

On Overlays and Minimization Diagrams

Vladlen Koltun
Computer Science Department
353 Serra Mall, Gates 464
Stanford University
Stanford, CA 94305, USA
vladlen@stanford.edu

Micha Sharir^{*}
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
michas@tau.ac.il

ABSTRACT

The overlay of $2 \leq m \leq d$ minimization diagrams of n surfaces in \mathbb{R}^d is isomorphic to a substructure of a suitably constructed minimization diagram of mn surfaces in \mathbb{R}^{d+m-1} . This elementary observation leads to a new bound on the complexity of the overlay of minimization diagrams of collections of d -variate semi-algebraic surfaces, a tight bound on the complexity of the overlay of minimization diagrams of collections of hyperplanes, and faster algorithms for constructing such overlays. Further algorithmic implications are discussed.

Categories and Subject Descriptors: F.2.2 [Nonnumerical Algorithms and Problems]: Geometrical problems and computations—*Computation on discrete structures, Geometrical problems and computations*; G.2 [Discrete Mathematics]: Combinatorics—*Counting problems*

General Terms: Algorithms, Theory

Keywords: Lower envelopes, Minimization diagrams, Overlays, Arrangements, Voronoi diagrams, Power diagrams, Hyperplanes

1. INTRODUCTION

The study of arrangements of linear and curved surfaces is a fundamental and notably active research field in both discrete and computational geometry. A central topic of study is the analysis of combinatorial geometric structures such as lower envelopes, minimization diagrams, and overlays of minimization diagrams. The wealth of algorithmic

^{*}Work by this author was supported by NSF Grants CCR-00-98246 and CCF-05-14079, by a grant from the Israeli Academy of Sciences for a Center of Excellence in Geometric Computing at Tel Aviv University, and by the Hermann Minkowski-MINERVA Center for Geometry at Tel Aviv University.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SCG'06, June 5–7, 2006, Sedona, Arizona, USA.

Copyright 2006 ACM 1-59593-340-9/06/0006 ...\$5.00.

and combinatorial applications of these particular structures has been amply described elsewhere [5], we only mention below some key results concerning their complexity. (See next section for precise definitions.)

It is easy to see that the complexity of an arrangement of n semi-algebraic d -variate functions of constant descriptive complexity is $O(n^{d+1})$ and that this bound is tight even for arrangements of (nonvertical) hyperplanes in \mathbb{R}^{d+1} (which are totally defined linear d -variate functions). The lower envelope of an arrangement of n hyperplanes in \mathbb{R}^{d+1} is a convex polyhedron with at most n facets and thus has maximum complexity $\Theta(n^{\lceil d/2 \rceil})$ [19, 22]. The lower envelope of an arrangement of n d -simplices in \mathbb{R}^{d+1} has maximum complexity $\Theta(n^d \alpha(n))$ [13], and the maximum complexity of the lower envelope of an arrangement of n semi-algebraic d -variate functions of constant descriptive complexity is $O(n^{d+\varepsilon})$ for any $\varepsilon > 0$ [20] (where the constant of proportionality depends on ε , a convention that we will keep implicit in the remainder of the paper), which is nearly tight in the worst case. Slightly sharper bounds, related to Davenport-Schinzel sequences, are available for univariate functions, and extend to a few special instances in higher dimensions too [21]. Thus the worst-case complexity of the lower envelope of an arrangement is generally about a factor of n smaller than the worst-case complexity of the whole arrangement. The *minimization diagram* of an arrangement is the projection of its lower envelope onto the d -space of the independent variables (see next section for more details). By definition, the complexity of a minimization diagram is equal to the complexity of the lower envelope.

For $d \leq 3$, it is known that the complexity of the *overlay* of two (or of any constant number of) minimization diagrams of n semi-algebraic d -variate functions of constant descriptive complexity is $O(n^{d+\varepsilon})$ for any $\varepsilon > 0$ [4, 17]. This matches the best known upper bounds on the complexity of a single minimization diagram and is nearly tight. (Again, slightly sharper bounds are known for $d = 1$.) For $d \geq 4$, however, the only known upper bound on the complexity of the overlay of, say, two minimization diagrams as above is the naive $O(n^{2d+\varepsilon})$ for any $\varepsilon > 0$. Improving this bound is a long standing open problem. This has recently been achieved for the special case of overlays of minimization diagrams of n simplices in \mathbb{R}^{d+1} , whose complexity was shown to be $O(n^{d+\varepsilon})$ for any $\varepsilon > 0$ [18]. The more general semi-algebraic case remained wide open.

In this paper we prove that the complexity of the overlay of minimization diagrams of $2 \leq m \leq d$ collections of n semi-algebraic d -variate functions of constant descriptive complexity is $O(n^{d+m-1+\varepsilon})$ for any $\varepsilon > 0$. We thus stand, for the special case $m = 2$, about one factor of n away from the known lower bound of $\Omega(n^d \alpha^2(n))$ [18]. On the algorithmic side, we show that we can enumerate the vertices, edges and 2-faces of the overlay of a pair of diagrams in time $O(n^{d+1+\varepsilon})$, for any $\varepsilon > 0$. This is again at most about one factor of n from the obvious lower bound. Applications of efficient overlay construction abound, see [5, 17, 18]; we highlight some further ones in Section 4.

An interesting kind of overlays is that of minimization diagrams of two or more collections of n hyperplanes in \mathbb{R}^{d+1} . Since the lower envelope of a collection of hyperplanes is a convex polyhedron, such minimization diagrams are exactly the regular polyhedral subdivisions of \mathbb{R}^d (see [22] for this terminology). We are thus concerned with the maximum possible complexity of the overlay of two (or more) such subdivisions, a rather basic question in the theory of polytopes. The reader is invited to verify that for $d = 2$ this complexity is $\Theta(n^2)$. For $d = 3$ the same lower bound holds, implying that the complexity of the overlay is $\Omega(n^2)$. Interestingly, it is also tight. An even stronger property is known [16]: The overlay of two (not necessarily regular) polyhedral subdivisions R and B of \mathbb{R}^3 has complexity $O(N_R + N_B + n_R n_B)$, where N_R, N_B are the complexities of R and B , and n_R, n_B are their respective numbers of three-dimensional cells. Since the number of three-dimensional cells of a regular polyhedral subdivision of \mathbb{R}^3 defined by a convex polyhedron with n facets is clearly n , and its complexity is $O(n^2)$, the bound of $O(n^2)$ on the complexity of the above overlay in \mathbb{R}^3 follows. However, this bound does not extend to the overlay of three polyhedral subdivisions of \mathbb{R}^3 .

In this paper we prove that the complexity of the overlay of m minimization diagrams of n hyperplanes in \mathbb{R}^{d+1} , for $2 \leq m \leq d$, is $O(n^{\lfloor (d+m)/2 \rfloor})$, thereby generalizing, for $m = 2$, the bound of [16] for regular subdivisions. We also show that this upper bound is tight. For $m = 2$, this demonstrates that the maximum complexity of the overlay of two regular polyhedral subdivisions of \mathbb{R}^d differs by a factor of n from the maximum complexity of one such subdivision for even d while being asymptotically equal to it for d odd! This surprising behavior is in contrast with the case for simplices [18] and with the conjectured state of affairs for semi-algebraic surfaces of constant descriptive complexity (but in these cases the bounds are much larger, so the difference between the two cases is not that striking).

All the above new results are derived by means of the following truly elementary observation:

The overlay of two (resp., m) minimization diagrams of n d -variate functions is isomorphic to a substructure of a suitably constructed minimization diagram of $2n$ ($d+1$)-variate functions (resp., mn ($d+m-1$)-variate functions).

It is remarkable that this basic property has not, to our knowledge, been discovered before.

The overlay of minimization diagrams of collections of hyperplanes has been previously used for the computation of the minimum-width spherical shell in dimensions two and three [1, 11, 12]. The above bound on the complexity of this

overlay immediately implies an algorithm for the minimum-width spherical shell problem in any dimension d , whose running time is $O(n^{\lfloor d/2 \rfloor + 1})$. An algorithm with the same running time was described by Chan [8], who improved by about a factor of n the running times achieved in [11] and (the conference version of) [1]. In fact, the above-stated main observation stemmed from contemplating Chan's mapping of all spherical shells enclosing n points in \mathbb{R}^d to a convex polytope in \mathbb{R}^{d+2} [8], which has led to the seed idea of the current paper.

It is our hope that this note will facilitate further advances in the study of overlays of minimization diagrams. We outline some algorithmic implications of the new results in Section 4, including a recent application of our results to the analysis of the ICP algorithm [15].

2. OVERLAYS AS MINIMIZATION DIAGRAMS

For simplicity of presentation, we first analyze the overlay of *two* minimization diagrams, and then, in the next section, generalize it to the overlay of arbitrarily many diagrams.

Let F be a collection of n d -variate functions. The lower envelope \mathcal{E}_F of F is defined to be their pointwise minimum:

$$\mathcal{E}_F(\mathbf{x}) := \min_{f \in F} f(\mathbf{x}), \quad \text{for } \mathbf{x} \in \mathbb{R}^d.$$

We assume for simplicity that at least one function of F is defined over every $\mathbf{x} \in \mathbb{R}^d$ and the above definition is thus valid for all $\mathbf{x} \in \mathbb{R}^d$; this assumption can easily be lifted. The minimization diagram \mathcal{M}_F of F is the orthogonal projection of \mathcal{E}_F onto \mathbb{R}^d . It is the subdivision of \mathbb{R}^d into maximal connected relatively open faces of dimensions ranging from 0 to d , over each of which \mathcal{E}_F is attained by a fixed subset of \mathcal{F} (of cardinality $d+1$ to 1, respectively, if one assumes general position—see below). Define \mathcal{E}_G and \mathcal{M}_G analogously for another collection G of n d -variate functions. The *overlay* $\mathcal{Q}(F, G)$ of \mathcal{M}_F and \mathcal{M}_G is their superposition in \mathbb{R}^d . It is a subdivision of \mathbb{R}^d into maximal connected relatively open faces of various dimensions, over each of which each of the envelopes \mathcal{E}_F and \mathcal{E}_G is attained by a fixed set of functions of F and of G , respectively. Without loss of generality, assume that the functions in $F \cup G$ are totally defined, continuous, and in general position with respect to each other. (See [5] for a definition and justification.) The overall number of faces of all dimensions of \mathcal{E}_F , \mathcal{M}_F , or $\mathcal{Q}(F, G)$ is called the combinatorial complexity (in short, complexity) of the respective structure. With a slight abuse of notation, throughout this paper we refer interchangeably to minimization diagrams of collections of functions and of collections of surfaces in Euclidean space that are the graphs of these functions.¹

For each $f \in F$ and $g \in G$, consider the following corresponding $(d+1)$ -variate functions:

$$\begin{aligned} f^*(\mathbf{x}, y) &:= f(\mathbf{x}), & \text{for } \mathbf{x} \in \mathbb{R}^d, & \quad y \in \mathbb{R}, \\ g^*(\mathbf{x}, y) &:= g(\mathbf{x}) + y, & \text{for } \mathbf{x} \in \mathbb{R}^d, & \quad y \in \mathbb{R}. \end{aligned}$$

¹We also define the upper envelope of F as the pointwise maximum of the functions of F . The orthogonal projection of this upper envelope onto \mathbb{R}^d is called the *maximization diagram* of F . The overlay of two maximization diagrams, or of a minimization and a maximization diagram, is defined in analogy to the overlay of minimization diagrams.

Let F^* (resp., G^*) be the collection of the functions f^* (resp., g^*), for all $f \in F$ (resp., $g \in G$). The following lemma is the key observation of this paper. It implies that the overlay of minimization diagrams of two collections of d -variate functions, each of cardinality n , arises as a substructure in the minimization diagram of a suitably constructed collection of $2n(d+1)$ -variate functions. Moreover, if the original d -variate functions are algebraic of degree k (or semi-algebraic of constant descriptive complexity), the constructed $(d+1)$ -variate functions are algebraic of degree k (or semi-algebraic of constant descriptive complexity) as well.

LEMMA 2.1. *There exists an injection from the set of faces of $\mathcal{Q}(F, G)$ to the set of faces of $\mathcal{M}_{F^* \cup G^*}$. That is, every face of the overlay $\mathcal{Q}(F, G)$ can be mapped to a unique corresponding face of the minimization diagram $\mathcal{M}_{F^* \cup G^*}$.*

PROOF. For a point $\mathbf{x} \in \mathbb{R}^d$, let $F_{\mathbf{x}}$ be the subset of F that attains \mathcal{E}_F over \mathbf{x} , at some value $\varphi_{\mathbf{x}}$. Let $G_{\mathbf{x}}$ be the subset of G that attains \mathcal{E}_G over \mathbf{x} , at some value $\gamma_{\mathbf{x}}$. Let $F_{\mathbf{x}}^*$ and $G_{\mathbf{x}}^*$ be the corresponding subsets of F^* and G^* . Note that $F_{\mathbf{x}}^* \cup G_{\mathbf{x}}^*$ attain $\mathcal{E}_{F^* \cup G^*}$ over the point $\mathbf{x}^* = (\mathbf{x}, \varphi_{\mathbf{x}} - \gamma_{\mathbf{x}})$ at value $\varphi_{\mathbf{x}}$, while all other functions of F^* and G^* lie strictly above this value over \mathbf{x}^* . Also, assume without loss of generality that $\gamma_{\mathbf{x}} < \varphi_{\mathbf{x}}$; then for all $\xi < \varphi_{\mathbf{x}} - \gamma_{\mathbf{x}}$ (resp., $\xi > \varphi_{\mathbf{x}} - \gamma_{\mathbf{x}}$) the envelope $\mathcal{E}_{F^* \cup G^*}$ over the point (\mathbf{x}, ξ) is attained by precisely the set $G_{\mathbf{x}}^*$ (resp., $F_{\mathbf{x}}^*$).

For a face s of $\mathcal{Q}(F, G)$, pick an arbitrary $\mathbf{x} \in s$. Define $F_{\mathbf{x}}$, $G_{\mathbf{x}}$, $\varphi_{\mathbf{x}}$, and $\gamma_{\mathbf{x}}$ as above. Let t denote the face of $\mathcal{M}_{F^* \cup G^*}$ that contains the point $\mathbf{x}^* = (\mathbf{x}, \varphi_{\mathbf{x}} - \gamma_{\mathbf{x}})$. The required correspondence between faces of $\mathcal{Q}(F, G)$ and of $\mathcal{M}_{F^* \cup G^*}$ is defined by mapping s to t . A schematic illustration of this mapping is given in Figure 1.




Figure 1: Turning a feature of the overlay of two minimization diagrams to a feature of the minimization diagram of extended functions.

To prove that this is an injection we need to show that (i) the face t is well defined, and is independent of the choice of $\mathbf{x} \in s$, and (ii) distinct faces s, s' of $\mathcal{Q}(F, G)$ are mapped to distinct faces t, t' of $\mathcal{M}_{F^* \cup G^*}$. We prove (ii), and then note that a similar and simpler argument establishes (i).

Assume for the sake of contradiction that distinct faces s, s' of $\mathcal{Q}(F, G)$ are mapped to the same face t of $\mathcal{M}_{F^* \cup G^*}$. That is, there are points $\mathbf{x}_1 \in s$ and $\mathbf{x}_2 \in s'$, such that the points \mathbf{x}_1^* and \mathbf{x}_2^* lie in the same face t of $\mathcal{M}_{F^* \cup G^*}$. Thus there is a continuous path π^* in \mathbb{R}^{d+1} that connects \mathbf{x}_1^* and \mathbf{x}_2^* and lies in its entirety in the face t . Let $F_t^* \subseteq F^*$ and $G_t^* \subseteq G^*$ be the functions that attain $\mathcal{E}_{F^* \cup G^*}$ over t and note that both sets are not empty. It follows that $F_t^* \cup G_t^*$

is exactly the set of functions that attain $\mathcal{E}_{F^* \cup G^*}$ over every point in π^* .

Consider the orthogonal projection π of the path π^* onto \mathbb{R}^d . π is a continuous path that connects \mathbf{x}_1 and \mathbf{x}_2 . Furthermore, we claim that the set of functions of F (resp., of G) that attain \mathcal{E}_F (resp., \mathcal{E}_G) is the same for every point along π and is precisely the set of functions that correspond to F_t^* (resp., G_t^*). Indeed, as observed above, the functions in F^* and G^* are constructed in such a fashion that for any $\mathbf{x} \in \mathbb{R}^d$, there is only one value ξ for which the envelope $\mathcal{E}_{F^* \cup G^*}$ is attained at the point $(\mathbf{x}, \xi) \in \mathbb{R}^{d+1}$ by both functions from F^* and functions from G^* , and these functions are exactly those that correspond to the functions of F and G that attain \mathcal{E}_F and \mathcal{E}_G , respectively, over \mathbf{x} ; since functions from both F^* and G^* attain $\mathcal{E}_{F^* \cup G^*}$ over π^* , the corresponding sets of functions from F and G are necessarily the sets that attain \mathcal{E}_F (resp., \mathcal{E}_G) over π .

Since \mathbf{x}_1 and \mathbf{x}_2 can be connected by a continuous path π , over every point of which the set of functions that attain \mathcal{E}_F and \mathcal{E}_G is the same, \mathbf{x}_1 and \mathbf{x}_2 lie in the same face of $\mathcal{Q}(F, G)$. We have reached a contradiction that proves the second claim (ii). The proof of (i) is simpler, and uses the same arguments in reverse. \square

We note several consequences of Lemma 2.1.

THEOREM 2.2. *The complexity of the overlay of minimization diagrams of two sets of n semi-algebraic d -variate functions of constant descriptive complexity is $O(n^{d+1+\varepsilon})$, for any $\varepsilon > 0$. Moreover, we can enumerate the vertices, edges and 2-faces of the overlay in time $O(n^{d+1+\varepsilon})$, for any $\varepsilon > 0$. (The constants of proportionality depend on d, ε , and on the descriptive complexity of the functions.)*

PROOF. The combinatorial bound follows from Lemma 2.1, which implies that the complexity of the overlay is at most the complexity of the minimization diagram of $2n$ semi-algebraic $(d+1)$ -variate functions of constant descriptive complexity, and the complexity of such minimization diagrams is known to be $O(n^{d+1+\varepsilon})$, for any $\varepsilon > 0$ [20]. Agarwal et al. [2] show that the vertices, edges and 2-faces of the minimization diagram of n semi-algebraic $(d+1)$ -variate functions of constant descriptive complexity can be enumerated in time $O(n^{d+1+\varepsilon})$, for any $\varepsilon > 0$. Combining this with Lemma 2.1 clearly implies the algorithmic part of the theorem. \square

The bound in Theorem 2.2 seems not to be tight, because the image of the embedding of $\mathcal{Q}(F, G)$ into $\mathcal{M}_{F^* \cup G^*}$ seems to consist of only a small portion of the minimization diagram, that is, of only one face at each slice through a fixed $\mathbf{x} \in \mathbb{R}^d$. Surprisingly, this objection does not apply to the case of hyperplanes, where the technique yields a worst-case tight bound:

THEOREM 2.3. *The complexity of the overlay of minimization diagrams of two sets of n hyperplanes in \mathbb{R}^{d+1} (regarded as graphs of linear d -variate functions) is $O(n^{\lfloor d/2 \rfloor + 1})$. This bound is tight in the worst case. The overlay can also be constructed in worst-case optimal time $O(n^{\lfloor d/2 \rfloor + 1})$.*

PROOF. The upper bound follows from Lemma 2.1, as in the proof of Theorem 2.2, this time plugging in the Upper Bound Theorem [19] that shows that the complexity of the lower envelope of $2n$ hyperplanes in \mathbb{R}^{d+2} is $O(n^{\lfloor d/2 \rfloor + 1})$.

For the lower bound construction, denote the coordinate axes in \mathbb{R}^{d+1} by x_1, \dots, x_{d+1} . Consider a collection of hyperplanes in \mathbb{R}^d (with coordinate axes $x_1, \dots, x_{d-1}, x_{d+1}$) whose lower envelope (with x_{d+1} as the dependent variable) has $\Theta(n^{\lfloor d/2 \rfloor})$ vertices. Such a collection can be constructed using duals of cyclic polytopes; see [22]. Take the product of each of these hyperplanes with the x_d -direction. The collection of these products forms the first family F of hyperplanes. Regarding the elements of F as graphs of linear functions with x_{d+1} as the dependent variable, the minimization diagram of F contains $\Theta(n^{\lfloor d/2 \rfloor})$ 1-faces (edges), each of which is an x_d -parallel line.

Let the second family G consist of n hyperplanes of the form $x_{d+1} = i(x_d + i)$, for $i = 1, \dots, n$. When we overlay \mathcal{M}_F with \mathcal{M}_G , each of the $n - 1$ projections of the $(d - 1)$ -flats of intersections of consecutive hyperplanes of G along \mathcal{E}_G intersects each of the $\Omega(n^{\lfloor d/2 \rfloor})$ 1-faces of \mathcal{M}_F in a point, creating $n - 1$ vertices of $\mathcal{Q}(F, G)$ on every 1-face of \mathcal{M}_F . This produces $\Omega(n^{\lfloor d/2 \rfloor + 1})$ vertices of $\mathcal{Q}(F, G)$.

Finally, the algorithmic result follows from standard convex hull algorithms [9]. \square

Remark. The maximum complexity of a single minimization diagram is $\Theta(n^{\lfloor d/2 \rfloor})$. Hence, as already mentioned, the bounds on the complexities of a single minimization diagram and of an overlay of two diagrams coincide (asymptotically) for d odd, and differ by a factor of n for d even.

Theorem 2.3 yields an interesting corollary, which follows from the well-known representation of power diagrams (which include Voronoi diagrams as special cases) as lower envelopes of hyperplanes; see [14].

COROLLARY 2.4. *The maximum possible complexity of the overlay of two (closest-point or farthest-point) power diagrams, each defined by n points in \mathbb{R}^d , is $\Theta(n^{\lfloor d/2 \rfloor + 1})$.*

3. OVERLAYS OF MANY DIAGRAMS

In some applications (see [15]) we need to consider the overlay of many minimization diagrams. Let F_1, F_2, \dots, F_m be m collections of n d -variate functions. For each $1 \leq i \leq m$, define \mathcal{E}_{F_i} and \mathcal{M}_{F_i} as in Section 2. Define $\mathcal{Q}(F_1, \dots, F_m)$ to be the superposition of all diagrams \mathcal{M}_{F_i} , for $i = 1, \dots, m$. For each $f \in F_i$, $1 \leq i \leq m$, we define a corresponding $(d + m - 1)$ -variate function as

$$f^*(\mathbf{x}, \mathbf{y}) := \begin{cases} f(\mathbf{x}) & i = 1 \\ f(\mathbf{x}) + \mathbf{y}_{i-1} & i > 1, \end{cases}$$

for $\mathbf{x} \in \mathbb{R}^d$ and $\mathbf{y} \in \mathbb{R}^{m-1}$. Let

$$\mathcal{F}^* := \{f^* : f \in F_i, 1 \leq i \leq m\}.$$

The following lemma is proved similarly to Lemma 2.1.

LEMMA 3.1. *There exists an injection from the set of faces of $\mathcal{Q}(F_1, \dots, F_m)$ to the set of faces of $\mathcal{M}_{\mathcal{F}^*}$.*

As above, this leads to the following result.

THEOREM 3.2. *The complexity of the overlay of minimization diagrams of m collections of n semi-algebraic d -variate functions of constant descriptive complexity is $O(n^{d+m-1+\varepsilon})$,*

for any $\varepsilon > 0$, for $2 \leq m \leq d$, and is $O(m^d n^{2d-1+\varepsilon})$, for any $\varepsilon > 0$, for $m > d$. Moreover, we can enumerate the vertices, edges and 2-faces of the overlay within the same asymptotic time bounds. (The constants of proportionality depend on d , ε , and on the descriptive complexity of the functions.)

PROOF. The bounds for $m \leq d$ follow from Lemma 3.1 and from [20, 21], as in the proof of Theorem 2.2. For $m > d$, we note that each vertex of the overlay is also a vertex of the overlay of some d of the given minimization diagrams. We thus iterate over all $O(m^d)$ d -tuples of diagrams, and apply the bound with $m = d$ to each tuple. Extending this argument to higher-dimensional features of the overlay is straight-forward. \square

As in Section 2, the case of hyperplanes yield improved bounds, which are tight for $d = O(1)$.

THEOREM 3.3. *The complexity of the overlay of minimization diagrams of m collections of n hyperplanes in \mathbb{R}^{d+1} (regarded as graphs of linear d -variate functions) is $O((mn)^{\lfloor (d+m)/2 \rfloor})$ and $\Omega(n^{\lfloor (d+m)/2 \rfloor})$ for $m \leq d$, and $O(\binom{m}{d} (dn)^d)$ and $\Omega\left(\left(n \lfloor \frac{m}{d} \rfloor\right)^d\right)$ for $m > d$.*

PROOF. The upper bound for $m \leq d$ follows by combining Lemma 3.1 with the Upper Bound Theorem [19]. For the upper bound for $m > d$ we note that every feature of the overlay lies in the intersection of at most d faces from the various minimization diagrams. Hence each vertex of the overlay also shows up as a vertex of the overlay of d of the minimization diagrams. Thus we can obtain all the vertices of the overlay by iterating over the overlays of d -tuples of diagrams. Extending this argument for higher dimensional features of the overlay is straightforward. This gives a considerably better upper bound for $m \gg d$, as stated in the theorem.

The first lower bound construction in an extension of the one in Theorem 2.3. Consider a collection of hyperplanes in \mathbb{R}^{d-m+2} (with coordinate axes $x_1, \dots, x_{d-m+1}, x_{d+1}$) whose lower envelope (with x_{d+1} as the dependent variable) has $\Omega(n^{\lfloor (d-m+2)/2 \rfloor})$ vertices. Take the product of each of these hyperplanes with the subspace spanned by the coordinates x_{d-m+2}, \dots, x_d . The collection of these products forms the family F_1 . For each $2 \leq i \leq m$, the family F_i consists of n hyperplanes of the form $x_{d+1} = j(x_{d-m+i} + j)$, for $j = 1, \dots, n$. It is easy to verify that the overlay of these collections has $\Omega\left(n^{\lfloor (d-m+2)/2 \rfloor} n^{m-1}\right) = \Omega\left(n^{\lfloor (d+m)/2 \rfloor}\right)$ vertices.

The second lower bound construction treats $d \lfloor \frac{m}{d} \rfloor$ of the collections as d “bundles” of $\lfloor \frac{m}{d} \rfloor$ collections. Thus a collection is referred to as $F_{i,j}$, where $1 \leq i \leq d$ and $1 \leq j \leq \lfloor \frac{m}{d} \rfloor$. The collection $F_{i,j}$ consists of hyperplanes of the form $x_{d+1} = k(x_i + k + jB)$, for $k = 1, \dots, n$ and a sufficiently large global constant B . The overlay of these $d \lfloor \frac{m}{d} \rfloor$ collections has $\Omega\left(\left(n \lfloor \frac{m}{d} \rfloor\right)^d\right)$ vertices. The bound can be tightened slightly by making use of the other $m - d \lfloor \frac{m}{d} \rfloor$ collections. \square

4. ALGORITHMIC IMPLICATIONS

Thinnest elliptic shell. It is well known that the center of the minimum-width spherical shell of a point set has to be a vertex of the overlay of its closest-point Voronoi diagram and its farthest-point Voronoi diagram, an observation that has been exploited in the design of efficient algorithms for finding such a spherical shell [1, 8, 11, 12]. Many further problems of this kind can be addressed using high-dimensional overlays. We give a simple illustration for a generalization of the minimum-width spherical shell problem in the plane.

Given a set P of n points in general position in the plane, we wish to find the thinnest (not necessarily axis-parallel) elliptic shell enclosing P . In certain applications, this can give a rather precise approximation of the point set and includes the smallest enclosing annulus [1] as a special case. Formally, we want to find parameters a, b, c , and θ that satisfy, for every point $(x_1, x_2) \in P$:

$$f_1 \leq (x'_1 - a)^2 + c(x'_2 - b)^2 \leq f_2,$$

$$\text{where } \begin{pmatrix} x'_1 \\ x'_2 \end{pmatrix} = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix},$$

such that $f_2 - f_1$ is minimized. (This is a somewhat arbitrary way of defining the thinness of an elliptic shell; it does however coincide with the usual notion of thinness when the elliptic shell in question becomes an annulus (i.e., when $c = 1$.)

It will be convenient to define $d = \tan \frac{\theta}{2}$ and substitute $\sin \theta = \frac{2d}{1+d^2}$ and $\cos \theta = \frac{1-d^2}{1+d^2}$. For each $\mathbf{x} = (x_1, x_2) \in P$ define a quadrivariate function

$$f_{\mathbf{x}}(a, b, c, d) = (x'_1 - a)^2 + c(x'_2 - b)^2, \quad (1)$$

where

$$\begin{pmatrix} x'_1 \\ x'_2 \end{pmatrix} = \begin{pmatrix} (1-d^2)/(1+d^2) & -2d/(1+d^2) \\ 2d/(1+d^2) & (1-d^2)/(1+d^2) \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}.$$

Let F be the collection of functions $f_{\mathbf{x}}$. We are seeking a, b, c , and d , for which

$$\max_{\mathbf{x} \in P} f_{\mathbf{x}}(a, b, c, d) - \min_{\mathbf{x} \in P} f_{\mathbf{x}}(a, b, c, d)$$

is minimized. In other words, the vertical distance between the upper envelope of F and the lower envelope of F is minimized at the point (a, b, c, d) . Let us rearrange (1) as

$$\begin{aligned} f_{\mathbf{x}}(a, b, c, d) &= (x'_1 - a)^2 + c(x'_2 - b)^2 \\ &= x_1^2 \left(\frac{d^4 + (4c-2)d^2 + 1}{(1+d^2)^2} \right) + x_2^2 \left(\frac{cd^4 + (4-2c)d^2 + c}{(1+d^2)^2} \right) \\ &\quad + 4x_1x_2 \left(\frac{d(c-1)(1-d^2)}{(1+d^2)^2} \right) - 2x_1 \left(\frac{2bcd - ad^2 + a}{1+d^2} \right) \\ &\quad + 2x_2 \left(\frac{bcd^2 - bc + 2ad}{1+d^2} \right) + (a^2 + b^2) \\ &= a_1w_1 + a_2w_2 + a_3w_3 + a_4w_4 + a_5w_5 + w_6, \end{aligned}$$

where

$$\begin{aligned} a_1 &= x_1^2 & w_1 &= \frac{d^4 + (4c-2)d^2 + 1}{(1+d^2)^2} \\ a_2 &= x_2^2 & w_2 &= \frac{cd^4 + (4-2c)d^2 + c}{(1+d^2)^2} \end{aligned}$$

$$\begin{aligned} a_3 &= 4x_1x_2 & w_3 &= \frac{d(c-1)(1-d^2)}{(1+d^2)^2} \\ a_4 &= -2x_1 & w_4 &= \frac{2bcd - ad^2 + a}{1+d^2} \\ a_5 &= 2x_2 & w_5 &= \frac{bcd^2 - bc + 2ad}{1+d^2} \\ & & w_6 &= a^2 + b^2. \end{aligned}$$

Thus $f_{\mathbf{x}}$ can also be viewed as a linear function $\varphi_{\mathbf{x}}(w_1, w_2, w_3, w_4, w_5, w_6)$ in six variables. (See, e.g., Agarwal et al. [3] for a similar linearization in the case of cylindrical shells.) The collection of n functions $\varphi_{\mathbf{x}}$ is denoted by Φ . The ‘‘valid’’ 6-tuples $(w_1, w_2, w_3, w_4, w_5, w_6)$ that correspond to quadruples (a, b, c, d) form a 4-manifold in \mathbb{R}^6 that we denote by Ψ . We are now looking for a 6-tuple $(w_1, w_2, w_3, w_4, w_5, w_6) \in \Psi$ that minimizes

$$\begin{aligned} \Delta(w_1, w_2, w_3, w_4, w_5, w_6) &= \\ &= \max_{\mathbf{x} \in P} \varphi_{\mathbf{x}}(w_1, w_2, w_3, w_4, w_5, w_6) - \\ &= \min_{\mathbf{x} \in P} \varphi_{\mathbf{x}}(w_1, w_2, w_3, w_4, w_5, w_6) \end{aligned}$$

among all possible $(w_1, w_2, w_3, w_4, w_5, w_6) \in \Psi$.

To this end, we use Theorem 2.3 (in \mathbb{R}^7) to construct the overlay \mathcal{O} of the minimization diagram and the maximization diagram of Φ , whose complexity is $O(n^4)$, in time $O(n^4)$. Since every cell of the overlay is a convex polyhedron, we can, in the same worst-case asymptotic time, decompose all cells of \mathcal{O} into simplices of various dimensions using the bottom-vertex simplicial decomposition [10]. Consider one simplex γ of this decomposition of dimension $2 \leq k \leq 6$. The minimization (resp., maximization) diagram of Φ is achieved by a fixed set of i (resp., j) hyperplanes of Φ over γ , such that $i + j = 8 - k$. Thus $\Delta(w_1, w_2, w_3, w_4, w_5, w_6)$ is a linear function over γ . We wish to find a point within $\gamma \cap \Psi$ that minimizes $\Delta(w_1, w_2, w_3, w_4, w_5, w_6)$ over this domain. Since γ and Ψ are both semi-algebraic of constant descriptive complexity, we can compute $\gamma \cap \Psi$ in constant time,² and then find the minimum point on it in additional constant time. Thus we can iterate over all simplices γ and find the global minimum of $\Delta(w_1, w_2, w_3, w_4, w_5, w_6)$ within Ψ in time $O(n^4)$. We have proved the following theorem:

THEOREM 4.1. *Given a set P of n points in general position in the plane, the thinnest (not necessarily axis-parallel) elliptic shell enclosing P , as defined above, can be found in time $O(n^4)$.*

Theorem 4.1 is just one concrete example of a variety of problems that can be tackled in a similar manner, in which we want to find a point that minimizes the difference between the upper and lower envelope of a collection of functions (or satisfies some other relation that involves the envelopes). If the given functions have reasonable linearization, then Theorem 2.3 can be used to solve the problem fairly efficiently. In other cases, Theorem 2.2 can be used, still yielding a considerable improvement over naive solutions.

Overlay of Voronoi diagrams and the ICP algorithm. Corollary 2.4 can be generalized, via Theorem 3.3, to yield the following.

²Of course, computations of this kind are quite involved, and use a battery of sophisticated algorithms in real algebraic geometry; see, e.g., [6].

COROLLARY 4.2. *The maximum possible complexity of the overlay of m (closest-point or farthest-point) power diagrams, each defined by n points in \mathbb{R}^d , is $O\left((mn)^{\lfloor (d+m)/2 \rfloor}\right)$ and $\Omega\left(n^{\lfloor (d+m)/2 \rfloor}\right)$ if $m \leq d$, and $O\left(\binom{m}{d}(dn)^d\right)$ and $\Omega\left(\left(n\lfloor \frac{m}{d} \rfloor\right)^d\right)$ for $m > d$.*

In particular, these bounds apply to the overlay of m (closest-point or farthest-point) Voronoi diagrams. This has, among other applications, the following interesting corollary.

COROLLARY 4.3. *Let A and B be two finite point sets in \mathbb{R}^d , with $|A| = m$ and $|B| = n$. For each $\mathbf{x} \in \mathbb{R}^d$ and $a \in A$, define $N_B(a + \mathbf{x})$ to be the nearest neighbor of $a + \mathbf{x}$ in B . Refer to the mapping $a \mapsto N_B(a + \mathbf{x})$ as the nearest-neighbor assignment at \mathbf{x} . Then the number of distinct nearest-neighbor assignments, over all possible translation \mathbf{x} , is $O(n^{\lfloor (d+m)/2 \rfloor})$ for $m \leq d$, and $O(m^d n^d)$ for $m \geq d$. Both bounds are tight in the worst case.*

PROOF. Note that $b = N_B(a + \mathbf{x})$ iff $a + \mathbf{x}$ lies in the Voronoi cell $V(b)$ in the closest-point Voronoi diagram $Vor(B)$, or iff $\mathbf{x} \in V(b - a)$ in $Vor(B - a)$, where $B - a = \{b - a \mid b \in B\}$. Hence the nearest-neighbor assignment is fixed in each cell of the overlay of the m Voronoi diagrams $Vor(B - a)$, for $a \in A$. The bounds are then immediate from Corollary 4.2 (treating d as a constant). Tightness also follows similar to the proof of Theorem 3.3. \square

Corollary 4.3 is used by Ezra et al. [15] to give a combinatorial bound on the number of iterations of the ICP algorithm, which has been proposed by Besl and McKay [7] and used over a decade as a heuristic approach to pattern matching, but has not been rigorously analyzed prior to the treatment in [15].

Sandwich regions. In conclusion, we note that one of the important applications of the overlay of minimization diagrams is to bound the complexity of the *sandwich region* between the upper envelope of one collection F of n semi-algebraic d -variate functions of constant descriptive complexity and the lower envelope of another such collection G [4, 17]. For this application our new bounds, though tight for the case of hyperplanes and within about one factor of n from the known lower bound for the case of semi-algebraic functions, are not strong enough to yield a notable improvement. Indeed, the sandwich region is easily seen to be a collection of cells in the arrangement of the graphs of the functions of $F \cup G$ in \mathbb{R}^{d+1} . Thus its complexity is at most the complexity of this arrangement, which is $O(n^{d+1+\epsilon})$ for any $\epsilon > 0$. This matches precisely the bound that can be derived through the application of Theorem 2.2. Any further tightening of our bounds on the complexity of the overlay will also yield an improvement for the complexity of the sandwich region.

Acknowledgements

We are grateful to Pankaj Agarwal for discussions concerning some of the applications outlined in Section 4, and to Esther Ezra for her kind permission to reproduce here the application of our results in [15].

5. REFERENCES

- [1] P. K. Agarwal, B. Aronov, S. Har-Peled, and M. Sharir. Approximation and exact algorithms for minimum-width annuli and shells. *Discrete Comput. Geom.*, 24(4):687–705, 2000.
- [2] P. K. Agarwal, B. Aronov, and M. Sharir. Computing envelopes in four dimensions with applications. *SIAM J. Comput.*, 26:1714–1732, 1997.
- [3] P. K. Agarwal, B. Aronov, and M. Sharir. Exact and approximation algorithms for minimum-width cylindrical shells. *Discrete Comput. Geom.*, 26:307–320, 2001.
- [4] P. K. Agarwal, O. Schwarzkopf, and M. Sharir. The overlay of lower envelopes and its applications. *Discrete Comput. Geom.*, 15:1–13, 1996.
- [5] 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.
- [6] S. Basu, R. Pollack, and M.-F. Roy. *Algorithms in Real Algebraic Geometry*. Springer Verlag, Berlin, 2003.
- [7] P. J. Besl and N. D. McKay. A method for registration of 3-d shapes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 14(2):239–256, 1992.
- [8] T. M. Chan. Approximating the diameter, width, smallest enclosing cylinder and minimum-width annulus. *Internat. J. Comput. Geom. Appl.*, 12(2):67–85, 2002.
- [9] B. Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete Comput. Geom.*, 10:377–409, 1993.
- [10] K. L. Clarkson. A randomized algorithm for closest-point queries. *SIAM J. Comput.*, 17:830–847, 1988.
- [11] C. A. Duncan, M. T. Goodrich, and E. A. Ramos. Efficient approximation and optimization algorithms for computational metrology. In *Proc. 8th ACM-SIAM Annu. Sympos. Discrete Algo.*, pages 121–130, 1997.
- [12] H. Ebara, N. Fukuyama, H. Nakano, and Y. Nakanishi. Roundness algorithms using the Voronoi diagrams. In *Abstracts 1st Canad. Conf. Comput. Geom.*, page 41, 1989.
- [13] H. Edelsbrunner. The upper envelope of piecewise linear functions: Tight complexity bounds in higher dimensions. *Discrete Comput. Geom.*, 4:337–343, 1989.
- [14] H. Edelsbrunner and R. Seidel. Voronoi diagrams and arrangements. *Discrete Comput. Geom.*, 1:25–44, 1986.
- [15] E. Ezra and M. Sharir. Bounds for the ICP algorithm. These proceedings, 2005.
- [16] L. J. Guibas, D. Halperin, J. Matoušek, and M. Sharir. On vertical decomposition of arrangements of hyperplanes in four dimensions. *Discrete Comput. Geom.*, 14:113–122, 1995.
- [17] V. Koltun and M. Sharir. The partition technique for overlays of envelopes. *SIAM J. Comput.*, 32:841–863, 2003.
- [18] V. Koltun and C. Wenk. Matching polyhedral terrains using overlays of envelopes. *Algorithmica*, 41:159–183, 2005.
- [19] P. McMullen. The maximal number of faces of a convex polytope. *Mathematika*, 17:179–184, 1970.
- [20] M. Sharir. Almost tight upper bounds for lower envelopes in higher dimensions. *Discrete Comput. Geom.*, 12:327–345, 1994.

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

[22] G. M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995. Revised edition 1998.