On the Overlay of Envelopes in Four Dimensions*

Vladlen Koltun† Micha Sharir‡

Abstract

We show that the complexity of the overlay of two envelopes of arrangements of n semi-algebraic surfaces or surface patches of constant description complexity in four dimensions is $O(n^{4-1/\lceil s/2\rceil+\varepsilon})$, for any $\varepsilon > 0$, where s is a constant related to the maximal degree of the surfaces. This is the first non-trivial (sub-quartic) bound for this problem, and for $s = 1, 2$ it almost matches the near-cubic lower bound. We discuss several applications of this result, including (i) an improved bound for the complexity of the region enclosed between two envelopes in four dimensions, (ii) an improved bound for the complexity of the space of all hyperplane transversals of a collection of simply-shaped convex sets in 4-space, (iii) an improved bound for the complexity of the space of all line transversals of a similar collection of sets in 3-space, and (iv) improved bounds for the complexity of the union of certain families of objects in four dimensions. The analysis technique we introduce is quite general, and has already proved useful in unrelated contexts.

1 Introduction

Let $\mathcal{R}$ be a family of n 'red' $(d-1)$-variate functions in $\mathbb{R}^d$, and $\mathcal{B}$ be a family of n 'blue' such functions. We assume each (partially or totally defined) function to be semi-algebraic of constant description complexity, that is, to be defined in terms of a constant number of polynomial equalities and inequalities of constant maximum degree. Let $\mathcal{E}_\mathcal{R}$ and $\mathcal{E}_\mathcal{B}$ denote the lower envelopes of $\mathcal{R}$ and of $\mathcal{B}$, respectively (upper envelopes are handled analogously and we will not treat them explicitly). Let $\mathcal{M}_\mathcal{R}$ and $\mathcal{M}_\mathcal{B}$ denote the *minimization diagrams* of $\mathcal{R}$ and of $\mathcal{B}$, respectively. These are the respective projections onto $\mathbb{R}^{d-1}$ of $\mathcal{E}_\mathcal{R}$ and $\mathcal{E}_\mathcal{B}$, which are subdivisions of $\mathbb{R}^{d-1}$. As is well known, the complexity of each of $\mathcal{E}_\mathcal{R}$, $\mathcal{E}_\mathcal{B}$, and thus also of $\mathcal{M}_\mathcal{R}$, $\mathcal{M}_\mathcal{B}$, is $O(n^{d-1+\varepsilon})$ [13].[1] The *overlay* $\mathcal{Q}_{\mathcal{R},\mathcal{B}}$ of $\mathcal{E}_\mathcal{R}$ and $\mathcal{E}_\mathcal{B}$ is defined as the superposition of $\mathcal{M}_\mathcal{R}$ and $\mathcal{M}_\mathcal{B}$. It is the subdivision of $\mathbb{R}^{d-1}$ into maximal connected relatively open faces, of any dimension, so that within each face the same subset of functions appear both on $\mathcal{E}_\mathcal{R}$ and on $\mathcal{E}_\mathcal{B}$. The *combinatorial complexity* (in short, complexity) of $\mathcal{Q}_{\mathcal{R},\mathcal{B}}$ is defined to be the number of faces of all dimensions in $\mathcal{Q}_{\mathcal{R},\mathcal{B}}$.

If $d = 2$, the complexity of the overlay is at most proportional to the sum of the complexities of the two respective envelopes, and is thus close to linear [13]. However, the most important algorithmic and combinatorial applications for overlays of envelopes arise in higher dimensions. This motivates the study of the complexity of overlays in higher dimensions.

A tight bound of $\Theta(n^{d-1}\alpha(n))$ (where $\alpha(n)$ is the extremely slowly growing functional inverse of Ackermann's function) is known for the complexity of the overlay $\mathcal{Q}_{\mathcal{R},\mathcal{B}}$ if $\mathcal{R}$ and $\mathcal{B}$ are collections of simplices [7]. In three dimensions, Agarwal, Schwarzkopf and Sharir [2] have proved that the complexity of the overlay $\mathcal{Q}_{\mathcal{R},\mathcal{B}}$ is $O(n^{2+\varepsilon})$, for the semi-algebraic case assumed above.

In four dimensions, it is easy to see that each feature of the overlay is defined by at most five surfaces (see below). A bound of $O(n^5)$ on the complexity of the overlay is thus trivial. It is also not hard to obtain a bound of $O(n^4\beta(n))$, where $\beta(n)$ is a slow-growing function related to Davenport-Schinzel sequences [13], using known results on the complexity of lower envelopes in two dimensions. The problem of improving this bound has been recognized as an important open problem [3], since such an improvement would have several algorithmic and combinatorial applications (such as those mentioned in the abstract).

In this paper, we obtain a bound of $O(n^{4-1/\lceil s/2\rceil+\varepsilon})$ for the complexity of such an overlay, where s is the maximal number of intersections between the (1-dimensional) projection of an intersection of three red

*Work on this paper has been supported by a grant from the Israel Science Fund (for a Center of Excellence in Geometric Computing). Work by Micha Sharir was also supported by NSF Grant CCR-97-32101, by a grant from the U.S.-Israel Binational Science Foundation and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.

†School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel. *vladlen@tau.ac.il*

‡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@math.tau.ac.il*

[1] Throughout this paper, a complexity bound of the form $f(n) = O(n^{q+\varepsilon})$ means that, for any arbitrarily small positive ε, there exists a constant c_ε, such that $f(n) \le c_\varepsilon n^{q+\varepsilon}$.

(resp., blue) surfaces onto 3-space, with a similarly-defined projection of the (2-dimensional) intersection of two blue (resp., red) surfaces. Our bound is near-cubic when $s = 1$ (as, for instance, in the case of linear surfaces) or when $s = 2$. This still leaves open the problem of obtaining a near-cubic bound for arbitrary constant values of s, or improving the near-cubic lower bound, which is given by the complexity of one lower envelope in 4-space. We tend to conjecture, though, that the correct bound is near-cubic.

Our result has several immediate applications. First, it can be used to obtain an improved bound on the complexity of the 'sandwich' region enclosed between two envelopes in four dimensions. It can also be applied to obtain an improved 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. Another application yields a sub-quartic bound on the complexity of the union of simply shaped (convex) objects in $\mathbb{R}^4$, if they have bounded curvature and nearly equal sizes.

The analysis technique we introduce seems to be quite general, and we have successfully applied it to obtain improved upper bounds for other geometric problems, as described in Section 6.

The rest of the paper is organized as follows. After introducing some preliminary definitions and observations in Section 2, we analyze the number of vertices and edges in the overlay in Sections 3 and 4, respectively, and obtain a system of inter-dependent recurrence relations. We solve this system in Section 5 and thus prove the main bound on the complexity of the overlay. The applications of this bound are described in Section 6, and we conclude in Section 7.

The proofs of all lemmas are omitted from this version of the paper due to space limitations.

2 Preliminaries

Observations and assumptions. The overlay $\mathcal{Q}_{\mathcal{R},\mathcal{B}}$ is a subdivision of $\mathbb{R}^3$ by algebraic surface patches. Any vertex of $\mathcal{Q}_{\mathcal{R},\mathcal{B}}$ is either a vertex of $\mathcal{M}_{\mathcal{R}}$, a vertex of $\mathcal{M}_{\mathcal{B}}$, an intersection of an edge of $\mathcal{M}_{\mathcal{R}}$ with a 2-face of $\mathcal{M}_{\mathcal{B}}$, or an intersection of an edge of $\mathcal{M}_{\mathcal{B}}$ with a 2-face of $\mathcal{M}_{\mathcal{R}}$. Note that an edge (2-face) of $\mathcal{M}_{\mathcal{R}}$ is the projection of a portion of an intersection curve (resp., intersection surface) of three (resp., of two) red surfaces, and similarly for $\mathcal{M}_{\mathcal{B}}$.

Before proceeding, we remark that it is sufficient to treat the case where $\mathcal{R}$ and $\mathcal{B}$ consist of totally-defined xyz-monotone functions, since all other cases can be reduced to it. Indeed, any semi-algebraic surface patch of constant description complexity can be decomposed into a constant number of xyz-monotone pieces. Each such piece can then be extended to be totally-defined, by means of near-vertical semi-infinite walls attached to its boundary. Formally, we replace each piece by 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. We will thus only deal with totally-defined xyz-monotone functions in the sequel.

We also assume that the surfaces in $\mathcal{R}$ and $\mathcal{B}$ are in *general position*. This excludes degenerate configurations, such as five surfaces meeting at a point. As argued in [8, 12], this assumption incurs no real loss of generality.

Definitions. Let v be a vertex of $\mathcal{Q}_{\mathcal{R},\mathcal{B}}$ that is formed by the intersection of an edge e of $\mathcal{M}_{\mathcal{R}}$ with a 2-face f of $\mathcal{M}_{\mathcal{B}}$. Let r_1, r_2, r_3 be the three red surfaces that attain $\mathcal{E}_{\mathcal{R}}$ over e, and let b_1, b_2 be the two blue surfaces that attain $\mathcal{E}_{\mathcal{B}}$ over f. In 4-space, v corresponds to a *vertical coincidence event*, in which $r_1 \cap r_2 \cap r_3$ and $b_1 \cap b_2$ are touched by a common vertical line λ, so that no red surface intersects λ below $r_1 \cap r_2 \cap r_3 \cap \lambda$ and no blue surface intersects λ below $b_1 \cap b_2 \cap \lambda$. By algebraicity and the general position assumption, a 5-tuple r_1, r_2, r_3, b_1, b_2 has at most $s = O(1)$ such events. We can extend the definition of vertical coincidence events as follows.

Definition 1. *Let r_1, r_2, r_3 be three red surfaces and let b_1, b_2, b_3 be three blue surfaces. The 5-tuple r_1, r_2, r_3, b_1, b_2 is said to define a $(3,2)$-event of level (i,j), if there exists a vertical line λ that intersects both $r_1 \cap r_2 \cap r_3$ and $b_1 \cap b_2$, such that i red surfaces pass below $r_1 \cap r_2 \cap r_3 \cap \lambda$ and j blue surfaces pass below $b_1 \cap b_2 \cap \lambda$. Symmetrically, r_1, r_2, b_1, b_2, b_3 are said to define a $(2,3)$-event of level (i,j), if there exists a vertical line λ that intersects both $r_1 \cap r_2$ and $b_1 \cap b_2 \cap b_3$, such that i red surfaces pass below $r_1 \cap r_2 \cap \lambda$ and j blue surfaces pass below $b_1 \cap b_2 \cap b_3 \cap \lambda$. Notice that $(3,2)$-events (resp., $(2,3)$-events) of level $(0,0)$ correspond to vertices of the overlay $\mathcal{Q}_{\mathcal{R},\mathcal{B}}$, where a red (resp., blue) edge crosses a blue (resp., red) 2-face. Also, any of the above 5-tuples define at most s $(3,2)$- or $(2,3)$-events, of any level.*

We say that an event is of level $(\leq i, \leq j)$ if it is an (a,b)-level event for some $a \leq i$, $b \leq j$. Denote the number of $(\leq i, \leq j)$-level $(3,2)$-events by $V^{32}_{i,j}(\mathcal{R},\mathcal{B})$, and define $V^{32}_{i,j}(n) \equiv \max V^{32}_{i,j}(\mathcal{R},\mathcal{B})$, where the maximum is taken over all red sets $\mathcal{R}$ and blue sets $\mathcal{B}$ of size n, as above. Define $V^{23}_{i,j}(n)$ symmetrically with respect to $(2,3)$-events, and let $V_{i,j}(n) \equiv V^{32}_{i,j}(n) + V^{23}_{i,j}(n)$.

Similar definitions pertain to 1-dimensional features of $\mathcal{Q}_{\mathcal{R},\mathcal{B}}$. Consider an intersection of a projection onto $\mathbb{R}^3$ of a 2-face of the blue arrangement with a similar projection of a 2-face of the red arrangement. This intersection is 1-dimensional, and may consist of up to a constant number of connected components (edges). (Each such edge is an (i, j)-level edge of the overlay, where i and j are the levels of the red and blue 2-faces that define it, in their respective arrangements.)

We can also consider intersections of projections of blue and red 2-intersection surfaces (each of which can contain numerous 2-faces). Each such intersection has at most a constant number of connected components, possibly containing many overlay edges at various levels. We refer to such connected components as $(2, 2)$-*events*. We say that a $(2, 2)$-event is at level $(\le i, \le j)$ if it contains a point that is at level $(\le i, \le j)$, where the level of a point on a $(2, 2)$-event is defined analogously to a level of $(3, 2)$-events (which, in fact, correspond to specific points on $(2, 2)$-events). We denote the number of $(\le i, \le j)$-level $(2, 2)$-events by $E_{i,j}(\mathcal{R}, \mathcal{B})$, and define $E_{i,j}(n) \equiv \max E_{i,j}(\mathcal{R}, \mathcal{B})$, as above.

THEOREM 2.1. *The combinatorial complexity of $\mathcal{Q}_{\mathcal{R},\mathcal{B}}$ is*

$$O\left(n^{4-\frac{1}{\lceil s/2\rceil}+\varepsilon}\right).$$

The proof of the theorem proceeds as follows. The next Section is devoted to recursively bounding the number of vertices of $\mathcal{Q}_{\mathcal{R},\mathcal{B}}$, in terms of higher-level vertices and edges. In the subsequent Section we derive a corresponding system of recurrence relations for the number of edges of $\mathcal{Q}_{\mathcal{R},\mathcal{B}}$. (Such a separate analysis is needed because $\mathcal{Q}_{\mathcal{R},\mathcal{B}}$ may contain edges that have no vertex incident to them.) Combining all these recurrences together in Section 5 yields the desired bound on the number of vertices and edges of the overlay, which is then easily extended to account for 2- and 3-dimensional faces as well.

3 Overlay Vertices

We will use a counting scheme for deriving a recurrence relation on the number of vertices of $\mathcal{Q}_{\mathcal{R},\mathcal{B}}$. Counting schemes, as a technique for analyzing the complexity of substructures in arrangements of surfaces, were introduced by Halperin and Sharir [8, 12] (see also [13]), and were further developed and extended in later works [1, 2, 5, 9, 11]. The general structure of our counting scheme is similar to that of [2], although the technical details are considerably more involved. It proceeds in two stages.

In the first stage, we estimate $V^{32}_{0,0}(n)$ and $V^{23}_{0,0}(n)$ in terms of $V^{32}_{0,k}(n)$ and $V^{23}_{k,0}(n)$, respectively, for some threshold parameter k, that we fix later on. In the second stage, we estimate $V^{32}_{0,k}(n)$ and $V^{23}_{k,0}(n)$ in terms of $V_{2k,2k}(n)$ and $E_{2k,2k}(n)$. Combining the results of the two stages, we get a bound on $V_{0,0}(n)$ in terms of $V_{2k,2k}(n)$ and $E_{2k,2k}(n)$. The Clarkson-Shor random sampling technique [6] is then applied to get a bound on $V_{2k,2k}(n)$ and $E_{2k,2k}(n)$ in terms of $V_{0,0}(n/k)$ and $E_{0,0}(n/k)$, respectively. Putting everything together, we get a recurrence relation on $V_{0,0}(n)$ that depends on both $V_{0,0}(n/k)$ and $E_{0,0}(n/k)$.

3.1 First Stage

The first stage of the counting scheme proceeds as follows. We fix an edge e of $\mathcal{E}_{\mathcal{R}}$, contained in the intersection of three red surfaces, take the vertical 2-dimensional surface ('curtain') Π_e, consisting of all x_4-parallel lines that pass through e, and consider the cross-section of $\mathcal{A}(\mathcal{B})$ within Π_e, which we denote by $\mathcal{A}(\mathcal{B}_e)$. The vertices of the lower envelope of $\mathcal{A}(\mathcal{B}_e)$ are the $(3, 2)$-events at level $(0, 0)$ that involve e, and the vertices of this arrangement at level at most k are the $(3, 2)$-events that involve e and are at level $(0, \le k)$.

Let t_e be the number of distinct blue surfaces that attain the lower envelope of $\mathcal{A}(\mathcal{B}_e)$. Assume first that $t_e \le k$. In this case, the number of vertices in the lower envelope, and thus the number of $(0, 0)$-level $(3, 2)$-events that involve e, is at most $\lambda_s(k)$, which we write as $O(k\beta(k))$, for an appropriate very slowly growing function $\beta(\cdot)$ [13]. Since there are at most $O(n^{3+\varepsilon})$ such edges e, there are overall $O(kn^{3+\varepsilon}\beta(k))$ $(0, 0)$-level $(3, 2)$-events that involve such edges.

Suppose now that $t_e > k$. In this case, a standard argument, as the one used in [2], implies that there are at least $\Omega(t_e k)$ vertices in levels 0 through k of $\mathcal{A}(\mathcal{B}_e)$. Therefore, there are at least $\Omega(t_e k)$ $(0, \le k)$-level $(3, 2)$-events that involve e, for any edge e as above. Combining this with the fact that there are at most $O(t_e\beta(t_e))$ $(0, 0)$-level $(3, 2)$-events that involve e, and summing over the edges e of $\mathcal{E}_{\mathcal{R}}$, we obtain

$$V^{32}_{0,0}(n) = O\left(\frac{\beta(n)}{k}V^{32}_{0,k}(n) + kn^{3+\varepsilon}\beta(k)\right). \quad (3.1)$$

We can repeat the above process for all edges e of $\mathcal{E}_{\mathcal{B}}$ as well, while considering a cross-section of the *red* arrangement with Π_e. In full symmetry with the above, this gives

$$V^{23}_{0,0}(n) = O\left(\frac{\beta(n)}{k}V^{23}_{k,0}(n) + kn^{3+\varepsilon}\beta(k)\right). \quad (3.2)$$

3.2 Second Stage

3.2.1 Preliminaries

In the second stage we wish to bound, say, $V_{0,k}^{32}(n)$ in terms of $V_{k,2k}(n)$ and $E_{k,2k}(n)$. To do this, we fix a pair of blue surfaces, consider the 3-dimensional arrangement of the red surfaces that lie 'over' their blue 2-intersection surface, and apply a charging scheme within this arrangement. Unlike the first stage, which dealt with 2-dimensional arrangements, it is much harder to carry out a charging scheme in three dimensions, and most of the technical difficulties with which the paper struggles stem from this fact.

In detail, we proceed as follows. We begin the second counting stage by bounding $V_{0,k}^{32}(n)$ in terms of $V_{k,2k}(n)$ and $E_{k,2k}(n)$, as follows. We use the same threshold parameter k as above. We fix a 2-face φ of $\mathcal{A}(\mathcal{B})$ at level at most k in this arrangement, take the vertical 3-dimensional curtain surface Π_φ, consisting, as above, of all x_4-parallel lines that pass through φ, and consider the cross-section $\mathcal{A}(\mathcal{R}_\varphi)$ of $\mathcal{A}(\mathcal{R})$ within Π_φ. The vertices of the lower envelope of $\mathcal{A}(\mathcal{R}_\varphi)$ are the $(3,2)$-events at level $(0,\leq k)$ that involve φ, and the vertices of this arrangement at level at most j are the $(3,2)$-events that involve φ and are at level $(\leq j, \leq k)$. Denote the 2-dimensional blue intersection surface that contains φ by χ. We will also consider the cross-sections $\mathcal{A}(\mathcal{R}_\chi)$ and $\mathcal{A}(\mathcal{B}_\chi)$, defined analogously within Π_χ. Note that the intersection surface χ is $x_1x_2x_3$-monotone, since all blue surfaces are. Similarly, $\mathcal{A}(\mathcal{R}_\chi)$ and $\mathcal{A}(\mathcal{B}_\chi)$ are arrangements of monotone 2-dimensional surfaces in $\mathbb{R}^3$.

Informally, we wish to show that for all vertices of the lower envelope of $\mathcal{A}(\mathcal{R}_\varphi)$ (that correspond to $(3,2)$-events at level $(0,\leq k)$ involving χ), we can either uniquely collect roughly $\zeta = k^{1/\lceil s/2\rceil}$ events at level $(\leq k, \leq 2k)$ involving χ, or uniquely charge $(2,2)$-events. We will choose k sufficiently large to ensure that $\zeta \geq 2$.

We partition the first k levels of $\mathcal{A}(\mathcal{R}_\chi)$ into $\lceil s/2\rceil + 1$ *layers*. Layer 0 is composed of the lower envelope. For $i = 1, \ldots, \lceil s/2\rceil$, layer i comprises levels ζ^{i-1} through $\zeta^i - 1$. Similarly, we partition levels $k+1$ through $2k-1$ of $\mathcal{A}(\mathcal{B}_\chi)$ into $\lceil s/2\rceil$ 'blue' layers, such that, for $i = 1, \ldots, \lceil s/2\rceil$, layer i comprises levels $k+\zeta^{i-1}$ through $k+\zeta^i-1$. 'Blue' layer 0 is composed of levels $(\leq k)$ of $\mathcal{A}(\mathcal{B}_\chi)$.

Intuitively, the reason for introducing layers is that we want to ensure that the collection process (described in detail below) will only collect/charge events at level $(\leq k, \leq 2k)$. This will be guaranteed by designing the process in such way that only events at red and blue layers at most $\lceil s/2\rceil$ are collected/charged.

3.2.2 The collection process

Let v be a vertex of $\mathcal{A}(\mathcal{R}_\chi)$ that corresponds to a $(3,2)$-event that belongs to red layer i and blue layer j. Let the three surfaces of $\mathcal{R}$ incident to v be r_1, r_2, r_3. Define the *index* of v to be the number of intersection points in $r_1 \cap r_2 \cap r_3$ in $\mathcal{A}(\mathcal{R}_\chi)$ that have a smaller x_1-coordinate than that of v. Notice that the index of v ranges between 0 and $s-1$.

We collect vertices for v as follows. Let $\zeta' \equiv \zeta^{i+1}/2$. Let the index of v be l. If $l \leq \lceil s/2\rceil - 1$ (resp., $l \geq \lceil s/2\rceil$), we trace an edge e of $\mathcal{A}(\mathcal{R}_\chi)$ that is incident to v, emanates from v in the negative (resp., the positive) x_1-direction, and lies above one of the surfaces incident to v (it is well known and easy to check that such a tracing direction always exists; see, e.g., [8, 12]). Below we only consider the case $l \leq \lceil s/2\rceil - 1$, as the case $l \geq \lceil s/2\rceil$ is handled in complete analogy.

Collecting ζ' vertices. If we pass ζ' vertices of $\mathcal{A}(\mathcal{R}_\chi)$ that lie in red layer $i+1$ (which correspond to $(3,2)$-events at red layer $i+1$), before encountering another point v_2 of intersection of the surfaces r_1, r_2, r_3 that lies along e, we collect these ζ' vertices. Similarly, if we pass above/below ζ' blue curves of the form $\chi \cap \psi$, for some blue surface ψ, while being in blue layer *at least* $i+1$, before encountering another point v_2 as above, we collect these ζ' events (which correspond to $(2,3)$-events at blue layer at least $i+1$). If we pass ζ' $(3,2)$-events *and* ζ' $(2,3)$-events as above, then we collect either the $(3,2)$-events or the $(2,3)$-events, depending on which "block" of ζ' events ends first.

It is easy to see that if we collect $(3,2)$-events, they necessarily lie at red layer $i+1$ and blue layer at most $\max(i,j)+1$. Analogously, if we collect $(2,3)$-events, they lie at red layer at most $i+1$ and blue layer between $i+1$ and $\max(i,j)+1$. If we manage to collect ζ' events as above, the vertex v is said to be a *heavy* vertex.

Passing the baton. If we encounter a vertex v_2 as above, we start the same collection process from v_2 (which lies in red layer i or $i+1$, and blue layer at most $\max(i,j)+1$), and the vertices collected by v_2 are said to be collected by both v and v_2. In this case, v is said to be a *light* vertex, and we say that it "passes the baton" to v_2.

Charging a $(2,2)$-event. The above rules apply only if we do not reach an x_1-extremum of e before collecting ζ' events or encountering v_2 as above. If we do reach such an extremum (which may be at infinity), we charge v to the $(2,2)$-event corresponding to the curve e, and terminate the collection process. v is then said to be an *extremal* vertex.

It is easy to see that if we charge a $(2,2)$-event, it necessarily lies at red layer at most $i+1$ and blue layer at most $\max(i,j)+1$.

Intuitively, one of the reasons for charging $(2,2)$-events at x_1-extrema is that we want to ensure that, if we pass the baton to a vertex v_2 as above, its index is lower than that of v. If we encounter an x_1-extremum while sliding, the direction of the sliding may reverse, and we will not be able to ensure this property. Our solution to this problem is to charge a $(2,2)$-event at the moment we reach such an x_1-extremum.

The reason for also charging $(2,2)$-events at infinity is that it is possible that there are 'very few' vertices that lie on the part of e that we slide on. In this case we might not be able to collect ζ' vertices or find a vertex v_2 as above before reaching infinity.

3.2.3 Analysis

It is easy to see that one of the three scenarios listed above always holds. Indeed, suppose we cannot collect ζ' vertices as in the first scenario, or pass the baton to a vertex v_2 as in the second. This can happen only if we slide on e *ad infinitum*, while always remaining at red layer at most $i+1$ and blue layer at most $\max(i,j)+1$. In this case, however, we reach an x_1-extremum at infinity, which lies at red layer at most $i+1$ and blue layer at most $\max(i,j)+1$, and we can charge the corresponding $(2,2)$-event as in the third scenario.

Lemma 3.1. *A $(3,2)$-event at red layer j is collected by at most $6s\zeta^{j-1}$ vertices. Furthermore, for any $j'<j$, a $(2,3)$-event at blue layer j is collected by at most $2s\zeta^{j'}$ vertices belonging to red layer at most j'.*

Lemma 3.2. *A vertex v of the lower envelope of $\mathcal{A}(\mathcal{R}_\varphi)$ is either heavy, or extremal, or its baton is passed to a heavy or extremal vertex during the sliding (possibly through a number of other light vertices in between), which corresponds to a $(3,2)$-event that belongs to layer at most $\lceil s/2\rceil-1$ in both arrangements.*

Lemma 3.3. *A $(2,2)$-event is charged by at most $O(k/\zeta)$ vertices that correspond to $(3,2)$-events that belong to red and blue layers at most $\lceil s/2\rceil-1$.*

Lemma 3.2 ensures that every vertex of the lower envelope of $\mathcal{A}(\mathcal{R}_\varphi)$ is either heavy or extremal or passes the baton to a heavy or an extremal vertex u that corresponds to a $(3,2)$-event that lies at some layer $j\le\lceil s/2\rceil-1$ in the red arrangement and layer $j'\le\lceil s/2\rceil-1$ in the blue arrangement. If u is heavy, then, by definition, either (a) u collects at least $\zeta^{j+1}/2$ $(3,2)$-events belonging to red layer $j+1$ and blue layer at most $\lceil s/2\rceil$, or (b) u collects at least $\zeta^{j+1}/2$ $(2,3)$-events, each belonging to red layer at most $j+1$ and blue layer j'', such that $j+1\le j''\le\lceil s/2\rceil$. If u is extremal, then, by definition (c) u charges a $(2,2)$-event belonging to red and blue layers at most $\lceil s/2\rceil$.

In case (a), Lemma 3.1 ensures that a $(3,2)$-event at red layer $j+1$ is collected by $O(\zeta^j)$ vertices. In case (b), Lemma 3.1 ensures for all $j+1\le j''\le\lceil s/2\rceil$, that a $(2,3)$-event at blue layer j'' is collected by $O(\zeta^j)$ vertices that lie in red layer at most j. In case (c), Lemma 3.3 ensures that each $(2,2)$-event is charged by $O(k/\zeta)$ vertices as above. We thus have

$$V^{32}_{0,k}(n)=O\left(\frac{1}{\zeta}V_{k,2k}(n)+\frac{k}{\zeta}E_{k,2k}(n)\right). \tag{3.3}$$

We can repeat the above process while reversing the roles of the red and the blue arrangements. This gives us a similar relation for $(2,3)$-events:

$$V^{23}_{k,0}(n)=O\left(\frac{1}{\zeta}V_{2k,k}(n)+\frac{k}{\zeta}E_{2k,k}(n)\right). \tag{3.4}$$

3.3 Wrap-up

We can estimate $V_{2k,2k}$ (which is an upper bound for both $V_{k,2k}$ and $V_{2k,k}$) in terms of $V_{0,0}$. Since every vertical coincidence event is defined by five surfaces, the standard random sampling argument of Clarkson and Shor [6] implies

$$V_{2k,2k}(n)=O\left(k^5V_{0,0}\left(\frac{n}{k}\right)\right). \tag{3.5}$$

Analogously, it holds for $(2,2)$-events that

$$E_{2k,2k}(n)=O\left(k^4E_{0,0}\left(\frac{n}{k}\right)\right). \tag{3.6}$$

Putting everything together, we get equation (3.7), given in Figure 1. In it, the second equality follows from (3.1) and (3.2), the third from (3.3) and (3.4), the fourth from the fact that both $V_{k,2k}(n)$ and $V_{2k,k}(n)$ (resp., $E_{k,2k}(n)$ and $E_{2k,k}(n)$) are not larger than $V_{2k,2k}(n)$ (resp., $E_{2k,2k}(n)$), and the fifth from (3.5) and (3.6).

We have thus recursively related the number of $(2,3)$- and $(3,2)$-events to the number of $(2,3)$-, $(3,2)$- and $(2,2)$-events in a sample of n/r surfaces. In the next Section we derive a corresponding system of recurrences for $(2,2)$-events. Combining all these relations together in Section 5 will yield the desired result.

4 Overlay Edges

In this section we analyze the number of $(2,2)$-events. As in the previous section, we design a counting scheme in order to obtain a recurrence relation. This counting scheme is one of the novel features of this paper, and no similar scheme, that charges curves rather than vertices, has been employed in previous works involving

$$
\begin{aligned}
V_{0,0}(n) &= V_{0,0}^{32}(n) + V_{0,0}^{23}(n) \\
&= O\left(\tfrac{\beta(n)}{k}\left(V_{0,k}^{32}(n) + V_{k,0}^{23}(n)\right) + kn^{3+\varepsilon}\beta(k)\right) \\
&= O\left(\tfrac{\beta(n)}{k}\left(\tfrac{1}{\zeta}V_{k,2k}(n) + \tfrac{k}{\zeta}E_{k,2k}(n) + \tfrac{1}{\zeta}V_{2k,k}(n) + \tfrac{k}{\zeta}E_{2k,k}(n)\right) + kn^{3+\varepsilon}\beta(k)\right) \\
&= O\left(\tfrac{\beta(n)}{k\zeta}V_{2k,2k}(n) + \tfrac{\beta(n)}{\zeta}E_{2k,2k}(n) + kn^{3+\varepsilon}\beta(k)\right) \\
&= O\left(\tfrac{\beta(n)}{k\zeta}k^5V_{0,0}\left(\tfrac{n}{k}\right) + \tfrac{\beta(n)}{\zeta}k^4E_{0,0}\left(\tfrac{n}{k}\right) + kn^{3+\varepsilon}\beta(k)\right) \\
&= O\left(\beta(n)k^{4-\frac{1}{\lceil s/2\rceil}}\left(V_{0,0}\left(\tfrac{n}{k}\right) + E_{0,0}\left(\tfrac{n}{k}\right)\right) + kn^{3+\varepsilon}\beta(k)\right) \qquad (3.7)
\end{aligned}
$$

Figure 1: The recurrence for $V_{0,0}(n)$.

substructures in arrangements. The counting scheme is different from the previous one in two significant ways.

First, it has only one stage. Intuitively, the reason is that each $(2,2)$-event is defined by four surfaces (two blue and two red ones), and a naive bound on the number of such events is $O(n^4)$. We wish to derive a sub-quartic bound, and thus need only to shave (less than) one order of magnitude off. A one-stage counting scheme is sufficient for this purpose.

The second difference is that we partition the set of $(2,2)$-events into a small number of groups, according to the *index* of the event. We prove a separate recurrence relation for the number of events in each group. These recurrences will depend on each other, as well as on the number of $(\le k)$-level $(2,3)$- and $(3,2)$-events. The notion of index employed here (and defined below) is different from the one used in the previous section, and is rather non-standard.

4.1 Preliminaries

Let us dispose of two easy cases before moving on to the main part of the analysis. Consider the (at most a constant number of) $(2,2)$-events defined by a fixed quadruple of surfaces b_1, b_2, r_1, r_2. Suppose one of these $(2,2)$-events is at level at most k and is incident to at least one $(2,3)$- or $(3,2)$-event. It is therefore incident to a $(\le k)$-level vertex of the overlay. We can charge all $(2,2)$-events defined by b_1, b_2, r_1, r_2 to this vertex. Each vertex can only be charged in this fashion at most a constant number of times. The overall number of $(2,2)$-events defined by b_1, b_2, r_1, r_2 as above is thus $O(V_{k,k}(n))$, and (3.5) implies that this is $O(k^5V_{0,0}(n/k))$.

Suppose now that one of these $(2,2)$-events is at level at most k, is not incident to a $(2,3)$- or $(3,2)$-event, and is not a closed curve (and thus "reaches" infinity). Consider a bounding hypersphere Ψ of the blue and red arrangements. (Informally, Ψ has the property that increasing its radius by an arbitrary amount does not change the combinatorial structure of its intersection with $\mathcal{A}(\mathcal{R})$ and $\mathcal{A}(\mathcal{B})$, and that all features of the overlay arise from events fully contained in Ψ.) Let $\mathcal{A}(\mathcal{R}_\Psi)$ and $\mathcal{A}(\mathcal{B}_\Psi)$ denote the 3-dimensional arrangements that are obtained by intersecting $\mathcal{A}(\mathcal{R})$ and $\mathcal{A}(\mathcal{B})$ with Ψ. Our assumption implies that the $(2,2)$-event that we consider corresponds to a $(\le k)$-level vertex of the overlay of $\mathcal{A}(\mathcal{R}_\Psi)$ and $\mathcal{A}(\mathcal{B}_\Psi)$ (in the appropriate embedding of Ψ into $\mathbb{R}^3$). The number of such vertices, and thus the number of $(2,2)$-events as above, is $O(k^2n^{2+\varepsilon})$, as follows from the analysis of overlays of envelopes and shallow levels in 3 dimensions [2].

In the remainder of this section we treat $(2,2)$-events defined by some b_1, b_2, r_1, r_2, such that all $(\le k)$-level $(2,2)$-events defined by b_1, b_2, r_1, r_2 are *closed* curves that are not incident to a $(2,3)$- or a $(3,2)$-event. We fix a 2-face φ that belongs to $\mathcal{E}_\mathcal{B}$, and consider the vertical 3-dimensional curtain surface Π_φ, as in the previous section. Lat $\mathcal{R}_\varphi$ be the cross-section of $\mathcal{R}$ within Π_φ, and let P_φ be the cross-section of P within Π_φ, for any $P \in \mathcal{R}$. $\mathcal{A}(\mathcal{R}_\varphi)$ is an arrangement of xy-monotone 2-dimensional surfaces in $\mathbb{R}^3$.

Let $\mathcal{S}(\mathcal{R}_\varphi)$ denote the set of curves, for which the following holds. Each curve c of $\mathcal{S}(\mathcal{R}_\varphi)$ is a $(0,0)$-level $(2,2)$-event that corresponds to an intersection curve (defined by some $r_1, r_2 \in \mathcal{R}_\varphi$) that lies completely in the lower envelope of $\mathcal{A}(\mathcal{R}_\varphi)$, such that all intersection curves (including c) that lie in level at most k of $\mathcal{A}(\mathcal{R}_\varphi)$ and are defined by r_1 and r_2, are closed and are not incident to a vertex.

Define $\mathcal{S}(\mathcal{R}, \mathcal{B}) \equiv \sum_\varphi |\mathcal{S}(\mathcal{R}_\varphi)|$, and $\mathcal{S}(n) \equiv \max S(\mathcal{R}, \mathcal{B})$, where the maximum is taken over all red sets $\mathcal{R}$ and blue sets $\mathcal{B}$ of size n as above. The above analysis implies

$$
(4.8)\ E_{0,0}(n) = O\left(S(n) + k^5V_{0,0}\left(\frac{n}{k}\right) + k^2n^{2+\varepsilon}\right).
$$

4.2 The Counting Scheme

We next define a new notion of *index* that we attach to intersection curves, as above. This definition is conceptually different from the one employed in the previous section and from previous similar-purpose definitions [12], in that this definition is not "local", meaning that the index given to a specific intersection curve is defined with respect to the whole arrangement $\mathcal{A}(\mathcal{R}_\varphi)$, and may change as the set $\mathcal{R}_\varphi$ decreases.

Specifically, consider the intersection $r_1 \cap r_2$, for any $r_1, r_2 \in \mathcal{R}_\varphi$. Consider all the connected components of this intersection that are events of $\mathcal{S}(\mathcal{R}_\varphi)$, and let j be their number. We set the index of the curve $r_1 \cap r_2$ to j. Clearly, j varies between 1 and q (the case $j = 0$ will not concern us), where q is the maximum possible number of connected components of an intersection $r_1 \cap r_2$ as above. By algebraicity and the general position assumption, q is constant. Let $\mathcal{S}^{(j)}(\mathcal{R}_\varphi)$ denote the subset of $\mathcal{S}(\mathcal{R}_\varphi)$ that contains events with index at least j. Define $\mathcal{S}^{(j)}(\mathcal{R}, \mathcal{B}) \equiv \sum_\varphi |\mathcal{S}^{(j)}(\mathcal{R}_\varphi)|$, and $\mathcal{S}^{(j)}(n) \equiv \max S^{(j)}(\mathcal{R}, \mathcal{B})$, where the maximum is taken over all red sets $\mathcal{R}$ and blue sets $\mathcal{B}$ of size n as above. Since the maximal index of a curve is q, we have $\mathcal{S}^{(q+1)}(n) = 0$. We also have, by definition, $S(n) = S^{(1)}(n)$.

We note that the index of $r_1 \cap r_2$ depends on the current set $\mathcal{R}$. When $\mathcal{R}$ is replaced by a smaller sample, as happens for example when applying the Clarkson-Shor bound, the index may increase, because more components of $r_1 \cap r_2$ may appear on the envelope. Clearly, the index can never decrease by such sampling.

The counting scheme below bounds $\mathcal{S}^{(j)}(n)$, for all $1 \le j \le q$. It proceeds by distinguishing between a number of possible scenarios, and by treating each in turn.

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

Case 2: There is a "shallow" connected component of the same intersection curve. Consider any pair of surfaces $P, Q \in \mathcal{R}$, such that there is a connected component c of the intersection $P_\varphi \cap Q_\varphi$ that is an event of $\mathcal{S}^{(j)}(\mathcal{R}_\varphi)$. Suppose there is a curve c' defined by $P_\varphi \cap Q_\varphi$ whose level is between 1 and k. In this case, we charge all events of $\mathcal{S}^{(j)}(\mathcal{R}_\varphi)$ defined by $P_\varphi \cap Q_\varphi$ to c'. It is easy to see that each such curve c' is charged in this fashion at most a constant number of times.

Since $\mathcal{S}^{(j)}(\mathcal{R}_\varphi) \subseteq \mathcal{S}(\mathcal{R}_\varphi)$, c is an event of $\mathcal{S}(\mathcal{R}_\varphi)$.

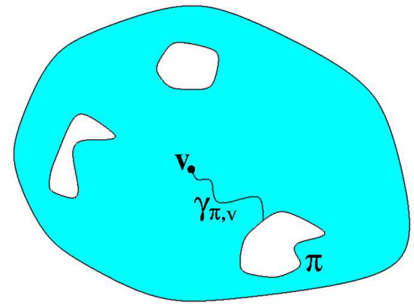

Figure 2: A schematic visualization of Case 3, showing a point v that is charged by a curve π, together with the connecting arc $\gamma_{\pi,v}$. The region Δ is shown shaded.

This implies that all connected components of $P_\varphi \cap Q_\varphi$ that lie at level at most k of $\mathcal{A}(\mathcal{R}_\varphi)$ (one of which is c') are closed and are not incident to a vertex. Since c belongs to $\mathcal{S}^{(j)}(\mathcal{R}_\varphi)$, the index of $P_\varphi \cap Q_\varphi$ is at least j, meaning that the number of connected components of the intersection $P_\varphi \cap Q_\varphi$ that lie on the lower envelope of $\mathcal{A}(\mathcal{R}_\varphi)$ is at least j.

Let Γ be the set of the (at most k) surfaces of $\mathcal{R}_\varphi$ that lie below c'. Set $\mathcal{R}'_\varphi \equiv \mathcal{R}_\varphi \setminus \Gamma$, and consider the arrangement $\mathcal{A}(\mathcal{R}'_\varphi)$. 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 it belongs to $\mathcal{S}^{(j+1)}(\mathcal{R}'_\varphi)$. A standard random sampling argument, as the ones used in [6, 12], now implies that the maximum number of events of $\mathcal{S}^{(j)}(n)$ in this case is $O(k^4 S^{(j+1)}(n/k))$. In the sequel we assume that there is no curve c' as above.

Case 3: There is a "shallow" vertex. By the analysis of Case 1, there are at least $(q+1)k$ intersection curves lying on any of P_φ, Q_φ, disregarding the components defined by $P_\varphi \cap Q_\varphi$. (Indeed, if r and r' are two red surfaces that appear on the envelope, then they must intersect over φ, as is easily checked.) Consider all the connected components (curves) of $P_\varphi \cap Q_\varphi$ that lie in the lower envelope of $\mathcal{A}(\mathcal{R}_\varphi)$. These curves partition 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 and a common projection onto φ. One of the regions in each pair (said to the the 'top' region) is completely above the other. One of the top regions, denoted by Δ, has to contain at least k intersection curves. Assume that Δ is part of Q_φ and consider an arbitrary curve π (defined by $P_\varphi \cap Q_\varphi$) that lies on its boundary. Let C be the set of intersection curves (each defined by $O_\varphi \cap Q_\varphi$, for some $O_\varphi \in \mathcal{R}_\varphi$) contained in Δ.

Suppose that there is at least one vertex v on Q_φ

within Δ that lies in level ($\leq k$), such that v can be connected to a point on π by an arc $\gamma_{\pi,v}$ (of finite description complexity) that lies on Q_φ within Δ, and is at level ($\leq k$), for all points in its interior. (This situation is visualized in Figure 2.) In this case, we charge all events of $\mathcal{S}^{(j)}(\mathcal{R}_\varphi)$ defined by $P_\varphi \cap Q_\varphi$ to an arbitrary such vertex. Note that each such vertex corresponds to a $(\leq k, 0)$-level $(3,2)$-event.

Similarly, suppose there is at least one curve that passes above/below the boundary of φ in (red) 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, v has to lie on a part of the boundary of Δ that is not defined by $P_\varphi \cap Q_\varphi$.) In this case, we charge to an arbitrary such point v, as above, noting that each such point corresponds to a $(\leq k, 0)$-level $(2,3)$-event.

LEMMA 4.1. *Each point v is charged by at most qk distinct curves of $\mathcal{S}^{(j)}(\mathcal{R}_\varphi)$ as described in Case 3.*

The proof, as usual, is omitted. Informally, this follows from the fact that, if v were charged also by some curve $P'_\varphi \cap Q_\varphi$, then P'_φ would have to be one of the surfaces that cross $\gamma_{\pi,v}$.

Combined with (3.5), Lemma 4.1 implies that the Maximum number of events of $\mathcal{S}^{(j)}(n)$ in this case is $O(kV_{k,0}(n)) = O(k^6 V_{0,0}(n/k))$. In the sequel we assume that there is no vertex v as above.

Case 4: There is a "shallow" curve that reaches infinity. We continue to use the setup introduced in Case 3, and suppose that there is a curve of C that lies in level at most k and is not closed. (Since we assume that the scenario treated in Case 3 does not hold, this can only occur if Δ is unbounded, and the curve reaches infinity.) In this case, we charge all events of $\mathcal{S}^{(j)}(\mathcal{R}_\varphi)$ defined by $P_\varphi \cap Q_\varphi$ to this curve. Arguing as in the previous subsection, we can show that the overall number of such curves is $O(k^2 n^{2+\varepsilon})$. The proof of Lemma 4.1 can easily be modified to show that each curve is charged at most qk times in this fashion. Thus, the maximum number of events of $\mathcal{S}^{(j)}(n)$ in this case is $O(k^3 n^{2+\varepsilon})$. In the sequel we assume that there is no curve of C as above.

Case 5. In this final case we distinguish between two subcases.

Subcase 5.A: There is a point on Δ that lies in level ($> k$). We can connect this point to a point on one of the boundary arcs π of Δ by an arc (of unknown finite description complexity) that lies on Q_φ within Δ, in its interior. Since one end-point of this arc lies in level ($> k$) and the other on the lower envelope, the arc intersects at least k distinct curves of C. Moreover, the first k distinct curves encountered when walking on the arc away from π, lie in level ($\leq k$), since π lies on the lower envelope. We charge π to these k curves of C.

Subcase 5.B: Δ lies entirely in level ($\leq k$). We charge π to k arbitrary curves of C.

In both subcases, we charge one event of $\mathcal{S}^{(j)}(\mathcal{R}_\varphi)$ defined by $P_\varphi \cap Q_\varphi$ (out of the at most q such events) to k curves of C that lie in level ($\leq k$). Let c be a curve of C that is charged in such fashion. We can show (proof omitted) that c is charged at most twice. Let Γ be the set of the (at most k) surfaces of $\mathcal{R}_\varphi$ that lie below c. Set $\mathcal{R}'_\varphi \equiv \mathcal{R}_\varphi \setminus \Gamma$. Since we assume that none of the scenarios treated in Cases 1–4 holds, it is easy to see that c is an event of $\mathcal{S}(\mathcal{R}'_\varphi)$. A standard random sampling argument, as the ones used in Case 2 and in [12], thus implies that the maximum number of events of $\mathcal{S}^{(j)}(n)$ in this case is $(q/k)O(k^4 S(n/k)) = O(k^3 S(n/k))$.

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

$$
\begin{aligned}
S^{(j)}(n) &= O\left(k^3 S\left(\frac{n}{k}\right)\right) + O\left(k^4 S^{(j+1)}\left(\frac{n}{k}\right)\right) \\
&\quad + O\left(k^6 V_{0,0}\left(\frac{n}{k}\right)\right) + O\left(k^2 n^{3+\varepsilon}\right). \qquad (4.9)
\end{aligned}
$$

5 Putting It All Together

The system of inter-dependent recurrences derived in the two preceding sections solves to

$$
V_{0,0}(n) = E_{0,0}(n) = O\left(n^{4-\frac{1}{\lceil s/2 \rceil}+\varepsilon}\right). \qquad (5.10)
$$

This is shown by induction, as in [12], choosing a different value of k for each recurrence. In more detail, we order the functions appearing in the recurrences as $(E_{0,0}, S^{(1)}, S^{(2)}, \ldots, S^{(q)}, V_{0,0})$ and denote this, for uniformity, as $(F_1, F_2, \ldots, F_{q+2})$. Each recurrence is roughly of the form

$$
\begin{aligned}
F_i(n) &= O\left(k_i^{\beta_1} F_{j_1}\left(\frac{n}{k_i}\right)\right) + O\left(k_i^{\beta_2} F_{j_2}\left(\frac{n}{k_i}\right)\right) \\
&\quad + \cdots + O\left(k_i^{\beta_r} F_{j_r}\left(\tfrac{n}{k_i}\right)\right) + O\left(f_i(n)\right) \qquad (5.11)
\end{aligned}
$$

(the recurrence for $V_{0,0}$ also involves the factor $\beta(n)$, which we can ignore by choosing the corresponding k to be much larger than $\beta(n)$). We represent this system symbolically by a directed graph G on the indices $\{1, 2, \ldots, q+2\}$, whose directed edges are $(i, j_1), \ldots, (i, j_r)$. We call an edge (i, j) a *forward* (resp., *backward*) edge if $j > i$ (resp., $j \leq i$). Let β be the maximum of the exponents β_t, taken over all corresponding edges (i, j_t) that are *backward* edges. Then one can show that the solution of this system is $F_i(n) = O(n^{\beta+\varepsilon})$, for any $\varepsilon > 0$ and for all i. Informally, larger exponents in terms that relate F_i to a function F_j with a *larger* index do not affect the overall bound, because their effect can be suppressed by the choice of much larger values for k_i with larger indices i.

Since in our case, under the order given above, $\beta = 4 - 1/\lceil s/2 \rceil$, we obtain the bound given in (5.10). This proves the asserted bound on the number of vertices of $\mathcal{Q}(\mathcal{R}, \mathcal{B})$ that correspond to $(3,2)$- and $(2,3)$-events, as well as on the number of 1-dimensional faces of $\mathcal{Q}(\mathcal{R}, \mathcal{B})$ that correspond to $(2,2)$-events and are not incident to a vertex, as each such face can be uniquely charged to a $(2,2)$-event.

The number of other vertices of the overlay is $O(n^{3+\varepsilon})$, since each such vertex is a projection of a vertex of either the red or the blue lower envelope. Similarly, the number of 1-dimensional faces that are not incident to a vertex and do not correspond to $(2,2)$-events is also $O(n^{3+\varepsilon})$, as each such face is a projection of an edge of one of the envelopes.

The number of 3-, 2-, and 1-dimensional faces that are incident to at least one vertex in the overlay is at most proportional to the number of vertices, since each such face can be charged to one of the vertices incident to it. By the assumption of general position, no vertex is charged more than a constant number of times.

We now bound the number of 2-dimensional faces of the overlay that are not incident to any vertex or 1-dimensional face. Each such 2-dimensional face is a projection of an intersection surface of two (red or blue) surfaces. There are thus $O(n^2)$ such faces.

Any 3-dimensional face that is incident to at least one lower-dimensional face can be charged to this lower-dimensional face, as in the case of vertices above. It is easy to see that, unless the overlay is empty, no 3-dimensional faces exist that are not incident to lower-dimensional faces.

This completes the proof of Theorem 2.1.

6 Applications

Our bound on the complexity of overlays in $\mathbb{R}^4$ has many applications. We mention here several of the more obvious ones. We omit all proofs and technical details due to space limitations. All the results listed below improve significantly upon the best previously known ones. Their proofs crucially rely on Theorem 2.1.

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

Theorem 6.1. *The combinatorial complexity of $\Sigma_{\mathcal{R},\mathcal{B}}$ is $O(n^{4-1/\lceil s/2 \rceil+\varepsilon})$, where s is as defined above.*

As a matter of fact, the proof of Theorem 6.1 implies the following stronger result; we refer the reader to [4, 9] for details concerning vertical decompositions in higher dimensions.

Corollary 6.1. *The combinatorial complexity of the first stage of the vertical decomposition of $\Sigma_{\mathcal{R},\mathcal{B}}$ is $O(n^{4-1/\lceil s/2 \rceil+\varepsilon})$.*

This corollary still leaves open the question of whether the complexity of the entire vertical decomposition of $\Sigma_{\mathcal{R},\mathcal{B}}$ is sub-quartic.

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 $\mathcal{C}$.

Theorem 6.2. *The combinatorial complexity of $T_3(\mathcal{C})$ is $O(n^{4-1/\lceil s/2 \rceil+\varepsilon})$, where s is as defined above.*

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 $\mathcal{C}$; i.e., the set of all lines that intersect every member of $\mathcal{C}$.

Theorem 6.3. *The combinatorial complexity of $T_1(\mathcal{C})$ is $O(n^{4-1/\lceil s/2 \rceil+\varepsilon})$, where s is as defined above.*

Remark 1. *The value of s in Theorem 6.2 has the following geometric interpretation: it is the maximum number of orientations u for which there exist two hyperplanes of orientation u, one tangent from below to three sets of $\mathcal{C}$ and the other tangent from above to two other sets of $\mathcal{C}$, or vice versa. An analogous definition holds for Theorem 6.3.*

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 maximum curvature 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 $\mathcal{U} = U_{\mathcal{C}}$ 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 6.4. *The combinatorial complexity of $\mathcal{U}$ is $O(n^{4-1/\lceil s/2 \rceil+\varepsilon})$, where s is the maximum number of* *intersections between the projections onto the $x_1x_2x_3$-hyperplane of the intersection curve of any three boundaries of sets in $\mathcal{C}$ and the 2-intersection surface of any two other boundaries.*

The analysis technique. The analysis technique introduced in this paper is general enough to have a wide variety of applications, and has already proved useful in unrelated contexts. In particular, we have applied it in a recent study of the combinatorial complexity of Voronoi diagrams of lines with a fixed number of orientations in three dimensions. With appropriate adaptations, the ideas of Sections 3.2 and 4 allowed us to obtain a substantially sub-cubic upper bound for this complexity [10].

The analysis in Section 4 seems to also be interesting on its own. For instance, it can be adjusted 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 $O(n^{1+\varepsilon})$. (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.)

7 Conclusion

We have obtained a sub-quartic bound on the complexity of overlays of envelopes of arrangements of surfaces in four dimensions, and have outlined some of its immediate applications. We strongly suspect that more applications of this result, and of the analysis technique used to obtain it, will become apparent in time, as has happened in the case of overlays in three dimensions [2].

At least two open problems remain. The first is improving the bound. We conjecture that the correct bound on the complexity of the overlay is near-cubic for arbitrary constant values of s. Another problem is motivated by the construction of lower envelopes in four dimensions. An improved algorithm for this task would be possible, given a sub-quartic bound on the complexity of the vertical decomposition of the overlay of two envelopes as above. Obtaining such a bound is still a challenging open problem, but the ideas introduced herein provide an essential step towards its solution.

Acknowledgements

Part of the work on this paper has been carried out while the authors were visiting ETH Zurich, 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, A. Efrat, and M. Sharir. Vertical decomposition of shallow levels in 3-dimensional arrangements and its applications. *SIAM J. Comput*, 29:912–953, 1999.

[2] P. K. Agarwal, O. Schwarzkopf, and M. Sharir. The overlay of lower envelopes and its applications. *Discrete Comput. Geom.*, 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] B. Chazelle, H. Edelsbrunner, L. J. Guibas, and M. Sharir. A singly-exponential stratification scheme for real semi-algebraic varieties and its applications. *Theoret. Comput. Sci.*, 84:77–105, 1991. Also in Proc. ICALP (1989), pp. 179–193.

[5] L. P. Chew, K. Kedem, M. Sharir, B. Tagansky, and E. Welzl. Voronoi diagrams of lines in 3-space under polyhedral convex distance functions. *J. Algorithms*, 29(2):238–255, 1998.

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

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

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

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

[10] V. Koltun and M. Sharir. Three dimensional Euclidean Voronoi diagrams of lines with a fixed number of orientations. *In preparation.*

[11] O. Schwarzkopf and M. Sharir. Vertical decomposition of a single cell in a three-dimensional arrangement of surfaces and its applications. *Discrete Comput. Geom.*, 18:269–288, 1997.

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

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