Matching Polyhedral Terrains Using Overlays of Envelopes (Extended Abstract)*

Vladlen Koltun[1] and Carola Wenk[2]

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

[2] University of Texas at San Antonio, Dept. of Computer Science, 6900 N. Loop 1604 West, San Antonio, TX 78249-0667, USA. carola@cs.utsa.edu

Abstract. We show that the combinatorial complexity of the overlay of the lower envelopes of two collections of d-variate piecewise linear functions of overall combinatorial complexity n is $\Omega(n^d\alpha^2(n))$ and $O(n^{d+\varepsilon})$ for any $\varepsilon > 0$ when $d \geq 2$, and $O(n^2\alpha(n)\log n)$ when $d = 2$. This extends and improves the analysis of de Berg et al. [9]. We also describe an algorithm that constructs the overlay in the same time.
We apply these results to obtain efficient general solutions to the problem of matching two polyhedral terrains in $\mathbb{R}^{d+1}$ under translation. For the perpendicular distance measure, which we adopt from functional analysis, we present a matching algorithm that runs in time $O(n^{2d+\varepsilon})$ for any $\varepsilon > 0$. For the directed and undirected Hausdorff distance measures, we present a matching algorithm that runs in time $O(n^{d^2+d+\varepsilon})$ for any $\varepsilon > 0$.

1 Introduction

Overlays of Envelopes. The *arrangement* $\mathcal{A}(\mathcal{F})$ of a collection $\mathcal{F}$ of graphs of d-variate functions (i.e., functions of d variables) in $\mathbb{R}^{d+1}$ is the subdivision of $\mathbb{R}^{d+1}$ induced by $\mathcal{F}$. The *lower envelope* $\mathcal{E}(\mathcal{F})$ of $\mathcal{A}(\mathcal{F})$ is the pointwise minimum of the functions of $\mathcal{F}$. For two collections $\mathcal{F}$ and $\mathcal{G}$ as above, the *sandwich region* $\mathcal{S}(\mathcal{F},\mathcal{G})$ consists of all points that lie below the lower envelope of $\mathcal{A}(\mathcal{F})$ and above the *upper envelope* of $\mathcal{A}(\mathcal{G})$ (defined as the pointwise maximum of the functions of $\mathcal{G}$). The *minimization diagram* $\mathcal{M}(\mathcal{F})$ of $\mathcal{E}(\mathcal{F})$ is the subdivision of $\mathbb{R}^d$ obtained by projecting $\mathcal{E}(\mathcal{F})$ onto the hyperplane $x_{d+1} = 0$. The *overlay* $\mathcal{O}(\mathcal{F},\mathcal{G})$ of envelopes $\mathcal{E}(\mathcal{F})$ and $\mathcal{E}(\mathcal{G})$ is the refined subdivision obtained by superimposing $\mathcal{M}(\mathcal{F})$ and $\mathcal{M}(\mathcal{G})$ in $\mathbb{R}^d$. The last definition can be naturally extended to the overlay $\mathcal{O}(\mathcal{F}_1,\mathcal{F}_2,\ldots,\mathcal{F}_k)$ of envelopes of arrangements of multiple collections $\mathcal{F}_1,\ldots,\mathcal{F}_k$. The (combinatorial) *complexity* of each structure introduced above is defined to be the overall number of its faces (of all dimensions).

The study of lower envelopes and related structures has a long and rich history in computational geometry, as they have innumerable applications to the

* A limited preliminary version of some of the results described in this paper has appeared in the second author's Ph.D. thesis [22].
** Supported in part by NSF Grant CCR-01-21555.

various problems in this field; see Sharir and Agarwal [20] for an overview. Edelsbrunner et al. [10, 11] and Pach and Sharir [18] have shown that the complexity of $\mathcal{E}(\mathcal{F})$ and of $\mathcal{S}(\mathcal{F},\mathcal{G})$, when $\mathcal{F}$ and $\mathcal{G}$ are collections of piecewise linear (possibly partially defined) functions in $\mathbb{R}^{d+1}$ of overall complexity n, is $O(n^d\alpha(n))$, where $\alpha(n)$ is the inverse Ackermann function. Agarwal et al. [2] have shown that when $\mathcal{F}$ and $\mathcal{G}$ consist of n semi-algebraic bivariate functions of constant description complexity, $\mathcal{O}(\mathcal{F},\mathcal{G})$ has complexity $O(n^{2+\varepsilon})$ for any $\varepsilon > 0$. Recently, Koltun and Sharir [13] have shown that for analogous collections $\mathcal{F}$ and $\mathcal{G}$ of trivariate functions, the complexity of $\mathcal{O}(\mathcal{F},\mathcal{G})$ is $O(n^{3+\varepsilon})$ for any $\varepsilon > 0$.

This paper deals with the special case of collections $\mathcal{F}$ and $\mathcal{G}$ of piecewise linear (possibly partially defined) d-variate functions of overall complexity n. A relevant result is that of de Berg et al. [9], who have studied the complexity of the vertical decomposition of an arrangement of a set of triangles in $\mathbb{R}^3$. Although their paper does not explicitly discuss overlays of envelopes, their analysis implies that the complexity of $\mathcal{O}(\mathcal{F},\mathcal{G})$ is $\Omega(n^2\alpha^2(n))$ and $O(n^2 2^{\alpha(n)}\log n)$ when $\mathcal{F}$ and $\mathcal{G}$ are as above and $d = 2$. We extend this analysis to show that the complexity of $\mathcal{O}(\mathcal{F},\mathcal{G})$ is $\Omega(n^d\alpha^2(n))$ and $O(n^{d+\varepsilon})$ for any $\varepsilon > 0$ when $d \geq 2$. This provides the first non-trivial upper bound on the complexity of the overlay of envelopes in dimensions $d > 3$. For $d = 2$ we prove a sharper bound of $O(n^2\alpha(n)\log n)$. We also show that the complexity of $\mathcal{O}(\mathcal{F}_1,\mathcal{F}_2,\ldots,\mathcal{F}_d)$ is $\Omega(n^d\alpha^d(n))$, for collections $\mathcal{F}_1,\ldots,\mathcal{F}_d$ as above. Finally, we describe an algorithm for constructing $\mathcal{O}(\mathcal{F},\mathcal{G})$, $\mathcal{S}(\mathcal{F},\mathcal{G})$ and $\mathcal{E}(\mathcal{F})$ in time that matches the respective complexity bounds.

Matching Terrains. The comparison of geometric objects is a task that naturally arises in many application areas, such as computer vision, computer aided design, robotics, medical imaging, etc. In many applications we are given a set of allowed transformations, and wish to *match* the shapes under these transformations, that is, to find an allowed transformation that, when applied to the first object, minimizes its distance (under some specific distance measure) to the second one. A natural transformation class is that of translations, which forms the focus of our work. See Alt and Guibas [4] for an overview of matching algorithms for various types of objects, distance measures and transformation classes.

Most matching algorithms in the existing literature either deal with two-dimensional problems or only consider point sets. Algorithms for matching shapes more complicated than points in dimensions higher than two have been presented for the first time only recently in the second author's thesis [22]. There it has been shown that a translation that minimizes the Hausdorff distance between two polyhedral sets of total complexity n in $\mathbb{R}^{d+1}$ can be computed in $O(n^{d^2+3d+2}\log^2 n)$ time for $d \geq 2$. The only other higher-dimensional result we recently learned about is the result by Agarwal et al. [1], who compute the minimum Hausdorff distance under translations for two sets of m and n L_2-balls in $\mathbb{R}^3$ in $O(m^2n^2(m+n)\log^3(mn))$ time.

Terrains are a natural subset of shapes that have particularly many applications in $\mathbb{R}^3$, especially for geographical data. However also in higher dimensions terrains are an important class of shapes since they are graphs of arbitrary d-variate functions. In Section 3 we present algorithms for matching polyhedral terrains in $\mathbb{R}^{d+1}$ under translations, for arbitrary $d \geq 1$. We present algorithms for these problems that reduce the matching task to the computation of certain overlays and sandwich regions of envelopes of collections of piecewise linear functions. We show that we can compute a translation of a terrain of complexity m which minimizes its *perpendicular distance* (which is an adaptation of the L_∞-Minkowski metric used in functional analysis) to a terrain of complexity n, in time $O((mn)^{d+\varepsilon})$ for any $\varepsilon > 0$. Sharper running time bounds are obtained for $d \leq 2$.

Assuming that the terrains are continuously defined over a convex domain, we provide an algorithm that matches two terrains of complexity n under the (directed or undirected) Hausdorff distance measure in time $O(n^{d^2+d+\varepsilon})$, for any $\varepsilon > 0$. Moreover, for the directed Hausdorff distance our algorithm applies even when we are matching a terrain with an arbitrary polyhedral set. For technical reasons, we assume that the metric in terms of which the Hausdorff distance is defined belongs to a certain class of convex polyhedral metrics of constant description complexity that includes for instance the L_∞- and L_1-metrics.

2 Overlays of Envelopes of Piecewise Linear Functions

2.1 Lower Bounds

In this section we describe two simple constructions of collections of n d-simplices in $\mathbb{R}^{d+1}$ for any $d \geq 2$ that define overlays of high complexity. Since d-simplices are special cases of piecewise linear d-variate functions, our lower bound naturally extends to the latter more general family of objects. When $d = 2$ both constructions are the same and are identical to the construction presented by de Berg et al. [9]. Throughout the remainder of the paper, denote the axes in the $(d+1)$-dimensional space by $x_1, \ldots, x_{d+1}$. Let x_{d+1} denote the vertical direction in $\mathbb{R}^{d+1}$, in terms of which the lower and upper envelopes are defined.

Theorem 1. *For $d \geq 2$, there are collections $\mathcal{F}$ and $\mathcal{G}$ of $O(n)$ d-simplices in $\mathbb{R}^{d+1}$, for which the complexity of $\mathcal{O}(\mathcal{F}, \mathcal{G})$ is $\Omega(n^d \alpha^2(n))$.*

Proof. For the sake of clarity, we describe the construction using infinite axis-parallel d-dimensional "strips". It can be trivially modified to use finite d-simplices in general position.

There exists a collection of n line segments in the plane, such that the complexity of the lower envelope of their arrangement is $\Omega(n\alpha(n))$, see [23]. Consider such a collection Γ in the $x_1 x_{d+1}$-plane. Without loss of generality, assume that the x_{d+1}-coordinates of the segments in Γ are strictly positive. Take the Cartesian product of each segment with the $(d-1)$-flat $x_1 = x_{d+1} = 0$. Let $\mathcal{F}$ be the resulting collection of d-dimensional strips.

Consider an analogously constructed collection C_1 of strips, this time orthogonal to the $x_2 x_{d+1}$-plane. For $2 \leq i \leq d-1$, consider also the collection $C_i = \bigcup_{j=1}^{n} s_j$ of strips, where s_j is the Cartesian product of the line segment $((2j, 0), (2j+1, 0))$, drawn in the $x_{i+1} x_{d+1}$-plane, with the $(d-1)$-flat $x_{i+1} = x_{d+1} = 0$. Define $\mathcal{G} = \bigcup_{i=1}^{d-1} C_i$.

When $d = 2$, overlaying $\mathcal{M}(\mathcal{F})$ and $\mathcal{M}(\mathcal{G})$ results in a grid of $\Omega(n\alpha(n)) \times \Omega(n\alpha(n))$ lines, thus producing $\Omega(n^2\alpha^2(n))$ vertices. In higher dimensions, overlaying $\mathcal{M}(\mathcal{F})$ and $\mathcal{M}(C_1)$ similarly produces $\Omega(n^2\alpha^2(n))$ infinite $(d-2)$-flats orthogonal to the x_1 and x_2 axes. The (partial) diagram $\mathcal{M}(C_2 \cup \ldots \cup C_{d-1})$, on the other hand, essentially contains a grid of $\Omega(n^{d-2})$ x_1x_2-parallel planes (each belonging to the boundary of the intersection of $d-2$ projections of strips, one strip from each of the groups $C_2, \ldots, C_{d-1}$). In the overlay $\mathcal{O}(\mathcal{F}, \mathcal{G})$, each of the latter planes intersects all of the former $(d-2)$-flats, resulting in $\Omega(n^d\alpha^2(n))$ vertices.

Theorem 2. *There are d collections of n d-simplices in $\mathbb{R}^{d+1}$, such that the complexity of the overlay of the d respective lower envelopes is $\Omega(n^d\alpha^d(n))$.*

Proof. For $1 \le i \le d$, let $\mathcal{F}_i$ be a collection of d-dimensional strips orthogonal to the x_ix_{d+1}-plane, constructed as follows. Consider, as above, a collection Γ of n segments, drawn in the x_ix_{d+1}-plane, such that the complexity of $\mathcal{E}(\Gamma)$ is $\Omega(n\alpha(n))$. We define $\mathcal{F}_i$ to be the collection of Cartesian products of the segments of Γ with the $(d-1)$-flat $x_i = x_{d+1} = 0$. The complexity of $\mathcal{O}(\mathcal{F}_1, \mathcal{F}_2, \ldots, \mathcal{F}_d)$ is easily seen to be as claimed.

Remark 1. Theorem 1 and the earlier construction of de Berg et al. [9] dispel a belief, expressed, e.g., in [3], that the analysis of Edelsbrunner et al. [11] implies a bound of $O(n^d\alpha(n))$ on the complexity of $\mathcal{O}(\mathcal{F}, \mathcal{G})$ when $\mathcal{F}$ and $\mathcal{G}$ are collections of piecewise linear functions of overall complexity n in $\mathbb{R}^{d+1}$.

2.2 Upper Bounds

We note that it is sufficient to analyze collections $\mathcal{F}$ and $\mathcal{G}$ of n d-simplices in general position, as such analysis easily carries over to arbitrary collections of piecewise linear functions of overall complexity n. We will thus confine ourselves to this setting. It is also easy to see that it is sufficient to count the vertices of $\mathcal{O}(\mathcal{F}, \mathcal{G})$, since all higher-dimensional faces of the overlay can be charged to its vertices, such that each vertex is charged at most a constant number of times.

Lemma 1. *Given a collection $\mathcal{F}$ of n d-simplices in $\mathbb{R}^{d+1}$, and a $(j+1)$-dimensional convex body P contained in the hyperplane $x_{d+1} = 0$, the combinatorial complexity of $\mathcal{E}(\mathcal{F}_{\partial P})$ is $O(n^{j+\varepsilon})$ for any $\varepsilon > 0$, where $\mathcal{F}_{\partial P}$ is the collection of cross-sections of the simplices of $\mathcal{F}$ within the x_{d+1}-vertical surface $\partial P \times x_{d+1}$ spanned by the boundary ∂P of P.*

Proof. Notice that $\mathcal{F}_{\partial P}$ is a collection of surfaces that are graphs of j-variate functions, partially defined over ∂P. Furthermore, every $(j+1)$-tuple of surfaces intersects in at most two points, and a vertical projection of a k-dimensional feature $(0 \le k \le j)$ of the arrangement $\mathcal{A}(\mathcal{F}_{\partial P})$ onto ∂P intersects an analogous projection of a $(j-k)$-dimensional feature of $\mathcal{A}(\mathcal{F}_{\partial P})$ in at most two points. In the full version of this paper [14] we show that the analysis of Sharir [19] for the complexity of the lower envelope of an arrangement of semi-algebraic surfaces carries over to our setting.

Theorem 3. *Given two collections $\mathcal{F}$ and $\mathcal{G}$ of n d-simplices in $\mathbb{R}^{d+1}$, the complexity of $\mathcal{O}(\mathcal{F},\mathcal{G})$ is $O(n^{d+\varepsilon})$ for any $\varepsilon > 0$.*

Proof. Our proof relies on the concept of *efficient hierarchical cuttings* [7, 16], and is based on the proof technique of de Berg et al. [9]. For the case $d = 2$ we provide a sharper upper bound of $O(n^2\alpha(n)\log n)$ (Theorem 4) using the analysis technique of Tagansky [21], which improves the result of de Berg et al. [9].

A $(1/r)$-*cutting* Ξ of a set Γ of n hyperplanes in $\mathbb{R}^d$ is a subdivision of the space into simplices, such that each simplex is intersected by at most n/r hyperplanes of Γ. The *size* of Ξ, denoted by $|\Xi|$, is defined to be the number of the simplices in the subdivision. A cutting Ξ' is said to C-*refine* a cutting Ξ if every simplex of Ξ' is completely contained in some simplex of Ξ, and every simplex of Ξ contains at most C simplices of Ξ'. Let C and ρ be appropriate constants. A sequence $\boldsymbol{\Xi} = \Xi_0, \Xi_1, \ldots, \Xi_k$ is called an *efficient hierarchical* $(1/r)$-*cutting* of Γ if Ξ_0 consists of the single degenerate 'simplex' $\mathbb{R}^d$, and for all $1 \le i \le k$, Ξ_i is a $(1/\rho^i)$-cutting of size $O(\rho^{di})$ of Γ that C-refines Ξ_{i-1}, and $\rho^{k-1} < r < \rho^k$. (Thus, $k = \lceil \log_\rho r \rceil$.) For any simplex s in Ξ_i, the simplex of Ξ_{i-1} that contains s is said to be the *parent* of s, denoted by parent(s).

To analyze the number of vertices of $\mathcal{O}(\mathcal{F},\mathcal{G})$ we first describe their combinatorial structure. For $d+1 = 3$, each vertex of $\mathcal{O}(\mathcal{F},\mathcal{G})$ is either a vertex of $\mathcal{M}(\mathcal{F})$ or $\mathcal{M}(\mathcal{G})$, or an intersection of an edge of $\mathcal{M}(\mathcal{F})$ with an edge of $\mathcal{M}(\mathcal{G})$ [2]. Similarly, it is easy to check that for any $d+1 \ge 3$, a vertex of $\mathcal{O}(\mathcal{F},\mathcal{G})$ is either a vertex of $\mathcal{M}(\mathcal{F})$ or $\mathcal{M}(\mathcal{G})$, or an intersection of a j-face of $\mathcal{M}(\mathcal{F})$ with a $(d-j)$-face of $\mathcal{M}(\mathcal{G})$, for some $1 \le j \le d-1$. We denote the vertices of the latter type as j-vertices. It is known that the number of vertices of $\mathcal{M}(\mathcal{F})$ and $\mathcal{M}(\mathcal{G})$ is $O(n^d\alpha(n))$ [10]. To bound the complexity of $\mathcal{O}(\mathcal{F},\mathcal{G})$ it remains to analyze the number of j-vertices, for $1 \le j \le d-1$.

Now, consider the $(d-1)$-faces of the projections of the simplices of $\mathcal{F}$ and $\mathcal{G}$ onto $x_{d+1} = 0$. For each such face, consider the $(d-1)$-hyperplane it spans. Let $\mathcal{H}$ be the collection of $2(d+1)n$ hyperplanes defined in this manner by all the simplices of $\mathcal{F}$ and $\mathcal{G}$. Construct an efficient hierarchical $(1/n)$-cutting of $\mathcal{H}$. By definition, each simplex at the last level of the cutting is crossed by at most a constant number of faces of the projections of $\mathcal{F}$ and $\mathcal{G}$. For convenience, we add one more refinement level to the hierarchical cutting such that no simplex belonging to this final level is cut by a face as above. We thus get a final hierarchy $\boldsymbol{\Xi} = \Xi_0, \Xi_1, \ldots, \Xi_k$, which satisfies the above definition of an efficient hierarchical cutting, with appropriate constants ρ and C, and the additional property that the simplices of Ξ_k are not intersected by the boundaries of the simplices in the projections of $\mathcal{F}$ and $\mathcal{G}$.

For any simplex s belonging to some level of $\boldsymbol{\Xi}$, define $\mathcal{F}_s^{\times}$ to be the set of simplices of $\mathcal{F}$ whose projections intersect the interior of s with their boundaries. Also let $\mathcal{F}_s^{\subset}$ be the set of simplices of $\mathcal{F}$ whose projections contain s, but do not contain parent(s). Define $\mathcal{G}_s^{\times}$ and $\mathcal{G}_s^{\subset}$ analogously with respect to $\mathcal{G}$, and put $\Gamma_s^{\times} = \mathcal{F}_s^{\times} \cup \mathcal{G}_s^{\times}$ and $\Gamma_s^{\subset} = \mathcal{F}_s^{\subset} \cup \mathcal{G}_s^{\subset}$. Also set $\mathcal{F}_s = \mathcal{F}_s^{\subset} \cup \mathcal{F}_s^{\times}$, and define $\mathcal{G}_s$ and Γ_s analogously.

Consider a j-vertex v of $\mathcal{O}(\mathcal{F}, \mathcal{G})$, for some $1 \leq j \leq d-1$. It is an intersection of a j-face of $\mathcal{M}(\mathcal{F})$ with a $(d-j)$-face of $\mathcal{M}(\mathcal{G})$, which are respectively defined by $d+1-j$ simplices of $\mathcal{F}$ and $j+1$ simplices of $\mathcal{G}$. The collection of these $(d+1-j)+(j+1) = d+2$ simplices is said to define v, and is denoted by def(v).

We claim that for every v as above, there exists a simplex s belonging to some level of Ξ, such that def(v) is contained in Γ_s and at least one of the simplices of def(v) belongs to $\Gamma_s^{\subset}$. Indeed, there is a simplex s_i at every level Ξ_i of Ξ that contains v. Every simplex s_i in the sequence $s_0, \ldots, s_k$ contains s_{i+1} (unless, of course, $i = k$). Since s_0 is the whole space $\mathbb{R}^d$, def(v) is completely contained in $\Gamma_{s_0}^{\times}$. On the other hand, $\Gamma_{s_k}^{\times}$ is by definition empty. Thus, there exists an i such that at least one simplex of def(v) is not contained in $\Gamma_{s_i}^{\times}$, but all of them are contained in $\Gamma_{parent(s_i)}^{\times}$ (and are thus not contained in $\Gamma_{parent(s_i)}^{\subset}$). All the simplices of def(v) that are not contained in $\Gamma_{s_i}^{\times}$ are thus contained in $\Gamma_{s_i}^{\subset}$, which proves our claim.

This claim implies that to count all j-vertices v as above it suffices to consider all simplices s of Ξ, and for each simplex to consider the vertices defined only by simplices from Γ_s, with at least one simplex coming from $\mathcal{F}_s^{\subset}$, without loss of generality. Let us consider a specific simplex s of Ξ and a specific value of j, and derive an upper bound on the number of j-vertices v that correspond to s in this fashion.

The j-face of $\mathcal{M}(\mathcal{F})$ that defines v lies on the projection of an intersection of $d+1-j$ simplices of $\mathcal{F}_s$. Our assumption implies that at least one of them belongs to $\mathcal{F}_s^{\subset}$. Consider some $(d-j)$-tuple of simplices of $\mathcal{F}_s$, and their $(j+1)$-dimensional intersection surface (which lies on a $(j+1)$-flat). The j-face of $\mathcal{E}(\mathcal{F})$ that defines v lies on the intersection of this surface, for some tuple as above, with the lower envelope $\mathcal{E}(\mathcal{F}_s^{\subset})$. Notice that the simplices of $\mathcal{F}_s^{\subset}$ are totally defined over s, and thus the envelope $\mathcal{E}(\mathcal{F}_s^{\subset})$ behaves over s as the lower envelope of a collection of hyperplanes, which is a convex polytope. The intersection of the above $(j+1)$-dimensional surface with $\mathcal{E}(\mathcal{F}_s^{\subset})$ is thus part of a j-dimensional convex polytope, defined as the intersection of $\mathcal{E}(\mathcal{F}_s^{\subset})$ with the $(j+1)$-flat containing the surface.

Consider the projection P of this polytope onto the hyperplane $x_{d+1} = 0$, and consider the cross-section of $\mathcal{M}(\mathcal{G}_s)$ within ∂P. It is the projection of the part of $\mathcal{E}(\mathcal{G}_s)$ that lies over ∂P, and Lemma 1 thus implies that its complexity is $O(|\mathcal{G}_s|^{j+\varepsilon})$ for any $\varepsilon > 0$. Any j-vertex v as above clearly corresponds to a vertex in this cross-section, for some $(d-j)$-tuple of simplices of $\mathcal{F}_s$ selected above. The number of such j-vertices v is thus $O\left(|\mathcal{F}_s|^{d-j}|\mathcal{G}_s|^{j+\varepsilon}\right)$, for any $\varepsilon > 0$. Notice now that the boundary of the projection of every simplex in Γ_s intersects the interior of parent(s), which implies that $|\Gamma_s| \leq n/\rho^{i_s-1}$, where i_s is the level of s in Ξ. The above quantity is thus bounded by

$$O\left(\left(\frac{n}{\rho^{i_s-1}}\right)^{d+\varepsilon}\right), \tag{1}$$

for any $\varepsilon > 0$. Summing over all simplices s, the number of j-vertices is

$$\sum_{s} O\left(\left(\frac{n}{\rho^{i_s-1}}\right)^{d+\varepsilon}\right) = \sum_{i=1}^{k}\sum_{|\Xi_i|} O\left(\left(\frac{n}{\rho^{i-1}}\right)^{d+\varepsilon}\right) = \sum_{i=1}^{k} O\left(\frac{n^{d+\varepsilon}\rho^{di}}{\rho^{(d+\varepsilon)(i-1)}}\right), \quad (2)$$

which equals $O(n^{d+\varepsilon})$ for any $\varepsilon > 0$. Noticing that the bound does not depend on j completes the proof.

Theorem 4. *Given two collections $\mathcal{F}$ and $\mathcal{G}$ of n triangles in three dimensions, the complexity of $\mathcal{O}(\mathcal{F},\mathcal{G})$ is $O(n^2\alpha(n)\log n)$.*

We skip the proof, and refer the reader to the full version of this paper [14].

2.3 Algorithms

Theorem 5. *Given two collections $\mathcal{F}$ and $\mathcal{G}$ of n d-simplices in $\mathbb{R}^{d+1}$, a complete combinatorial representation of $\mathcal{O}(\mathcal{F},\mathcal{G})$ (resp., $\mathcal{E}(\mathcal{F})$ and $\mathcal{S}(\mathcal{F},\mathcal{G})$) can be constructed in randomized expected time $O(n^{d+\varepsilon})$ for any $\varepsilon > 0$ (resp., time $O(n^d\alpha(n))$). When $d = 2$ the running time is $O(n^2\alpha(n)\log n)$.*

Proof. As in the proof of Theorem 3, consider the $(d-1)$-faces of the projections of the simplices of $\mathcal{F}$ and $\mathcal{G}$ onto $x_{d+1} = 0$. For each face, consider the $(d-1)$-hyperplane it spans. Let $\mathcal{H}$ be the collection of $2(d+1)n$ hyperplanes defined in this manner by all the simplices of $\mathcal{F}$ and $\mathcal{G}$. Consider the refinement of $\mathcal{O}(\mathcal{F},\mathcal{G})$ with these hyperplanes, as in [10, 18]. The cross-section of $\mathcal{O}(\mathcal{F},\mathcal{G})$ within some $h \in \mathcal{H}$ is actually the overlay $\mathcal{O}(\mathcal{F}_h,\mathcal{G}_h)$, where $\mathcal{F}_h$ (resp., $\mathcal{G}_h$) is the collection of cross-sections of the simplices of $\mathcal{F}$ (resp., of $\mathcal{G}$) with the x_{d+1}-vertical hyperplane spanned by h. Theorem 3 implies that the complexity of $\mathcal{O}(\mathcal{F}_h,\mathcal{G}_h)$ is $O(n^{d-1+\varepsilon})$, for every $\varepsilon > 0$. Therefore, refining $\mathcal{O}(\mathcal{F},\mathcal{G})$ (which is a subdivision of $\mathbb{R}^d$) with the $O(n)$ hyperplanes of $\mathcal{H}$ does not asymptotically increase the complexity of the subdivision, which remains $O(n^{d+\varepsilon})$ for every $\varepsilon > 0$. Each cell in the resulting refined subdivision is convex. It can thus be easily decomposed into simplices using the bottom-vertex simplicial decomposition [8, 15].

This representation of $\mathcal{O}(\mathcal{F},\mathcal{G})$ allows us to construct this overlay using a standard randomized incremental approach that utilizes a conflict graph. In fact, our setting fits into standard abstract frameworks, see e.g. [6, Section 5.2]. The construction proceeds by choosing a random permutation of the simplices of $\mathcal{F} \cup \mathcal{G}$. (We will refer to these simplices, extended by hyperplanes as above, as "objects" throughout the rest of this paragraph, to avoid confusion with the simplices of the decomposition.) We first construct in constant time the decomposition of the "overlay" of just the first object. We then add the objects one by one according to the random order. With every addition of an object, we insert it into the overlay and update the decomposition and the conflict graph. The conflict graph stores for every simplex in the decomposition a list of objects (that have not yet been added) that intersect it. Additionally, it stores for every such object a list of simplices that it intersects. This allows knowing which simplices are affected by the addition of a particular object. The restructuring of all affected simplices and their conflict lists is a standard procedure and we omit its rather routine details. By standard arguments [6, Section 5.2], the expected running time of the construction algorithm is

$$O\left(n \sum_{r=1}^{2n} \frac{f(r)}{r^2}\right), \tag{3}$$

where $f(r)$ denotes the maximal complexity of the overlay of envelopes of two sets of r simplices overall. Theorem 3 shows that $f(r) = O(r^{d+\varepsilon})$ for any $\varepsilon > 0$. The running time of the algorithm is thus $O(n^{d+\varepsilon})$ for any $\varepsilon > 0$. For $d = 2$ the running time becomes $O(n^2\alpha(n)\log n)$ due to Theorem 4.

One envelope $\mathcal{E}(\mathcal{F})$ and the sandwich region $\mathcal{S}(\mathcal{F}, \mathcal{G})$ can be constructed analogously in time $O(n^d\alpha(n))$, using the fact that these structures can also be refined into convex subdivisions that can be decomposed using the bottom-vertex decomposition.

Remark 2. We note that a $O(n^2\alpha(n)\log n)$ randomized incremental algorithm for constructing $\mathcal{E}(\mathcal{F})$ when $d = 2$ has been described by Boissonnat and Dobrindt [5]. Their goal was to obtain an on-line algorithm and their construction followed a different approach that uses a two-level *history graph* instead of the conflict graph. Also, a randomized divide-and-conquer algorithm for constructing $\mathcal{E}(\mathcal{F})$ in time $O(n^{d+\varepsilon})$ for any $\varepsilon > 0$ and $d \geq 2$ has been described by Sharir and Agarwal [20, Section 7.2.2].

3 Matching Terrains

In this section we apply the above results for overlays and sandwich regions to matching terrains in an arbitrary fixed dimension. A (k-dimensional) *terrain* F in $\mathbb{R}^{d+1}$ is the graph $F = \{(x, f(x)) \,|\, x \in D_f\}$ of a k-variate function $f : D_f \longrightarrow \mathbb{R}$, $0 \leq k \leq d$, where the domain D_f is a k-dimensional subset of $\mathbb{R}^d$. F is a *polyhedral terrain* if D_f is a polyhedral subset of $\mathbb{R}^d$, and f is a linear function over each polyhedron in D_f. Hence, a polyhedral terrain is a polyhedral set with the property that every x_{d+1}-vertical line intersects the terrain in at most one point. We assume in the following that a polyhedral set always consists of a collection of simplices. As long as the terrains are given as collections of convex polytopes this assumption is not restrictive, since each convex polytope of complexity n can be easily partitioned into $O(n)$ simplices [8, 15]. Hence we can associate with each terrain F a simplicial partition M_f of its domain D_f, such that f is linear over each simplex in M_f.

3.1 Perpendicular Distance

Let two polyhedral terrains $F = \{(x, f(x)) \mid x \in D_f\}$ and $G = \{(x, g(x)) \mid x \in D_g\}$ in $\mathbb{R}^{d+1}$ of complexity m and n, respectively, be given. Since each terrain intersects every vertical line at most once, it is natural to consider the height difference between vertically adjacent points of F and G as a distance measure. We therefore consider the *perpendicular distance* (also called *uniform metric* or *Chebyshev metric*) $\boldsymbol{\delta}_{\perp}(F,G) = \sup_{x \in D_f} |f(x) - g(x)|$, where we assume that $D_f \subseteq D_g$. Notice that the perpendicular distance is the standard L_∞-Minkowski metric for the functions f and g.

We consider a translation $t' = (t_1, \ldots, t_d, t_{d+1}) \in \mathbb{R}^{d+1}$ to be composed of a translation $t = (t_1, \ldots, t_d) \in \mathbb{R}^d$ and a translation $t_{d+1} \in \mathbb{R}$, hence $t' = (t, t_{d+1})$. Using this notation we have $F + t' = \{(x, f(x-t) + t_{d+1}) \mid x \in D_f + t\}$. We wish to compute a translation $(t^*, t^*_{d+1}) \in \mathbb{R}^{d+1}$, where $t^* \in \mathbb{R}^d$, $t^*_{d+1} \in \mathbb{R}$, such that $D_f \subseteq D_g$ and the perpendicular distance between F and G is minimized, hence

$$\boldsymbol{\delta}_{\perp}(F + (t^*, t^*_{d+1}), G) = \min_{\substack{t \in \mathbb{R}^d \\ D_f + t \subseteq D_g}} \min_{t_{d+1} \in \mathbb{R}} \boldsymbol{\delta}_{\perp}(F + (t, t_{d+1}), G) \ . \tag{4}$$

Reformulating $\boldsymbol{\delta}_{\perp}(F + t', G)$ produces

$$\boldsymbol{\delta}_{\perp}(F + t', G) = \max_{x \in D_f + t \subseteq D_g} |f(x-t) - g(x) + t_{d+1}| \tag{5}$$

$$= \max\{ \max_{x \in D_f \subseteq D_g - t} h_x(t) - t_{d+1} \ , \ t_{d+1} - \min_{x \in D_f \subseteq D_g - t} h_x(t)) \ \} \tag{6}$$

with $h_x(t) = g(x+t) - f(x)$. Observe that the condition $D_f \subseteq D_g - t$ is equivalent to $t \in D_h$, where $D_h = \overline{\overline{D_g} \oplus (-D_f)}$. Here and throughout the rest of the paper, $A \oplus B = \bigcup_{a \in A} \bigcup_{b \in B} \{a + b\}$ denotes the *Minkowski sum* (or *vector sum*) of two sets $A, B \subseteq \mathbb{R}^d$, while $\overline{A}$ denotes the complement of a set $A \subseteq \mathbb{R}^d$. We define two functions:

$$\overline{h} : D_h \longrightarrow \mathbb{R} \ ; \ t \mapsto \max_{x \in D_f} h_x(t) \tag{7}$$

$$\underline{h} : D_h \longrightarrow \mathbb{R} \ ; \ t \mapsto \min_{x \in D_f} h_x(t) \ . \tag{8}$$

Let $\overline{H}$ and $\underline{H}$ be the respective graphs of $\overline{h}$ and $\underline{h}$. These graphs are polyhedral terrains, and are, respectively, the upper and lower envelopes of the functions h_x, for all $x \in D_f$. Let $-F$ denote the set $-F = \{((x, -f(x)) \mid x \in D_f\}$. The next lemma follows directly from the definitions of $\overline{H}$ and $\underline{H}$.

Lemma 2. *Let $F, G, \overline{H}, \underline{H}$ be as defined above. Then $\overline{H}$ (resp., $\underline{H}$) is the upper (resp., the lower) envelope of $G \oplus (-F)$ restricted to the region above D_h.*

Reformulating (4) we have

$$\boldsymbol{\delta}_{\perp}(F + (t^*, t^*_{d+1}), G) = \min_{t \in D_h} \min_{t_{d+1} \in \mathbb{R}} \max\{\overline{h}(t) - t_{d+1}, t_{d+1} - \underline{h}(t)\}, \tag{9}$$

which implies that the translation $(t^*, t^*_{d+1}) \in \mathbb{R}^{d+1}$ we are seeking is such that:

$$\frac{\overline{h}(t^*) - \underline{h}(t^*)}{2} = \min_{t \in D_h} \frac{\overline{h}(t) - \underline{h}(t)}{2} \tag{10}$$

$$\text{and} \quad t^*_{d+1} = (\overline{h}(t^*) + \underline{h}(t^*))/2. \tag{11}$$

This leads to the following algorithm: We compute $D_h = \overline{D_g \oplus (-D_f)}$ by computing the Minkowski sums for each pair of simplices, one from D_g and one from $(-D_f)$, and computing the complement of their union with a brute-force arrangement approach, in $O((mn)^d)$ time. See the full version of this paper [14] for the technical details. The lower envelope $\underline{\mathcal{E}}$ and the upper envelope $\overline{\mathcal{E}}$ of $G \oplus (-F)$ are envelopes of nm pairs of simplices. Let M be the simplicial partition of the domain of $\{(t, (\overline{h}(t) - \underline{h}(t))/2) \mid t \in D_h\}$. Lemma 2 implies that M is the overlay of $\overline{\mathcal{E}}$ and $\underline{\mathcal{E}}$ restricted to D_h. We compute the overlay of $\overline{\mathcal{E}}$ and $\underline{\mathcal{E}}$, additionally superimposed with the $O(mn)$ hyperplanes defining D_h. Notice that in the proof of Theorem 5 the overlay of the envelopes is additionally overlaid with a set of hyperplanes, and for the correctness of the proof only the number of the hyperplanes matters. Thus, Theorem 5 implies that the overlay can be constructed in $O((mn)^{d+\varepsilon})$ time for any $\varepsilon > 0$. The function $(\overline{h}(t) - \underline{h}(t))/2$ is linear within each simplex of M. Therefore, the global minimum t^* that minimizes this function, as described in (10), is necessarily reached at a vertex of M_{h^-}. Hence it suffices to iterate over all vertices in M_{h^-}, which takes time proportional to their number, which is $O((mn)^{d+\varepsilon})$ for any $\varepsilon > 0$. We thus obtain t^*, which we can plug into (11) to get t^*_{d+1}.

Overall, the described algorithm runs in time $O((mn)^{d+\varepsilon})$, for $d \geq 2$. For $d = 1$ we can compute $\overline{\mathcal{E}}$ and $\underline{\mathcal{E}}$ in time $O(mn \log(mn))$ [12]. The computation of the overlay and the clipping can be done with a simple sweep in $O(mn\alpha(mn))$ time. For $d = 2$ we can construct $\overline{\mathcal{E}}$ and $\underline{\mathcal{E}}$ in time $O((mn)^2\alpha(mn))$ [11] and construct the overlay in time $O((mn)^2\alpha(mn)\log(mn))$ using Theorem 5. The following theorem summarizes the results of this section.

Theorem 6. *Let F and G be two polyhedral terrains in $\mathbb{R}^{d+1}$, with complexities m and n respectively. We can decide whether there exists a translation of the domain of F to within the interior of the domain of G, and we can compute a translation that minimizes the perpendicular distance between F and G in time $O((mn)^{d+\varepsilon})$ for any $\varepsilon > 0$. For $d = 1$ the running time is $O(mn \log(mn))$ and for $d = 2$ the running time is $O((mn)^2\alpha(mn)\log(mn))$.*

3.2 Hausdorff Distance

Given two polyhedral terrains F and G in $\mathbb{R}^{d+1}$, we consider the task to compute the translation that brings F into the smallest possible distance to G, according to the (directed or undirected) Hausdorff distance measure. We accomplish this by first solving the corresponding decision problem: For F and G as above, decide whether there exists a translation that brings F into Hausdorff distance at most δ of G, for a given parameter $\delta > 0$. For technical reasons, we assume that one of the terrains we have to match, say G, is continuously defined over a convex domain. In the full version of the paper [14] we prove our results for a more general class of terrains called δ-terrains (where δ is the threshold parameter in the decision procedure).

Let ρ be a metric in $\mathbb{R}^{d+1}$, and let $A, B \subseteq \mathbb{R}^d$ be two compact sets. The *directed Hausdorff distance* $\boldsymbol{\delta'}_{\mathrm{H}}(A, B)$ is defined as $\boldsymbol{\delta'}_{\mathrm{H}}(A, B) = \max_{x \in A} \min_{y \in B}$

$\rho(x, y)$. The *(undirected) Hausdorff distance* $\boldsymbol{\delta}_{\mathrm{H}}(A, B)$ is defined as $\boldsymbol{\delta}_{\mathrm{H}}(A, B) = \max\{\boldsymbol{\delta'}_{\mathrm{H}}(A, B), \boldsymbol{\delta'}_{\mathrm{H}}(B, A)\}$. The Hausdorff distance is a natural way to extend a metric ρ to the class of compact sets. Note that $\boldsymbol{\delta}_{\mathrm{H}}$ is indeed a metric, while $\boldsymbol{\delta'}_{\mathrm{H}}$ is not, since it is not symmetric. Nonetheless, the directed Hausdorff distance is often used in partial matching applications, where the task is to find a subset of the shape B that resembles the shape A the most.

Our algorithms assume that ρ is a convex polyhedral metric of constant description complexity that has the property that the $(d+1)$-dimensional unit ball vertically projects onto the d-dimensional unit ball (defined in terms of ρ). This is for example the case for the commonly used L_∞- and L_1-metrics. We call a metric that satisfies these assumptions *projectable*. For $d = 1$ our approach also works for the Euclidean metric.

Unfortunately, due to space limitations we have to omit all technical details of our treatment of matching under the Hausdorff distance. We state the main results without proof, which proceed by reducing the described decision problems to testing whether a certain sandwich region in an appropriately defined arrangement is empty. Details can be found in the full version of this paper [14].

Theorem 7. *Let F be a polyhedral set and G be a convex-domain polyhedral terrain in $\mathbb{R}^{d+1}$, with respective complexities m and n. We can test whether there exists a translation in $\mathbb{R}^{d+1}$ that brings F into directed Hausdorff distance at most δ of G in the following time:*

- $O(mn\log(mn))$ *when* $d = 1$ *and the underlying metric is projectable, and* $O(mn2^{\alpha(mn)}\log(mn))$ *when the metric is Euclidean.*
- $O(m^d n^{d^2} \alpha^{d+1}(m+n))$ *when* $d \geq 2$ *and the metric is projectable.*

For the undirected Hausdorff distance we obtain similar runtimes which are symmetric in m and n. In order to solve the optimization problem we apply the technique of parametric searching [17].

Theorem 8. *Let F and G be convex-domain polyhedral terrains in $\mathbb{R}^{d+1}$, with respective complexities m and n. Put $N = m + n$. We can compute a translation that minimizes the directed (resp., undirected) Hausdorff distance between F and G in time $O(m^{d+\varepsilon} n^{d^2+\varepsilon})$ (resp., $O(N^{d^2+d+\varepsilon})$), for any $\varepsilon > 0$ and $d \geq 1$, assuming that the underlying point metric is projectable.*

Acknowledgements

We are particularly grateful to Danny Halperin for pointing us to the work of de Berg et al. [9]. We would also like to thank Pankaj K. Agarwal and Micha Sharir for helpful discussions.

References

1. P. K. Agarwal, S. Har-Peled, M. Sharir, and Y. Wang. Hausdorff distance under translation for points, disks, and balls. In *Proc. Symp. Comp. Geom.*, 2003.

2. P. K. Agarwal, O. Schwarzkopf, and M. Sharir. The overlay of lower envelopes and its applications. *Discrete and Computational Geometry*, 15:1–13, 1996.
3. P. K. Agarwal and M. Sharir. Arrangements and their applications. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 49–119. Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000.
4. H. Alt and L. Guibas. Discrete geometric shapes: Matching, interpolation, and approximation – a survey. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*. Elsevier Science Publishers B.V. North-Holland, 2000.
5. J.-D. Boissonnat and K. Dobrindt. On-line construction of the upper envelope of triangles and surface patches in three dimensions. *Comput. Geom. Theory Appl.*, 5:303–320, 1996.
6. J.-D. Boissonnat and M. Yvinec. *Algorithmic Geometry*. Cambridge University Press, UK, 1998.
7. B. Chazelle. Cutting hyperplanes for divide-and-conquer. *Discrete Comput. Geom.*, 9(2):145–158, 1993.
8. K. L. Clarkson. A randomized algorithm for closest-point queries. *SIAM Journal on Computing*, 17:830–847, 1988.
9. M. de Berg, L. J. Guibas, and D. Halperin. Vertical decompositions for triangles in 3-space. *Discrete Comput. Geom.*, 15:35–61, 1996.
10. H. Edelsbrunner. The upper envelope of piecewise linear functions: Tight complexity bounds in higher dimensions. *Discrete and Comp. Geom.*, 4:337–343, 1989.
11. H. Edelsbrunner, L. Guibas, and M. Sharir. The upper envelope of piecewise linear functions: algorithms and applications. *Discr. Comp. Geom.*, 4:311–336, 1989.
12. J. Hershberger. Finding the upper envelope of n line segments in $O(n \log n)$ time. *Inform. Process. Lett.*, 33:169–174, 1989.
13. V. Koltun and M. Sharir. The partition technique for overlays of envelopes. *SIAM Journal on Computing*, to appear.
14. V. Koltun and C. Wenk. Matching polyhedral terrains using overlays of envelopes. http://www.cs.berkeley.edu/~vladlen/linear-overlays-full.zip.
15. C. Lee. Subdivisions and triangulations of polytopes. In J. E. Goodman and J. O'Rourke, editors, *Discr. Comp. Geom.*, pages 271–290. CRC Press, 1997.
16. J. Matoušek. Range searching with efficient hierarchical cuttings. *Discrete Comput. Geom.*, 10(2):157–182, 1993.
17. N. Megiddo. Applying parallel computation algorithms in the design of serial algorithms. *J. ACM*, 30(4):852–865, 1983.
18. J. Pach and M. Sharir. The upper envelope of piecewise linear functions and the boundary of a region enclosed by convex plates: combinatorial analysis. *Discrete Comput. Geom.*, 4:291–309, 1989.
19. M. Sharir. Almost tight upper bounds for lower envelopes in higher dimensions. *Discrete and Computational Geometry*, 12:327–345, 1994.
20. M. Sharir and P. K. Agarwal. *Davenport-Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, New York, 1995.
21. B. Tagansky. A new technique for analyzing substructures in arrangements of piecewise linear surfaces. *Discrete and Computational Geometry*, 16:455–479, 1996.
22. C. Wenk. *Shape Matching in Higher Dimensions*. Ph.D. thesis, Free University Berlin, Berlin, Germany, 2002.
23. A. Wiernik and M. Sharir. Planar realizations of nonlinear Davenport-Schinzel sequences by segments. *Discrete Comput. Geom.*, 3:15–47, 1988.