The Overlay of Lower Envelopes and its Applications*

Pankaj K. Agarwal† Otfried Schwarzkopf‡ Micha Sharir§

March 16, 1996

Abstract

Let $\mathcal{F}$ and $\mathcal{G}$ be two collections of a total of n (possibly partially-defined) bivariate algebraic functions of constant maximum degree. The minimization diagrams of $\mathcal{F}$, $\mathcal{G}$ are the planar maps obtained by the xy-projections of the lower envelopes of $\mathcal{F}$, $\mathcal{G}$, respectively. We show that the combinatorial complexity of the overlay of the minimization diagrams of $\mathcal{F}$ and of $\mathcal{G}$ is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. This result has several applications: (i) a near-quadratic upper bound on the complexity of the region in 3-space enclosed between the lower envelope of one such collection of functions and the upper envelope of another collection; (ii) an efficient and simple divide-and-conquer algorithm for constructing lower envelopes in three dimensions; and (iii) a near-quadratic upper bound on the complexity of the space of all plane transversals of a collection of simply-shaped convex sets in three dimensions.

1 Introduction

Let $\mathcal{F} = \{f_1, \ldots, f_n\}$ be a collection of n d-variate, possibly partially defined, functions, all algebraic of some constant maximum degree b (and if they are partially defined, the domain of definition of each f_i is also described by a constant number of polynomial equalities and inequalities of maximum degree b). Abusing the notation slightly, we will not distinguish

*Work on this paper by the first author has been supported by National Science Foundation Grant CCR-93–01259 and an NYI award. Work on this paper by the second author has been supported by the Netherlands' Organization for Scientific Research (NWO) and partially supported by ESPRIT Basic Research Action No. 6546 (project PROMotion). Work on this paper by the third author has been supported by NSF Grant CCR-91-22103, by a Max-Planck Research Award, and by grants from the U.S.-Israeli Binational Science Foundation, the Fund for Basic Research administered by the Israeli Academy of Sciences, and the G.I.F., the German-Israeli Foundation for Scientific Research and Development.

†Department of Computer Science, Box 1029, Duke University, Durham, NC 27708-0129, USA

‡Department of Computer Science, Utrecht University, P.O. Box 80.089, 3508 TB Utrecht, the Netherlands

§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

between a function and its graph. The *lower envelope* $E_{\mathcal{F}}$ of $\mathcal{F}$ is defined as

$$E_{\mathcal{F}}(\mathbf{x}) = \min_i f_i(\mathbf{x}) ,$$

where the minimum is taken over all functions of $\mathcal{F}$ that are defined at $\mathbf{x}$. Similarly, we define the *upper envelope* $E'_{\mathcal{F}}$ of $\mathcal{F}$ as

$$E'_{\mathcal{F}}(\mathbf{x}) = \max_i f_i(\mathbf{x}) .$$

The *minimization diagram* $M_{\mathcal{F}}$ of $\mathcal{F}$ is the decomposition of $\mathbb{R}^d$ into maximal connected relatively open cells, of any dimension, so that within each cell the same subset of functions appear on the envelope $E_{\mathcal{F}}$. If the functions of $\mathcal{F}$ are partially defined, we also require that, over each cell c, each of the polynomials defining the domain of definition of any function that attains $E_{\mathcal{F}}$ over c has a fixed sign. Informally, this means that if a function f attains $E_{\mathcal{F}}$ over a cell c, then either c is fully contained in the boundary of the domain of f or is disjoint from that boundary. The *combinatorial complexity* of $M_{\mathcal{F}}$ and of $E_{\mathcal{F}}$ is the number of cells of all dimensions in $M_{\mathcal{F}}$. The *maximization diagram* and its combinatorial complexity are defined in an analogous manner.

Recently there has been significant progress in the analysis of the combinatorial complexity of lower envelopes of multivariate functions [15, 19]. In particular, it was shown in [19] that the maximum complexity of $M_{\mathcal{F}}$ is $O(n^{d+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε, d, and b. This result almost settles a major open problem and has already led to many applications [1, 15, 19]. In some applications, however, one has to consider the interaction between the lower envelope of one collection of functions and the upper envelope of another collection. A major application of this type, which has motivated the work on the present paper, is the analysis of the combinatorial complexity of the space of *k-transversals* of a collection $\mathcal{C}$ of n compact convex sets in d dimensions; a k-transversal is a k-flat that intersects all the sets of $\mathcal{C}$ (see [11, 12, 13]). Using an appropriate coordinate system for representing the space of k-flats in $\mathbb{R}^d$ (as is well known, the dimension of that space is $N = (k+1)(d-k)$), one can show that the space of k-transversals of $\mathcal{C}$ can be represented as the region enclosed between the upper envelope of one collection of functions and the lower envelope of another collection, where each function in the first (resp. second) collection represents all k-flats that are tangent to one of the given sets from below (resp. from above). Hence, the study of spaces of transversals calls for combinatorial (as well as algorithmic) analysis of the region enclosed between two envelopes in higher dimensions [11, 13]. Edelsbrunner et al. [11] showed that the complexity of the region between the two envelopes of d-variate, partially-defined, linear functions is $O(n^d \alpha(n))$, where $\alpha(n)$ is the inverse Ackermann function, which in turn yields a near-optimal bound on the complexity of the space of hyperplane transversals for convex polytopes. No such bound is known for nonlinear functions, even for $d = 2$.

In this paper we provide such an analysis for the case of bivariate functions. We show that the combinatorial complexity of the region enclosed between the lower envelope of a

collection of n low-degree bivariate algebraic functions and the upper envelope of another collection of n such functions, is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends, as in the case of a single envelope, on ε, on the maximum degree of the given functions, and, in case of partial functions, of their domain boundaries. In other words, the worst-case complexity of the region in question is asymptotically no worse than that of a single envelope.

The proof uses techniques that resemble those used in the proofs given in [15, 19], but requires several additional tricks. The basic result that we derive in this paper, which we consider to be interesting in its own right, is the analysis of the combinatorial complexity of the overlay of the minimization diagrams of the lower envelopes of two collections of bivariate functions. Notice that this problem is easy for the case of univariate functions, because the complexity of the overlay of the x-projections of two envelopes of univariate functions is proportional to the sum of the complexities of the individual envelopes. This is, however, not true for envelopes of bivariate functions; see Figure 1. Nevertheless, we show that the complexity of the overlay of the minimization diagrams of two collections of a total of n functions in 3-space is only $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. This result not only implies the asserted bound on the complexity of the region enclosed between two envelopes, but also has several other useful applications, among which is a simple, deterministic, divide-and-conquer algorithm for computing lower envelopes, which we believe to be conceptually much simpler than the competing techniques of [4, 9, 19].

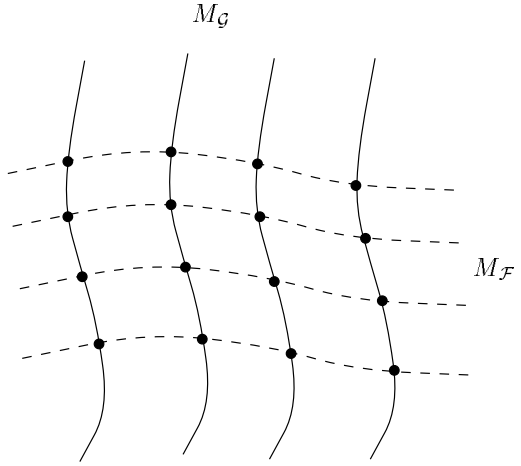

Figure 1: Overlay of two envelopes with quadratic complexity

The paper is organized as follows. In Section 2 we prove the main result concerning the overlay of the projections of two envelopes in 3-space. In section 3 we apply the result to obtain:

- an efficient and simple divide-and-conquer algorithm for constructing lower envelopes in 3 dimensions; and
- a near-quadratic upper bound on the complexity of the region enclosed between a lower envelope and an upper envelope;
- a near-quadratic upper bound on the complexity of the space of all plane transversals of an arbitrary collection of simply-shaped convex sets in 3 dimensions.

2 Complexity of the Overlay of Two Envelopes

Let $\mathcal{F}$ and $\mathcal{G}$ be two given families of a total of n bivariate functions, satisfying the following condition:

($\star$) Each $f \in \mathcal{F} \cup \mathcal{G}$ is a continuous, totally or partially defined, bivariate algebraic function of constant maximum degree b; if f is only partially defined, the domain of definition of f is bounded by a constant number of algebraic arcs of constant maximum degree b.

Let M denote the planar map obtained by superimposing $M_{\mathcal{F}}$ and $M_{\mathcal{G}}$. We refer to M as the *overlay* of $M_{\mathcal{F}}$ and of $M_{\mathcal{G}}$. We first prove an upper bound on M for the case when the functions in $\mathcal{F} \cup \mathcal{G}$ are totally defined, and then extend the proof to partially defined functions.

Theorem 2.1 *Let $\mathcal{F}$ and $\mathcal{G}$ be two collections of n totally defined bivariate functions, satisfying condition ($\star$). Then the combinatorial complexity of the overlay of the minimization diagrams of $\mathcal{F}$ and $\mathcal{G}$, as defined above, is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$ (where the constant of proportionality depends on ε and on the maximum degree b).*

Proof: We use a two-stage counting argument to obtain a recurrence for the complexity of the overlay. For the sake of simplicity, we assume that the functions in $\mathcal{F} \cup \mathcal{G}$ are in *general position*. This excludes degenerate configurations where four function graphs meet at a point, a pair of graphs are tangent to each other, a singular point on one graph lies on an intersection curve between two other graphs, etc.) Similar conditions were assumed in the papers [15, 19]. We refer the reader to these papers for more details, and for an argument that no real loss of generality is made by assuming general position. An appropriate variant of this argument shows that our proof can also be extended to collections $\mathcal{F}, \mathcal{G}$ not in general position.

Our general position assumption implies that over each face of $M_{\mathcal{F}}$ the envelope is attained by a single function (or by no function at all), that over each edge the envelope is attained by two functions simultaneously, and that over each vertex of $M_{\mathcal{F}}$ the envelope is attained by three functions simultaneously. By Euler's formula for planar maps, the

complexity of the overlay M is proportional to the number of vertices of M. Each vertex of M is a vertex of $M_{\mathcal{F}}$, a vertex of $M_{\mathcal{G}}$, or an intersection point of an edge of $M_{\mathcal{F}}$ and an edge of $M_{\mathcal{G}}$. Since the total number of vertices in $M_{\mathcal{F}}$ and $M_{\mathcal{G}}$ is $O(n^{2+\varepsilon})$, as proved in [15, 19], it suffices to bound the number of intersection points between the edges of $M_{\mathcal{F}}$ and the edges of $M_{\mathcal{G}}$.

We call an intersection between an edge of $M_{\mathcal{F}}$ and an edge of $M_{\mathcal{G}}$ an *crossing* in M. For the purpose of analysis, we generalize the notion of a crossing, as follows. Let $\mathcal{A}(\mathcal{F})$ denote the arrangement of $\mathcal{F}$, namely the three-dimensional space decomposition induced by the graphs of the functions of $\mathcal{F}$ (see [10] for a more detailed definition). The *level* of a point w in $\mathcal{A}(\mathcal{F})$ is defined as the number of surfaces of $\mathcal{F}$ that lie vertically below w (note that 0-level points are precisely those that lie on or below the lower envelope $E_{\mathcal{F}}$). Let e be an edge of $\mathcal{A}(\mathcal{F})$. Clearly, the level of all points on e is the same, so we define the level of e to be the level of any point on e. We define the arrangement $\mathcal{A}(\mathcal{G})$, and the level of a point or of an edge in this arrangement, in an analogous manner for the collection $\mathcal{G}$. Let e be an edge of $\mathcal{A}(\mathcal{F})$, and let e' be an edge of $\mathcal{A}(\mathcal{G})$, such that the xy-projections of e and e' cross each other at a point σ. Let ξ, ξ' be the levels of the respective edges e, e'. Then we say that (e, e', σ) is an *edge-crossing* in $(\mathcal{A}(\mathcal{F}), \mathcal{A}(\mathcal{G}))$ at *level* (ξ, ξ'). If the point σ is not important, or is clear from the context, we just use (e, e') to denote the edge crossing (e, e', σ) (by our assumptions, for any pair of edges (e, e'), as above, there is only a constant number of points σ that appear in edge-crossings of the form (e, e', σ)). Note that the original crossings in M correspond to edge-crossings at level $(0, 0)$. (We slightly confuse the notation here, because crossings in M involve arcs in the xy-plane, whereas edge-crossings in $(\mathcal{A}(\mathcal{F}), \mathcal{A}(\mathcal{G}))$ involve arcs of these arrangements in 3-space; this abuse will also take place in what follows.) Let $C_{p,q}(\mathcal{F}, \mathcal{G})$ denote the number of edge-crossings in $(\mathcal{A}(\mathcal{F}), \mathcal{A}(\mathcal{G}))$ whose level is (p', q') for some $p' \leq p$, $q' \leq q$, and let

$$C_{p,q}(n) = \max C_{p,q}(\mathcal{F}, \mathcal{G}),$$

where the maximum is taken over all collections $\mathcal{F}$ and $\mathcal{G}$, as above, such that $|\mathcal{F}| + |\mathcal{G}| = n$. The goal is thus to obtain a sharp upper bound for $C_{0,0}(n)$.

Let k be a threshold parameter, whose value will be specified later on. Let e be an edge in the graph of $E_{\mathcal{F}}$, and let V_e be the vertical 2-manifold obtained as the union of all z-vertical lines passing through points of e. The intersection of the graph of each function $g \in \mathcal{G}$ with V_e is an algebraic arc of constant maximum degree, so each pair of these arcs intersect in at most some constant number, s, of points (where s depends only on the maximum degree of the functions of $\mathcal{F} \cup \mathcal{G}$ and of their graph boundaries, but not on e). Let $\mathcal{A}^{(e)}(\mathcal{G})$ denote the cross-section of $\mathcal{A}(\mathcal{G})$ with V_e, and let $C_{0,q}(e, \mathcal{G})$ denote the number of edge-crossings of the form (e, e') whose level is $(0, q')$, for any $q' \leq q$. See Figure 2 for an illustration. A simple but crucial observation is:

Lemma 2.2 *Let e' be an edge of $\mathcal{A}(\mathcal{G})$. Then (e, e', σ) is an edge-crossing in $(\mathcal{A}(\mathcal{F}), \mathcal{A}(\mathcal{G}))$ at level $(0, \xi)$, if and only if the point of $e' \cap V_e$ that lies on the z-vertical line through σ is a vertex at level ξ in $\mathcal{A}^{(e)}(\mathcal{G})$, and vice-versa.*

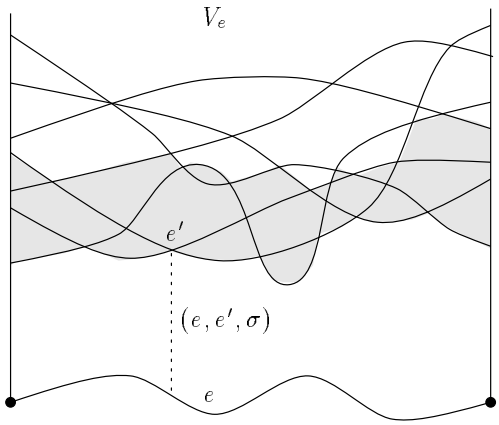

Figure 2: The arrangement $\mathcal{A}^{(e)}(\mathcal{G})$; the shaded region consist of points at level ≤ 3.

This lemma implies that each crossing of e in M corresponds to a vertex in the cross-section $E_{\mathcal{G}}^{(e)}$ of the lower envelope $E_{\mathcal{G}}$ within V_e. Let $\mathcal{G}^{(e)} \subseteq \mathcal{G}$ be the subset of functions of $\mathcal{G}$ that appear on $E_{\mathcal{G}}^{(e)}$, and let $t = |\mathcal{G}^{(e)}|$. By the standard Davenport-Schinzel theory [2, 16], $C_{0,0}(e, \mathcal{G}) \leq \lambda_s(t)$, where s is an appropriate constant (depending on the maximum degree of the surfaces in $\mathcal{F} \cup \mathcal{G}$), and where $\lambda_s(t)$ is the maximum length of a (t, s)-Davenport-Schinzel sequence.

If $t \leq k$ then, by Lemma 2.2, there are at most $\lambda_s(k)$ edge-crossings at level $(0, 0)$ involving e. Since the number of edges in $M_{\mathcal{F}}$ is $O(n^{2+\varepsilon})$, the overall number of crossings involving such edges is at most $O(\lambda_s(k) n^{2+\varepsilon})$.

Next, assume that $t > k$. Let g, g' be a pair of distinct functions in $\mathcal{G}^{(e)}$. By continuity, g and g' must intersect within V_e at least once. Thus each function $g \in \mathcal{G}^{(e)}$ must cross at least $t - 1$ other functions of $\mathcal{G}$ within V_e, that is, each function $g \in \mathcal{G}^{(e)}$ is incident to at least $t - 1$ vertices of $\mathcal{A}^{(e)}(\mathcal{G})$. Since the graph of g contains points at level 0 in this cross section, it follows that g is incident to at least k vertices of $\mathcal{A}^{(e)}(\mathcal{G})$ at level $\leq k$. The number of vertices of $\mathcal{A}^{(e)}(\mathcal{G})$ at level $\leq k$ is therefore $\Omega(tk)$, which, by Lemma 2.2, implies

$$C_{0,k}(e, \mathcal{G}) = \Omega(tk) = \Omega\left(k \frac{t}{\lambda_s(t)} \cdot C_{0,0}(e, \mathcal{G})\right) \geq \frac{k}{\beta(n)} \cdot C_{0,0}(e, \mathcal{G}), \tag{1}$$

where $\beta(n) = \Theta(\lambda_s(n)/n)$ is an extremely slowly growing function of n [2, 16].

Summing (1) over all edges of $M_{\mathcal{F}}$ that cross more than k edges of $M_{\mathcal{G}}$, adding the bound for the other edges of $M_{\mathcal{F}}$, and observing that each edge crossing in $C_{0,k}(\mathcal{F}, \mathcal{G})$ is counted in this manner exactly once, we obtain:

$$C_{0,0}(\mathcal{F}, \mathcal{G}) = \sum_{e \in M_{\mathcal{F}}} C_{0,0}(e, \mathcal{G}) \leq \frac{\beta(n)}{k} C_{0,k}(\mathcal{F}, \mathcal{G}) + O(kn^{2+\varepsilon}),$$

which implies:

$$C_{0,0}(n) \leq \frac{\beta(n)}{k} C_{0,k}(n) + O(kn^{2+\varepsilon}) . \tag{2}$$

We next bound $C_{0,k}(n)$ in terms of $C_{k,k}(n)$. Let e' be an edge of $\mathcal{A}(\mathcal{G})$ at some level $\xi' \leq k$, let $V_{e'}$ be the vertical 2-manifold erected from e', defined as above, and consider the cross section $\mathcal{A}^{(e')}(\mathcal{F})$ of $\mathcal{A}(\mathcal{F})$ within $V_{e'}$. Let t denote the number of functions of $\mathcal{F}$ that appear on the lower envelope of $\mathcal{A}_{\mathcal{F}}^{(e')}$. If $t \leq k$, then e' contributes at most $\lambda_s(k)$ edge-crossings to $C_{0,k}(\mathcal{F}, G)$. Since there are only $O(k^{1-\varepsilon}n^{2+\varepsilon})$ edges of $\mathcal{A}(\mathcal{G})$ at level at most k (see, e.g., [18, 20]), the number of crossings as above is $O(\lambda_s(k)k^{1-\varepsilon}n^{2+\varepsilon}) = O(k^2n^{2+\varepsilon})$, for any $\varepsilon > 0$.

We thus assume that $t > k$. We can now repeat, within $V_{e'}$, the preceding analysis, replacing $\mathcal{G}$ by $\mathcal{F}$, so as to conclude that the number of edge crossings of the form (e, e') at level (ξ, ξ'), for all $\xi \leq k$, is $\Omega(tk)$. Following the same arguments as above, and noting that each such crossing (e, e') is counted in this manner at most once, we easily obtain the recurrence:

$$C_{0,k}(n) \leq \frac{\beta(n)}{k} C_{k,k}(n) + O(k^2n^{2+\varepsilon}) . \tag{3}$$

Next, we estimate $C_{k,k}(n)$, by using the probabilistic technique of Clarkson and Shor [8] (see also [18]). Set $r = \lceil n/k \rceil$. Choose a random subset $R \subseteq \mathcal{F}$ of size r, where each subset of size r is chosen with equal probability. Similarly, choose a random subset $S \subseteq \mathcal{G}$ of size r. Let (e, e', σ) be an edge-crossing in $(\mathcal{A}(\mathcal{F}), \mathcal{A}(\mathcal{G}))$ at level (ξ, ξ'), and let ℓ_σ be the vertical line passing through σ. This edge-crossing appears in $(\mathcal{A}(R), \mathcal{A}(S))$ at level $(0, 0)$ if and only if the following conditions hold.

(i) The two functions whose intersection curve contains e are chosen in R,

(ii) the two functions whose intersection curve contains e' are chosen in S,

(iii) none of the ξ functions whose graphs intersect ℓ_σ below e is chosen in R, and

(iv) none of the ξ' functions whose graphs intersect ℓ_σ below e' is chosen in S.

The probability that (e, e', σ) is an edge-crossing in $(\mathcal{A}(\mathcal{F}), \mathcal{A}(\mathcal{G}))$ at level $(0, 0)$ is thus

$$\frac{\binom{n-\xi-2}{r-2}\binom{n-\xi'-2}{r-2}}{\binom{n}{r}\binom{n}{r}}.$$

Following the same analysis as in [8], we can show that, for $\xi, \xi' \leq k$ and for the specific choice of r, this probability is at least $1/(ck^4)$, for some absolute constant c. Summing this over all edge-crossings counted in $C_{k,k}(\mathcal{F}, \mathcal{G})$, we thus obtain, as in [8], that the expected number of edge-crossings in $(\mathcal{A}(R), \mathcal{A}(S))$ at level $(0, 0)$ is

$$\mathbf{E}[C_{0,0}(R, S)] \geq \frac{1}{ck^4} C_{k,k}(\mathcal{F}, \mathcal{G}) .$$

Hence, we obtain

$$C_{k,k}(n) = O(k^4) \cdot C_{0,0}\Big(\Big\lceil \frac{n}{k} \Big\rceil\Big) \ . \tag{4}$$

Combining (2), (3), and (4), we thus obtain

$$\begin{aligned} C_{0,0}(n) &= O((k + k\beta(n))n^{2+\varepsilon}) + \frac{\beta^2(n)}{k^2} \cdot O(k^4) \cdot C_{0,0}\Big(\Big\lceil \frac{n}{k} \Big\rceil\Big) \\ &= O(k\beta(n)n^{2+\varepsilon}) + O(k^2\beta^2(n)) \cdot C_{0,0}\Big(\Big\lceil \frac{n}{k} \Big\rceil\Big) \ . \end{aligned} \tag{5}$$

The solution of this recurrence is $O(n^{2+\delta})$, for any $\delta > \varepsilon$. This is shown by induction, choosing $k = \beta^{1+2/\delta}(n)$ and using the fact that $\beta(n)$ is an extremely slowly growing function of n. This concludes the proof of Theorem 2.1. □

Next, we extend the above proof to partially defined functions. We call an edge-crossing (e, e') a *boundary edge-crossing* if e or e' is contained in the boundary of some function graph in $\mathcal{F}$ or in $\mathcal{G}$, respectively.

Lemma 2.3 *The overall number of boundary edge-crossings (e, γ), where e is an edge of $E_{\mathcal{F}}$ and γ is a boundary edge in $\mathcal{A}(\mathcal{G})$, is $O(n\lambda_{s'}(n))$, where s' is an appropriate constant depending on the maximum degree of the functions in $\mathcal{F} \cup \mathcal{G}$.*

Proof: Let γ be a boundary arc of a function graph in $\mathcal{G}$ (there are a total of $O(n)$ such arcs), and let V_γ denote the vertical 2-manifold erected from γ, defined exactly as V_e in the proof of Theorem 2.1. Let $E_{\mathcal{F}}^{(\gamma)}$ denote the cross section $E_{\mathcal{F}} \cap V_\gamma$. As in Lemma 2.2, each boundary edge-crossing involving γ in the present lemma corresponds to a breakpoint of $E_{\mathcal{F}}^{(\gamma)}$, and, by the standard Davenport-Schinzel theory, the number of such breakpoints is $O(\lambda_{s'}(n))$, for an appropriate constant s'. This is easily seen to imply the lemma. □

We now establish the recurrence (2) for partially defined functions, in the same way as above, but with the following additional modifications. Again, we assume that the functions in $\mathcal{F} \cup \mathcal{G}$ are in general position. Let k be a threshold parameter. We split each edge e of $E_{\mathcal{F}}$ at a point σ if there exist an integer $l \leq k$ and a boundary edge $e' \in \mathcal{A}(\mathcal{G})$ such that (e, e', σ^*) is a boundary edge-crossing at level $(0, l)$, where σ^* is the xy-projection of σ. This step ensures that there is no boundary edge-crossing at level $(0, l)$ for any $l \leq k$. By Lemma 2.3, we introduce a total of at most $O(n\lambda_{s'}(n))$ new vertices, over all edges e of $E_{\mathcal{F}}$, so the number of edges in $M_{\mathcal{F}}$ is still $O(n^{2+\varepsilon})$.

Fix an edge e of the (refined) lower envelope $E_{\mathcal{F}}$. Define V_e, $\mathcal{A}^{(e)}(\mathcal{G})$, and $E_{\mathcal{G}}^{(e)}$ as above. By construction, the level of any point on $\partial g \cap V_e$, for any $g \in \mathcal{G}$, is greater than k. Let $\mathcal{G}^{(e)}$ be the set of connected components of $g \cap V_e$, for $g \in \mathcal{G}$, that appear on $E_{\mathcal{G}}^{(e)}$, and let $t = |\mathcal{G}^{(e)}|$. The case $t \leq k$ is handled exactly as in the preceding proof, so assume that $t > k$. We claim that any arc $\gamma \in \mathcal{G}^{(e)}$ is incident to at least k vertices of $\mathcal{A}^{(e)}(\mathcal{G})$ whose levels are $\leq k$. There are two cases to consider:

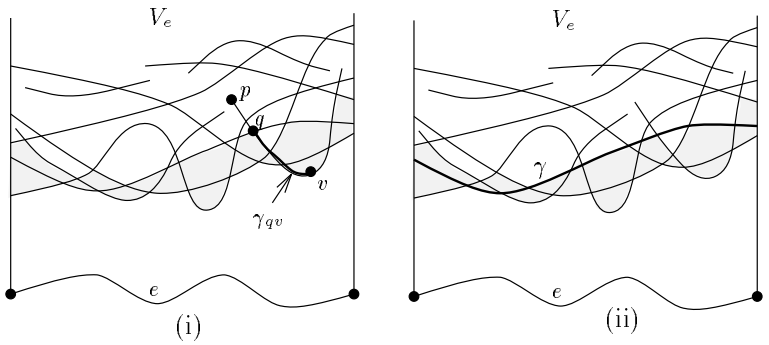

Figure 3: Illustration of the claim (for $k = 3$): (i) γ has a point p at level $\geq k$; (ii) the level of all points on γ is $< k$.

(i) γ has a point p whose level is $\geq k$. Let v be a point on $\gamma \cap E_{\mathcal{G}}^{(e)}$; without loss of generality, assume that v lies to the right of p. Let q be the rightmost point on γ to the left of v whose level is k, and let γ_{qv} denote the portion of γ between q and v. The level of all points on γ_{qv} is at most k. Let $g' \in \mathcal{G}$ be any of the k function graphs lying below q. Obviously, g' cannot lie below v. Moreover, by construction, no point of $\partial g'$ can lie below γ_{qv}, so g' has to intersect γ_{qv} (see Figure 3(i)). Since there are k function graphs of $\mathcal{G}$ lying below q, the arc γ_{qv} contains at least k vertices of $\mathcal{A}^{(e)}(\mathcal{G})$, and the level of each of them is at most k.

(ii) The level of all points on γ is $< k$. In this case, by construction, the endpoints of γ must lie on the vertical boundary edges of V_e. Let γ' be another arc in $\mathcal{G}^{(e)}$. If the endpoints of γ' also lie on the vertical boundary edges of V_e, then, as argued in the proof of Theorem 2.1, γ and γ' intersect within V_e. Otherwise, γ' has an endpoint p that lies inside V_e. By construction, the level of p is $> k$. Since the level of all points in γ is $\leq k$, the endpoint p lies above γ, implying again, that γ and γ' intersect (see Figure 3 (ii)). Since $t > k$, we obtain at least k such intersections with γ, and the level of each of these intersection points is $\leq k$.

This completes the proof of the claim. Hence $\mathcal{A}^{(e)}(\mathcal{G})$ has $\Omega(tk)$ vertices at level $\leq k$. Following the same argument as in the proof of Theorem 2.1, we can obtain the recurrence (2) in this case too.

A similar argument can be applied to obtain the recurrence (3) for the case of partial functions. The only difference is that we now split an edge $e' \in \mathcal{A}(\mathcal{G})$ at a point p if there are integers $j, j' \leq k$ and a boundary edge e of some function in $\mathcal{F}$ such that (e, e', p^*) is an edge-crossing at level (j, j'), where p^* is the xy-projection of p. Using the proof of Lemma 2.3, in conjunction with the Clarkson-Shor technique, we can show that the number of newly added vertices in $\mathcal{A}(\mathcal{G})$ is $O(k^2 n \lambda_{s'}(\lceil n/k \rceil))$, for an appropriate constant s'. Hence the number of edges of $\mathcal{A}(\mathcal{G})$ at level $\leq k$ remains $O(k^{1-\varepsilon} n^{2+\varepsilon})$, as in the preceding analysis.

The recurrence (3) now follows by the same analysis as above. This is turn yields the same final recurrence (5) for $C_{0,0}(n)$, whose solution, as above, is $O(n^{2+\varepsilon})$. We thus obtain the main result of the paper:

Theorem 2.4 *Let $\mathcal{F}$ and $\mathcal{G}$ be two collections of n, possibly partially defined, bivariate functions, satisfying condition $(\star)$. Then the combinatorial complexity of the overlay of the minimization diagrams of $\mathcal{F}$ and $\mathcal{G}$, as defined above, is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$ (where the constant of proportionality depends on ε and on the maximum degree b).*

3 Applications

3.1 Computing lower envelopes in 3-space

Let $\mathcal{F}$ be a collection of n bivariate functions satisfying condition $(\star)$. Our goal is to construct the lower envelope $E_{\mathcal{F}}$ of $\mathcal{F}$. This is equivalent to constructing the minimization diagram $M_{\mathcal{F}}$, as defined above, so that each face ϕ of $M_{\mathcal{F}}$ is labeled with the unique function of $\mathcal{F}$ (if exists) attaining $E_{\mathcal{F}}$ over ϕ. Several algorithms for this construction have recently been designed (see [4, 9, 19]), but they are either rather complicated or require the use of randomization. Here we present a simple deterministic algorithm based on the divide-and-conquer approach, which is similar to Atallah's algorithm [3] for computing the minimization diagram of univariate functions.

The algorithm partitions $\mathcal{F}$ into two subcollections, $\mathcal{F}_1$, $\mathcal{F}_2$, of roughly $n/2$ functions each, constructs recursively the minimization diagrams $M_{\mathcal{F}_1}$, $M_{\mathcal{F}_2}$, and then merges these diagrams to obtain the final minimization diagram $M_{\mathcal{F}}$.

The merge step is done as follows. We first compute the superposition M of $M_{\mathcal{F}_1}$ and $M_{\mathcal{F}_2}$. This can be done, e.g., by applying a standard sweep-line procedure, whose running time is $O((|M_{\mathcal{F}_1}| + |M_{\mathcal{F}_2}| + |M|)\log n)$; by Theorem 2.4, this is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. We can implement the sweep so that it also constructs the *vertical decomposition* of M. The *vertical decomposition* of M is a refinement of M, obtained by drawing a vertical segment upwards and downwards (in the y-direction) from each vertex of M and from each point on any edge of M that has y-vertical tangency, and by extending each segment until it hits another edge of M or to infinity if no such edge exists. The number of resulting cells, usually referred to as 'pseudo-trapezoids', is proportional to the complexity of M, namely it is also $O(n^{2+\varepsilon})$.

Let c be a pseudo-trapezoid in this vertical decomposition. Note that, over c, the envelope $E_{\mathcal{F}_1}$ is attained by a single function $f_1 \in \mathcal{F}_1$ (or by no function at all), and $E_{\mathcal{F}_2}$ is attained by a single function $f_2 \in \mathcal{F}_2$ (or by no function at all). Hence, the envelope $E_{\mathcal{F}}$ is equal over c to $\min\{f_1, f_2\}$ if both functions exists, or to one of these functions if the other does not exist, or, if both functions do not exist, $E_{\mathcal{F}}$ is undefined over c. In

any case, we can compute $E_{\mathcal{F}}$ over c in constant time.[1] We repeat this computation over all pseudo-trapezoids of M, in overall $O(n^{2+\varepsilon})$ time, and thus obtain the entire envelope $E_{\mathcal{F}}$. We still need to apply a final clean-up stage, in which the computed portions of $E_{\mathcal{F}}$ are properly glued together, removing, as appropriate, any redundant data concerning the behavior of $E_{\mathcal{F}}$ over edges of the pseudo-trapezoids of M. This stage also produces the final minimization diagram $M_{\mathcal{F}}$, with its faces labeled in the required manner. We omit the routine details of this step, and note that it also takes only $O(n^{2+\varepsilon})$ time. It follows that the cost of the entire divide-and-conquer process is also $O(n^{2+\varepsilon})$. In conclusion, we thus have:

Theorem 3.1 *The lower envelope of a collection of n bivariate functions satisfying condition $(\star)$, can be computed, in an appropriate model of computation, by a deterministic divide-and-conquer algorithm, in time $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε and on the maximum algebraic degree of the given functions (and of their domain boundaries).*

3.2 Complexity of the region enclosed between two envelopes in 3-space

Let $\mathcal{T}$ and $\mathcal{B}$ be two given families of a total of n, possibly partially-defined, bivariate functions satisfying condition $(\star)$. We denote by $\mathcal{L}_{\mathcal{T}}$ the lower envelope of the 'top' family $\mathcal{T}$, and by $\mathcal{U}_{\mathcal{B}}$ the upper envelope of the 'bottom' family $\mathcal{B}$. We consider the region $K = \{(x, y, z) \mid \mathcal{U}_{\mathcal{B}}(x, y) \leq z \leq \mathcal{L}_{\mathcal{T}}(x, y)\}$ of points lying between the two envelopes, and our goal is to derive an $O(n^{2+\varepsilon})$ bound on the combinatorial complexity of K.

We establish this bound as follows. Let $M_{\mathcal{T}}$, $M_{\mathcal{B}}$ denote the minimization and maximization diagrams of the envelopes $\mathcal{L}_{\mathcal{T}}$, $\mathcal{U}_{\mathcal{B}}$, respectively. By Theorem 2.4, the combinatorial complexity of the overlay M of these two planar maps is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. Construct the vertical decomposition of M, as defined above. As noted, the number of pseudo-trapezoids of this decomposition is proportional to the complexity of M, i.e., it is $O(n^{2+\varepsilon})$. Observe that, for each resulting pseudo-trapezoid τ, there is a single function $f \in \mathcal{T}$ and a single function $g \in \mathcal{B}$ such that $\mathcal{L}_{\mathcal{T}} \equiv f$ and $\mathcal{U}_{\mathcal{B}} \equiv g$ over τ (if the given functions are only partially defined then either f or g or both may not exist at all, in which case the corresponding envelope(s) are undefined over τ). This implies that the portion of K that projects into τ has constant complexity—it is defined by the interaction between f, g, and the functions defining the (at most 4) edges of τ. Since the number of pseudo-trapezoids is $O(n^{2+\varepsilon})$, we immediately obtain:

[1] We are implicitly assuming an appropriate model of computation, in which computing the pointwise minimum of two given functions, as well as various primitive operations involving edges of the minimization diagrams, can be performed in constant time. For example, we can use precise rational arithmetic to perform each of these operations in constant time, using standard techniques from computational real algebraic geometry; see, e.g., [17].

Theorem 3.2 *The combinatorial complexity of the region enclosed between a lower envelope and an upper envelope of two respective collections of n bivariate functions satisfying condition $(\star)$, is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε and on the maximum algebraic degree of the given functions (and of their domain boundaries).*

It is also easy to construct the desired region K, in a manner that resembles the divide-and-conquer algorithm presented above. That is, we compute $\mathcal{L}_{\mathcal{T}}$ and $\mathcal{U}_{\mathcal{B}}$ separately, in time $O(n^{2+\varepsilon})$, using the algorithm of the preceding subsection. Next we compute the overlay of the minimization diagram $M_{\mathcal{T}}$ and of the maximization diagram $M_{\mathcal{B}}$, using the same sweep technique described above, and decompose the resulting map into pseudo-trapezoids. Finally we compute the portions of K over each pseudo-trapezoid separately, and 'glue' together the resulting pieces to obtain the whole K. It is easily verified that the overall complexity of the algorithm is $O(n^{2+\varepsilon})$, so we have:

Theorem 3.3 *The region enclosed between two envelopes in 3-space, as above, can be computed in (deterministic) time $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$.*

3.3 Complexity of the space of plane transversals

In this subsection we obtain new bounds on the combinatorial complexity of the space of plane transversals of a collection of simply-shaped convex sets in 3-space. Let $\mathcal{C} = \{C_1, \ldots, C_n\}$ be a collection of n compact convex sets in 3-space. A plane π is a *transversal* of $\mathcal{C}$ if it intersects every set in $\mathcal{C}$. The space of all plane transversals of $\mathcal{C}$ is denoted by $T(\mathcal{C})$.

It is more convenient to represent $T(\mathcal{C})$ in the dual space, where each nonvertical plane $z = \xi x + \eta y + \zeta$ is mapped to a point (ξ, η, ζ), and each point (u, v, w) is mapped to a plane $z = -ux - vy + w$. Note that a plane $z = \xi x + \eta y + \zeta$ intersects a compact convex set C if and only if $\phi_C(\xi, \eta) \leq \zeta \leq \psi_C(\xi, \eta)$, where $\phi_C(\xi, \eta)$, $\psi_C(\xi, \eta)$ are defined so that the plane $z = \xi x + \eta y + \phi_C(\xi, \eta)$ (resp. $z = \xi x + \eta y + \psi_C(\xi, \eta)$) is tangent to C from below (resp. from above). Thus, in the dual space, the set of all plane transversals of $\mathcal{C}$ is the set

$$\{(\xi, \eta, \zeta) \mid \max_{C \in \mathcal{C}} \phi_C(\xi, \eta) \leq \zeta \leq \min_{C \in \mathcal{C}} \psi_C(\xi, \eta)\} .$$

That is, $T(\mathcal{C})$ is, in the dual space, the region enclosed between a lower envelope and an upper envelope of two respective collections of functions.

We can therefore apply Theorem 3.2 to this case, but we first have to ensure that the functions ϕ_C and ψ_C satisfy the assumptions of that theorem. This will be the case if we assume that each $C \in \mathcal{C}$ has *constant description complexity*, that is, it is defined by a constant number of algebraic equalities and inequalities of constant maximum degree. In

this case one can easily show that the functions ϕ_C and ψ_C do indeed satisfy condition $(\star)$.[2] We thus have:

Theorem 3.4 *The complexity of the space of plane transversals of a collection of n compact convex sets in 3-space, each of constant description complexity, is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$.*

Remarks: (1) Convexity is not essential here, since we can replace each set in $\mathcal{C}$ by its convex hull, without affecting the transversality of any plane.

(2) If the sets in $\mathcal{C}$ do not have constant description complexity, the complexity of $T(\mathcal{C})$ can be arbitrarily large. However, if one assumes, in addition, that the sets are *separated*, in the sense that no three of these sets is crossed by a common line, then it is shown by Cappell et al. [5] that, for such a collection $\mathcal{C}$, the complexity of $T(\mathcal{C})$ is $O(n^2)$. This bound, in this restricted case, is slightly better than the bound derived above. The result of [5] applies in higher dimensions too: Under an appropriate assumption of separation of the sets in $\mathcal{C}$, the complexity of $T(\mathcal{C})$ is $O(n^{d-1})$. Other related results on transversals can be found in a recent paper by Goodman et al. [13], and in a survey paper [14] by the same authors.

Acknowledgments

We wish to thank Boris Aronov, Leo Guibas, and Olivier Devillers for useful discussions concerning the problems studied in this paper. Part of the work on the paper has been carried out in the Mathematisches Forschungsinstitut in Oberwolfach, and we would like to thank the institute for its hospitality.

References

[1] P. Agarwal and M. Sharir, On the number of views of polyhedral terrains, *Proc. 5th Canadian Conf. Computational Geometry*, 1993, 55–61.

[2] P. Agarwal, M. Sharir and P. Shor, Sharp upper and lower bounds for the length of general Davenport Schinzel sequences, *J. Combin. Theory, Ser. A.* 52 (1989), 228–274.

[3] M. Atallah, Some dynamic computational geometry problems, *Computers and Mathematics with Applications* 11 (1985), 1171–1181.

[4] J.D. Boissonnat and K. Dobrindt, On-line randomized construction of the upper envelope of triangles and surface patches in $\mathbb{R}^3$, Tech. Rept. 1878, INRIA, Sophia-Antipolis, France, 1993.

[2] Actually, ϕ_C and ψ_C may only be piecewise-algebraic, so we may have to replace them by a constant number of appropriate partially defined algebraic functions, and apply Theorem 3.2 to the resulting new collections of functions.

[5] S. Cappell, J.E. Goodman, J. Pach, R. Pollack, M. Sharir and R. Wenger, Common tangents and common transversals, *Advances in Mathematics* 106 (1994), 198–215.

[6] 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.

[7] 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.

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

[9] M. de Berg, K. Dobrindt and O. Schwarzkopf, On lazy randomized incremental construction, to appear in *26th Annu. ACM Symp. Theory of Computing,* 1994, pp. 105–114.

[10] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer–Verlag, Berlin, 1987.

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

[12] H. Edelsbrunner and M. Sharir, The maximum number of ways to stab n convex nonintersecting objects in the plane is $2n-2$, *Discrete Comput. Geom.* 5 (1990), 35–42.

[13] J. Goodman, R. Pollack, and R. Wenger, Geometric transversal theory, in: *New Trends in Discrete and Computational Geometry* (J. Pach, ed.), Springer-Verlag, New York–Berlin–Heidelberg, 1993, pp. 163–198.

[14] J. Goodman, R. Pollack, and R. Wenger, Bounding the number of geometric permutations induced by k-transversals, *Proc. 10th Annu. Symp. Computational Geometry*, 1994, pp. 192–197.

[15] 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.

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

[17] J. Heintz, T. Recio and M.-F. Roy, Algorithms in real algebraic geometry and applications to computational geometry, in *Discrete and Computational Geometry: Papers from the DIMACS Special Year* (J.E. Goodman, R. Pollack and W. Steiger, Eds.), AMS Press, Providence, RI 1991, pp. 137–163.

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

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

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