The Partition Technique for Overlays of Envelopes*

Vladlen Koltun† Micha Sharir‡

Abstract

We obtain a near-tight bound of $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$, on the complexity of the overlay of the minimization diagrams of two collections of surfaces in four dimensions. This settles a long-standing problem in the theory of arrangements, most recently cited by Agarwal and Sharir [3, Open Problem 2], and substantially improves and simplifies a result previously published by the authors [17].

Our bound is obtained by introducing a new approach to the analysis of combinatorial structures arising in geometric arrangements of surfaces. This approach, which we call the 'partition technique', is based on k-fold divide and conquer, in which a given collection $\mathcal{F}$ of n surfaces is partitioned into k subcollections $\mathcal{F}_i$ of n/k surfaces each, and the complexity of the relevant combinatorial structure in $\mathcal{F}$ is recursively related to the complexities of the corresponding structures in each of the $\mathcal{F}_i$'s. We introduce this approach by applying it first to obtain a new *simple* proof for the known near-quadratic bound on the complexity of an overlay of two minimization diagrams of collections of surfaces in $\mathbb{R}^3$, thereby simplifying the previously available proof [2].

The main new bound on overlays has numerous algorithmic and combinatorial applications, some of which are presented in this paper.

1 Introduction

In this paper we obtain combinatorial bounds on overlays of minimization diagrams, by introducing a new approach to the analysis of combinatorial structures in arrangements of surfaces. Let us start with the basic definitions. (For a thorough treatment of the topic we are about to briefly introduce, the reader is referred to [20, Chapter 7].)

Let $\mathcal{F}$ be a family of n d-variate (not necessarily continuous or totally defined) functions of *constant description complexity*, that is, the graph of each function is a semi-algebraic set in $\mathbb{R}^{d+1}$ defined by a constant number of polynomial equalities and inequalities of constant maximum degree. The lower envelope $E_{\mathcal{F}}$ of $\mathcal{F}$ is the pointwise minimum of the functions

*This work was supported by a grant from the Israel Science Fund (for a Center of Excellence in Geometric Computing), and is part of the first author's Ph.D. dissertation, prepared under the supervision of the second author in Tel Aviv University. Work by Micha Sharir was also supported by NSF Grants CCR-97-32101 and CCR-00-98246, by a grant from the U.S.-Israel Binational Science Foundation, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University. A preliminary version of this paper appeared in the Proceedings of the 43nd IEEE Symposium on Foundations of Computer Science, 2002.

†Computer Science Division, University of California, Berkeley, CA 94720-1776, USA. *vladlen@cs.berkeley.edu*

‡School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA. *sharir@cs.tau.ac.il*

of $\mathcal{F}$:

$$E_{\mathcal{F}}(\mathbf{x}) = \min_{f \in \mathcal{F}} f(\mathbf{x}), \qquad \text{for } \mathbf{x} \in \mathbb{R}^d.$$

$E_{\mathcal{F}}$ is itself a d-variate function whose graph is a semi-algebraic set in $\mathbb{R}^{d+1}$. The projection onto d-space of this graph is called the *minimization diagram* of $\mathcal{F}$, and is denoted by $M_{\mathcal{F}}$. This is a subdivision of $\mathbb{R}^d$ into maximal connected relatively open cells of dimensions that range between 0 and d, so that, for each cell τ, a fixed subset of $\mathcal{F}$ attains $E_{\mathcal{F}}$ over all points $\mathbf{x} \in \tau$ (and no other function attains the envelope over any point in τ). The complexity of $M_{\mathcal{F}}$ (and of $E_{\mathcal{F}}$) is the number of cells (of all dimensions) of $M_{\mathcal{F}}$.

It has been shown by Halperin and Sharir [12, 19] that the complexity of $M_{\mathcal{F}}$ (and of $E_{\mathcal{F}}$) is $O(n^{d+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε, d, and the maximum degree of the polynomials defining the functions of $\mathcal{F}$. A (slightly) super-$\Omega(n^d)$ lower bound is known, so the bound of [12, 19] is almost tight in the worst case. Its proof is fairly involved, and is based on a counting (or charging) scheme, where vertices of the envelope are charged to sets of 'nearby' vertices of the arrangement $\mathcal{A}(\mathcal{F})$ of the function graphs.

This technique, and various refinements and extensions thereof, have been successful in bounding the complexity of lower envelopes and of several related structures, such as a single cell in a d-dimensional arrangement [6, 13]. However, the technique has had only partial success in analyzing the *overlay* of minimization diagrams. The overlay of two (or several) minimization diagrams of d-variate functions, as above, is the superposition of these diagrams in $\mathbb{R}^d$; specifically, it is the arrangement in d-space of the union of the curves or surfaces (of various dimensions) that constitute the individual diagrams. The *complexity* of the overlay is the complexity of this arrangement, namely, the number of its cells of all possible dimensions. The overlay of minimization diagrams became an important concept in the theory of arrangements after it was demonstrated that a successful analysis of the complexity of the overlay can lead to simple divide and conquer algorithms for the computation of lower envelopes and related structures [11]. Moreover, overlays arise naturally in many applications, as will be described below in more detail.

For $d = 1$ (the case of univariate functions), each minimization diagram is simply a partition of the x-axis into a finite number of intervals; if each pair of functions intersect in at most s points, then the size of the minimization diagram of n functions is known to be at most $\lambda_s(n)$ for totally-defined continuous functions, or $\lambda_{s+2}(n)$ for partially-defined continuous functions, where $\lambda_s(n)$ is the maximum length of (n, s)-Davenport-Schinzel sequences [20], which is near-linear in n for any fixed s. The overlay of two (or more) minimization diagrams is simply the partition obtained by merging the breakpoints of the diagrams into a single sequence. The complexity of the overlay is thus proportional to the sum of the complexities of the individual diagrams; in particular, it is near-linear in the number of functions.

For $d = 2$, it was shown by Agarwal et al. [2] that the overlay of two minimization diagrams, each defined for some set of n bivariate functions of constant description complexity, is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. That is, the asymptotic bounds on the complexity of the overlay and on the complexity of a single diagram are the same, which is somewhat counter-intuitive. The proof uses a refined and more involved variant of the counting scheme mentioned above.

The prevailing conjecture is that the complexity of the overlay of two minimization diagrams of d-variate functions of constant description complexity is $O(n^{d+\varepsilon})$, for any $\varepsilon > 0$ (the same asymptotic bound as that for the complexity of a single envelope). This has been an open problem for all $d \geq 3$. Recently [17], the authors have made a small step towards establishing the conjecture for $d = 3$, obtaining bounds of the form $O(n^{4-\frac{1}{s}+\varepsilon})$, where s is a constant integer parameter that depends on the shape and the (constant) maximum degree of the given functions. The proof in [17] is based on the counting scheme, and is highly complicated.

In this paper, we settle the conjecture affirmatively for $d = 3$, and prove the following theorem:

Theorem 1.1. *(a) The complexity of the overlay of two minimization diagrams of a total of n trivariate functions of constant description complexity is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$.*
(b) The complexity of the overlay of $k \geq 3$ minimization diagrams, each of n/k trivariate functions of constant description complexity, is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$.

This is achieved by introducing the *partition technique*, a new approach to this problem, which we hope will also prove useful in the analysis of the complexity of other substructures in arrangements. The technique is based on *k-fold divide and conquer*, in which each of the given collections is partitioned into k subcollections of n/k functions each, where k is some parameter, and the complexity of the entire overlay is expressed in terms of the complexities of various 'sub-overlays' and related substructures. The analysis exploits and extends ideas used by Har-Peled [14] for the analysis of the complexity of substructures in the overlay of planar arrangements.

We introduce our approach by presenting a new *simple* proof for the complexity of the overlay of the minimization diagrams of two collections of bivariate functions. This compares favorably with the previously available proof [2].

So far, the partition technique faces technical difficulties when applied to the analysis of other structures (like vertical decompositions), or to overlays in higher dimensions. Nevertheless, we feel hopeful that it will develop further, so as to be able to tackle these problems, and to find many additional applications.

Our result has several applications, which we enumerate in Section 5 (some of which were already noted in our previous work [17]). It can be used to obtain an improved near-cubic bound, which is nearly tight in the worst case, on the complexity of the region enclosed between two envelopes in four dimensions. Another application is an improved near-cubic bound on the complexity of the space of all hyperplane transversals of a collection of simply-shaped convex sets in 4-space, and on the complexity of the space of all line transversals of a similar collection of convex sets in 3-space. Using these bounds, one can adapt randomized incremental techniques, proposed in [1, 5], to construct the boundary of these transversal spaces in expected near-cubic time. We also obtain an improved near-cubic bound on the number of geometric permutations in a collection of disjoint convex bodies in $\mathbb{R}^3$. Our new bound can also be used to obtain a near-cubic bound on the complexity of the union of certain families of fat convex objects of nearly equal size in 4-space.

Parts of our analysis are of independent interest, and we adapt one to show that the complexity of the lower envelope of an arrangement of n totally defined semi-algebraic surfaces of constant description complexity in $\mathbb{R}^3$, that does not contain any vertices, is $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$.

The rest of the paper is organized as follows. We begin by introducing the partition technique. It is introduced by example, in Section 2, where it is used to re-derive the known near-quadratic bound for the complexity of the overlay of the minimization diagrams of two collections of bivariate functions. In Section 3 we provide a useful technical tool needed for applying the partition technique in three dimensions. Finally, we use the partition technique to prove a near-cubic bound on the complexity of the overlay of minimization diagrams of trivariate functions in Section 4. Several applications of this new bound are described in Section 5.

2 The Overlay of Bivariate Minimization Diagrams

We start by introducing the partition technique. This introduction is done by example, on the analysis of overlays of minimization diagrams of two collections of bivariate functions.

Let $\mathcal{F}$ and $\mathcal{G}$ be two collections, each consisting of n bivariate functions of constant description complexity. We prove that the complexity of the overlay of the minimization diagrams $M_\mathcal{F}$ of $\mathcal{F}$ and $M_\mathcal{G}$ of $\mathcal{G}$ is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. Denote the overlay by $Q(\mathcal{F}, \mathcal{G})$, and define the *bichromatic complexity* $C(\mathcal{F}, \mathcal{G})$ of the overlay to be the number of intersections between edges of $M_\mathcal{F}$ and edges of $M_\mathcal{G}$. Clearly, since each overlay is a planar map, the actual complexity of $Q(\mathcal{F}, \mathcal{G})$ is proportional to $C(\mathcal{F}, \mathcal{G}) + |M_\mathcal{F}| + |M_\mathcal{G}|$, where $|M_\mathcal{F}|$ (resp., $|M_\mathcal{G}|$) is the complexity of the minimization diagram $M_\mathcal{F}$ (resp., of $M_\mathcal{G}$).

Throughout the sequel, we will assume that all the functions in $\mathcal{F}$ and $\mathcal{G}$ are continuous and totally defined. This involves no real loss of generality, because one can always partition each function graph in $\mathcal{F}$ and $\mathcal{G}$ into a constant number of continuous patches, which can then be extended to be totally defined ones without decreasing the complexity of the overlay. Indeed, a continuous partially defined function graph can be extended to be totally defined, by means of near-vertical semi-infinite walls attached to its boundary. Formally, we replace each graph by the boundary of its Minkowski sum with a steeply-sloped vertical cone; it is easy to see that, if the slope of the cone is large enough, this can only increase the complexity of the overlay. This extension can be analogously performed for collections $\mathcal{F}$ and $\mathcal{G}$ of functions in any dimension.

Partition $\mathcal{G}$ into k groups, $\mathcal{G}_1, \ldots, \mathcal{G}_k$, each of n/k functions, for some threshold parameter k that we will determine later. Fix an edge e of $M_\mathcal{F}$, and consider the vertical 2-dimensional *wall* $V^{(e)}$ erected over e; this is the union of all z-parallel lines that pass through points of e. Restrict the functions of $\mathcal{G}$ over e, to obtain a collection $\mathcal{G}^{(e)}$ of univariate functions of constant description complexity. It is partitioned in an obvious way into k subcollections $\mathcal{G}_1^{(e)}, \ldots, \mathcal{G}_k^{(e)}$.

We shall now see that the complexity of the lower envelope of $\mathcal{G}^{(e)}$ is roughly proportional to the sum of the complexities of the lower envelopes of all the subsets $\mathcal{G}_i^{(e)}$, disregarding a small additive term and a near-constant multiplicative factor.

Let s denote the (constant) maximum number of intersections of the xy-projections of an intersection curve of two function graphs in $\mathcal{F}$, and of an intersection curve of two function graphs in $\mathcal{G}$. In particular, the number of intersections between any pair of (graphs of) functions in $\mathcal{G}^{(e)}$ is at most s.

Note that the lower envelope $E_{\mathcal{G}^{(e)}}$ of $\mathcal{G}^{(e)}$ is the lower envelope of the lower envelopes

$E_{\mathcal{G}_i^{(e)}}$ of the subcollections $\mathcal{G}_i^{(e)}$, for $i = 1, \ldots, k$. Define the complexity $|E^{(e)}|$ of a univariate envelope $E^{(e)}$ over a connected arc e to be the number of vertices (breakpoints) of $E^{(e)}$ over points in the relative interior of e. Using an easy modification of an observation due to Har-Peled [14], the complexity of $E_{\mathcal{G}^{(e)}}$ is

$$|E_{\mathcal{G}^{(e)}}| = O\left(\frac{\lambda_s(k)}{k}\sum_{i=1}^{k}\left(1 + |E_{\mathcal{G}_i^{(e)}}|\right)\right). \tag{1}$$

Indeed, Har-Peled's proof merges the breakpoints of the $E_{\mathcal{G}_i^{(e)}}$'s into a single sequence, and partitions this sequence into blocks of k breakpoints each. Within each block, at most $2k$ functions can appear on the merged envelope, and they contribute at most $O(\lambda_s(k))$ breakpoints to $E_{\mathcal{G}^{(e)}}$. If an envelope $E_{\mathcal{G}_i^{(e)}}$ does not have any breakpoint over (the relative interior of) e, we still need to consider the single function that it contributes to the overall envelope. The term 1, added to each term $|E_{\mathcal{G}_i^{(e)}}|$ in the above bound, takes care of these extreme cases.

Putting $\beta_s(k) = \lambda_s(k)/k$, and summing these bounds over all edges e of $M_{\mathcal{F}}$, we obtain that the bichromatic complexity $C(\mathcal{F}, \mathcal{G})$ of the overlay of $M_{\mathcal{F}}$ and $M_{\mathcal{G}}$ satisfies

$$C(\mathcal{F}, \mathcal{G}) = O\left(|M_{\mathcal{F}}|\lambda_s(k) + \beta_s(k)\sum_{i=1}^{k} C(\mathcal{F}, \mathcal{G}_i)\right). \tag{2}$$

Informally, we have just shown that, ignoring an additive term dominated by $|M_{\mathcal{F}}|$ and a negligible multiplicative factor of $\beta_s(k)$, the complexity of the overlay of $M_{\mathcal{F}}$ with $M_{\mathcal{G}}$ is roughly proportional to the sum of the complexities of the overlays of $M_{\mathcal{F}}$ with all the minimization diagrams $M_{\mathcal{G}_i}$. This is the essence of the partition technique for overlays, and it allows us to easily finish off the analysis as follows.

We reverse the roles of $\mathcal{F}$ and $\mathcal{G}$, as follows. Fix a subset $\mathcal{G}_i$, and consider an edge e' of $M_{\mathcal{G}_i}$. Partition $\mathcal{F}$ into k groups, $\mathcal{F}_1, \ldots, \mathcal{F}_k$, each of n/k functions, and consider the vertical 2-dimensional wall $V^{(e')}$ erected over e'. Restrict the functions of $\mathcal{F}$ over e', to obtain a collection $\mathcal{F}^{(e')}$ of univariate functions of constant description complexity, which is partitioned in an obvious way into $\mathcal{F}_1^{(e')}, \ldots, \mathcal{F}_k^{(e')}$.

As above, the lower envelope of the individual lower envelopes $E_{\mathcal{F}_j^{(e')}}$ over e' is the lower envelope $E_{\mathcal{F}^{(e')}}$. Using once again the technique of Har-Peled [14], the complexity of $E_{\mathcal{F}^{(e')}}$ is

$$|E_{\mathcal{F}^{(e')}}| = O\left(\beta_s(k)\sum_{j=1}^{k}\left(1 + |E_{\mathcal{F}_j^{(e')}}|\right)\right).$$

Summing these bounds over all edges e' of $M_{\mathcal{G}_i}$, we obtain that the bichromatic complexity $C(\mathcal{F}, \mathcal{G}_i)$ of the overlay of $M_{\mathcal{F}}$ and $M_{\mathcal{G}_i}$ satisfies

$$C(\mathcal{F}, \mathcal{G}_i) = O\left(|M_{\mathcal{G}_i}|\lambda_s(k) + \beta_s(k)\sum_{j=1}^{k} C(\mathcal{F}_j, \mathcal{G}_i)\right).$$

This is essentially the same equation as (2), with $\mathcal{G}_i$ in the role of $\mathcal{F}$ and $\mathcal{F}$ in the role of $\mathcal{G}$, and it was obtained using an identical mechanism. If we now simply substitute this equation into (2), we obtain

$$C(\mathcal{F},\mathcal{G}) = O\left(|M_{\mathcal{F}}|\lambda_s(k) + \beta_s(k)\sum_{i=1}^{k}\left(|M_{\mathcal{G}_i}|\lambda_s(k) + \beta_s(k)\sum_{j=1}^{k} C\left(\mathcal{F}_j,\mathcal{G}_i\right)\right)\right). \qquad (3)$$

Let $C(n)$ denote the maximum complexity of the overlay of the minimization diagrams of two collections of n bivariate functions, each of the same constant description complexity,[1] and recall that the complexity of a lower envelope of n bivariate functions of constant description complexity is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$ [12, 19]. This simplifies (3) into the recurrence

$$C(n) = O\left(n^{2+\varepsilon}\lambda_s(k) + k^2\beta_s^2(k)C\left(\frac{n}{k}\right)\right),$$

the solution of which is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$ (see [13, 19, 20] for demonstrations of solutions of similar recurrence relations, which are obtained by choosing a suitable parameter k as a function of ε). We have thus shown the following theorem.

Theorem 2.1. *The complexity of the overlay of the minimization diagrams of two collections of n bivariate functions of constant description complexity is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε.*

3 The Partition Technique in Three Dimensions

In this section we take a crucial preparatory step towards a bound on the complexity of overlays of collections of trivariate functions. The partition technique, as exposed in the previous section, calls for relating the complexity of the minimization diagram of a collection $\mathcal{F} = \bigcup_{i=1}^{k}\mathcal{F}_i$ of functions to the sum of the complexities of the minimization diagrams of all the sub-collections $\mathcal{F}_i$. This is precisely what was established in equation (1), in the case of univariate functions, where it was essentially shown that

$$|M_{\mathcal{F}}| = O\left(\beta_s(k)\sum_{i=1}^{k}\left(1 + |M_{\mathcal{F}_i}|\right)\right),$$

when $\mathcal{F}$ is a collection of univariate functions.

Utilizing the partition technique in three dimensions requires a parallel relation to be established for the case of bivariate functions. Such a relation, which we believe to be of independent interest, is proved below. Although it is sufficient for our purposes, it is not as intuitive as its counterpart in the univariate case.

Theorem 3.1. *Let $\mathcal{F}_1,\ldots,\mathcal{F}_k$ be k sets of bivariate functions of constant description complexity, and put $\mathcal{F} = \bigcup_{i=1}^{k}\mathcal{F}_i$. Then*

$$|M_{\mathcal{F}}| = O\left(k^{2+\varepsilon} + k^{1+\varepsilon}\sum_{i=1}^{k}|M_{\mathcal{F}_i}| + k^{\varepsilon}\sum_{i=1}^{k}\sum_{j=i+1}^{k} C(\mathcal{F}_i,\mathcal{F}_j)\right),$$

for any $\varepsilon > 0$, where $C(\mathcal{F}_i,\mathcal{F}_j)$ is, as above, the bichromatic complexity of the overlay $Q(\mathcal{F}_i,\mathcal{F}_j)$.

[1] Having a fixed constant description complexity means that a function graph is defined by a fixed maximum number of polynomials of a fixed maximum degree.

Proof. $E_{\mathcal{F}}$ is the lower envelope of $E_{\mathcal{F}_1}, \ldots, E_{\mathcal{F}_k}$. Take, as above, the minimization diagrams $M_{\mathcal{F}_1}, \ldots, M_{\mathcal{F}_k}$, and overlay them, to obtain a planar subdivision, which is the arrangement of the edges of these individual minimization diagrams. We may, of course, interpret this overlay as the combination of the overlays of pairs of these minimization diagrams (for example, each vertex of the overlay, which is not a vertex of one of the diagrams $\mathcal{M}_{\mathcal{F}_i}$, is also a vertex of one of these pairwise overlays). However, in order to derive the relation asserted in the theorem, we treat the overlay in a more 'economical' manner, which can be regarded as a 2-dimensional extension of the technique of Har-Peled [14].

Specifically, let N_i denote the number of edges of $M_{\mathcal{F}_i}$, for $i = 1, \ldots, k$. Clearly, $N_i = O(|M_{\mathcal{F}_i}|)$. Put $N = \sum_{i=1}^{k} N_i$, and let V denote the number of crossings between these N arcs. Note that, by definition,

$$V = \sum_{i=1}^{k} \sum_{j=i+1}^{k} C(\mathcal{F}_i, \mathcal{F}_j).$$

It is therefore sufficient to show that

$$|M_{\mathcal{F}}| = O\left(k^{2+\varepsilon} + k^{1+\varepsilon} N + k^{\varepsilon} V\right),$$

for any $\varepsilon > 0$, and we establish this bound as follows.

Put $r = \lceil N/k \rceil$, and construct a $(1/r)$-cutting Ξ of the arrangement of the above N edges. This is a decomposition of the plane into cells, each of constant description complexity, such that each cell is crossed by at most N/r arcs of the arrangement. The size of a cutting is said to be its number of cells. As shown, e.g., by de Berg and Schwarzkopf [9], there exists a cutting Ξ of size

$$O\left(r + \frac{Vr^2}{N^2}\right) = O\left(1 + \frac{N}{k} + \frac{V}{k^2}\right).$$

Let τ be a cell of Ξ. It is crossed by at most $N/r \le k$ edges. Let m_i denote the number of edges of $M_{\mathcal{F}_i}$ that cross τ; we have $\sum_{i=1}^{k} m_i \le k$. It is easily seen that the number of functions of $\mathcal{F}_i$ that can attain $E_{\mathcal{F}_i}$ over τ is at most $m_i + 1$. Indeed, construct a spanning tree T of the adjacency graph of the faces of $M_{\mathcal{F}_i} \cap \tau$ (T exists since the adjacency graph is clearly connected). Each edge of T corresponds to an edge of $M_{\mathcal{F}_i}$ that crosses τ, so T has at most m_i edges, and thus at most $m_i + 1$ nodes, corresponding to at most $m_i + 1$ faces of $M_{\mathcal{F}_i}$ that cross τ; this is easily seen to imply the claim.

We have thus shown that the number of functions that can attain $E_{\mathcal{F}}$ over τ is at most $\sum_{i=1}^{k}(m_i + 1) \le 2k$. The complexity of $E_{\mathcal{F}}$ over τ is thus $O(k^{2+\varepsilon})$, for any $\varepsilon > 0$ [12, 19]. Summing this bound over all cells τ of Ξ, the overall complexity of $E_{\mathcal{F}}$ is

$$O(k^{2+\varepsilon}|\Xi|) = O\left(k^{2+\varepsilon} \cdot \left(1 + \frac{N}{k} + \frac{V}{k^2}\right)\right) = O(k^{2+\varepsilon} + k^{1+\varepsilon} N + k^{\varepsilon} V),$$

for any $\varepsilon > 0$, as asserted. □

Remark. An obvious open problem is to extend Theorem 3.1 to the case of trivariate functions. Here we have a collection of 2-dimensional surface patches in $\mathbb{R}^3$, which are the faces of the individual minimization diagrams, and we want to construct a $(1/r)$-cutting for this collection. The crucial ingredient in the preceding proof is a sharp bound for the size of such a cutting, which becomes a considerably harder task in the trivariate case. Specifically, the only known general-purpose method for constructing cuttings of curved surfaces in three (and higher) dimensions uses the *vertical decomposition* of a sample Σ of the given surfaces (see [20]). The size of such a vertical decomposition depends on the number of *visibility events* in Σ, which are triples of the form (e, e', s), where each of e, e' is an intersection curve of two surfaces in Σ or a boundary edge or the silhouette of a single surface, and s is a vertical segment that connects a point on e to a point on e' and does not meet any other surface of Σ. Obtaining sharp bounds on the number of visibility events in a sample of faces of the k given minimization diagrams appears to be a fairly involved problem, which makes an extension of Theorem 3.1 to the trivariate case a difficult task.

4 The Overlay of Trivariate Minimization Diagrams

Armed with the extension given in Theorem 3.1, we apply the partition technique to prove the main result of the paper, which yields a near-cubic bound for the complexity of overlays of minimization diagrams of trivariate functions. The general approach is the same as the one demonstrated in the proof of Theorem 2.1, but the technical details are unfortunately more complicated.

4.1 Preliminaries

Let $\mathcal{F}$ and $\mathcal{G}$ be two collections, each consisting of n totally-defined trivariate functions of constant description complexity. Consider the overlay $Q(\mathcal{F}, \mathcal{G})$ of the minimization diagrams $M_{\mathcal{F}}$ (of $\mathcal{F}$) and $M_{\mathcal{G}}$ (of $\mathcal{G}$). The combinatorial complexity of $Q(\mathcal{F}, \mathcal{G})$ counts the number of cells of all dimensions in the overlay.

Each vertex (0-dimensional cell) of $Q(\mathcal{F}, \mathcal{G})$ is either a vertex of $M_{\mathcal{F}}$ or of $M_{\mathcal{G}}$, or a crossing between an edge of one diagram and a 2-face of the other. Denote by $C_{32}(\mathcal{F}, \mathcal{G})$ the number of crossings between edges of $M_{\mathcal{F}}$ and 2-faces of $M_{\mathcal{G}}$ (the subscripts 3 and 2 indicate that we consider interactions between features of $M_{\mathcal{F}}$ defined by three functions and features of $M_{\mathcal{G}}$ defined by two functions). Our analysis will concentrate on $C_{32}(\mathcal{F}, \mathcal{G})$, and the complementary count, $C_{23}(\mathcal{F}, \mathcal{G})$, of the number of crossings between 2-faces of $M_{\mathcal{F}}$ and edges of $M_{\mathcal{G}}$, will be handled in a fully symmetric manner. Vertices counted in $C_{32}(\mathcal{F}, \mathcal{G})$ will sometimes be referred to as $(3, 2)$-vertices, and vertices in $C_{23}(\mathcal{F}, \mathcal{G})$ can similarly be called $(2, 3)$-vertices.

Each edge (1-dimensional cell) of $Q(\mathcal{F}, \mathcal{G})$ is either (a portion of) an edge of $M_{\mathcal{F}}$ or of $M_{\mathcal{G}}$, or a maximal connected portion of an intersection curve between a 2-face of $M_{\mathcal{F}}$ and a 2-face of $M_{\mathcal{G}}$, which does not meet any other 2-face of either diagram. Any edge of the latter kind that has an endpoint can be charged to that endpoint, which is a vertex of the overlay. Under an appropriate general position assumption, each vertex is charged in this manner only $O(1)$ times, so we will not have to be concerned explicitly with such edges. Any other edge is either unbounded, or a closed bounded connected curve without a vertex. An unbounded edge e, formed by the intersection of two 2-faces φ_1 of $M_{\mathcal{F}}$ and φ_2 of $M_{\mathcal{G}}$, can be charged to a crossing between φ_1 and φ_2 at infinity (or, alternatively, at some sufficiently large sphere or cube, enclosing all bounded features of $Q(\mathcal{F}, \mathcal{G})$). The number of such crossings is equal to the complexity of the overlay of $M_{\mathcal{F}}$ and $M_{\mathcal{G}}$ at infinity, which is the overlay of two minimization diagrams of *bivariate* functions of constant description complexity, and its complexity, denoted as $C_\infty(\mathcal{F},\mathcal{G})$, is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$ (see [2] and the analysis in Section 2).

We are thus left with edges of Q that are closed and bounded connected components of 'bichromatic' intersection curves. Computing the number of such edges turns out to be the most involved part of our analysis. We denote their number by $C_{22}(\mathcal{F},\mathcal{G})$.

There is no need to consider separately 2-faces of $Q(\mathcal{F},\mathcal{G})$. Any 2-face φ of $Q(\mathcal{F},\mathcal{G})$ can be charged to a bounding edge, except for those 2-faces that have no boundary. The number of 2-faces of this latter kind is only $O(n^2)$, as is easily seen.

Levels. We can extend the above definitions as follows. The *level* of a point $\mathbf{x}$ in the arrangement $\mathcal{A}(\mathcal{F})$ of $\mathcal{F}$ is the number of graphs of functions in $\mathcal{F}$ that lie vertically below $\mathbf{x}$, and similarly for $\mathcal{A}(\mathcal{G})$. Consider a projection of an a-level edge (resp., 2-face) of $\mathcal{A}(\mathcal{F})$, and a projection of a b-level 2-face (resp., edge) of $\mathcal{A}(\mathcal{G})$. A crossing between these two projections is said to be an *(a, b)-level overlay vertex* (note, though, that these vertices do not show up at all in the actual overlay, unless $a = b = 0$). The number of such vertices is denoted by $C_{32}^{(a,b)}(\mathcal{F},\mathcal{G})$ (resp., $C_{23}^{(a,b)}(\mathcal{F},\mathcal{G})$). Denote by $C_{32}^{(\le k)}(\mathcal{F},\mathcal{G})$ (resp., $C_{23}^{(\le k)}(\mathcal{F},\mathcal{G})$) the number of all such vertices for which $a + b \le k$. $C_{22}^{(\le k)}(\mathcal{F},\mathcal{G})$ is defined analogously. (Note that, since our functions are totally defined and since we count here curves that have no vertex, the level of a curve is well defined, and is the same for all points on the curve.) Obviously, $C_{32}(\mathcal{F},\mathcal{G}) = C_{32}^{(0,0)}(\mathcal{F},\mathcal{G})$, and similarly for $C_{23}(\mathcal{F},\mathcal{G})$ and $C_{22}(\mathcal{F},\mathcal{G})$.

Denote by $C_{32}^{(\le k)}(n)$ the maximum value of $C_{32}^{(\le k)}(\mathcal{F},\mathcal{G})$, over all collections $\mathcal{F},\mathcal{G}$, each consisting of n trivariate functions of the same constant description complexity. The quantities $C_{23}^{(\le k)}(n)$, $C_{22}^{(\le k)}(n)$ are defined analogously. For the case $k = 0$, where features of the actual overlay are counted, we use the notations $C_{32}(n)$, $C_{23}(n)$, $C_{22}(n)$, respectively. Since $\mathcal{F}$ and $\mathcal{G}$ are assumed to belong to the same general class of functions, we have $C_{32}(n) = C_{23}(n)$ and $C_{32}^{(\le k)}(n) = C_{23}^{(\le k)}(n)$, so we will only address the quantities $C_{32}(n)$ and $C_{32}^{(\le k)}(n)$, and not their symmetric counterparts, in what follows.

We can estimate $C_{32}^{(\le k)}(n)$ in terms of $C_{32}(n)$. Since every crossing counted in $C_{32}^{(\le k)}(\mathcal{F},\mathcal{G})$ is defined by five surfaces, the standard random sampling argument of Clarkson and Shor [8] implies

$$C_{32}^{(\le k)}(n) = O\left(k^5 C_{32}\left(\frac{n}{k}\right)\right). \tag{4}$$

Similarly, we get for overlay edges (each defined by four surfaces) that

$$C_{22}^{(\le k)}(n) = O\left(k^4 C_{22}\left(\frac{n}{k}\right)\right). \tag{5}$$

Remark. A more general problem involving overlays of trivariate functions is that of the overlay of *three* or more minimization diagrams, rather than just two. The reason is that overlays of three minimization diagrams contain a new kind of features: vertices formed by the intersection of three 2-faces, one from each diagram. Overlays of two diagrams are special in that they do not give rise to such vertices. Nevertheless, as will be shown below, our analysis will lead us straight into the consideration of such triple overlays. As a byproduct, we will also obtain near-cubic bounds on their complexity.

4.2 Overlay Vertices (Counting $C_{32}(n)$)

As in Section 2, we apply a 2-stage analysis. We fix a parameter k, and, in the first stage, we partition $\mathcal{G}$ into k subgroups $\mathcal{G}_1, \ldots, \mathcal{G}_k$, each consisting of n/k functions.

Fix an edge e of $M_{\mathcal{F}}$, and erect a 2-dimensional wall $V^{(e)}$ over e, consisting of all x_4-parallel lines that pass through e. Restrict the functions of $\mathcal{G}$ over e, to obtain a collection $\mathcal{G}^{(e)}$ of univariate functions of constant description complexity, which is partitioned in an obvious way into k subcollections $\mathcal{G}_1^{(e)}, \ldots, \mathcal{G}_k^{(e)}$.

Let s denote the maximum number of intersections between the xy-projections of an intersection curve of three function graphs in $\mathcal{F}$, and of an intersection 2-surface of two function graphs in $\mathcal{G}$. Since $\mathcal{F}$ and $\mathcal{G}$ are assumed to belong to the same general class of functions, we may assume that s also bounds the number of intersections between the xy-projections of any intersection curve of three function graphs in $\mathcal{G}$, and of an intersection 2-surface of two function graphs in $\mathcal{F}$.

Consider the lower envelopes $E_{\mathcal{G}_i^{(e)}}$ of $\mathcal{G}_i^{(e)}$, for $i = 1, \ldots, k$. Note that the lower envelope of the $E_{\mathcal{G}_i^{(e)}}$'s is the restriction $E_{\mathcal{G}^{(e)}}$ of the lower envelope $E_{\mathcal{G}}$ over e, and that the number of breakpoints of $E_{\mathcal{G}^{(e)}}$ is the number of crossings between e and the 2-faces of $M_{\mathcal{G}}$ (and similarly for each $E_{\mathcal{G}_i^{(e)}}$). Hence, applying the analysis of Section 2, the complexity of $E_{\mathcal{G}^{(e)}}$ is

$$|E_{\mathcal{G}^{(e)}}| = O\left(\beta_s(k) \sum_{i=1}^{k} \left(1 + |E_{\mathcal{G}_i^{(e)}}|\right)\right).$$

Summing this bound over all edges e of $M_{\mathcal{F}}$, we obtain the recurrence

$$C_{32}(\mathcal{F}, \mathcal{G}) = O\left(|M_{\mathcal{F}}|\lambda_s(k) + \beta_s(k) \sum_{i=1}^{k} C_{32}(\mathcal{F}, \mathcal{G}_i)\right). \tag{6}$$

In the (considerably more involved) second stage, we partition $\mathcal{F}$ into k groups $\mathcal{F}_1, \ldots, \mathcal{F}_k$, each consisting of n/k functions. Fix a subset $\mathcal{G}_b$, and a 2-face φ of $M_{\mathcal{G}_b}$. Let $V^{(\varphi)}$ denote the 3-dimensional wall erected over φ (it is the union of all x_4-parallel lines passing through points of φ). Restrict the functions of $\mathcal{F}$ over φ, to obtain a collection $\mathcal{F}^{(\varphi)}$ of bivariate functions of constant description complexity, which is partitioned in an obvious way into $\mathcal{F}_1^{(\varphi)}, \ldots, \mathcal{F}_k^{(\varphi)}$.

Note that the number of crossings between edges of $M_{\mathcal{F}}$ and φ is equal to the number of vertices of the lower envelope $E_{\mathcal{F}^{(\varphi)}}$ (over the relative interior of φ). Using Theorem 3.1, we can estimate $|E_{\mathcal{F}^{(\varphi)}}|$, or, rather, $|M_{\mathcal{F}^{(\varphi)}}|$, as follows.

$$|M_{\mathcal{F}^{(\varphi)}}| = O\left(k^{2+\varepsilon} + k^{1+\varepsilon} \sum_{i=1}^{k} |M_{\mathcal{F}_i^{(\varphi)}}| + k^{\varepsilon} \sum_{i=1}^{k} \sum_{j=i+1}^{k} C\left(\mathcal{F}_i^{(\varphi)}, \mathcal{F}_j^{(\varphi)}\right)\right).$$

We sum this bound over all 2-faces φ of $M_{\mathcal{G}_b}$, to obtain an upper bound for $C_{32}(\mathcal{F}, \mathcal{G}_b)$. The terms $k^{2+\varepsilon}$ add up to $O(k^{2+\varepsilon}|M_{\mathcal{G}_b}|)$. The sum $\sum_{\varphi} |M_{\mathcal{F}_i^{(\varphi)}}|$, for any fixed i, is equal, by definition, to

$$O\left(C_{32}(\mathcal{F}_i,\mathcal{G}_b)+C_{23}(\mathcal{F}_i,\mathcal{G}_b)+C_{22}(\mathcal{F}_i,\mathcal{G}_b)+C_{\infty}(\mathcal{F}_i,\mathcal{G}_b)\right).$$

Indeed, vertices of $M_{\mathcal{F}_i^{(\varphi)}}$ that lie in the relative interior of φ are counted in $C_{32}(\mathcal{F}_i,\mathcal{G}_b)$. Closed and bounded connected edges of $M_{\mathcal{F}_i^{(\varphi)}}$ that have no vertex and are fully contained in the relative interior of φ are counted in $C_{22}(\mathcal{F}_i,\mathcal{G}_b)$. Edges with no vertex that reach the boundary of φ induce, at their boundary crossings, $(2,3)$-vertices, and are thus counted in $C_{23}(\mathcal{F}_i,\mathcal{G}_b)$. Edges that reach infinity (which can happen when φ is unbounded) are counted in $C_{\infty}(\mathcal{F}_i,\mathcal{G}_b)$. Finally, edges with a vertex can be charged to that vertex.

The sum $\sum_{\varphi} C(\mathcal{F}_i^{(\varphi)},\mathcal{F}_j^{(\varphi)})$, for any fixed i,j, is equal to the number of vertices of the *triple overlay* of $M_{\mathcal{F}_i}$, $M_{\mathcal{F}_j}$, and $M_{\mathcal{G}_b}$, which are intersections of three 2-faces, one of each diagram; we refer to such vertices as *trichromatic.* We denote this triple overlay by $Q^*(\mathcal{F}_i,\mathcal{F}_j,\mathcal{G}_b)$, and denote by $C_{222}(\mathcal{F}_i,\mathcal{F}_j,\mathcal{G}_b)$ the number of trichromatic vertices of the overlay. We define the notations $C_{222}^{(\le k)}(\mathcal{F}_i,\mathcal{F}_j,\mathcal{G}_b)$, $C_{222}^{(\le k)}(n)$, and $C_{222}(n)$ in complete analogy with the definition of the similar quantities given above.

Note that we started our analysis by considering the overlay $Q(\mathcal{F},\mathcal{G})$ of only two minimization diagrams, and there we only needed to consider 'bichromatic' vertices, formed by the intersection of edges of one diagram with 2-faces of the other. As mentioned in the remark above, when we overlay three or more diagrams, we also encounter trichromatic vertices, formed by the intersection of three 2-faces, one of each diagram.

Combining the analysis just given with (6), and using the fact that $|M_{\mathcal{F}}| = O(n^{3+\varepsilon})$, for any $\varepsilon > 0$ [19], we obtain, for any $\varepsilon > 0$,

$$\begin{aligned} C_{32}(n) &= O\Bigg(\lambda_s(k)n^{3+\varepsilon}+\beta_s(k)\sum_{b=1}^{k}\Bigg[k^{2+\varepsilon}|M_{\mathcal{G}_b}|+ \\ &\qquad k^{1+\varepsilon}\sum_{i=1}^{k}\left(C_{32}\left(\frac{n}{k}\right)+C_{22}\left(\frac{n}{k}\right)+\left(\frac{n}{k}\right)^{2+\varepsilon}\right)+k^{\varepsilon}\sum_{i=1}^{k}\sum_{j=i+1}^{k}C_{222}\left(\frac{n}{k}\right)\Bigg]\Bigg) \\ &= O\left(\lambda_s(k)n^{3+\varepsilon}+k^{3+\varepsilon}C_{32}\left(\frac{n}{k}\right)+k^{3+\varepsilon}C_{22}\left(\frac{n}{k}\right)+k^{3+\varepsilon}C_{222}\left(\frac{n}{k}\right)\right). \end{aligned} \tag{7}$$

Note that the final value of ε in this relation has to be taken slightly larger than the one we started with, to accommodate the extra factor $\beta_s(n)$. However, the recurrence still holds for any, arbitrarily small $\varepsilon > 0$.

4.3 Trichromatic Vertices (Counting $C_{222}(n)$)

Let $\mathcal{F},\mathcal{G},\mathcal{H}$ be three collections, each consisting of n trivariate functions of constant description complexity, as above. We want to estimate the number $C_{222}(\mathcal{F},\mathcal{G},\mathcal{H})$ of triple intersections of 2-faces in the triple overlay $Q^*(\mathcal{F},\mathcal{G},\mathcal{H})$ of $M_{\mathcal{F}},M_{\mathcal{G}},M_{\mathcal{H}}$. Here we use a 3-stage analysis, based on a variant of the analysis of the overlay of bivariate functions, given in Section 2. Interestingly, this part of our analysis of overlays of trivariate minimization diagrams is the simplest. Informally, this is because, in each of the three stages below, we only have to consider overlays of univariate functions, as in Section 2.

We fix a parameter k, and partition $\mathcal{H}$ into k groups $\mathcal{H}_1,\ldots,\mathcal{H}_k$ of n/k functions each. Fix a 2-face φ_1 of $M_{\mathcal{F}}$ and a 2-face φ_2 of $M_{\mathcal{G}}$, and let e be a connected component of the intersection curve $\varphi_1 \cap \varphi_2$ in $\mathbb{R}^3$. Note that the number of such edges e that have at least one endpoint is proportional to $C_{32}(\mathcal{F}, \mathcal{G}) + C_{23}(\mathcal{F}, \mathcal{G})$, because each endpoint of e is a crossing between an edge of one diagram and a 2-face of the other. Any other edge is either unbounded or a closed and bounded Jordan curve, and their number, as analyzed above, is $O(n^{2+\varepsilon} + C_{22}(n))$. Thus the number of edges e under consideration is $O(n^{2+\varepsilon} + C_{32}(n) + C_{22}(n))$.

Restrict the functions in $\mathcal{H}$ to the vertical wall $V^{(e)}$, as defined above, to obtain a collection $\mathcal{H}^{(e)}$ of univariate functions of constant description complexity, which is partitioned in an obvious way into k subcollections $\mathcal{H}_1^{(e)}, \ldots, \mathcal{H}_k^{(e)}$. As above, we have

$$|M_{\mathcal{H}^{(e)}}| = O\left(\beta_s(k) \sum_{i=1}^{k} \left(1 + |M_{\mathcal{H}_i^{(e)}}|\right)\right),$$

and, summing this bound over all such edges e, we obtain, by definition,

$$C_{222}(\mathcal{F}, \mathcal{G}, \mathcal{H}) = O\left(\left(n^{2+\varepsilon} + C_{22}(n) + C_{32}(n)\right) \lambda_s(k) + \beta_s(k) \sum_{i=1}^{k} C_{222}(\mathcal{F}, \mathcal{G}, \mathcal{H}_i)\right).$$

We repeat the same counting stage twice more. In the second stage, we fix $\mathcal{F}$ and one subcollection $\mathcal{H}_c$, partition $\mathcal{G}$ into k subgroups $\mathcal{G}_1, \ldots, \mathcal{G}_k$, and conclude, as above, that

$$C_{222}(\mathcal{F}, \mathcal{G}, \mathcal{H}_c) = O\left(\left(n^{2+\varepsilon} + C_{22}(n) + C_{32}(n)\right) \lambda_s(k) + \beta_s(k) \sum_{j=1}^{k} C_{222}(\mathcal{F}, \mathcal{G}_j, \mathcal{H}_c)\right).$$

(Note that the 'overhead' term depends on $C_{22}(\mathcal{F}, \mathcal{H}_c) + C_{32}(\mathcal{F}, \mathcal{H}_c) + C_{23}(\mathcal{F}, \mathcal{H}_c) + C_{\infty}(\mathcal{F}, \mathcal{H}_c)$, which still involves $\Theta(n)$ functions in $\mathcal{F}$.)

Similarly, in the third stage, we fix two subsets $\mathcal{G}_b$ and $\mathcal{H}_c$, partition $\mathcal{F}$ into k subgroups $\mathcal{F}_1, \ldots, \mathcal{F}_k$, and obtain

$$C_{222}(\mathcal{F}, \mathcal{G}_b, \mathcal{H}_c) = O\left(\left(\left(\frac{n}{k}\right)^{2+\varepsilon} + C_{22}\left(\frac{n}{k}\right) + C_{32}\left(\frac{n}{k}\right)\right) \lambda_s(k) + \beta_s(k) \sum_{\ell=1}^{k} C_{222}\left(\mathcal{F}_\ell, \mathcal{G}_b, \mathcal{H}_c\right)\right).$$

Substituting these estimates into one another, we obtain

$$C_{222}(n) = O\Bigg(\left(n^{2+\varepsilon} + C_{22}(n) + C_{32}(n)\right) \lambda_s(k) + \lambda_s(k) \Bigg[\left(n^{2+\varepsilon} + C_{22}(n) + C_{32}(n)\right) \lambda_s(k) + \lambda_s(k) \left[\left(\left(\frac{n}{k}\right)^{2+\varepsilon} + C_{22}\left(\frac{n}{k}\right) + C_{32}\left(\frac{n}{k}\right)\right) \lambda_s(k) + \lambda_s(k) C_{222}\left(\frac{n}{k}\right)\right]\Bigg]\Bigg),$$

or

$$C_{222}(n) = O\left(\lambda_s^2(k)\left[n^{2+\varepsilon} + C_{22}(n) + C_{32}(n)\right] + \lambda_s^3(k)\left[C_{22}\left(\frac{n}{k}\right) + C_{32}\left(\frac{n}{k}\right) + C_{222}\left(\frac{n}{k}\right)\right]\right). \tag{8}$$

4.4 Overlay Edges (Counting $C_{22}(n)$)

In this section we analyze the quantity $C_{22}(\mathcal{F}, \mathcal{G})$, which is the number of 'bichromatic' overlay edges that are closed bounded Jordan curves, each formed as a connected component of an intersection of a face of $M_{\mathcal{F}}$ and a face of $M_{\mathcal{G}}$, which are not adjacent to any overlay vertex. Overlay edges are structurally quite different from overlay vertices. To bound their number, we do not use the partition technique, but employ the well-known technique of counting schemes, introduced by Halperin and Sharir [12, 19] (see also [20]), and already mentioned in the introduction. This part of the analysis is essentially identical to the corresponding part in the precursor paper [17].

Our counting scheme for $C_{22}(\mathcal{F}, \mathcal{G})$ is another novel technical feature of this paper, and no similar scheme, that charges curves rather than vertices, has been employed in previous works involving substructures in arrangements. In this counting scheme, we partition the set of overlay edges into a small number of groups, according to the *index* of the edge (a notion to be defined shortly, and quite different from the standard notion of indices, as used in the analysis of substructures in arrangements [20]). We establish a separate recurrence relation for the number of edges in each group. These recurrences will depend on each other, as well as on $C_{32}(n)$.

Fix some threshold parameter k. Consider the (at most a constant number of) connected components of a bichromatic overlay intersection curve defined by a fixed quadruple of surfaces $f_1, f_2 \in \mathcal{F}$, $g_1, g_2 \in \mathcal{G}$. Suppose one of these components contains a point at level at most k and is either incident to an $(\le k)$-level overlay vertex or extends to infinity. The discussion in Section 4.1 easily implies that the number of such overlay edges (that is, edges counted in $C_{22}(n)$ for which a sibling component of the same intersection curve satisfies the above properties), is, for any $\varepsilon > 0$,

$$O\left(k^5 C_{32}\left(\frac{n}{k}\right) + k^2 n^{2+\varepsilon}\right). \tag{9}$$

In the remainder of this section we treat bichromatic overlay edges defined by some quadruple f_1, f_2, g_1, g_2, such that all $(\le k)$-level edges defined by f_1, f_2, g_1, g_2 are closed and bounded Jordan curves that are not incident to any vertex. (The level of such a curve is well-defined: all points on the curve have the same level, assuming, as above, that the functions in $\mathcal{F}$ and $\mathcal{G}$ are totally defined.) We fix a 2-face φ that belongs to $E_{\mathcal{G}}$, and is a connected portion of the intersection of the graphs of two functions $g_1, g_2 \in \mathcal{G}$, and consider the vertical 3-dimensional wall $V^{(\varphi)}$ erected over φ, as in the previous sections. Let $\mathcal{F}^{(\varphi)}$ be the cross-section of $\mathcal{F}$ within $V^{(\varphi)}$, and let f_φ be the cross-section of the graph of a function f within $V^{(\varphi)}$, for each $f \in \mathcal{F}$. $\mathcal{A}(\mathcal{F}^{(\varphi)})$ can be regarded as an arrangement of xy-monotone 2-dimensional surfaces in $\mathbb{R}^3$.

Let $\Gamma(\mathcal{F}^{(\varphi)})$ denote the set of edges, for which the following holds. Each edge c of $\Gamma(\mathcal{F}^{(\varphi)})$ is a 0-level edge that lies completely in the lower envelope of $\mathcal{A}(\mathcal{F}^{(\varphi)})$ and is a connected component of an intersection curve $f_1 \cap f_2$, for some $f_1, f_2 \in \mathcal{F}^{(\varphi)}$, such that all the connected components (including c) of this intersection that lie at level at most k of $\mathcal{A}(\mathcal{F}^{(\varphi)})$, are closed and bounded Jordan curves, and are not incident to any vertex. Note that any bichromatic overlay edge of the kind we are after, corresponds to a curve in $\Gamma(\mathcal{F}^{(\varphi)})$ for some 2-face φ of $E_{\mathcal{G}}$.

Define $S(\mathcal{F}, \mathcal{G}) \equiv \sum_\varphi |\Gamma(\mathcal{F}^{(\varphi)})|$, and $S(n) \equiv \max S(\mathcal{F}, \mathcal{G})$, where the maximum is taken over all sets $\mathcal{F}$ and $\mathcal{G}$ of size n as above. The above analysis implies, for any $\varepsilon > 0$,

$$C_{22}(n) = O\left(S(n) + k^5 C_{32}\left(\frac{n}{k}\right) + k^2 n^{2+\varepsilon}\right). \qquad (10)$$

We next define the notion of *index* that we attach to intersection curves. For readers who have encountered previous similar-purpose definitions [19], we remark that our definition is conceptually different, in that it is not "local", meaning that the index given to a specific intersection curve is defined with respect to the whole arrangement $\mathcal{A}(\mathcal{F}^{(\varphi)})$, and may change as the set $\mathcal{F}^{(\varphi)}$ decreases.

Specifically, consider the intersection $f_1 \cap f_2$, for any $f_1, f_2 \in \mathcal{F}^{(\varphi)}$. Consider all the connected components of this intersection that are edges of $\Gamma(\mathcal{F}^{(\varphi)})$, and let j be their number. We set the index of the curve $f_1 \cap f_2$ to j. We will say that j is also the index of all the edges of $\Gamma(\mathcal{F}^{(\varphi)})$ that are connected components of $f_1 \cap f_2$.

Let q be the maximum possible number of connected components of an intersection $f_1 \cap f_2$ as above. By the constant description complexity and the general position assumptions, q is constant. Clearly, j varies between 1 and q. The case $j = 0$ means that either no component of $f_1 \cap f_2$ shows up on the envelope, or some such components do show up, but there exists a component at level at most k that has a vertex or reaches infinity. We will not be concerned with edges of index 0, because either we do not have to count them at all, or else we can bound their number using (9).

Let $\Gamma^{(j)}(\mathcal{F}^{(\varphi)})$ denote the subset of $\Gamma(\mathcal{F}^{(\varphi)})$ that contains edges with index at least j. Define $S^{(j)}(\mathcal{F}, \mathcal{G}) \equiv \sum_{\varphi} |\Gamma^{(j)}(\mathcal{F}^{(\varphi)})|$, and $S^{(j)}(n) \equiv \max S^{(j)}(\mathcal{F}, \mathcal{G})$, where the maximum is taken over all sets $\mathcal{F}$ and $\mathcal{G}$ of size n as above. Since the maximal index of an edge is q, we have $S^{(q+1)}(n) = 0$. We also have, by definition, $S(n) = S^{(1)}(n)$.

We note that the index of $f_1 \cap f_2$ depends on the current set $\mathcal{F}$. When $\mathcal{F}$ is replaced by a smaller sample, as happens for example when applying the Clarkson-Shor bound, the index may increase, because either (i) more components of $f_1 \cap f_2$ appear on the envelope, or (ii) all vertices of $\mathcal{A}(\mathcal{F}^{(\varphi)})$ that lied on components of $f_1 \cap f_2$ at level at most k disappear, since all the third surfaces incident to these vertices have been removed from $\mathcal{F}$. (Note that the level of a curve can only decrease when functions are removed from $\mathcal{F}$.) In this latter case, the index jumps from 0 (no component of $f_1 \cap f_2$, even those on the envelope, qualified to belong to $\Gamma(\mathcal{F}^{(\varphi)})$ before $\mathcal{F}$ was reduced) to some j, equal to the number of components that lie on the envelope after the reduction.

The index of $f_1 \cap f_2$ can decrease in only one way, as follows. There may exist an intersection curve $f_1 \cap f_2$ that had a positive index, but, after $\mathcal{F}$ has been reduced, new components of $f_1 \cap f_2$ reach the ($\leq k$)-level of $\mathcal{A}(\mathcal{F}^{(\varphi)})$, and do contain vertices, or do reach infinity. The index of $f_1 \cap f_2$ then drops to 0. However, the number of such curves in the reduced arrangement can be estimated using the bound in (9) (applied to the reduced arrangement). To conclude, with the exception of these drops to zero, the index of a curve can only increase when the size of $\mathcal{F}$ is reduced.

The counting scheme below bounds $S^{(j)}(n)$, for all $1 \leq j \leq q$. It proceeds by distinguishing between five possible scenarios (Cases 1–5), and treating each in turn.

Case 1: $\Gamma^{(j)}(\mathcal{F}^{(\varphi)})$ is small. Suppose first that at most $(q+1)k+2 = O(k)$ surfaces of $\mathcal{F}$ attain the lower envelope of $\mathcal{A}(\mathcal{F}^{(\varphi)})$. In this case, we use the naive bound $|\Gamma^{(j)}(\mathcal{F}^{(\varphi)})| = O(k^2)$. Since there are $O(n^{3+\varepsilon})$ possible faces φ, the maximum number of edges of this kind that are counted in $S^{(j)}(n)$ is $O(k^2 n^{3+\varepsilon})$ (for all j). In the sequel we only consider faces φ, such that more than $(q+1)k+2$ surfaces of $\mathcal{F}$ attain the lower envelope of $\mathcal{A}(\mathcal{F}^{(\varphi)})$.

Case 2: There is a "shallow" connected component of the same intersection curve. Consider any pair of surfaces $P, Q \in \mathcal{F}$, such that there is a connected component c of the intersection $P_\varphi \cap Q_\varphi$ that is an edge of $\Gamma^{(j)}(\mathcal{F}^{(\varphi)})$. The component c is, by definition, also an edge of $\Gamma(\mathcal{F}^{(\varphi)})$, which implies that all the connected components of $P_\varphi \cap Q_\varphi$ that lie at level at most k of $\mathcal{A}(\mathcal{F}^{(\varphi)})$ are closed and bounded Jordan curves, and are not incident to a vertex. Also, since the index of $P_\varphi \cap Q_\varphi$ is at least j, the number of connected components of $P_\varphi \cap Q_\varphi$ that lie on the lower envelope of $\mathcal{A}(\mathcal{F}^{(\varphi)})$ is at least j. Suppose there is a connected component c' defined by $P_\varphi \cap Q_\varphi$ whose level is between 1 and k. (As noted, the level of a curve that satisfies the above properties is well-defined.) In this case, we charge all edges of $\Gamma^{(j)}(\mathcal{F}^{(\varphi)})$ defined by $P_\varphi \cap Q_\varphi$ to c'. It is easy to see that each such edge c' is charged in this fashion at most a constant number of times (that is, at most $q-1$ times).

Let $\Gamma(c')$ be the set of the (at most k) surfaces of $\mathcal{F}^{(\varphi)}$ that lie below c'. Set $\tilde{\mathcal{F}}^{(\varphi)} \equiv \mathcal{F}^{(\varphi)} \setminus \Gamma(c')$, and consider the arrangement $\mathcal{A}(\tilde{\mathcal{F}}^{(\varphi)})$. Clearly, c' lies on its lower envelope. Moreover, there are at least $j+1$ connected components of the intersection $P_\varphi \cap Q_\varphi$ that lie on this lower envelope. Thus, the index of c' is now at least $j+1$, and c' belongs to $\Gamma^{(j+1)}(\tilde{\mathcal{F}}^{(\varphi)})$. More accurately, c' belongs to $\Gamma^{(j+1)}(\tilde{\mathcal{F}}^{(\varphi)})$ unless a new component of $P_\varphi \cap Q_\varphi$, that either reaches infinity or contains a vertex of $\mathcal{A}(\tilde{\mathcal{F}}^{(\varphi)})$, has 'descended' to the first k levels in $\mathcal{A}(\tilde{\mathcal{F}}^{(\varphi)})$, thereby dropping the index of c' to 0. (Note that this analysis also applies to any subset $\mathcal{F}' \subseteq \tilde{\mathcal{F}}^{(\varphi)}$.) A standard random sampling argument, as the ones used in Section 4.1, now implies, in combination with (9), that the maximum number of edges of this kind that are counted in $S^{(j)}(\mathcal{F}, \mathcal{G})$ is

$$O\left(k^4 S^{(j+1)}\left(\frac{n}{k}\right) + k^4\left[k^5 C_{32}\left(\frac{n}{k^2}\right) + n^{2+\varepsilon}\right]\right) = O\left(k^4 S^{(j+1)}\left(\frac{n}{k}\right) + k^9 C_{32}\left(\frac{n}{k^2}\right) + k^4 n^{2+\varepsilon}\right).$$

In the sequel we assume that there is no connected component c' as above.

Case 3: There is a "shallow" vertex. Since Case 1 is ruled out, there are at least $(q+1)k$ surfaces, other than P_φ, Q_φ, that attain the lower envelope of $\mathcal{A}(\mathcal{F}^{(\varphi)})$ over φ. Consider all the connected components (edges) of $P_\varphi \cap Q_\varphi$ that lie in the lower envelope of $\mathcal{A}(\mathcal{F}^{(\varphi)})$. These edges partition both P_φ and Q_φ into at most $(q+1)$ pairs of relatively open regions, where each pair consists of two regions, one on P_φ and one on Q_φ, that have a common boundary (over the relative interior of φ) and a common projection onto φ. One of these projections, denoted by Δ_0, has to contain at least k subregions where k other distinct surfaces attain the envelope. Each of these surfaces, T_φ, lies strictly above $P_\varphi \cap Q_\varphi$ over all points of $\partial\Delta_0$ defined by $P_\varphi \cap Q_\varphi$. Hence, T_φ intersects both P_φ and Q_φ over Δ_0, and the projection of each component of either intersection onto φ is not incident to those boundary components of Δ_0 that are defined by $P_\varphi \cap Q_\varphi$. (We emphasize that the region Δ_0 is *fixed* in all the remaining charging steps, during the present case and in Cases 4 and 5 below.)

Consider an arbitrary edge π (defined by $P_\varphi \cap Q_\varphi$) whose projection onto φ lies on the boundary of Δ_0, and assume, without loss of generality, that P_φ lies below Q_φ when approaching (the projection of) π from within Δ_0. Let Δ denote the portion of Q_φ (the surface that is higher near π) that projects onto Δ_0 (see Figure 1(a)), and let C be the set of edges (each being a connected component of $O_\varphi \cap Q_\varphi$, for some $O_\varphi \in \mathcal{F}^{(\varphi)}$) contained in Δ. By what we have just argued, $|C| \geq k$.

Suppose that there is at least one vertex v on Q_φ within Δ that lies at level at most k, such that v can be connected to a point on π by an arc $\gamma_{\pi,v}$ (of some finite though not necessarily constant description complexity) that lies on Q_φ within Δ, and is at level $(\leq k)$, for all points in its relative interior. In this case, we charge all the edges of $\Gamma^{(j)}(\mathcal{F}^{(\varphi)})$ defined by $P_\varphi \cap Q_\varphi$ to an arbitrary such vertex (see Figure 1(b)). Note that each such vertex corresponds to an $(\leq k, 0)$-level overlay vertex.

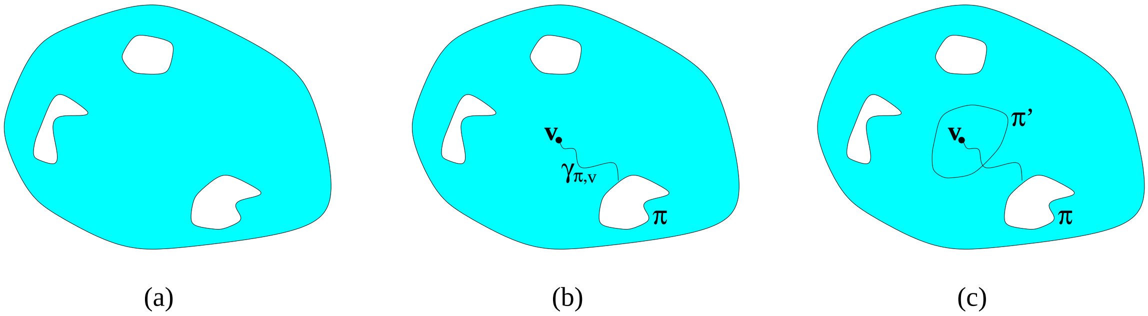

Figure 1: The region Δ (shaded) is illustrated in (a). (b) shows a point v that is charged by an edge π, together with the connecting arc $\gamma_{\pi,v}$. The interior of the arc $\gamma_{\pi,v}$ lies fully within Δ, and thus cannot intersect any connected component of $P_\varphi \cap Q_\varphi$ that lies on the lower envelope. (c) depicts the edge π' whose assumed existence leads to the contradiction in the proof of Lemma 4.1.

Similarly, suppose there is at least one edge that passes above/below the boundary of φ at level $\leq k$, while lying on Q_φ within Δ, at a certain point v, such that v can be connected to π by an arc $\gamma_{\pi,v}$ as above. By construction, this implies that part of $\partial\Delta$ is covertical with the boundary of φ, and this is the part of $\partial\Delta$ that contains v. (In general, notice that Δ is always 'bounded' by φ, in the sense that the projection Δ_0 of Δ is contained in φ. In the case under consideration, Δ_0 touches the boundary of φ, and thus part of $\partial\Delta_0$ lies on $\partial\varphi$.) In this case, we charge all the edges in $\Gamma^{(j)}(\mathcal{F}^{(\varphi)})$ defined by $P_\varphi \cap Q_\varphi$ to an arbitrary such point v, as above, noting that each such point again corresponds to an $(\leq k, 0)$-level overlay vertex. (Note that, in the preceding case, the charged vertex was a $(3,2)$-vertex, whereas now it is a $(2,3)$-vertex.)

An observation that will prove useful in the proof of the following lemma is that an arc $\gamma_{\pi,v}$ cannot cross any connected component of $P_\varphi \cap Q_\varphi$ that lies in the lower envelope of $\mathcal{A}(\mathcal{F}^{(\varphi)})$, since the interior of $\gamma_{\pi,v}$ is required to lie inside the region Δ.

Lemma 4.1. *Each point v is charged by at most qk distinct edges of $\Gamma^{(j)}(\mathcal{F}^{(\varphi)})$ as prescribed in Case 3.*

Proof. The proof is visualized in Figure 1. Suppose a point v that lies on Q_φ is charged by an edge π. By construction, π is a connected component of an intersection of Q_φ with another surface P_φ. The crucial observation is that P_φ has to lie below v. Indeed, suppose P_φ lies above v. By construction, there exists an arc $\gamma_{\pi,v}$ that connects v to a point on π, and lies fully within level $\leq k$ in its interior. Also by construction, P_φ lies below Q_φ when we approach π on $\gamma_{\pi,v}$, and our assumption states that P_φ lies above Q_φ when we approach v on $\gamma_{\pi,v}$. Thus, P_φ has to intersect the interior of $\gamma_{\pi,v}$, which implies the existence of an edge π' defined by $P_\varphi \cap Q_\varphi$, distinct from π, within level $\leq k$. As observed just before the statement of the lemma, $\gamma_{\pi,v}$ cannot intersect π' at level 0. Thus, the level of π' is between 1 and k. Such situation has however been ruled out in Case 2, leading us to a contradiction.

We have thus shown that v can only be charged by connected components of intersections of Q_φ with surfaces that lie below v. Since at most k surfaces lie below v, and each defines at most q such connected components (edges) along Q_φ, v is charged by at most qk distinct edges that lie on Q_φ. Since v lies on at most 3 surfaces of $\mathcal{F}^{(\varphi)}$, it can be charged by at most $3qk$ edges overall. □

Combined with (4), and with the fact that q is a constant, Lemma 4.1 implies that the maximum number of edges of this kind that are counted in $S^{(j)}(\mathcal{F},\mathcal{G})$ is $O(kC_{32}^{(\leq k)}(n)) = O(k^6 C_{32}(n/k))$. In the sequel we assume that there is no vertex v as above (for the specific region Δ which we now keep fixed).

Case 4: There is a "shallow" edge that reaches infinity. We continue to use the setup introduced in Case 3, and suppose that there is an edge c of C that lies at level at most k and is not a closed and bounded Jordan curve, and c can be reached from π along an arc $\gamma_{\pi,c}$, as above, that stays at level $\leq k$. (Since we assume that the scenario treated in Case 3 does not hold, this can only occur if Δ is unbounded, c has no vertices, and it reaches infinity. This, in turn, can only happen when φ is unbounded.) In this case, we charge all the edges of $\Gamma^{(j)}(\mathcal{F}^{(\varphi)})$ defined by $P_\varphi \cap Q_\varphi$ to the edge c. Arguing as above, we can show that the overall number of such edges is $O(k^2 n^{2+\varepsilon})$. The proof of Lemma 4.1 can easily be modified to show that each edge is charged at most $2qk$ times in this fashion. (In the modified proof we use the fact that the set of surfaces that lie below a point on c is the same for all points of c.) Thus, the maximum number of edges of this kind that are counted in $S^{(j)}(\mathcal{F},\mathcal{G})$ is $O(k^3 n^{2+\varepsilon})$. In the sequel we assume that there is no edge $c \in C$ as above (for our fixed Δ).

Case 5. In this final case we distinguish between two subcases. In both, we charge one edge π of $\Gamma^{(j)}(\mathcal{F}^{(\varphi)})$ defined by $P_\varphi \cap Q_\varphi$ (out of the at most q such edges), which bounds the region Δ under consideration, to k edges $c \in C$ that lie at level $\leq k$, and can be reached from π along an $(\leq k)$-level arc $\gamma_{\pi,c}$, as above.

Subcase 5.A: There is a point on Δ that lies at level $> k$. We can connect this point to a point on one of the boundary arcs π of Δ, defined by $P_\varphi \cap Q_\varphi$, by an arc (of some finite though not necessarily constant description complexity) that lies on Q_φ within Δ, in its interior. Since one end-point of this arc lies at level $> k$ and the other on the lower envelope, the arc intersects at least k distinct edges of C. Moreover, the first k distinct edges encountered when walking along the arc away from π, lie at level $\leq k$, since π lies on the lower envelope. We charge π to these first k edges of C.

Subcase 5.B: Δ lies entirely at level $\leq k$. We charge π to k arbitrary edges of C.

We emphasize that in both subcases there exists an arc $\gamma_{\pi,c}$, for each edge c charged by π, that connects a point on c to a point on π, lies on Q_φ within Δ, and its interior lies fully at level $\leq k$. In Subcase 5.A, $\gamma_{\pi,c}$ is the appropriate prefix of the arc used to identify the k edges that are charged, while in Subcase 5.B, $\gamma_{\pi,c}$ exists since Δ lies entirely at level $\leq k$ and is connected. Therefore, since Cases 1–4 are assumed not to occur, it follows that each of these edges c is a closed bounded Jordan curve not incident to any vertex. Indeed, had vertices incident to c existed, one of them would lie at level $\leq k$ and could be connected to a point on π as prescribed in Case 3. We have assumed however that there are no such vertices. The exclusion of Case 4 similarly implies that c is a closed bounded Jordan curve. (Note that there may nevertheless exist edges $c \in C$ at level $\leq k$ that contain a vertex or are unbounded, but are *unreachable* from π along a 'low-level' arc $\gamma_{\pi,c}$, as above.)

Let c be an edge of C that is charged in the above fashion. Lemma 4.2 below states that c is charged at most twice. Let $\Gamma(c)$ be the set of the (at most k) surfaces of $\mathcal{F}^{(\varphi)}$ that lie below c, and set $\tilde{\mathcal{F}}^{(\varphi)} \equiv \mathcal{F}^{(\varphi)} \setminus \Gamma(c)$. Since we assume that none of the scenarios treated in Cases 1–4 holds, it is easy to see that c is an edge of $\Gamma(\tilde{\mathcal{F}}^{(\varphi)})$, unless there exists a component c' 'sibling' to c, that meets the first k levels of $\mathcal{A}(\tilde{\mathcal{F}}^{(\varphi)})$ and either reaches infinity or contains a vertex of $\mathcal{A}(\tilde{\mathcal{F}}^{(\varphi)})$; the component c' has either descended to the first k levels in $\mathcal{A}(\tilde{\mathcal{F}}^{(\varphi)})$, thereby dropping the index of c to 0, or has already existed within the first k levels of $\mathcal{A}(\mathcal{F}^{(\varphi)})$, but was 'hidden' from the charging curve π, as described in the preceding paragraph. (These properties also hold for any subset of $\tilde{\mathcal{F}}^{(\varphi)}$.) As in Case 2, a standard random sampling argument, as the ones used in Section 4.1, now implies, in combination with (9), that the maximum number of edges of this kind that are counted in $S^{(j)}(\mathcal{F}, \mathcal{G})$ is

$$\frac{q}{k} \cdot O\left(k^4 S\left(\frac{n}{k}\right) + k^4\left[k^5 C_{32}\left(\frac{n}{k^2}\right) + n^{2+\varepsilon}\right]\right) = O\left(k^3 S\left(\frac{n}{k}\right) + k^8 C_{32}\left(\frac{n}{k^2}\right) + k^3 n^{2+\varepsilon}\right).$$

We now give the lemma, referred to in the beginning of the paragraph, that ensures that each edge $c \in C$ is charged in the above fashion at most twice.

Lemma 4.2. *Each curve is charged by at most two distinct edges of* $\Gamma^{(j)}(\mathcal{F}^{(\varphi)})$ *as prescribed in Case 5.*

Proof. By construction, an edge $c \subseteq Q_\varphi \cap O_\varphi$ can only be charged by edges that lie either on Q_φ or on O_φ. Assume, for the sake of contradiction, that c is charged by two edges, $\pi \subseteq Q_\varphi \cap P_\varphi$ and $\theta \subseteq Q_\varphi \cap T_\varphi$, for some $P_\varphi, T_\varphi \in \mathcal{F}^{(\varphi)}$ (see Figure 2). Notice that the level of c is at most k and that c can be connected to π and θ by arcs $\gamma_{\pi,c}$ and $\gamma_{\theta,c}$, respectively, as described above (see Figure 2(a) for an illustration). As argued above, this implies that c is a closed bounded Jordan curve that is incident to no vertex.

We can assume, without loss of generality, that θ and c lie on the same side of the closed curve π. The arguments in the proof of Lemma 4.1 imply that P_φ has to lie completely below c. Consider the arc $\gamma_{\theta,c}$ connecting a point on c to a point on θ, as above. P_φ lies below Q_φ when we approach c on $\gamma_{\theta,c}$, but, since θ lies on the lower envelope, P_φ lies above Q_φ when we approach θ on $\gamma_{\theta,c}$. Thus, P_φ has to intersect the relative interior of $\gamma_{\theta,c}$, which implies the existence of a closed curve π' defined by $P_\varphi \cap Q_\varphi$, within level $\leq k$, such that θ and c lie on different sides of π'. π' is therefore distinct from π. If its level is between 1 and k, we reach a contradiction, since such situations have been ruled out in Case 2. π' therefore lies on the lower envelope.

Consider the part $\gamma'_{\theta,c}$ of $\gamma_{\theta,c}$ that lies between (the last intersection of $\gamma_{\theta,c}$ with) π' and θ (as shown in Figure 2(b)). By construction, T_φ lies below Q_φ when we approach θ on

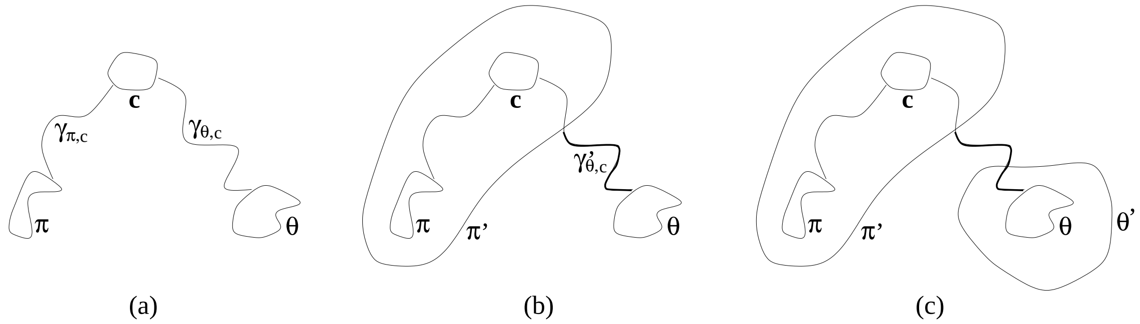

Figure 2: A schematic visualization of the proof of Lemma 4.2. The general setup is introduced in (a); (b) illustrates the edge π' and the arc $\gamma'_{\theta,c}$ (with the latter thickened), and (c) illustrates the edge θ'.

$\gamma'_{\theta,c}$, but, since π' lies on the lower envelope, T_φ lies above Q_φ when we approach π' on $\gamma'_{\theta,c}$. Thus, T_φ has to intersect the relative interior of $\gamma'_{\theta,c}$, which implies the existence of a closed curve θ' defined by $T_\varphi \cap Q_\varphi$, distinct from θ, within level $\le k$ (as illustrated in Figure 2(c)). As observed just before the statement of Lemma 4.1, $\gamma_{\theta,c}$ cannot intersect θ' at level 0. Thus, the level of θ' is between 1 and k. Such situation has however been ruled out in Case 2, leading to a contradiction.

We have shown that c can be charged by a connected component of an intersection of Q_φ with only one other surface. A symmetric statement holds for O_φ. $c \subseteq Q_\varphi \cap O_\varphi$ can thus be charged at most twice. $\square$

We can now write the following relations, for all $1 \le j \le q$. (Note that for $j = q$ the second term on the right side is not present.)

$$S^{(j)}(n) = O\left[k^3 S\left(\frac{n}{k}\right) + k^4 S^{(j+1)}\left(\frac{n}{k}\right) + k^6 C_{32}\left(\frac{n}{k}\right) + k^9 C_{32}\left(\frac{n}{k^2}\right) + k^4 n^{3+\varepsilon}\right]. \quad (11)$$

4.5 Putting It All Together

We claim that the system of inter-dependent recurrences derived in this section, given in (7), (8), (10), (11), solves to

$$C_{32}(n) = O(n^{3+\varepsilon}), \qquad C_{22}(n) = O(n^{3+\varepsilon}), \qquad C_{222}(n) = O(n^{3+\varepsilon}),$$

for any $\varepsilon > 0$. This is shown by induction, as in [19], choosing a different value of k for each recurrence. In more detail, we order the functions appearing in the recurrences as $(C_{22}, S^{(1)}, S^{(2)}, \ldots, S^{(q)}, C_{222}, C_{32})$, and denote this, for uniformity, as $(F_1, F_2, \ldots, F_{q+3})$. Each recurrence is roughly of the form

$$F_i(n) = O\left(k_i^{\beta_1} F_{j_1}\left(\frac{n}{k_i^{\alpha_1}}\right)\right) + O\left(k_i^{\beta_2} F_{j_2}\left(\frac{n}{k_i^{\alpha_2}}\right)\right) + \cdots + O\left(k_i^{\beta_r} F_{j_r}\left(\frac{n}{k_i^{\alpha_r}}\right)\right) + O\left(f_i(n)\right).$$

We represent this system symbolically by a directed graph G on the indices $\{1, 2, \ldots, q+3\}$, whose directed edges are $(i, j_1), \ldots, (i, j_r)$, for all i. We call an edge (i, j) a *forward* (resp., *backward*) edge if $j > i$ (resp., $j \le i$). Let γ be the maximum of the ratios β_t/α_t, taken over all corresponding edges (i, j_t) that are *backward* edges, and assume also that $f_i(n) = O(n^{\gamma+\varepsilon})$, for each i and for any $\varepsilon > 0$. Then one can show that the solution of this system is $F_i(n) = O(n^{\gamma+\varepsilon})$, for any $\varepsilon > 0$ and for all i. Informally, larger exponent ratios in terms that relate F_i to a function F_j with a *larger* index do not affect the overall bound, because (almost all of) their effect can be suppressed by the choice of appropriate values for the k_i's, which decrease exponentially with i.

Since in our case, under the order given above, $\gamma = 3$, we obtain the bound asserted above.[2] This completes the proof of Theorem 1.1.

5 Applications

Our bound on the complexity of overlays in $\mathbb{R}^4$ has many applications. We mention several of the more obvious ones. All the results listed below improve significantly upon the best previously known ones. Their proofs crucially rely on Theorem 1.1. Some of the more standard details in the proofs are omitted.

The region 'sandwiched' between two envelopes. Let $\mathcal{F}$ and $\mathcal{G}$ be two families of n trivariate functions of constant description complexity, as above. Let $\Sigma_{\mathcal{F},\mathcal{G}}$ denote the *sandwich region* consisting of all points that lie above the upper envelope $E_{\mathcal{G}}$ of $\mathcal{G}$ and below the lower envelope $E_{\mathcal{F}}$ of $\mathcal{F}$. That is, $\Sigma_{\mathcal{F},\mathcal{G}}$ is the set of all quadruples (x_1, x_2, x_3, x_4), such that $g(x_1, x_2, x_3, x_4) \le x_4 \le f(x_1, x_2, x_3, x_4)$, for each $f \in \mathcal{F}$, $g \in \mathcal{G}$.

Theorem 5.1. *The combinatorial complexity of the sandwich region $\Sigma_{\mathcal{F},\mathcal{G}}$ is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$.*

Proof. Consider for example the number of vertices of $\Sigma_{\mathcal{F},\mathcal{G}}$. Any such vertex is either (i) a vertex of $E_{\mathcal{F}}$ or of $E_{\mathcal{G}}$ (there are $O(n^{3+\varepsilon})$ such vertices, for any $\varepsilon > 0$), or (ii) an intersection between an edge e of $E_{\mathcal{F}}$ and a facet φ of $E_{\mathcal{G}}$, or (iii) an intersection between an edge e of $E_{\mathcal{G}}$ and a facet φ of $E_{\mathcal{F}}$, or (iv) an intersection between a 2-face f of $E_{\mathcal{F}}$ and a 2-face g of $E_{\mathcal{G}}$. Consider the overlay $Q(\mathcal{F}, \mathcal{G})$ of the minimization diagram $M_{\mathcal{F}}$ of $\mathcal{F}$ and the *maximization diagram* $M_{\mathcal{G}}$ of $\mathcal{G}$ (defined in complete analogy to the definition of minimization diagrams). In cases (ii) and (iii), the projections of e and of φ in $Q(\mathcal{F}, \mathcal{G})$ have a nonempty intersection. That is, there exists a connected portion of e that appears as a feature of $Q(\mathcal{F}, \mathcal{G})$, and the cells of $Q(\mathcal{F}, \mathcal{G})$ that it bounds are portions of the projection of φ. Similarly, in case (iv), the projections of f and of g intersect in a curve that is a feature (or a union of features) of $Q(\mathcal{F}, \mathcal{G})$. We can thus charge vertices of $\Sigma_{\mathcal{F},\mathcal{G}}$ to features of $Q(\mathcal{F}, \mathcal{G})$ in a unique manner, which clearly implies the claim. □

Note that the bound in Theorem 5.1 is nearly tight in the worst case. As a matter of fact, the proof of Theorem 5.1 implies the following stronger result; we refer the reader to [7, 16] for details concerning vertical decompositions in four dimensions.

Corollary 5.2. *The combinatorial complexity of the first stage of the vertical decomposition of the sandwich region $\Sigma_{\mathcal{F},\mathcal{G}}$ is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$.*

[2] Technically, γ is not quite 3, because of the factors $\beta_s(k)$ and k^{ε} that are also present in our recurrences. However, any $\gamma > 3$ can be used as a bound for the exponent of the solution, so $O(n^{3+\varepsilon})$ is a solution of the system.

This corollary still leaves open the question of whether the complexity of the entire vertical decomposition of $\Sigma_{\mathcal{F},\mathcal{G}}$ is near-cubic, or at least sub-quartic. This problem is still open even when one collection is empty, i.e., the problem concerning the vertical decomposition of the region below the lower envelope of a collection of trivariate functions.

The space of hyperplane transversals in 4-space. Let $\mathcal{C}$ be a collection of n convex sets in $\mathbb{R}^4$, each being semi-algebraic of constant description complexity. Let $T_3(\mathcal{C})$ denote the space of all hyperplane transversals of $\mathcal{C}$; i.e., the set of all hyperplanes that intersect every member of $\mathcal{C}$. Using a standard duality transformation [10], we map hyperplanes to points and points to hyperplanes, so that the incidence and the above/below relationships between points and hyperplanes are preserved. (This transformation excludes hyperplanes parallel to the x_4-axis, which can be handled separately, in a much simpler manner.) Then each $c \in \mathcal{C}$ is mapped into two totally-defined trivariate functions f_c^+, f_c^-, such that a hyperplane $x_4 = h_1x_1 + h_2x_2 + h_3x_3 + h_4$ intersects c if and only if $f_c^-(h_1, h_2, h_3) \le h_4 \le f_c^+(h_1, h_2, h_3)$. See [2] for more details. Hence, $T_3(\mathcal{C})$ is the region sandwiched between the upper envelope of $\{f_c^- | c \in \mathcal{C}\}$ and the lower envelope of $\{f_c^+ | c \in \mathcal{C}\}$. Using Theorem 5.1, we thus obtain:

Theorem 5.3. *The combinatorial complexity of the space $T_3(\mathcal{C})$ of all hyperplane transversals of a set $\mathcal{C}$ of n convex sets of constant description complexity in $\mathbb{R}^4$ is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$.*

Remark. Note that each vertex of $T_3(\mathcal{C})$ is dual to a hyperplane transversal that is tangent to four members of $\mathcal{C}$. Similar geometric interpretations hold for other features of $\partial T_3(\mathcal{C})$.

The space of line transversals in 3-space. Let $\mathcal{C}$ be a collection of n convex sets in $\mathbb{R}^3$, each being semi-algebraic of constant description complexity. Let $T_1(\mathcal{C})$ denote the space of all line transversals of $\mathcal{C}$. We can map each line l in $\mathbb{R}^3$, given by the equations $y = a_1x + a_2$, $z = a_3x + a_4$, to the point $l^* = (a_1, a_2, a_3, a_4) \in \mathbb{R}^4$. (This excludes lines parallel to the yz-plane, which can be handled separately, in a much simpler manner.) As above (see [2] for details), each $c \in \mathcal{C}$ is mapped to a pair of partially-defined trivariate functions f_c^+, f_c^-, such that f_c^+ and f_c^- have the same domain of definition, and a line l, with $l^* = (a_1, a_2, a_3, a_4)$, is a transversal of c if and only if the functions f_c^+, f_c^- are defined at (a_1, a_2, a_3) and $f_c^-(a_1, a_2, a_3) \le a_4 \le f_c^+(a_1, a_2, a_3)$. Hence, this problem too reduces to a sandwich region in four dimensions, and Theorem 5.1 implies:

Theorem 5.4. *The combinatorial complexity of the space $T_1(\mathcal{C})$ of all line transversals of a set $\mathcal{C}$ of n convex sets of constant description complexity in $\mathbb{R}^3$ is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$.*

Remark. As above, vertices of $T_1(\mathcal{C})$ are dual to lines that are tangent to four members of $\mathcal{C}$.

This implies the following theorem concerning the number of *geometric permutations*. Such a permutation is the order in which a collection of disjoint convex bodies can be stabbed by a line transversal.

Corollary 5.5. *The number of geometric permutations in a collection $\mathcal{C}$ of n pairwise disjoint convex sets of constant description complexity in $\mathbb{R}^3$ is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$.*

This improves the general known upper bound of $O(n^4)$ [21] (for the special case of sets with constant description complexity), but is not known to be tight, since the only known lower bound is $\Omega(n^2)$ [15].

Efficient construction of transversal spaces. Theorems 5.3 and 5.4 do not address the problem of efficient construction of the respective spaces $T_3(\mathcal{C})$, $T_1(\mathcal{C})$. We next show that the boundaries of these spaces can be constructed efficiently, in time $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$. To avoid (the routine, though somewhat technical) details involving the representation of the combinatorial structure of the boundary as a 3-dimensional cell complex, we restrict the algorithm to produce, for every pair $c_1, c_2 \in \mathcal{C}$, just the portion of $\partial T_3(\mathcal{C})$ (or $\partial T_1(\mathcal{C})$) that consists of all the points representing hyperplanes (or lines) that are tangent to c_1, c_2 and intersect all the other sets in $\mathcal{C}$. The representation of such a portion, which consists of some faces of a 2-dimensional arrangement (see below), is easy to define and compute. Our construction technique follows the approach used in [1, 5] and is easy to adopt to constructing a complete representation of $\partial T_3(\mathcal{C})$ (or $\partial T_1(\mathcal{C})$).

In more detail, let us consider the case of $T_3(\mathcal{C})$. Fix a pair of sets $c_1, c_2 \in \mathcal{C}$, and consider any pair of functions from $\{f^+_{c_1}, f^-_{c_1}\} \times \{f^+_{c_2}, f^-_{c_2}\}$, say $f^+_{c_1}, f^+_{c_2}$. The intersection of their graphs is a 2-dimensional surface φ. For each $c \in \mathcal{C}^* \equiv \mathcal{C} \setminus \{c_1, c_2\}$, let

$$K_c = \{h \in \varphi \mid f^-_c(h_1, h_2, h_3) \le h_4 \le f^+_c(h_1, h_2, h_3)\}.$$

We need to construct $\bigcap_{c \in \mathcal{C}^*} K_c$. We do it using the randomized incremental technique of [1, 5]. That is, we insert the sets K_c, for $c \in \mathcal{C}^*$, in some random order, and update the intersection after the insertion of each new set. We omit further details, which can be easily adapted from the algorithms just cited. The analysis given in [5], combined with Theorem 5.3, can easily be adjusted to the case at hand, implying that the overall expected running time of this algorithm, when applied to all pairs $c_1, c_2 \in \mathcal{C}$, is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$. Applying a fully analogous procedure for the case of line transversals in 3-space, we obtain:

Theorem 5.6. *(a) The boundary of the space of hyperplane transversals $T_3(\mathcal{C})$ in four dimensions, as defined above, can be computed in $O(n^{3+\varepsilon})$ randomized expected time, for any $\varepsilon > 0$.*
(b) The boundary of the space of line transversals $T_1(\mathcal{C})$ in three dimensions, as defined above, can be computed in $O(n^{3+\varepsilon})$ randomized expected time, for any $\varepsilon > 0$.

Union of objects in 4-space. Let $\mathcal{C}$ be a collection of n convex sets in $\mathbb{R}^4$, each being semi-algebraic of constant description complexity, such that (i) The mean curvature [18] of any $c \in \mathcal{C}$ is at most some constant κ, and (ii) For any pair of sets $c_1, c_2 \in \mathcal{C}$, the ratio between their diameters is at most some constant α. (We refer to such sets as being of 'nearly equal size'.)

Let $\mathcal{U}$ denote the union of $\mathcal{C}$. The combinatorial complexity of $\mathcal{U}$ is the number of faces of all dimensions of the arrangement of the boundaries ∂c of the sets $c \in \mathcal{C}$, which appear on $\partial \mathcal{U}$.

Theorem 5.7. *The combinatorial complexity of the union $\mathcal{U}$ of n convex sets of constant description complexity in $\mathbb{R}^4$ that satisfy properties (i) and (ii) is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$.*

Proof. We may assume that the diameter of any $c \in \mathcal{C}$ is between 1 and α. This, plus the bounded mean curvature assumption, implies the following two properties. Let G be an infinite axis-parallel grid in $\mathbb{R}^4$, where each cell of G is a hypercube of side length b, for some sufficiently small constant $b < 1$. Then (a) Each $c \in \mathcal{C}$ intersects only $O(1)$ cells of G; (b) Let c be a set in $\mathcal{C}$ such that ∂c intersects a cell τ of G. Let $\Delta(c, \tau)$ denote the set of all directions d, such that $\partial c \cap \tau$ is monotone orthogonally to d. That is, $\partial c \cap \tau$ can be regarded as the graph of a (partially-defined) function, where the dependent variable is in direction d. (Clearly, $\Delta(c, \tau)$ is centrally symmetric: $d \in \Delta(c, \tau) \Leftrightarrow -d \in \Delta(c, \tau)$.) Then the measure of $\Delta(c, \tau)$ is at least 7/8 of the measure of the entire sphere of directions (provided b is sufficiently small).

This is easy to establish and a similar analysis in three dimensions is provided in [4]. To verify property (b) note that the bounded mean curvature assumption implies that for any pair of points $x, y \in \partial c$ on a surface $c \in \mathcal{C}$

$$d_S(N_c(x), N_c(y)) \leq \kappa' \|x - y\|,$$

where κ' is a constant dependant on κ, $N_c(x)$ denotes the direction of the outward normal to ∂c at x and d_S is the geodesic distance along the unit sphere of directions $\mathbb{S}^3$. Choose b sufficiently small, so that $2\kappa' b < \delta$, where the value of δ will be determined shortly. Fix some $x_0 \in \partial c \cap \tau$. Let d be any direction forming an angle θ with $N_c(x_0)$. Suppose that there exists a line λ parallel to d that intersects $\partial c \cap \tau$ at two points u, v. Then $N_c(u) \cdot d$ and $N_c(v) \cdot d$ have different signs, so that, say, $N_c(u) \cdot d < 0$. Assuming θ to be smaller than $\pi/2 - \delta$, it follows that $d_S(N_c(x_0), N_c(u)) \geq \pi/2 - \theta > \delta$. On the other hand, the bounded mean curvature assumption implies that $d_S(N_c(x_0), N_c(u)) \leq \kappa' \|u - x_0\| \leq 2\kappa' b \leq \delta$, a contradiction. This implies that the set of directions d that satisfy the property in (b) contains the two spherical caps centered at $\pm N_c(x_0)$ and having geodesic radius $\pi/2 - \delta$. Choosing δ (and b) sufficiently small, the validity of property (b) follows.

Property (b) implies that there exists a set Δ of $O(1)$ directions, such that, for any quadruple $c_1, c_2, c_3, c_4 \in \mathcal{C}$ whose boundaries all cross a cell τ, we have $\Delta \cap \bigcap_{i=1}^4 \Delta(c_i, \tau) \neq \emptyset$.

Now, fix a cell τ. If τ is fully contained in some set $c \in \mathcal{C}$ then we ignore τ — it contributes nothing to $\partial \mathcal{U}$. Otherwise, we consider the set $\mathcal{C}_\tau = \{c \in \mathcal{C} \mid \partial c \cap \tau \neq \emptyset\}$, and further partition it into the subsets $\mathcal{C}^+_{\tau,d}$, $\mathcal{C}^-_{\tau,d}$, for $d \in \Delta$, where $\mathcal{C}^+_{\tau,d}$ (resp., $\mathcal{C}^-_{\tau,d}$) consists of all sets $c \in \mathcal{C}_\tau$ for which $d \in \Delta(c, \tau)$, and such that, if we move slightly from any point on $\partial c \cap \tau$ in the direction $+d$ (resp., $-d$), we enter c. (A set c may belong to more than one of these subcollections.)

The preceding discussion implies the following property: (b') Let $v \in \tau$ be a vertex of $\mathcal{U}$, incident to the boundaries of four sets $c_1, c_2, c_3, c_4 \in \mathcal{C}$. Then there exists a direction $d \in \Delta$ such that $c_1, c_2, c_3, c_4 \in \mathcal{C}^+_{\tau,d} \cup \mathcal{C}^-_{\tau,d}$.

Property (b') implies that the number of vertices of $\partial \mathcal{U} \cap \tau$ can be upper bounded by the sum, over $d \in \Delta$, of the number of vertices of the sandwich region between the upper envelope of the boundaries of the sets in $\mathcal{C}^+_{\tau,d}$ and the lower envelope of the boundaries of the sets in $\mathcal{C}^-_{\tau,d}$, where both boundaries are clipped to within τ. Using Theorem 5.1, plus the facts that $|\Delta| = O(1)$ and that $\sum_\tau |\mathcal{C}_\tau| = O(n)$ (which follows from property (a)), the bound on the complexity of $\mathcal{U}$ follows. $\square$

Arrangements with no vertices. The analysis in Section 4.4, which uses quite a non-standard counting scheme, is of interest on its own, and can be adapted to other settings. In particular, it can be easily adjusted to show the following.

Theorem 5.8. *The complexity of the lower envelope of an arrangement of* n *totally defined semi-algebraic surfaces of constant description complexity in* $\mathbb{R}^3$*, that does not contain any vertices, is* $O(n^{1+\varepsilon})$*, for any* $\varepsilon > 0$.

Note that if the surfaces are not totally defined, the complexity of the lower envelope may still be quadratic. An easy lower bound construction is provided by a family of $n/2$ nearly x-parallel lines and another family of $n/2$ nearly y-parallel lines, that together make up a grid-structure when viewed from below.

6 Conclusion

We have obtained several results concerning overlays of minimization diagrams using a novel approach. We feel that this approach might find applications for related problems, like the analysis of vertical decompositions of arrangements [16]. Although the partition technique seems quite general, our initial steps in applying it in more general contexts have encountered some technical difficulties, which we hope to be able to resolve in the future. We also hope to be able to apply the technique to settle the conjecture concerning the complexity of overlays of minimization diagrams in all dimensions.

In general, it would be interesting to analyze the partition technique from a more 'philosophical' point of view, and to understand in particular the underlying reason for why it was so successful in the analysis of overlays, where the technique of counting schemes has provided only partial results, but does not seem to be easily applicable to the related problem of the analysis of single cells, where a near-optimal solution using counting schemes exists [6]. The two techniques seem to be related, at least in the fact that the final recurrences that are derived by both techniques have very similar structure. This is apparent for instance in the case of overlays for bivariate functions, in which both techniques provide the same near-optimal solution. The difference, in this case at least, is that the way to obtain these recurrences via the partition technique is arguably much simpler.

Acknowledgements

Part of the work on this paper has been carried out while the authors were visiting ETH Zürich, hosted by the European Graduate Program on Combinatorics, Geometry and Computation. The authors are grateful to the program, and in particular to Emo Welzl, for the hospitality. Emo Welzl has also contributed valuable ideas through helpful discussions.

References

[1] P. K. Agarwal, B. Aronov, and M. Sharir. Computing envelopes in four dimensions with applications. *SIAM Journal on Computing*, 26:1714–1732, 1997.

[2] P. K. Agarwal, O. Schwarzkopf, and M. Sharir. The overlay of lower envelopes and its applications. *Discrete and Computational Geometry*, 15:1–13, 1996.

[3] P. K. Agarwal and M. Sharir. Arrangements and their applications. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 49–119. Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000.

[4] P. K. Agarwal and M. Sharir. Pipes, cigars, and kreplach: The union of Minkowski sums in three dimensions. *Discrete and Computational Geometry*, 24:645–685, 2000.

[5] B. Aronov, M. Sharir, and B. Tagansky. The union of convex polyhedra in three dimensions. *SIAM Journal on Computing*, 26:1670–1688, 1997.

[6] S. Basu. On the combinatorial and topological complexity of a single cell. In *Proc. 39th IEEE Symposium on Foundations of Computer Science*, pages 606–616, 1998.

[7] B. Chazelle, H. Edelsbrunner, L. J. Guibas, and M. Sharir. A singly-exponential stratification scheme for real semi-algebraic varieties and its applications. *Theoretical Computer Science*, 84:77–105, 1991. Also in *Proc. International Colloquium on Automata, Languages and Programming*, pages 179–193, 1989.

[8] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete and Computational Geometry*, 4:387–421, 1989.

[9] M. de Berg and O. Schwarzkopf. Cuttings and applications. *International Journal of Computational Geometry and Applications*, 5:343–355, 1995.

[10] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*, volume 10 of *EATCS Monographs on Theoretical Computer Science*. Springer-Verlag, Heidelberg, West Germany, 1987.

[11] H. Edelsbrunner, L. J. Guibas, and M. Sharir. The upper envelope of piecewise linear functions: algorithms and applications. *Discrete and Computational Geometry*, 4:311–336, 1989.

[12] D. Halperin and M. Sharir. New bounds for lower envelopes in three dimensions, with applications to visibility in terrains. *Discrete and Computational Geometry*, 12:313–326, 1994.

[13] D. Halperin and M. Sharir. Almost tight upper bounds for the single cell and zone problems in three dimensions. *Discrete and Computational Geometry*, 14:385–410, 1995.

[14] S. Har-Peled. The complexity of many cells in the overlay of many arrangements. M.Sc. dissertation, School of Computer Science, Tel Aviv University, 1995. See also: Multicolor combination lemmas, *Computational Geometry Theory and Applications*, 12 (1999), 155–176.

[15] M. Katchalski, T. Lewis, and A. Liu. The different ways of stabbing disjoint convex sets. *Discrete and Computational Geometry*, 7:197–206, 1992.

[16] V. Koltun. Almost tight upper bounds for vertical decompositions in four dimensions. In *Proc. 42nd IEEE Symposium on Foundations of Computer Science*, pages 56–65, 2001.

[17] V. Koltun and M. Sharir. On the overlay of envelopes in four dimensions. In *Proc. 13th ACM-SIAM Symposium on Discrete Algorithms*, pages 810–819, 2002.

[18] A. Pressley. *Elementary Differential Geometry*. Springer-Verlag, London, 2001.

[19] M. Sharir. Almost tight upper bounds for lower envelopes in higher dimensions. *Discrete and Computational Geometry*, 12:327–345, 1994.

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

[21] R. Wenger. Upper bounds on geometric permutations for convex sets. *Discrete and Computational Geometry*, 5:27–33, 1990.