We show that the combinatorial complexity of the vertical decomposition of a single cell in the arrangement $\mathcal{A}(\Sigma)$ is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε and on the maximum degree of the surfaces and of their boundaries. As an application, we obtain a near-quadratic motion planning algorithm for general systems with three degrees of freedom.

1 Introduction

Let $\Sigma = \{\sigma_1, \ldots, \sigma_n\}$ be a collection of n algebraic surface patches in $\mathbb{R}^3$ of constant maximum degree b, such that the boundary of each surface consists of a constant number of algebraic arcs, each of degree at most b as well. Let $\mathcal{A}(\Sigma)$ denote the arrangement of Σ. (We assume that the reader is familiar with arrangements—see for instance the recent book [17] for details concerning arrangements of surfaces in higher dimensions.) Let ω be a fixed point, not lying on any surface of Σ. We denote by $\mathcal{C}_\omega(\Sigma)$ the 3-dimensional cell of $\mathcal{A}(\Sigma)$ containing ω. The combinatorial complexity of $\mathcal{C}_\omega(\Sigma)$ is the number of vertices, edges, and 2-faces of $\mathcal{A}(\Sigma)$ appearing on the boundary of that cell. For simplicity, we will measure this complexity only by the number of vertices of the cell. It is well known that the number of all other boundary features of $\mathcal{C}_\omega(\Sigma)$ is proportional to the number of vertices (assuming general position—see below), plus an additive term of $O(n^2)$.

Recently, it has been shown that the combinatorial complexity of $\mathcal{C}_\omega(\Sigma)$ is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε and on the maximum degree b of the surfaces and of their boundaries [11]. The corresponding algorithmic problem, however, of computing $\mathcal{C}_\omega(\Sigma)$ in near-quadratic time, has been open, with the exception of several solutions for special classes of surfaces [4, 5, 12]. The main motivation for this algorithmic problem comes from motion planning, and is explained in detail in the papers just cited, and in the recent survey paper [13].

An algorithm for constructing $\mathcal{C}_\omega(\Sigma)$ can be obtained using the *vertical decomposition* of such a cell [11, 12, 13]. This is a standard decomposition scheme, described in detail in several recent works [7, 8, 17], that partitions cells in arrangements of algebraic surfaces into subcells of constant description complexity (see below), provided the maximum degree of the surfaces is also constant.

For the sake of completeness, we also give a brief informal description of the vertical decomposition. We first assume that each surface patch in Σ is xy-monotone. This can always be enforced by splitting each such patch into $O(1)$ xy-monotone subpatches. In the first decomposition stage, we erect within $\mathcal{C}$ a vertical 'wall' up and/or down from each edge of $\mathcal{C}$ (both surface boundary edges and intersection edges of pairs of surfaces). Each such wall consists of maximal verti-

*Work on this paper by the first author has been supported by the Netherlands' Organization for Scientific Research (NWO), and partially by Pohang University of Science and Technology Grant P95015, 1995. Work on this paper by the second author has been supported by NSF Grants CCR-93-11127 and CCR-94-24398, by a Max-Planck Research Award, and by grants from the U.S.-Israeli Binational Science Foundation, the Israel Science Fund administered by the Israeli Academy of Sciences, and the G.I.F., the German-Israeli Foundation for Scientific Research and Development.

†Dept. of Computer Science, Pohang University of Science and Technology, San 31, Hyoja-Dong, Pohang 790–784, South Korea

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

cal segments contained in (the closure of) $\mathcal{C}$ and passing through the points of the edge. The collection of these walls partition $\mathcal{C}$ into subcells, each having the property that it has a unique 'top' facet and a unique 'bottom' facet (one or both of these facets may be undefined when the subcell is unbounded; all other facets of the subcell lie on the vertical walls). However, the complexity of each subcell may still be arbitrarily large. Thus, in the second decomposition stage, we take each subcell $\mathcal{C}'$, project it onto the xy-plane, and apply to the projection a similar but 2-dimensional vertical decomposition, erecting a y-vertical segment from each vertex of the projected subcell and from each point of local x-extremum on its edges. This yields a collection of trapezoidal-like subcells, and we then lift each of them to 3-space, to obtain a decomposition of $\mathcal{C}'$ into prism-like subcells, each having 'constant description complexity', meaning that each of them is a semialgebraic set defined by a constant number of polynomials of constant maximum degree (which depends on b). Repeating this second stage for all subcells $\mathcal{C}'$ produced in the first stage, we obtain the desired vertical decomposition of $\mathcal{C}$. More details can be found elsewhere [7, 8, 17].

Using this decomposition scheme, one can then apply, for instance, a lazy randomized incremental algorithm [6] to construct the vertical decomposition of $\mathcal{C}_\omega(\Sigma)$, by adding the surfaces one after the other in random order, and by updating the decomposition as the surfaces are added. The efficiency of this algorithm crucially depends on the size (number of subcells) of the decomposition (of the cells $\mathcal{C}_\omega(\Sigma')$, for any subset $\Sigma' \subseteq \Sigma$). A near-quadratic bound on the size of the vertical decomposition of a single cell implies that the (expected) complexity of the above algorithm is also near-quadratic.

In this paper we show that the complexity of the vertical decomposition of a single cell in a 3-dimensional arrangement, as above, is indeed $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends, as above, on ε and on the maximum degree b of the surfaces and of their boundaries. The proof technique borrows ideas from several recent papers [1, 10, 16, 18] that have analyzed several related problems.

It is instructive to note that if all our surfaces are xy-monotone without boundaries (in other words, they are graphs of continuous totally-defined algebraic bivariate functions), then the near-quadratic bound on the complexity of the vertical decomposition of a single cell is an immediate consequence of the recent results of Agarwal et al. [1], which give a near-quadratic bound for the complexity of the vertical decomposition of the region enclosed between the lower envelope of one collection of such surfaces and the upper envelope of another such collection; in this special case, our single cell is a portion of such a 'sandwiched' region. In the general case, though, the topological structure of a single cell can be much more complex, and this makes the analysis considerably harder.

As a corollary of our bound, we obtain that a single cell in a 3-dimensional arrangement of surfaces, as above, can be constructed in randomized expected $O(n^{2+\varepsilon})$ time, for any $\varepsilon > 0$. This in turn implies that motion planning for fairly general systems with three degrees of freedom can be performed in near quadratic time. This solves one of the major open problems in the area. These applications of our bound are briefly presented in Section 3.

2 Complexity of the Vertical Decomposition of a Single Cell

Let Σ and ω be as in the introduction. For the purpose of our analysis, we will require the surface patches to be xy-monotone. This involves no real loss of generality, because, as already mentioned in the introduction, we can partition each of the surfaces into a constant number of xy-monotone portions (where the constant depends on the maximum degree b). We also assume that the surfaces are in *general position*, in the standard sense considered in, for instance, [17]. One can show that this involves no real loss of generality, but we omit this discussion in this abstract.

As is well known [17], the complexity of the vertical decomposition of $\mathcal{C} = \mathcal{C}_\omega(\Sigma)$ is proportional (up to an additive near-quadratic term) to the number of *vertically-visible* pairs of edges of $\mathcal{C}$. These are ordered pairs (e, e') of edges of $\mathcal{C}$ such that there exists a vertical segment g whose bottom endpoint lies on e, whose top endpoint lies on e', and whose relative interior is contained in $\mathcal{C}$. More precisely, the relevant quantity for measuring the complexity of the vertical decomposition is the number of *vertical visibility configurations* of the form (e, e', g), where e, e' and g are as above. However, the assumptions concerning the surfaces of Σ are easily seen to imply that, under the general position assumption, the maximum number of vertical visibility configurations that correspond to any fixed pair (e, e') of vertically-visible edges is at most some constant s (which depends on the maximum degree b of the surfaces and of their boundaries).

If e or e' is a portion of the boundary of a surface of Σ, we call (e, e', g) an *outer* vertical visibility configuration; otherwise (e, e', g) is an *inner* configuration. We will later show that the overall number of outer configurations is $O(n\lambda_{s+2}(n))$, where $\lambda_s(n)$ is the maximum length of an (n, s) Davenport-Schinzel sequence [2, 14, 17]. Hence, in what follows, we will only consider inner vertical visibility configurations. For convenience, we will not mention the qualifier 'inner' from now on.

For technical reasons, we extend the notion of vertical visibility configurations as follows. Let e and e' be two edges of $\mathcal{A}(\Sigma)$ such that there exists a vertical segment g whose bottom endpoint lies on e and whose top endpoint lies on e'. We say that (e, e', g) is a *vertical edge-crossing* at *level* ξ if

(i) the subset $\Sigma' \subseteq \Sigma$ of surfaces that intersect the relative interior of g has cardinality ξ, and

(ii) g is fully contained in $\mathcal{C}_\omega(\Sigma \setminus \Sigma')$.

Note that the four surfaces incident to e and e' cannot intersect the relative interior of g. Thus, vertical edge-crossings at level 0 are precisely the vertical visibility configurations. We denote by $C_q(\Sigma; \omega)$ the number of vertical edge-crossings (with respect to the cell $\mathcal{C}_\omega(\Sigma)$) of level at most q. We also denote by $C_q(n)$ the maximum possible value of $C_q(\Sigma; \omega)$, over all collections Σ of n surfaces as above, and over all points ω not lying on any surface.

The notion of levels is also extended to vertices and edges of $\mathcal{A}(\Sigma)$: We say that a vertex v (resp. an edge e) of $\mathcal{A}(\Sigma)$ is at level ξ (with respect to the cell $\mathcal{C}_\omega$) if there exists a subset Σ' of ξ surfaces, so that v is a vertex of (resp. e is contained in an edge of) $\mathcal{C}_\omega(\Sigma \setminus \Sigma')$, and if ξ is the smallest number with that property. Again, the actual vertices and edges of $\mathcal{C}_\omega(\Sigma)$ are precisely the vertices and edges at level 0.

Let k be a threshold parameter, whose value will be specified later on. Our goal is to prove a bound on $C_0(n)$ that has roughly the form

$$C_0(n) \leq \frac{1}{k^2} C_k(n) + O(k^\alpha n^{2+\varepsilon}), \qquad (1)$$

where α is some fixed exponent, from which we can deduce the near quadratic bound on $C_0(n)$, by using Clarkson and Shor's technique [9] to bound $C_k(n)$ by $O(k^4 C_0(n/k))$, and by solving the resulting recurrence for C_0. The exact inequality that we will derive will be somewhat weaker than (1), but it will still yield the desired bound on $C_0(n)$.

The idea of proving a bound like (1) is to identify about $k^2 C_0(\Sigma; \omega)$ distinct edge-crossings at level at most k in the arrangement $\mathcal{A}(\Sigma)$. We do (something close to) this, using a two-stage counting argument, similar to that used by Agarwal et al. [1].

Preliminaries. Let e be an edge of $\mathcal{A}(\Sigma)$ and let V_e be the vertical 2-manifold obtained as the union of all z-vertical rays whose bottom endpoints lie on e. The intersection of each surface $\sigma \in \Sigma$ with V_e is a (not necessarily connected) algebraic arc of constant maximum degree (and with a constant number of connected components), so each pair of these arcs intersect in at most some constant number, s, of points (where s depends only on the maximum degree b of the given surfaces and of their boundaries; it is the same parameter s mentioned at the beginning of this section). We denote the set of these arcs by Σ_e, and their arrangement on V_e by $\mathcal{A}(\Sigma_e)$.

Completely analogously, we define the vertical 2-manifold V^e obtained as the union of all downward directed z-vertical rays whose top endpoints lie on e. (Imagine V^e as a 'curtain' hanging down from e, while V_e is a curtain standing on e.) We denote the set of arcs formed by the intersections of the surfaces of Σ with V^e by Σ^e, and their arrangement in V^e by $\mathcal{A}(\Sigma^e)$. We define the *level* of a point p in V_e (resp. in V^e) to be the number of arcs in Σ_e (resp. in Σ^e) that lie below (resp. above) p.

A simple but crucial observation is:

Lemma 2.1 *Let e be an edge of $\mathcal{A}(\Sigma)$ with $\mathcal{C}$ locally above e. Then (e, e', g) is an edge-crossing at level ξ, if and only if the point of $e' \cap V_e$ that lies on the z-vertical line through g is a vertex of $\mathcal{A}(\Sigma_e)$ at level ξ.*

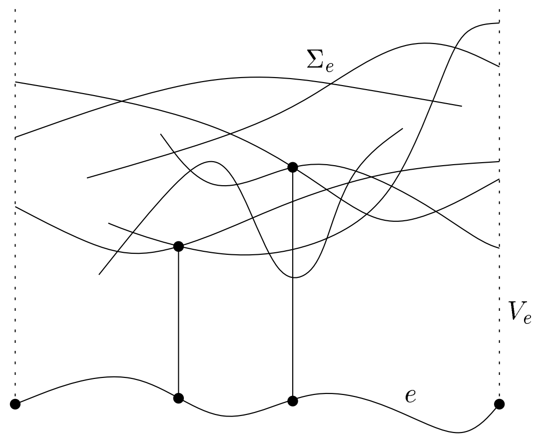

Figure 1: The arrangement $\mathcal{A}(\Sigma_e)$; with a vertical visibility configuration and a vertical edge-crossing at level 3.

See Figure 1 for an illustration. This lemma implies that each vertical visibility configuration with bottom edge e corresponds to a vertex in the lower envelope of the arcs in Σ_e. (Of course, a similar and symmetric statement holds for $\Sigma^{e'}$.)

Now that we have introduced this terminology and observations, we can dispose, as promised, of outer visibility configurations:

Lemma 2.2 *The number of outer vertical visibility configurations is $O(n\lambda_{s+2}(n))$.*

Proof: Let δ be an arc bounding some surface in Σ. By Lemma 2.1 each outer vertical visibility configuration having δ as its bottom edge is represented by some vertex (breakpoint) of the lower envelope of Σ_δ within V_δ. By the standard Davenport-Schinzel theory [2, 14, 17], the number of such breakpoints is $O(\lambda_{s+2}(n))$ (recall that Σ_δ consists of $O(n)$ connected arcs, each pair of which intersect in at most s points). We repeat this analysis for each of the $O(n)$ boundary arcs of the surfaces of Σ, and also apply a symmetric analysis within the "hanging curtains" V^δ. This implies the assertion of the lemma. □

It will be convenient for our analysis to assume that the arrangements $\mathcal{A}(\Sigma_e)$ and $\mathcal{A}(\Sigma^{e'})$ do not contain any arc endpoints at level $\leq 3k$, except on the relative boundaries of V_e and $V^{e'}$. We can achieve this by splitting the edges of $\mathcal{A}(\Sigma)$ into what we call *split edges*, as follows.

Let γ be an arc in Σ_e. If γ has an endpoint w within the relative interior of V_e, which lies in the first $3k$ levels of $\mathcal{A}(\Sigma_e)$, we erect a vertical line through w and split e and V_e at that line into two portions. We repeat this splitting for all edges e of $\mathcal{A}(\Sigma)$ and for each $\gamma \in \Sigma_e$, whenever it is applicable. We apply a symmetric procedure in all the corresponding downward-directed curtains $V^{e'}$. Furthermore, we split all edges e at points where their projection onto the xy-plane has a tangent parallel to the y-axis. This will guarantee that all split edges are x-monotone.

Lemma 2.3 *The overall number of such edge-splittings is $O(k^2 n\lambda_{s+2}(n/k))$.*

Proof: The intersection curve of two surfaces has only a constant number of points where the projection on the xy-plane has a tangent parallel to the y-axis, so the total numer of such points is $O(n^2)$.

We bound the number of splits induced by an endpoint at level at most $3k$ using a similar argument to that of Lemma 2.2. Let δ be an arc bounding some surface in Σ. It is easily seen that each edge-split induced by δ corresponds to a vertex of $\mathcal{A}(\Sigma_\delta)$ in V_δ, or to a vertex of $\mathcal{A}(\Sigma^\delta)$ in V^δ, at level $\leq 3k$, and the overall number of such vertices, over all boundary arcs δ, is known to be $O(k^2 n\lambda_{s+2}(n/k))$ [9, 15]. □

Let's make one final definition before we start with the real proof. For a vertical edge-crossing $\chi = (e, e', g)$, we have four distinct surfaces $\sigma_1, \sigma_2, \sigma_3, \sigma_4 \in \Sigma$, such that $e \subset \sigma_1 \cap \sigma_2$ and $e' \subset \sigma_3 \cap \sigma_4$. Let ℓ be the vertical line through g, and let ℓ' be a copy of ℓ shifted infinitesimally along e in decreasing x-direction. Then $\sigma_3 \cap \ell'$ and $\sigma_4 \cap \ell'$ are two distinct points. We put $\sigma(\chi) = \sigma_3$ if $\sigma_3 \cap \ell'$ lies below $\sigma_4 \cap \ell'$, otherwise $\sigma(\chi) = \sigma_4$.

First Stage. In this stage we identify a set $\mathcal{R}$ of special vertical edge-crossings in $\mathcal{A}(\Sigma)$.

Consider first a (split) edge e of $\mathcal{C}$ with $\mathcal{C}$ locally above e. We partition it into two subedges as follows: We start from the right endpoint of e (recall that all split edges are x-monotone) and move along e to the left until we encounter the $(k+1)$-st distinct surface of Σ directly above the point. We denote the portion of e traversed by this process by e_r, and the remaining part of e by e_l. (It can happen that we encounter the left endpoint of e before seeing more than k distinct surfaces—in that case $e_r = e$ and e_l is empty.) By Lemma 2.1, every edge-crossing at level 0 with bottom point on e_r corresponds to a vertex of the lower envelope of $\mathcal{A}(\Sigma_{e_r})$ on V_{e_r}. Since there are only k surfaces appearing on the lower envelope over e_r, its complexity is at most $O(\lambda_{s+2}(k))$ [2, 14]. Since the number of edges bounding $\mathcal{C}$ is $O(n^{2+\varepsilon})$ [11], and they can be split into $O(k^2 n\lambda_{s+2}(n/k))$ additional split edges, the overall number of vertical visibility configurations involving the right subedge e_r of any split edge of $\mathcal{C}$ is at most $O(\lambda_{s+2}(k)(n^{2+\varepsilon} + k^2 n\lambda_{s+2}(n/k)))$, for any $\varepsilon > 0$. In the following, we will therefore restrict our attention to the vertical visibility configurations that appear above the left subedges e_l of the edges of $\mathcal{C}$.

Consider a pair of surfaces $\sigma_1, \sigma_2 \in \Sigma$, and consider their intersection curve $\sigma_1 \cap \sigma_2$. This curve consists of a constant number of x-monotone connected pieces. Let γ be one such piece. Let V_γ be the union of all V_e, for all (split) edges e of the arrangement that are contained in γ. Let γ' be the subset of γ that consists of the left subedges e_l of all edges $e \subseteq \gamma$ of $\mathcal{C}$ such that $\mathcal{C}$ lies locally above e. Let $\Sigma(\gamma)$ be the set of surfaces in Σ that appear on the lower envelope on V_γ restricted over γ', and let $t = t_\gamma = |\Sigma(\gamma)|$. The number of breakpoints of the lower envelope above γ' is at most $a\lambda_{s+2}(t)$ [2, 14], where a is an appropriate constant (depending on the maximum degree of the surfaces; it arises because, as above, an intersection $\sigma \cap V_\gamma$ may consist of more than one connected arc).

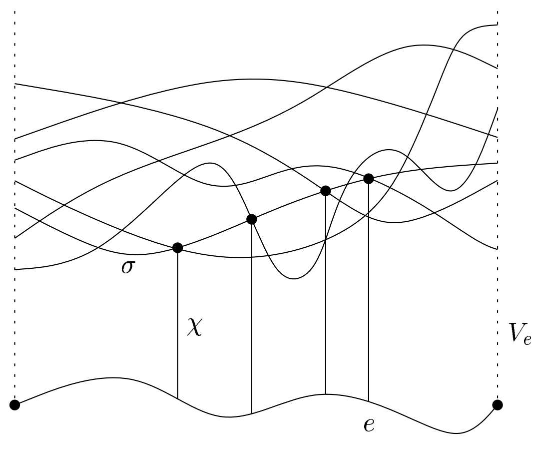

Figure 2: The setup in the construction of $\mathcal{R}$

Consider now a surface $\sigma \in \Sigma(\gamma)$. It appears on the lower envelope over the left subedge of some (split) edge $e \subseteq \gamma$ with $\mathcal{C}$ locally above e, and therefore there are at least k surfaces $\sigma' \in \Sigma$ that appear over e to the right of σ. By continuity and by our construction, either σ and such a surface σ' intersect within V_e at least once, or each of them has a point at level $> k$. We will now collect k vertices on $\sigma \cap V_e$ as follows: We start at some point where σ appears on the lower envelope on V_e (over the left subedge e_l), and follow $\sigma \cap V_e$ in increasing x-direction (recall that all split edges are x-monotone). We will pass, before we reach the end of e, at least k vertices v, at which we encounter a new, distinct surface in Σ, because we must either encounter all the k surfaces that appear above e_r or reach the k-level. All these vertices are at level $\le k$, since when we first reach the k-level, we must have passed all the k surfaces lying below the point. For every such vertex v, let χ_v be the vertical edge-crossing with bottom edge e corresponding to v by Lemma 2.1. Note that $\sigma(\chi_v) = \sigma$. See Figure 2 for an illustration. We let $\mathcal{R}(\gamma, \sigma)$ denote the collection of these edge-crossings χ_v. Note that we collect these k vertices starting from *only one occurrence* of σ along the entire curve γ, so $|\mathcal{R}(\gamma,\sigma)| = k$. This will be used in deriving property (R3) below. Put $\mathcal{R}(\gamma) = \bigcup_{\sigma\in\Sigma(\gamma)} \mathcal{R}(\gamma,\sigma)$, and $\mathcal{R} = \bigcup_\gamma \mathcal{R}(\gamma)$, over all x-monotone pieces γ of intersection curves.

As observed before, the number of visibility configurations above γ' is at most $a\lambda_{s+2}(t)$. On the other hand, we have that

$$|\mathcal{R}(\gamma)| = \left| \bigcup_{\sigma\in\Sigma(\gamma)} \mathcal{R}(\gamma,\sigma) \right| = kt = \tag{2}$$

$$= \Omega\Big(k \tfrac{t}{\lambda_{s+2}(t)} \cdot a\lambda_{s+2}(t)\Big) \ge \tfrac{k}{\beta(n)} \cdot a\lambda_{s+2}(t_\gamma)\,,$$

where $\beta(n) = \Theta(\lambda_{s+2}(n)/n)$ is an extremely slowly growing function of n [2, 14]. Summing (2) over all x-monotone pieces γ of intersection curves, and observing that any edge-crossing in $\mathcal{R}$ is counted in this sum exactly once, we obtain:

$$C_0(\Sigma;\omega) \le \frac{\beta(n)}{k}|\mathcal{R}| + O(\lambda_{s+2}(k)(n^{2+\varepsilon} + k^2 n \lambda_{s+2}(n/k))), \tag{3}$$

for any $\varepsilon > 0$ (where the second term in the right-hand side bounds the number of vertical visibility configurations over the portions e_r). We thus obtain the following lemma.

Lemma 2.4 *Given a set of surfaces Σ and a point ω as above, there is a set $\mathcal{R}$ of vertical edge-crossings such that $|\mathcal{R}|$ satisfies (3) and such that the following conditions hold.*

- *(R1) Each $\chi \in \mathcal{R}$ is at level at most k.*
- *(R2) For each $\chi = (e, e', g) \in \mathcal{R}$, the edge e is an edge of $\mathcal{C}$ such that $\mathcal{C}$ lies locally above e.*
- *(R3) For any three surfaces $\sigma_1, \sigma_2, \sigma_3 \in \Sigma$, there are at most k vertical edge-crossings $\chi = (e, e', g)$ in $\mathcal{R}$ with $e \subseteq \sigma_1 \cap \sigma_2$ and $\sigma(\chi) = \sigma_3$.*
- *(R4) For each $\chi = (e, e', g) \in \mathcal{R}$, and for any surface β that intersects the relative interior of g, there is an intersection point v of β and $\sigma = \sigma(\chi)$ on V_e, so that the portion of $\beta \cap V_e$ between v and g does not meet σ, and the portion of $\sigma \cap V_e$ between v and g does not meet the other surface σ' incident to the top endpoint of g. In addition, if g' denotes the vertical segment connecting e and v, then the "trapezoidal-like" region Δ formed within V_e by g, g', the part of e between g and g', and the part of $\sigma \cap V_e$ between g and g', is intersected by at most k surfaces of Σ (see Figure 3).*

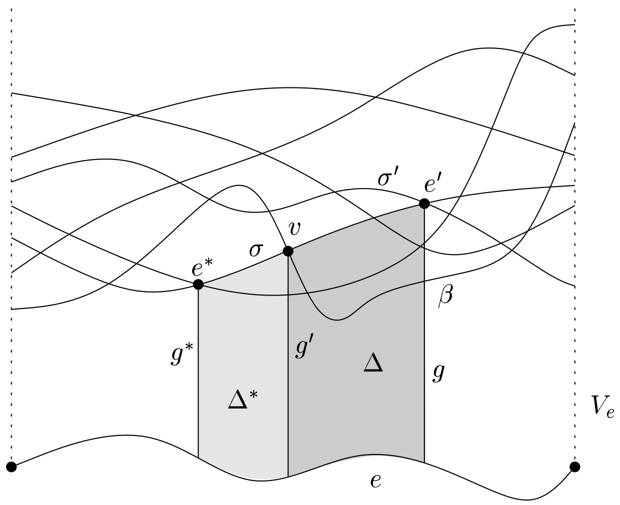

Figure 3: Condition (R4)

Proof: All that remains to be shown is condition (R4): Since $\chi \in \mathcal{R}$ and $\sigma(\chi) = \sigma$, there exists a 0-level edge-crossing of the form (e, e^*, g^*), where e^* is incident to σ, such that, if we follow $\sigma \cap V_e$ from the top endpoint of g^* to the top endpoint of g, we encounter at most $k-1$ distinct surfaces of Σ, and do not encounter σ'. Let Δ^* denote the trapezoidal-like region formed within V_e by g, g^*, the part of e between g and g^*, and the part ζ of $\sigma \cap V_e$ between g and g^*. Then Δ^* can be intersected by at most k surfaces of Σ. To see this, we note that, as just argued, at most k surfaces intersect ζ, and that no surface can intersect e (which is a portion of an edge of $\mathcal{A}(\Sigma)$) or g^*. We claim that no boundary arc of any surface of Σ can cross Δ^*. Indeed, let u be the leftmost point of intersection of a boundary arc with Δ^*. Then the number of surfaces that vertically separate u from e is at most k (any such surface must cross ζ), so that, by construction, such a crossing would have caused e to be split below it. It now follows that any surface crossing Δ^* must also cross ζ, and condition (R4) is thus immediate. ◻

Second Stage. We next bound $|\mathcal{R}|$ in terms of $C_k(n)$. Let e' be a split edge of $\mathcal{A}(\Sigma)$ at level $\le k$. Let $t = t_{e'}$ denote the number of surfaces of Σ that appear on the first $2k$ (top) levels of $\mathcal{A}(\Sigma^{e'})$. If $t \le 4k$, then e' contributes at most $O(k^2)$ edge-crossings to $\mathcal{R}$. We claim that there are only $O(k^{1-\varepsilon}n^{2+\varepsilon})$ split edges e' of $\mathcal{A}(\Sigma)$ at level $\le k$. To see this, charge each such edge e' to one of its endpoints, and observe that the endpoint is either a vertex of $\mathcal{A}(\Sigma)$ at level at most k or a splitting point of the edge of $\mathcal{A}(\Sigma)$ containing e'. By the analysis in [15, 17], the number of endpoints of the first type is $O(k^{1-\varepsilon}n^{2+\varepsilon})$, for any $\varepsilon > 0$. The number of endpoints of the second type is, by Lemma 2.3, $O(k^2 n\lambda_{s+2}(n/k))$, which is subsumed by the first bound. Hence, the overall number of edge-crossings in $\mathcal{R}$ within all the curtains $V^{e'}$ for which $t_{e'} \le 4k$, is $O(k^2 \cdot k^{1-\varepsilon}n^{2+\varepsilon}) = O(k^3 n^{2+\varepsilon})$, for any $\varepsilon > 0$. We can thus assume that $t > 4k$.

We want to repeat, within $V^{e'}$, the analysis of the first step. However, we face a complication that the number of edge-crossings that are counted in $\mathcal{R}$ within $V^{e'}$ could be as large as $\Omega(tk)$, as it is possible that *all* vertices of level $\le k$ in $\mathcal{A}(\Sigma^{e'})$ correspond to edge-crossings in $\mathcal{R}$. Our goal in this second stage is to bound $|\mathcal{R}|$ by something close to $\frac{1}{k}C_k(n)$, but the technique of the first counting stage will not imply this when $V^{e'}$ is 'full' of edge-crossings in $\mathcal{R}$. To overcome this problem, we will first bound the number of 'excessive' edge-crossings in $\mathcal{R}$ within $V^{e'}$, using a different approach, and only then bound the number of remaining crossings, using the same approach as in the first stage.

We first identify a class of vertical edge-crossings in $\mathcal{R}$ (the 'excessive' edge-crossings), whose number we will be able to bound independently. We define an edge-crossing (e, e', g) at level at most k to be *covered* if it satisfies the following condition.

> Let σ_1, $\sigma_2 \in \Sigma$ be the two surfaces incident to e. There is a surface $\beta \in \Sigma$ that intersects the relative interior of g, and either σ_1 or σ_2 (say, for definiteness, σ_1) crosses β within $V^{e'}$, either to the left or to the right of g, at some point w. Moreover, if g^* denotes the vertical segment connecting w to e', then the trapezoidal-like portion of $V^{e'}$ bounded by g, g^*, the portion of e' between g and g^*, and the portion of σ_1 between g and g^*, is crossed by at most $2k$ surfaces of Σ.

See Figure 4 for an illustration of this definition; note that the definition encompasses several different subcases, as is illustrated in the figure.

We can now establish the two central lemmas.

Lemma 2.5 *The number of uncovered edge-crossings in $\mathcal{R}$ within $V^{e'}$ is $O(t_{e'}\beta^2(t_{e'}))$.*

Proof: In the first step of the proof, we partition the arcs of $\Sigma^{e'}$ into 'small' subarcs, as follows. Recall that, by construction, each endpoint of any arc in $\Sigma^{e'}$ must either lie on the relative boundary of $V^{e'}$, or else be at level $> 3k$. Replace $\Sigma^{e'}$ by the subcollection Σ^* of only those $t = t_{e'}$ arcs that appear within the first $2k$ levels of $\mathcal{A}(\Sigma^{e'})$. The total number of vertices of $\mathcal{A}(\Sigma^*)$ at level $\le 2k$ is $O(k^2\lambda_{s+2}(t/k))$ [15, 17], and so there must exist a level $k \le k^* \le 2k$ that contains only $O(k\lambda_{s+2}(t/k)) = O(t\beta(t/k))$ vertices. Hence, the portion $\mathcal{A}^+(\Sigma^*)$ of $\mathcal{A}(\Sigma^*)$ that lies at level $\le k^*$ is formed by $O(t\beta(t/k))$ connected subarcs. Moreover, since the number of vertices in $\mathcal{A}^+(\Sigma^*)$ is $O(kt\beta(t/k))$, we can partition further each of these subarcs into smaller connected pieces, so that each piece is incident to at most k vertices of $\mathcal{A}^+(\Sigma^*)$, and so that the overall number of these smaller subarcs is still $O(t\beta(t/k))$. Let Σ^+ denote the resulting collection of the new subarcs. Note that all vertices of $\mathcal{A}(\Sigma^{e'})$ at level $\le k^*$ are also vertices of $\mathcal{A}(\Sigma^+)$.

Define the *level* $\ell(\delta)$ of a subarc $\delta \in \Sigma^+$ to be the smallest level in $\mathcal{A}(\Sigma^+)$ of any point on δ. Note that, by construction, $\ell(\delta)$ is always at most k^*, and it is also equal to the smallest level in $\mathcal{A}(\Sigma^{e'})$ of any point on δ. We have:

Claim: Let (e, e', g) be an uncovered edge-crossing in $\mathcal{R}$ within $V^{e'}$, and suppose that the relative interior of g is crossed by $h \le k$ surfaces. Let σ_1 and σ_2 be the two surfaces incident to e, and let δ_1, $\delta_2 \in \Sigma^+$ be the respective subarcs of $\sigma_1 \cap V^{e'}$, $\sigma_2 \cap V^{e'}$ incident to the bottom endpoint v of g. Then $\ell(\delta_1) = \ell(\delta_2) = h$.

Proof of Claim: Note that, by definition, both $\ell(\delta_1)$ and $\ell(\delta_2)$ are at most h. Suppose to the contrary that, say $\ell(\delta_1) < h$. Then there is a point $v^* \in \delta_1$ whose level h^* is strictly smaller than h; see Figure 5. This implies that one of the surfaces, call it β, that crosses the relative interior of g cannot cross the relative interior of the vertical segment g^* that connects v^* to e'. But then β must cross δ_1 at some point w between v and v^* (β cannot cross e' and since δ_1 has no point of level $> k^*$, there cannot be a boundary point above δ_1 within $V^{e'}$). It is now easy to see that (e, e', g) is a covered edge-crossing. Indeed, let Δ denote the trapezoidal-like region formed within $V^{e'}$ by g, the vertical segment connecting w to e', the portion of δ_1 between v and w, and e'. Any surface of Σ intersecting Δ must either intersect the interior of g or must form a vertex on δ_1 between v and w. There are at most k surfaces of the first kind, and at most $k-2$ of the second (recall that δ_1 has at most k vertices in $\mathcal{A}(\Sigma^{e'})$), and so Δ is crossed by less than $2k$ surfaces. This contradiction completes the proof of the claim. ◻

Let Σ_h^+ denote the set of all subarcs of Σ^+ whose level is h, for $h = 1, \ldots, k$, and put $t_h = |\Sigma_h^+|$. If (e, e', g) is as in the Claim, then the corresponding vertex v is a vertex of the upper envelope of Σ_h^+ within $V^{e'}$: Indeed,

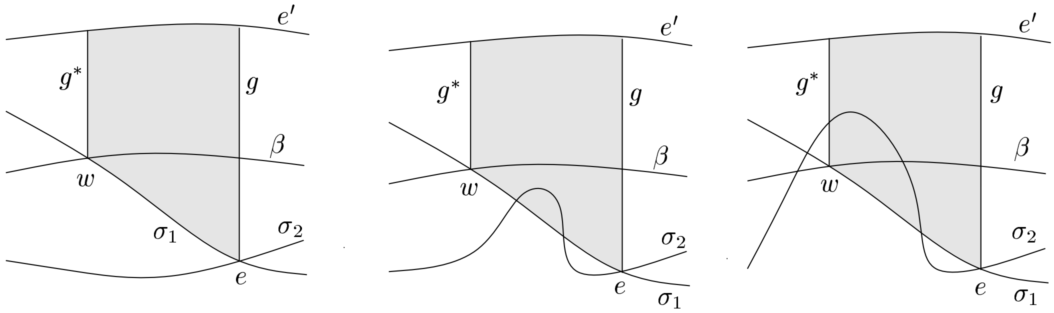

Figure 4: Several types of covered edge-crossings

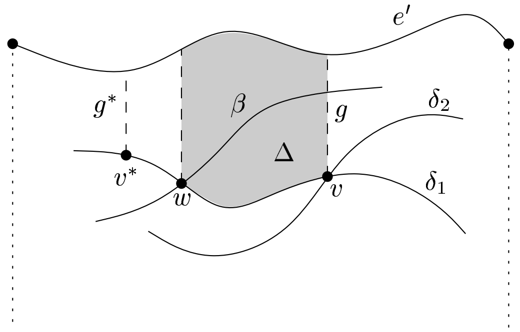

Figure 5: The proof of the bound on the number of uncovered edge-crossings

the Claim implies that v is a vertex of $\mathcal{A}(\Sigma_h^+)$, and no subarc of Σ_h^+ can pass above v, because the level of any such subarc must be strictly smaller than h. Hence, the number of uncovered edge-crossings in $\mathcal{R}$ within $V^{e'}$ is

$$\sum_{h=1}^{k} \lambda_{s+2}(t_h) \le \lambda_{s+2}\left(\sum_{h=1}^{k} t_h\right) = $$
$$= \quad \lambda_{s+2}\left(O(t\beta(t/k))\right) = O(t\beta^2(t)) . \qquad \Box$$

Lemma 2.6 *The total number of covered edge-crossings in $\mathcal{R}$, in the entire arrangement, is $O(k^4 n^{2+\varepsilon})$, for any $\varepsilon > 0$.*

Proof: Let $\mathcal{R}'$ be the set of all covered edge-crossings in $\mathcal{R}$, but with certain multiplicities removed: For every four surfaces $\sigma_1, \sigma_2, \sigma_3, \sigma_4 \in \Sigma$, we choose only one edge-crossing (e, e', g) with $e \subseteq \sigma_1 \cap \sigma_2$ and $e' \subseteq \sigma_3 \cap \sigma_4$ for $\mathcal{R}'$, and for every three surfaces $\sigma_1, \sigma_2, \sigma_3 \in \Sigma$, we choose only one edge-crossing $\chi = (e, e', g)$ with $e \subseteq \sigma_1 \cap \sigma_2$ and $\sigma(\chi) = \sigma_3$ for $\mathcal{R}'$. Exploiting condition (R3) and the fact that there are at most s edge-crossings defined by the same four surfaces, it is now sufficient to show that $|\mathcal{R}'| = O(k^3 n^{2+\varepsilon})$, for any $\varepsilon > 0$.

Let $\chi = (e, e', g)$ be a covered edge-crossing in $\mathcal{R}'$. Let σ_1, σ_2, and $\beta \in \Sigma$ be as in the definition of a covered edge-crossing, and let's assume that σ_1 crosses β within $V^{e'}$ to the left of g at some point w. By condition (R4), there is an intersection point v of β and σ on V_e, where $\sigma = \sigma(\chi)$ is one of the two surfaces defining e', the portion of $\beta \cap V_e$ between v and g does not meet σ, and the portion of $\sigma \cap V_e$ between v and g does not meet the other surface σ' incident to the top endpoint of g. Let Δ be the trapezoidal-like region within V_e, as defined in condition (R4), and let Δ' be the trapezoidal-like region within $V^{e'}$, as in the definition of covered edge-crossings. By the preceding arguments, at most $3k$ surfaces cross either Δ or Δ'. See Figure 6 for an illustration.

Let K be a random sample of n/k surfaces of Σ. As argued in [9], the probability that K contains the five surfaces σ_1, σ_2, β, σ and σ', and does not contain any of the other surfaces crossing $\Delta \cup \Delta'$, is at least c/k^5, for some absolute constant $c > 0$. Let $\mathcal{R}'_K$ be the set of edge-crossings $\chi \in \mathcal{R}'$ that appear in $\mathcal{A}(K)$, in the sense that the above choice of surfaces in K materializes. Since the expected size of $\mathcal{R}'_K$ is at least c/k^5 times the size of $\mathcal{R}'$, it suffices to prove that $|\mathcal{R}'_K| = O(n^{2+\varepsilon}/k^2)$, for any $\varepsilon > 0$.

So let $\mathcal{C}(K) = \mathcal{C}_\omega(K)$ denote the cell in $\mathcal{A}(K)$ that contains ω, and consider an edge-crossing $\chi = (e, e', g) \in \mathcal{R}'_K$. Clearly, the segment g crosses in $\mathcal{A}(K)$ only the surface β, and its portion below β is fully contained in $\mathcal{C}(K)$. For technical reasons, we distinguish between the two (top and bottom) sides of each surface in Σ; we appeal to the intuition of the reader, and refer to [3, 4, 11] for a formal definition. An important property that this distinction has is that a curve drawn on the top side of a surface is not considered to cross a curve drawn on the bottom side.

We will construct, for every $\chi \in \mathcal{R}'_K$, a path $\pi = \pi(\chi)$ on $\partial\mathcal{C}(K)$ as the concatenation of the following three subpaths:

- $\pi_1(\chi)$: the subarc of $\beta \cap V_e$ connecting v to $u = g \cap \beta$;
- $\pi_2(\chi)$: a subarc of $\beta \cap V^{e'}$ extending from u towards w, stopping as soon as it hits either σ_1 or σ_2 in a point w' (w' may or may not be equal to w, see the two cases in Figure 6);
- $\pi_3(\chi)$: this subpath extends from w' along $\sigma_1 \cap V^{e'}$

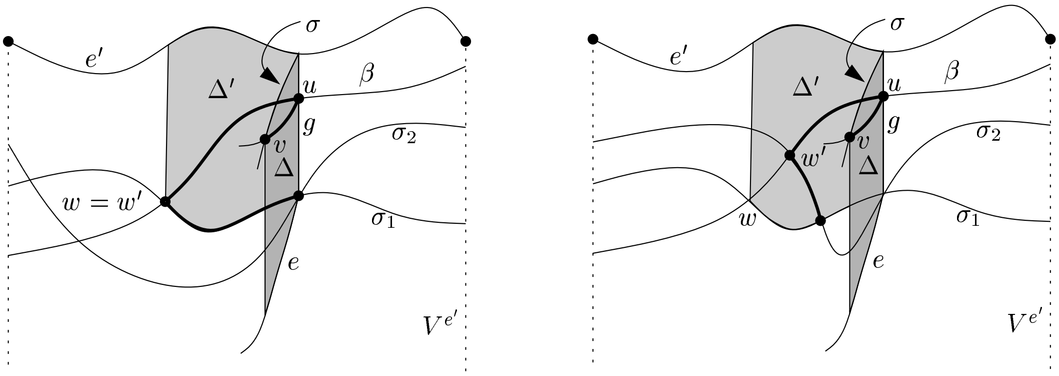

Figure 6: Two cases of covered edge-crossings $\chi = (e, e', g) \in \mathcal{R}$, and the corresponding paths $\pi(\chi)$ (shown in fat)

or $\sigma_2 \cap V^{e'}$ (depending on which surface contains w') towards g, and stops as soon as it hits the other surface defining e (this can happen either at g or earlier, see Figure 6).

The first two portions $\pi_1(\chi)$ and $\pi_2(\chi)$ are drawn on the *bottom* side of β, and the last portion $\pi_3(\chi)$ is drawn on the *top* side of σ_1.

We draw a path $\pi(\chi)$ for each edge-crossing $\chi \in \mathcal{R}'_K$, in all curtains $V^{e'}$.

Claim: Two distinct paths $\pi(\chi)$, $\pi(\chi^*)$ are disjoint.

Proof of Claim: Suppose to the contrary that $\pi(\chi)$ and $\pi(\chi^*)$ contain a common point, for two distinct edge-crossings $\chi = (e, e', g)$ and $\chi^* = (e^*, (e')^*, g^*)$ in $\mathcal{R}'_K$. There are several cases.

$\pi_1(\chi)$ intersects $\pi_1(\chi^*)$: By our construction, the path $\pi_1(\chi)$ lies above an edge of $\mathcal{A}(K)$ (the one containing e), starts at an edge of $\mathcal{A}(K)$ and stops as soon as it passes directly below another edge (the edge e'); a similar property holds for $\pi_1(\chi^*)$. It follows that if $\pi_1(\chi)$ and $\pi_1(\chi^*)$ have a common point (including the sharing of an endpoint), then $e' = (e')^*$, $e = e^*$, and the two paths must coincide completely. This implies that χ and χ^* are defined by the same four surfaces, which is a contradiction to the definition of $\mathcal{R}'$.

$\pi_2(\chi)$ intersects $\pi_2(\chi^*)$: The path $\pi_2(\chi)$ passes directly below an edge of $\mathcal{A}(K)$ (the one containing e'), ends on an edge of $\mathcal{A}(K)$, and starts directly above an edge of $\mathcal{A}(K)$ (the one containing e). It may pass above other edges $\bar{e}$ of $\mathcal{A}(K)$, but then $\bar{e}$ is defined by the same two surfaces as e. It follows that if $\pi_2(\chi)$ and $\pi_2(\chi^*)$ have a common point, then $e' = (e')^*$, and hence $\sigma(\chi) = \sigma(\chi^*)$. Furthermore, e and e^* must be defined by the same two surfaces of Σ. Again, this is impossible by our definition of $\mathcal{R}'$.

$\pi_3(\chi)$ intersects $\pi_3(\chi^*)$: The path $\pi_3(\chi)$ lies below an edge of $\mathcal{A}(K)$ (the one containing e'), and exactly one surface of K (namely, β) passes between the path and the edge. Hence, if $\pi_3(\chi)$ and $\pi_3(\chi^*)$ have a common point, then $e' = (e')^*$, and $\pi_3(\chi) = \pi_3(\chi^*)$ (since they both lie in $V^{e'}$ and extend in both directions until they hit an edge of $\mathcal{A}(K)$). This implies that χ and χ^* are again defined by the same four surfaces, which is impossible.

Finally, $\pi_1(\chi)$ intersects $\pi_2(\chi^*)$ (or, symmetrically, $\pi_1(\chi^*)$ intersects $\pi_2(\chi)$): A point $p \in \pi_2(\chi^*)$ has to lie below an edge of $\mathcal{A}(K)$. The only point with this property on $\pi_1(\chi)$ is its endpoint u. On the other hand, u lies directly above an edge of $\mathcal{A}(K)$. The only such point on $\pi_2(\chi^*)$ is *its* endpoint, and hence $e = e^*$ and $e' = (e')^*$, a contradiction.

This completes the proof of the Claim. $\square$

We continue with the proof of the lemma. We now have a system G of pairwise openly-disjoint paths drawn on $\partial\mathcal{C}_\omega(K)$, and our next goal is to bound their number, using Euler's formula for planar graphs, in a manner similar to, though somewhat more complex than, the technique used by Tagansky [18]. This is done as follows.

Fix a face f of $\partial\mathcal{C}_\omega(K)$ (which lies on either the top side or the bottom side of some surface), and clip all paths that cross f to within f (note that either all these clippings retain the first two portions of each such path, if f lies on the bottom side of a surface, or they all retain the third portions of these paths, if f lies on the top side). Let G_f denote the resulting collection of clipped paths. We regard G_f as a plane drawing of a graph, whose nodes are the edges of f and whose arcs are the clipped paths. Since f is (homeomorphic to) a planar region, we do indeed obtain a plane drawing of a planar graph, and we can apply Euler's formula to conclude that the number of arcs in G_f is at most three times the number of edges of f, plus the number of faces of G_f of degree 2. Applying this analysis to each (sided) face of $\mathcal{C}_\omega(K)$, and summing up these bounds, we conclude that the overall number of clipped subpaths is proportional to the complexity of $\mathcal{C}_\omega(K)$, which is $O((n/k)^{2+\varepsilon})$, for any $\varepsilon > 0$ [11], plus the overall number of graph-faces of degree 2.

To get a better handle on those degree-2 faces, we go over all faces f of $\mathcal{C}_\omega(K)$, take each 'run' of adjacent degree-2 faces within f, and delete all their incident subpaths, except for the first and the last one. Clearly, the number of remaining subpaths is $O((n/k)^{2+\varepsilon})$, for any $\varepsilon > 0$.

Now take a full path $\pi(\chi)$. If either of its two clipped subpaths has survived after the above trimming, we charge $\pi(\chi)$ to that subpath. Since this charging is unique, the number of paths $\pi(\chi)$ of this kind is $O((n/k)^{2+\varepsilon})$, for any $\varepsilon > 0$. Suppose then that both clipped subpaths of $\pi(\chi)$ have been trimmed. This is easily seen to imply that there is a sequence of 'parallel' paths, all of which connect between the same pair of edges of $\mathcal{C}_\omega(K)$, such that $\pi(\chi)$ is a middle element of the sequence. Recall that, in the notations used above, one of the terminal edges of $\pi(\chi)$ is incident to the two surfaces σ_1, σ_2 that meet also at the bottom endpoint of g (recall that the endpoint of $\pi(\chi)$ needs not coincide with the endpoint of g), and the other terminal edge is incident to β and to $\sigma(\chi)$. It follows that there are at least three different paths $\pi(\chi')$ in the above sequence, such that the corresponding edge-crossings χ' share the three surfaces σ_1, σ_2, and $\sigma = \sigma(\chi')$. By our definition of $\mathcal{R}'$, this is impossible, and therefore every path $\pi(\chi)$ is uniquely charged to one of its subpaths.

All these arguments readily imply that the total number of paths $\pi(\chi)$ that are drawn on $\partial\mathcal{C}_\omega(K)$ is $O((n/k)^{2+\varepsilon})$, for any $\varepsilon > 0$. This implies that $|\mathcal{R}'_K| = O(n^{2+\varepsilon}/k^2)$, and by our conclusions above, this completes the proof of the lemma. □

We now show how these bounds imply our main result. Put $t^* = \sum_{e'} t_{e'}$, where the sum extends over all split edges e' at level at most k, for which $t_{e'} > 4k$. Lemmas 2.5 and 2.6 imply that

$$|\mathcal{R}| = O(k^4 n^{2+\varepsilon}) + \sum_{e'} O(t_{e'}\beta^2(t_{e'})) = \\ = O(k^4 n^{2+\varepsilon}) + O(t^*\beta^2(n)) .$$

On the other hand, for each edge e' as above, the number of vertices appearing in the first $3k$ levels of $\mathcal{A}(\Sigma^{e'})$ is $\Omega(t_{e'}k)$. Indeed, the number of surfaces whose intersection arcs are fully contained within the first $3k$ levels of $\mathcal{A}(\Sigma^{e'})$ is at most $3k$, because, for every such surface σ, the endpoints of the curve $\sigma \cap V^{e'}$ must be among the first $3k$ curves below each endpoint of e' (the curve cannot have an endpoint in the interior of $V^{e'}$, because such an endpoint would have caused e' to be further split). Since $t_{e'} > 4k$, at least $t_{e'} - 3k > \frac{1}{4}t_{e'}$ of these curves have a point at level $> 3k$, and thus each of them must contain at least k vertices of $\mathcal{A}(\Sigma^{e'})$ at level $\leq 3k$ (because, by definition, it also shows up among the top $2k$ levels). Since each such vertex induces an edge-crossing at level at most $3k$ (with respect to our cell $\mathcal{C}$), it follows that

$$C_{3k}(\Sigma;\omega) = \Omega(k\sum_{e'} t_{e'}) = \Omega(t^* k) .$$

Hence, we have

$$|\mathcal{R}| = O(k^4 n^{2+\varepsilon}) + O\left(\frac{\beta^2(n)}{k}\cdot kt^*\right) = \\ = O(k^4 n^{2+\varepsilon}) + O\left(\frac{\beta^2(n)}{k}\cdot C_{3k}(n)\right) . \quad (4)$$

Next, we estimate $C_{3k}(n)$, by using the probabilistic technique of Clarkson and Shor [9] (see also [15]). Since each edge-crossing (e, e', g) is defined by four surfaces (two surfaces incident to e and two incident to e'), the Clarkson-Shor technique is easily seen to imply that $C_{3k}(n) = O(k^4 C_0(n/k))$. Combining (3), (4), and the Clarkson-Shor bound, we readily obtain

$$\begin{aligned} C_0(n) &= O((k\beta(k) + k^3\beta(n))n^{2+\varepsilon} + \\ &\quad + k^3\beta(k)n\lambda_{s+2}(n/k)) + \\ &\quad + \frac{\beta^3(n)}{k^2}\cdot O(k^4 C_0(n/k)) \\ &= O(k^3\beta(n)n^{2+\varepsilon}) + O(k^2\beta^3(n))\cdot C_0(n/k) . \end{aligned}$$

The solution of this recurrence is $O(n^{2+\delta})$, for any $\delta > \varepsilon$. This is shown by induction on n, choosing $k = \beta^{1+3/\delta}(n)$ and using the fact that $\beta(n)$ is an extremely slowly growing function of n (see also [1]).

In conclusion, we have thus obtained the main result of the paper:

Theorem 2.7 *The complexity of the vertical decomposition of a single cell in an arrangement of n algebraic surface patches in $\mathbb{R}^3$, such that the degrees of the surfaces and of their boundary curves are all bounded by some constant b, is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε and on b.*

3 Applications

3.1 Constructing a Single Cell

Theorem 3.1 *Given an arrangement of n algebraic surface patches in $\mathbb{R}^3$, such that the degrees of the surfaces and of their boundary curves are all bounded by some constant b, one can construct, in an appropriate model of computation, the cell of the arrangement containing a given point, in randomized expected time $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε and on b.*

Proof: As noted in the introduction, this can be accomplished, in a rather routine manner, by applying the lazy randomized algorithm of de Berg et al. [6]. We omit the details in this version. □

3.2 Motion Planning for Systems with Three Degrees of Freedom

As noted in the introduction, the main application of Theorems 2.7 and 3.1 is to motion planning for arbitrary systems with three degrees of freedom. This application is described in detail in the survey paper [13]. The results of our paper imply the following:

Theorem 3.2 *Let B be a robot system with three degrees of freedom, such that the free configuration space of B can be described as a Boolean combination of n polynomial equalities and inequalities, of constant maximum degree b, in the three parameters that define the degrees of freedom of B. Then, given any two free placements Z_1, Z_2 of B, one can determine, in randomized expected time $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, whether there exists a collision-free motion of B from Z_1 to Z_2, and, if so, produce such a motion. (The constant of proportionality in this bound depends on ε and on b.)*

Acknowledgments

We wish to thank Danny Halperin for useful discussions concerning the problems studied in this paper. Part of the work on the paper has been carried out in the Mathematical Research Institute of Tel Aviv University, which the first author has visited in the spring of 1995.

References

[1] P.K. Agarwal, O. Schwarzkopf and M. Sharir, The overlay of lower envelopes in three dimensions and its applications, *Discrete Comput. Geom.* 15 (1996), 1–13.

[2] 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.

[3] B. Aronov and M. Sharir, Triangles in space, or: Building (and analyzing) castles in the air, *Combinatorica* 10 (2) (1990), 137–173.

[4] B. Aronov and M. Sharir, Castles in the air revisited, *Discrete Comput. Geom.* 12 (1994), 119–150.

[5] B. Aronov and M. Sharir, On translational motion planning in three dimensions, *Proc. 10th ACM Symp. on Computational Geometry,* 1994, pp. 21–30.

[6] M. de Berg, K. Dobrindt and O. Schwarzkopf, On lazy randomized incremental construction, *Discrete Comput. Geom.* 14 (1995), 261–286.

[7] B. Chazelle, H. Edelsbrunner, L. Guibas and M. Sharir, A singly exponential stratification scheme for real semi–algebraic varieties and its applications, *Proc. 16th Int. Colloq. on Automata, Languages and Programming*, 1989, pp. 179–193. (Also in *Theoretical Computer Science* 84 (1991), 77–105.)

[8] 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.

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

[10] 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.

[11] D. Halperin and M. Sharir, Almost tight upper bounds for the single cell and zone problems in three dimensions, *Discrete Comput. Geom.* 14 (1995), 385–410.

[12] 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. (Also to appear in *Discrete Comput. Geom.*)

[13] D. Halperin and M. Sharir, Arrangements and their applications in robotics: Recent developments, in *The Algorithmic Foundations of Robotics* (K. Goldberg, D. Halperin, J.C. Latombe and R. Wilson, Eds.), A.K. Peters, Boston MA, 1995, pp. 495–511.

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

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

[16] M. Sharir, Almost tight upper bounds for lower envelopes in higher dimensions, *Discrete Comput. Geom.* 12 (1994), 327–345.

[17] M. Sharir and P. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.

[18] B. Tagansky, A new technique for analyzing substructures in arrangements, *Proc. 11th ACM Symp. on Computational Geometry,* 1995, pp. 200–210.