

Matching Polyhedral Terrains Using Overlays of Envelopes[‡]

Vladlen Koltun[†]

Carola Wenk[§]

May 24, 2004

Abstract

For a collection \mathcal{F} of d -variate piecewise linear functions of overall combinatorial complexity n , the lower envelope $\mathcal{E}(\mathcal{F})$ of \mathcal{F} is the pointwise minimum of these functions. The minimization diagram $\mathcal{M}(\mathcal{F})$ is the subdivision of \mathbb{R}^d obtained by vertically (i.e., in direction x_{d+1}) projecting $\mathcal{E}(\mathcal{F})$. The overlay $\mathcal{O}(\mathcal{F}, \mathcal{G})$ of two such subdivisions $\mathcal{M}(\mathcal{F})$ and $\mathcal{M}(\mathcal{G})$ is their superposition. We extend and improve the analysis of de Berg et al. [17] by showing that the combinatorial 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 $d \geq 2$, and $O(n^2 \alpha(n) \log n)$ when $d = 2$. We also describe an algorithm that constructs $\mathcal{O}(\mathcal{F}, \mathcal{G})$ in this time.

We apply these results to obtain efficient general solutions to the problem of matching two polyhedral terrains in higher dimensions under translation. That is, given two piecewise-linear terrains of combinatorial complexity n in \mathbb{R}^{d+1} , we wish to find a translation of the first terrain that minimizes its distance to the second, according to some distance measure. For the perpendicular distance measure, which we adopt from functional analysis since it is natural for measuring the similarity of terrains, we present a matching algorithm that runs in time $O(n^{2d+\varepsilon})$ for any $\varepsilon > 0$. Sharper running time bounds are shown for $d \leq 2$. 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. Before providing the necessary background to our work, let us define the basic terms that will be of use throughout the paper. 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 $\mathcal{E}(\mathcal{F})$

*A limited preliminary version of some of the results described in this paper has appeared in the second author's Ph.D. thesis [36].

[†]An extended abstract of this paper has appeared in [28]

[‡]Computer Science Division, University of California, Berkeley, CA 94720-1776, USA. vladlen@cs.berkeley.edu. Supported in part by NSF Grant CCR-01-21555.

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

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 $\mathcal{O}(\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_k)$ of envelopes of arrangements of multiple collections $\mathcal{F}_1, \dots, \mathcal{F}_k$. 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 various problems in this field; see Sharir and Agarwal [34] for an overview. The special setting of lower envelopes of piecewise linear functions has received considerable attention as it has multiple applications and is easier to analyze than the general setting of semi-algebraic functions; see Edelsbrunner et al. [19] and the references within. In particular, Edelsbrunner et al. [18, 19] and Pach and Sharir [32] 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.

The overlay of envelopes is a central structure in the theory of arrangements, and numerous problems in computational geometry can be reduced to algorithmic and combinatorial problems on overlays [2, 27]. 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 [27] 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. [17], who have studied the complexity of the vertical decomposition of an arrangement of a set of triangles in \mathbb{R}^3 . The vertical decomposition is a decomposition of an arrangement into cells of constant description complexity (that is, defined by a constant number of polynomial equations and inequalities of constant maximum degree), each of which is contained in a cell of $\mathcal{A}(\Gamma)$. This decomposition is obtained by refining the arrangement using auxiliary vertical surfaces.

Although the paper [17] 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 $\mathcal{O}(\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_d)$ is $\Omega(n^d \alpha^d(n))$, for collections $\mathcal{F}_1, \dots, \mathcal{F}_d$ 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. In the second part of the paper we investigate matching problems on polyhedral terrains, as introduced below. 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. The upper bounds and algorithms that we obtain in the first part of the paper allow us to present efficient algorithms for matching terrains. We also present results on terrain matching that do not directly rely on

our study of overlays.

We begin by introducing the considered problems. The comparison of geometric objects is a task that naturally arises in many application areas, such as for example 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 [10] 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 [10]. For example, a lot of attention has been directed to matching polygons in the plane [6, 8, 14, 16]. A translation that minimizes the Hausdorff distance between two polygons of complexity n in the plane, for example, can be computed in $O(n^4 \log^3 n)$ time, see [6]. An exception of sorts is the algorithm for computing the Hausdorff distance between two sets of k -dimensional simplices in \mathbb{R}^{d+1} for arbitrary fixed dimension d , described by Alt et al. [9] and Godau [20]. It runs in $O(n^{k+1+\varepsilon})$ randomized expected time, where n is the complexity of both sets of simplices, and $\varepsilon > 0$ is an arbitrarily small constant. This is however not a matching algorithm since it computes the distance between the two sets but does not allow to search for some transformation that minimizes it. 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 [36]. 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 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 L_2 -disks in the plane in $O(mn(m+n) \log^3(mn))$ time, which is a linear factor faster than the general result of Agarwal et al. [6], and for two sets of L_2 -balls in three dimensions in $O(m^2 n^2 (m+n) \log^3(mn))$ time. A different strategy is not to insist on finding an optimal transformation, but to find an approximately optimal transformation with a considerably faster algorithm. Approximation algorithms are not the focus of this paper, but we refer the interested reader to Alt et al. [8, 7] for reference-point based algorithms, and to Goodrich et al. [21] for algorithms based on a pinning strategy for two sets of points.

Terrains are a natural subset of shapes that have particularly many applications in $d+1 = 3$ dimensions, 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. One distance measure we consider is the *perpendicular distance*, which is an adaptation of the L_∞ Minkowski metric used in functional analysis. It naturally measures the similarity of two terrains (or functions) by assessing the largest height difference between any pair of vertically adjacent points on the terrains. We show that we can compute a translation of a terrain of complexity m which minimizes its perpendicular distance to a terrain of complexity n , in time $O((mn)^{d+\varepsilon})$ for any $\varepsilon > 0$. Our algorithm operates under the assumption that the domain of the first terrain, under some translation, is contained in the domain of the second. The algorithm reduces the matching problem to the computation of an overlay of envelopes. Sharper running time bounds are obtained for $d \leq 2$.

Another distance measure we consider is the popular Hausdorff distance [25], with its directed and undirected variants. Assuming that the terrains are continuously defined over a convex domain (our algorithms actually apply to a more general class of terrains, see Section 3 for details), 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 $\varepsilon > 0$. This is achieved by first reducing the corresponding decision problem of determining whether the Hausdorff distance of two terrains under translation is at most δ , for $\delta > 0$, to the problem of testing whether a certain sandwich region is empty, which can be solved by constructing this sandwich region. We then use parametric searching [31] to turn this decision algorithm into an optimization procedure. 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, see Section 3 for definitions.

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.

We first construct (in Theorem 2.1) two such collections \mathcal{F} and \mathcal{G} , such that the complexity of $\mathcal{O}(\mathcal{F}, \mathcal{G})$ is $\Omega(n^d \alpha^2(n))$. We then construct (in Theorem 2.2) d such collections \mathcal{F}_i , for $1 \leq i \leq d$, such that $\mathcal{O}(\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_d)$ has complexity $\Omega(n^d \alpha^d(n))$. Overlays of more than two minimization diagrams were studied, e.g., by Koltun and Sharir [27]. Note that when $d \geq 3$ an upper bound on the complexity of the overlay of two minimization diagrams does not extend to overlays of more than two diagrams.

When $d = 2$ both constructions are the same and are identical to the construction presented by de Berg et al. [17] (see Figure 1). Throughout the remainder of the paper, denote the axes in the $(d + 1)$ -dimensional space by x_1, \dots, 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 2.1. *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.

By the result of Wiernik and Sharir [37], there exists a collection of n line segments in the plane, such that the complexity of the lower envelope of their arrangement is $\Omega(n\alpha(n))$. 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.




Figure 1: Schematic illustration of the overlay in the lower bound construction for $d = 2$.

Consider an analogously constructed collection C_1 of strips, this time orthogonal to the x_2x_{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 $\mathcal{G} = \bigcup_{i=1}^{d-1} C_i$.

Let us verify that $\mathcal{O}(\mathcal{F}, \mathcal{G})$ has $\Omega(n^d \alpha^2(n))$ vertices. 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 $\mathcal{M}(C_2 \cup \dots \cup C_{d-1})$, 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, \dots, 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. \square

Theorem 2.2. *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 $1 \leq i \leq d$, let \mathcal{F}_i be a collection of d -dimensional strips orthogonal to the $x_i x_{d+1}$ -plane, constructed as follows. Consider, as above, a collection Γ of n segments, drawn in the $x_i x_{d+1}$ -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 $\mathcal{O}(\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_d)$ is easily seen to be as claimed. \square

Corollary 2.3. *For any $2 \leq k \leq d$, there are k collections of n d -simplices in \mathbb{R}^{d+1} , such that the complexity of the overlay of the k respective lower envelopes is $\Omega(n^d \alpha^k(n))$.*

Proof. A simple combination of the constructions in Theorems 2.1 and 2.2. \square

Remark. Theorem 2.1 and the earlier construction of de Berg et al. [17] dispel a belief, expressed, e.g., in [5], that the analysis of Edelsbrunner et al. [19] 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

In this section we prove a bound of $O(n^{d+\varepsilon})$ for any $\varepsilon > 0$ on the complexity of $\mathcal{O}(\mathcal{F}, \mathcal{G})$, for $d \geq 2$ (Theorem 2.6). Our proof relies on the concept of hierarchical cuttings [13, 30], and is based on the proof technique of de Berg et al. [17]. For the case $d = 2$ we provide a sharper upper bound of $O(n^2 \alpha(n) \log n)$ (Theorem 2.7) using the analysis technique of Tagansky [35]. This improves the result of de Berg et al. [17].

Before stating and proving the described theorems, 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.

We begin with two lemmas that will play a part in proving Theorem 2.6.

Lemma 2.4. *Let Γ be a collection of surfaces that are graphs of j -variate (possibly partially defined) functions (not necessarily of constant description complexity) in \mathbb{R}^{j+1} , such that Γ satisfies the following assumptions for some global constant c :*

- *The surfaces of Γ are in general position.*
- *Let $v(\Gamma)$ be the number of vertices of the lower envelope of the arrangement $\mathcal{A}(\Gamma)$. Then the combinatorial complexity of this lower envelope is bounded by $O(v(\Gamma) + n^j)$.*
- *Every $(j + 1)$ -tuple of surfaces of Γ intersects in at most c points.*
- *A vertical projection of a k -dimensional feature ($0 \leq k \leq j$) of $\mathcal{A}(\Gamma)$ onto the hyperplane $x_{j+1} = 0$ intersects an analogous projection of a $(j - k)$ -dimensional feature of $\mathcal{A}(\Gamma)$ in at most c points.*
- *For any feature φ of $\mathcal{A}(\Gamma)$ and any $1 \leq k \leq j + 1$, the number of local x_k -maxima is at most c .*
- *For any feature φ of $\mathcal{A}(\Gamma)$ and any x_{j+1} -vertical line ℓ , $\varphi \cap \ell$ consists of at most one point (i.e., φ is a graph of an appropriately defined function).*
- *For every k -dimensional feature φ , $0 \leq k < j$, of the arrangement $\mathcal{A}(\Gamma)$, let Γ^φ be the collection of surfaces $\{\gamma \cap W_\varphi \mid \gamma \in \Gamma, W_\varphi = \varphi \times x_{j+1}\}$. In other words, Γ^φ is obtained by restricting the surfaces of Γ to within the x_{j+1} -vertical wall spanned by φ . Then $|\Gamma^\varphi| = O(n)$ and the above conditions hold for the collection Γ^φ as well.*

Then the combinatorial complexity of the lower envelope of the arrangement $\mathcal{A}(\Gamma)$ is $O(n^{j+\varepsilon})$ for any $\varepsilon > 0$.

Proof. The lemma is a direct consequence of the proof technique of Sharir [33] for the complexity of the lower envelope of an arrangement of a collection of semi-algebraic surfaces of constant description complexity. Although the surfaces of Γ may not have bounded description complexity, the conditions listed in the statement of the lemma are sufficient for Sharir's proof to work. \square

Lemma 2.5. *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 .*

The reader is invited to consider the case $j = 1$, when ∂P is a convex curve (i.e., the boundary of a two-dimensional convex set) and $\mathcal{F}_{\partial P}$ is a collection of curves on the two-dimensional vertical surface spanned by ∂P .

Proof. Notice that since ∂P is a closed surface, the x_{d+1} -vertical surface $\Delta = \partial P \times x_{d+1}$ has cylindrical topology. For technical reasons, however, we assume that it is homeomorphic to \mathbb{R}^{j+1} . This can be accomplished by a mapping that cuts Δ along a vertical j -dimensional hyperplane and unfolds it appropriately. Such mapping can easily be applied without disturbing the above topological properties of any collection $\mathcal{F}_{\partial P}$ contained in Δ as above.

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 ($0 \leq k \leq j$) 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. It is easy to verify that $\mathcal{F}_{\partial P}$ satisfies the assumptions of Lemma 2.4, which proves the current lemma. \square

The proof of Theorem 2.6 below relies on the concept of *efficient hierarchical cuttings*, which were shown to exist by Chazelle [13] (see also Matoušek [30]). 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 $\Xi = \Xi_0, \Xi_1, \dots, \Xi_k$ 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 $1 \leq i \leq k$, 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 $\text{parent}(s)$. With hierarchical cuttings at hand, we state and prove the main result of this section.

Theorem 2.6. *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. 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 $d+1 \geq 3$, 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 $1 \leq j \leq d-1$. We denote the vertices of the latter type as j -vertices. (E.g., for $d+1=3$, one type of vertices of $\mathcal{O}(\mathcal{F}, \mathcal{G})$ is referred to as 1-vertices.) It is known that the number of vertices of $\mathcal{M}(\mathcal{F})$ and $\mathcal{M}(\mathcal{G})$ is $O(n^d \alpha(n))$ [18]. To bound the complexity of $\mathcal{O}(\mathcal{F}, \mathcal{G})$ it remains to analyze the number of j -vertices, for $1 \leq j \leq d-1$.

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} . (The projection of each of the $2n$ simplices defines $d+1$ hyperplanes.) 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 $\Xi = \Xi_0, \Xi_1, \dots, \Xi_k$, 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 Ξ , 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^c be the set of simplices of \mathcal{F} whose projections contain s , but do not contain $\text{parent}(s)$. Define \mathcal{G}_s^\times and \mathcal{G}_s^c analogously with respect to \mathcal{G} , and put $\Gamma_s^\times = \mathcal{F}_s^\times \cup \mathcal{G}_s^\times$ and $\Gamma_s^c = \mathcal{F}_s^c \cup \mathcal{G}_s^c$. Also set $\mathcal{F}_s = \mathcal{F}_s^c \cup \mathcal{F}_s^\times$, 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 $\text{def}(v)$.

We claim that for every v as above, there exists a simplex s belonging to some level of Ξ , such that $\text{def}(v)$ is contained in Γ_s and at least one of the simplices of $\text{def}(v)$ belongs to Γ_s^c . 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, \dots, s_k contains s_{i+1} (unless, of course, $i = k$). Since s_0 is the whole space \mathbb{R}^d , $\text{def}(v)$ 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 $\text{def}(v)$ is not contained in $\Gamma_{s_i}^\times$, but all of them are contained in $\Gamma_{\text{parent}(s_i)}^\times$ (and are thus not contained in $\Gamma_{\text{parent}(s_i)}^c$). All the simplices of $\text{def}(v)$ that are not contained in $\Gamma_{s_i}^\times$ are thus contained in $\Gamma_{s_i}^c$, 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^c , 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^c . 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 $\mathcal{E}(\mathcal{F}_s^c)$. Notice that the simplices of \mathcal{F}_s^c are totally defined over s , and thus the envelope $\mathcal{E}(\mathcal{F}_s^c)$ 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 $\mathcal{E}(\mathcal{F}_s^c)$ is thus part of a j -dimensional convex polytope, defined as the intersection of $\mathcal{E}(\mathcal{F}_s^c)$ 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 2.5 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 $\text{parent}(s)$, 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

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

for any $\varepsilon > 0$. Summing over all simplices s , the number of j -vertices of $\mathcal{O}(\mathcal{F}, \mathcal{G})$ is

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

for any $\varepsilon > 0$. Noticing that the bound does not depend on j completes the proof. \square

Remark. We have devised an alternative proof for Theorem 2.6 using the technique of counting schemes [22, 33]. While it produces the same upper bound, it is somewhat more involved than the above proof using hierarchical cuttings. Another advantage of the above proof is that the factor of ε in the exponent of the resulting bound is not inherent in the proof technique. Instead, it comes from the bound on the complexity of the lower envelope of surfaces in higher dimensions [33]. If the latter bound is ever sharpened (say, by replacing the n^ε -factor by a polylogarithmic one), the bound in Theorem 2.6 will be automatically refined.

Theorem 2.7. *Given two collections \mathcal{F} and \mathcal{G} of n triangles in three dimensions, the complexity of $\mathcal{O}(\mathcal{F}, \mathcal{G})$ is $O(n^2 \alpha(n) \log n)$.*

Proof. Since the overlay $\mathcal{O}(\mathcal{F}, \mathcal{G})$ is a planar map, it is sufficient to bound the number of its vertices. Every such vertex is either a vertex of one of the minimization diagrams $\mathcal{M}(\mathcal{F})$, $\mathcal{M}(\mathcal{G})$, or an intersection of an edge of $\mathcal{M}(\mathcal{F})$ with an edge of $\mathcal{M}(\mathcal{G})$ [2]. There are $O(n^2 \alpha(n))$ vertices in $\mathcal{M}(\mathcal{F})$ and $\mathcal{M}(\mathcal{G})$ [32], so we are left with intersections of pairs of edges. An edge of $\mathcal{M}(\mathcal{F})$ is a projection of an edge of $\mathcal{E}(\mathcal{F})$, which necessarily lies either on the boundary of a triangle of \mathcal{F} or on the intersection of two triangles of \mathcal{F} . A standard argument shows that the number of vertices of $\mathcal{O}(\mathcal{F}, \mathcal{G})$ that lie on edges of $\mathcal{M}(\mathcal{F})$ of the former type is $O(n^2 \alpha(n))$ [2, 17, 32]. The situation is analogous for vertices that lie on edges of $\mathcal{M}(\mathcal{G})$ of this type. It thus remains to count vertices that lie on intersections of edges defined by intersections of pairs of triangles.

Define the *level* of a point $\mathbf{x} \in \mathbb{R}^3$ in the arrangement $\mathcal{A}(\mathcal{F})$ (resp., $\mathcal{A}(\mathcal{G})$) to be the number of triangles of \mathcal{F} (resp., of \mathcal{G}) that lie vertically below \mathbf{x} . Consider a crossing point v between a projection of an edge of $\mathcal{A}(\mathcal{F})$ defined by two triangles of \mathcal{F} and a projection of an edge of $\mathcal{A}(\mathcal{G})$ defined by two triangles of \mathcal{G} , such that the point that lies on an edge of $\mathcal{A}(\mathcal{F})$ (resp., of $\mathcal{A}(\mathcal{G})$) and projects to v has level a (resp., b). Put $c = a + b$. v is said to be a *c-level overlay vertex*. (Note that our definition of *c-level overlay vertices* is non-standard in that it does not encompass vertices incident to projections of boundaries of triangles of \mathcal{F} or \mathcal{G} .) It is easy to see that 0-level overlay vertices are exactly the vertices of $\mathcal{O}(\mathcal{F}, \mathcal{G})$ that we need to consider. To bound their number we use the analysis technique introduced by Tagansky [35]. It is a general method of bounding the number of 0-level configurations by charging them to 1-level configurations, and applies in a wide variety of settings. The following lemma formulates the technique's implication in our particular scenario. The lemma is directly implied by Tagansky's analysis [35], and is thus given without proof.

Lemma 2.8. *Assume that each 0-level overlay vertex, aside from a collection of $O(n^2\alpha(n))$ "excess" vertices, can be charged to four 1-level overlay vertices, such that every 1-level overlay vertex is charged at most twice. Then the complexity of $\mathcal{O}(\mathcal{F}, \mathcal{G})$ is $O(n^2\alpha(n) \log n)$.*

The rest of the proof of the theorem is devoted to showing how the charging described in the lemma can be performed. Consider a 0-level overlay vertex v . Denote the envelope edges whose projections define v by e_f and e_g , where $e_f \subseteq F_1 \cap F_2$ and $e_g \subseteq G_1 \cap G_2$, for $F_1, F_2 \in \mathcal{F}$ and $G_1, G_2 \in \mathcal{G}$.

Consider the vertical wall W_f (resp., W_g) spanned by e_f (resp., by e_g), defined as the union of all vertical lines passing through e_f (resp., e_g). Let $\tilde{\mathcal{G}}$ (resp., $\tilde{\mathcal{F}}$) be the collection of segments $G_i \cap W_f$, for all $G_i \in \mathcal{G}_f$ (resp., $F_i \cap W_g$, for all $F_i \in \mathcal{F}$). It is clear that the point $v_g \in e_g$ (resp., $v_f \in e_f$) that projects to v lies on a vertex of $\mathcal{E}(\tilde{\mathcal{G}})$ (resp., of $\mathcal{E}(\tilde{\mathcal{F}})$). Consider sliding a point u on one of the two 1-level edges of $\mathcal{A}(\tilde{\mathcal{G}})$ (resp., of $\mathcal{A}(\tilde{\mathcal{F}})$) adjacent to v_g (resp., v_f), away from v_g (resp., v_f). Denote this edge by a . One of the following events has to occur during the sliding, see Figure 2:

- (i) u reaches the end-point of a .
- (ii) u passes above an end-point of some segment of $\tilde{\mathcal{G}}$ (resp., of $\tilde{\mathcal{F}}$).
- (iii) The projection of u onto e_f (resp., onto e_g) reaches an end-point of e_f (resp., of e_g).
- (iv) The projection of u onto e_f (resp., onto e_g) passes above an end-point of a segment of $\tilde{\mathcal{F}}$ (resp., of $\tilde{\mathcal{G}}$).
- (v) u reaches a 1-level vertex of $\mathcal{A}(\tilde{\mathcal{G}})$ (resp., of $\mathcal{A}(\tilde{\mathcal{F}})$).

When any of these events occurs, the sliding is stopped. We claim that the projection of u does not pass any vertex of $\mathcal{O}(\mathcal{F}, \mathcal{G})$ or a 1-level overlay vertex before an event of one of the above types occurs. Indeed, a vertex of $\mathcal{O}(\mathcal{F}, \mathcal{G})$ is either either a vertex of $\mathcal{M}(\mathcal{F})$ or of $\mathcal{M}(\mathcal{G})$ or an intersection of an edge of $\mathcal{M}(\mathcal{F})$ with an edge of $\mathcal{M}(\mathcal{G})$. u cannot pass a vertex of the second of these kinds as that would be a violation of the general position assumption. A vertex of the first kind (a vertex of $\mathcal{M}(\mathcal{F})$) is precisely an event of type (iii). A vertex of the third kind can only occur as an event of type (ii). Finally, a 1-level overlay vertex can only occur as an event of type (v), which proves our claim.




Figure 2: The five types of events that can happen during the sliding inside the wall W_f spanned by e_f .

We perform the above sliding in both possible directions away from v_g (resp., v_f). Overall, we have described four sliding processes. It is easy to see that if an event of type (v) is the first to occur during a particular sliding process then it corresponds to a 1-level overlay vertex. If all the four slidings terminate at an event of type (v), we charge the original 0-level overlay vertex v to the four 1-level overlay vertices corresponding to these type-(v) events. Otherwise, if any of the four slidings reach an event of one of the types (i–iv), v is branded an “excess” vertex. We omit the easy argument showing that the number of “excess” vertices is $O(n^2\alpha(n))$, and concentrate on showing that each 1-level overlay vertex is charged at most twice.

Consider a 1-level overlay vertex v' , and let e'_f (resp., e'_g) be the edge of $\mathcal{A}(\mathcal{F})$ (resp., of $\mathcal{A}(\mathcal{G})$) defining it. Let v'_f (resp., v'_g) be the point on e'_f (resp., e'_g) that projects to v' . Assume, without loss of generality, that v'_f has level 0 in $\mathcal{A}(\mathcal{F})$, while v'_g has level 1 in $\mathcal{A}(\mathcal{G})$. Thus there exists a triangle $G_3 \in \mathcal{G}$ that lies below v'_g .

There are four directions of sliding along which v' can be charged, described below. We can slide back along each of these directions, away from v' and towards the potential charging vertex. v' is charged along a specific such direction of backward sliding if and only if we reach a 0-level overlay vertex v along this direction without any event of type (i–iv) happening first. We will now show that this can be the case along at most two directions of backward sliding.

Let us first describe the four possible directions along which v' can be charged. Two directions are obtained by considering the cross-section of $\mathcal{A}(\mathcal{F})$ within the vertical wall W_g spanned by e'_g , and sliding along the two *bottom edges* adjacent to v'_f in this cross-section (see Figure 3(a)). Considering bottom edges is sufficient since by sliding away from v' along a top edge we clearly cannot reach a 0-level overlay vertex without an event of type (i–iv) happening on the way. The other two directions are symmetrically obtained by considering the two bottom edges adjacent to v'_g in the cross-section of $\mathcal{A}(\mathcal{G})$ within the vertical wall spanned by e'_f .




Figure 3: Notation related to the process of sliding a point u backwards from v'_f towards v_f is introduced in (a). Since G_3 lies below e'_g when the sliding begins, but does not lie below e'_g when u reaches v_f , the projection of u onto e'_g has to either slide above the boundary of G_3 , as shown in (b), or pass the intersection point $e'_g \cap G_3$, as shown in (c).

We claim that v' cannot be charged along the former two directions of sliding. Indeed, consider sliding a point u away from v'_f along one of the two adjacent bottom edges, inside the wall W_g spanned by e'_g . In the beginning of the sliding, when $u = v'_f$, the triangle G_3 lies below the the projection of u onto e'_g . If v' is charged along the currently considered sliding direction, the sliding of u has to terminate at a vertex v_f of $\mathcal{E}(\mathcal{F})$, such that no triangle of \mathcal{G} , including G_3 , lies below the projection of v_f onto e'_g . Thus, the projection of u onto e'_g has to slide above the boundary of G_3 (Figure 3(b)) or pass the intersection point $e'_g \cap G_3$ (Figure 3(c)). In the latter case, however, the edge e'_g ends at the intersection point $e'_g \cap G_3$ and the former case is exactly an events of type (iv), which means that v' cannot be charged along the considered sliding direction.

We have shown that no 1-level overlay vertex v' can be charged along more than two sliding directions, and thus any such vertex is charged at most twice. This completes the proof of Theorem 2.7. \square

2.3 Algorithms

This section presents an algorithm for constructing $\mathcal{O}(\mathcal{F}, \mathcal{G})$. It can be adapted to constructing a single envelope $\mathcal{E}(\mathcal{F})$ or the sandwich region $\mathcal{S}(\mathcal{F}, \mathcal{G})$, which will be useful in the terrain matching algorithms described in Section 3.

Theorem 2.9. *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 of the construction of $\mathcal{O}(\mathcal{F}, \mathcal{G})$ is $O(n^2 \alpha(n) \log n)$.*

Proof. As in the proof of Theorem 2.6, 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 [18, 32]. 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 2.6 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 $\varepsilon > 0$. It is easy to see that each cell in the resulting refined subdivision is convex. It can thus be easily decomposed into simplices using the bottom-vertex simplicial decomposition [15, 29].

The ability to represent $\mathcal{O}(\mathcal{F}, \mathcal{G})$ as a convex subdivision that is decomposed with the bottom-vertex decomposition allows us to construct this overlay using a randomized incremental algorithm that utilizes a conflict graph. A general description of this standard approach is given, e.g., in [12, 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 [12, Section 5.2], the expected running time of the construction algorithm is

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

where $f(r)$ denotes the maximal complexity of the overlay of envelopes of two sets of r simplices overall. Theorem 2.6 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 2.7.

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. \square

Remark. 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 [11]. 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 [34, 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. It turns out that the above structures occur naturally in the two matching problems we consider: For the matching of terrains under translations using the Hausdorff distance, properties of terrains allow us to substitute the brute-force approach of constructing a complete arrangement in the transformation space [36] by the construction of only a certain sandwich region in this space, see below for definitions. In the case of the perpendicular distance, which is a natural distance measure for terrains as defined below, we reduce the matching problem to overlaying two envelopes that arise directly from the definition of the distance measure. The results for the perpendicular distance are given in Subsection 3.1, and the results for the directed and undirected Hausdorff distance are given in Subsection 3.2. We begin by defining our notation for terrains.

A (k -dimensional) *terrain* F in \mathbb{R}^{d+1} is the graph $F = \{(x, f(x)) \mid x \in D_f\}$ of a k -variate function $f : D_f \rightarrow \mathbb{R}$, $0 \leq k \leq d$, 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 piecewise linear (x_1, \dots, x_d) -monotone surface, or in other words 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 [15, 29]. 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*):

$$\delta_{\perp}(F, G) := \sup_{x \in D_f} |f(x) - g(x)|,$$

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 . Our task is to translate F in \mathbb{R}^{d+1} such that $D_f \subseteq D_g$ and the perpendicular distance between F and G is minimized. We assume for now that at least one such translation exists; we will later see how this can be verified.

We consider a translation $t' = (t_1, \dots, t_d, t_{d+1}) \in \mathbb{R}^{d+1}$ to be composed of a translation $t = (t_1, \dots, t_d) \in \mathbb{R}^d$ 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

$$\delta_{\perp}(F + (t^*, t_{d+1}^*), G) = \min_{\substack{t \in \mathbb{R}^d \\ D_f + t \subseteq D_g}} \min_{t_{d+1} \in \mathbb{R}} \delta_{\perp}(F + (t, t_{d+1}), G) \quad (1)$$

Reformulating $\delta_{\perp}(F + t', G)$ produces

$$\begin{aligned} \delta_{\perp}(F + t', G) &= \max_{x \in D_f + t \subseteq D_g} |f(x - t) - g(x) + t_{d+1}| \\ &= \max\left\{ \max_{x \in D_f \subseteq D_g - t} (g(x + t) - f(x)) - t_{d+1}, t_{d+1} - \min_{x \in D_f \subseteq D_g - t} (g(x + t) - f(x)) \right\} \\ &= \max\left\{ \max_{x \in D_f \subseteq D_g - t} h_x(t) - t_{d+1}, t_{d+1} - \min_{x \in D_f \subseteq D_g - t} h_x(t) \right\} \end{aligned}$$

with $h_x(t) := g(x + t) - f(x)$. Then for each fixed $t \in \mathbb{R}^d$, $\max_{x \in D_f} h_x(t)$ is the pointwise maximum of the functions h_x , for all $x \in D_f$. Observe that the condition $D_f \subseteq D_g - t$ is equivalent to $t \in D_h$, where $D_h := \overline{D_g} \oplus (-D_f)$. 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:

$$\begin{aligned} \overline{h} : D_h &\longrightarrow \mathbb{R}; & t &\mapsto \max_{x \in D_f} h_x(t) \\ \underline{h} : D_h &\longrightarrow \mathbb{R}; & t &\mapsto \min_{x \in D_f} h_x(t). \end{aligned}$$

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 $-F := \{(x, -f(x)) \mid x \in D_f\}$.

Lemma 3.1. *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 .*

Proof.

$$G \oplus (-F) = \bigcup_{z \in D_g} \bigcup_{x \in D_f} \{(z - x, g(z) - f(x))\} = \bigcup_{t \in \mathbb{R}^d} \bigcup_{\substack{x \in D_f \\ x + t \in D_g}} \{(t, g(x + t) - f(x))\}$$

The lemma now follows directly from the definitions of \overline{H} and \underline{H} . □

Reformulating (1) we thus have

$$\delta_{\perp}(F + (t^*, t_{d+1}^*), G) = \min_{t \in D_h} \min_{t_{d+1} \in \mathbb{R}} \max\{\bar{h}(t) - t_{d+1}, t_{d+1} - \underline{h}(t)\}, \quad (2)$$

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

$$\frac{\bar{h}(t^*) - \underline{h}(t^*)}{2} = \min_{t \in D_h} \frac{\bar{h}(t) - \underline{h}(t)}{2} \quad (3)$$

$$\text{and } t_{d+1}^* = (\bar{h}(t^*) + \underline{h}(t^*)) / 2. \quad (4)$$

This leads to the following algorithm:

1. Compute $D_h = \overline{D_g \oplus (-D_f)}$. We represent D_g by a collection of boundary $(d-1)$ -simplices that make up ∂D_g . With each simplex we associate information specifying to which of its sides lies D_g . The complement $\overline{D_g}$ is represented by the same collection of simplices with the side information reversed. We also represent $(-D_f)$ in a similar manner. To compute the Minkowski sum $\overline{D_g} \oplus (-D_f)$ we compute, for each simplex γ in the representation of $\overline{D_g}$ and each simplex ϕ in the representation of $(-D_f)$, the Minkowski sum $\gamma \oplus \phi$. Overall we process $O(mn)$ pairs, and the computation for one pair takes constant time as it reduces to computing the convex hull of a constant number of points (the vector sums of the vertices of γ and ϕ). Now consider the arrangement \mathcal{A} of the $O(mn)$ hyperplanes spanned by the boundary faces of the individual Minkowski sums. These sums, and consequently the hyperplanes they span, inherit the side information from the original simplices. The set $\overline{D_g} \oplus (-D_f)$ is a union of a number of cells in \mathcal{A} . These cells can be detected and marked by a simple process that traverses the arrangement and uses the side information associated with the hyperplanes. We omit the details. The complement $\overline{D_g} \oplus (-D_f) = D_h$ is simply the union of the unmarked cells in \mathcal{A} . This arrangement can be computed and traversed in time $O((mn)^d)$, and this is also the time taken by the overall algorithm for constructing D_h .
2. Compute the overlay of the upper envelope $\bar{\mathcal{E}}$ and the lower envelope $\underline{\mathcal{E}}$ of $G \oplus (-F)$, and clip it to within D_h . Since f and g are piecewise linear, $G \oplus (-F)$ is polyhedral. Hence both envelopes are also piecewise linear, or in other words polyhedral terrains. We view $G \oplus (-F)$ as the union of the individual Minkowski sums $\gamma \oplus \phi$ for all d -simplices ϕ in $(-F)$ and γ in G . However, we do not have to compute the whole union $G \oplus (-F)$, since we only need the upper and lower envelopes of the collection of the individual Minkowski sums. $\gamma \oplus \phi$ has constant complexity and can be computed in constant time. Its surface can be partitioned into $O(1)$ d -simplices. There are nm different (γ, ϕ) -pairs, hence altogether we have $O(mn)$ d -simplices in \mathbb{R}^{d+1} . Their lower and upper envelopes have complexity $O((mn)^d \alpha(mn))$ [18].

Define $h^- : D_h \rightarrow \mathbb{R}$ as $h^- := (\bar{h} - \underline{h}) / 2$. Let $H^- := \{(x, h^-(x)) \mid x \in D_h\}$ be its graph (a polyhedral terrain), and M_{h^-} the simplicial partition of its domain. Lemma 3.1 implies that M_{h^-} is the overlay of $\bar{\mathcal{E}}$ and $\underline{\mathcal{E}}$ restricted to D_h . We thus compute the overlay of the minimization diagrams of a lower and an upper envelope of a set of $O(mn)$ d -simplices, additionally superimposed with the $O(mn)$ hyperplanes defining D_h . Notice that in the proofs of Theorem 2.6 and Theorem 2.9 the overlay of the envelopes is additionally overlaid with a set of hyperplanes, and for the correctness of

the proofs only the number of the hyperplanes matters. Thus, Theorem 2.6 implies that the complexity of this overlay is $O((mn)^{d+\varepsilon})$ for any $\varepsilon > 0$, and Theorem 2.9 states that it can be constructed within a matching time bound.

3. For every vertex v in M_{h^-} , compute $(\bar{h}(v) - \underline{h}(v))/2$. The function $(\bar{h}(t) - \underline{h}(t))/2$ is linear within each cell C of M_{h^-} . Therefore, the global minimum t^* that minimizes this function, as described in (3), is necessarily reached at a vertex of M_{h^-} . In order to compute t^* , it thus suffices to iterate over all vertices in M_{h^-} . A complete combinatorial representation of M_{h^-} was computed in Step 2 above, and can be used to examine the vertices in time proportional to their number, which is $O((mn)^{d+\varepsilon})$ for any $\varepsilon > 0$. We thus obtain t^* , which we can plug into (4) to get t_{d+1}^* . This yields the complete description of the minimizing transformation (t^*, t_{d+1}^*) .

Notice that if $D_h = \emptyset$ then there exists no translation such that the translated domain of F is contained in the domain of G . This is detected at the first stage of the algorithm. Overall, the described algorithm runs in time $O((mn)^{d+\varepsilon})$, for $d \geq 2$. For $d = 1$ we can compute $\bar{\mathcal{E}}$ and $\underline{\mathcal{E}}$ in time $O(mn \log(mn))$ [24]. In this case, each domain is a set of $O(mn\alpha(mn))$ intervals on a line, such that the overlay has the same complexity and can be computed with a simple sweep in $O(mn\alpha(mn))$ time. It can be clipped to within D_h by an additional sweep. Hence for $d = 1$ the algorithm runs in time $O(mn \log(mn))$. Since for $d = 2$ we can construct $\bar{\mathcal{E}}$ and $\underline{\mathcal{E}}$ in time $O((mn)^2\alpha(mn))$ [19] and construct the overlay in time $O((mn)^2\alpha(mn) \log(mn))$ (Theorem 2.9), the running time for $d = 2$ is $O((mn)^2\alpha(mn) \log(mn))$. The following theorem summarizes the results of this section.

Theorem 3.2. *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

In this subsection we consider matching terrains according to the directed and the undirected Hausdorff distance. Specifically, given two polyhedral terrains F and G in \mathbb{R}^{d+1} , we present algorithms that 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 $\delta > 0$. Stated equivalently, the decision problem is to determine whether the minimum Hausdorff distance between F and G , under translations, is at most δ . We then plug this decision procedure into the general technique of *parametric searching* [31] to derive an optimization algorithm that computes the minimum Hausdorff distance. This last step is discussed in Section 3.2.5, prior to which we concentrate on the design of the decision algorithm.

For technical reasons, we assume throughout this subsection that one of the terrains we have to match, say G , is continuously defined over a convex domain (we will call such a terrain a *convex-domain terrain*). Actually, it is easy to verify that the decision algorithms we derive, although stated under the assumption that G is a convex-domain terrain, hold

also for a more general class of terrains called δ -terrains (where δ is the threshold parameter in the decision procedure), defined below. We show in Lemma 3.4 that a convex-domain terrain G is a δ -terrain for any $\delta > 0$.

Definition 3.1. For a given $\delta > 0$, and a given metric ρ , a terrain $G = \{(x, g(x)) \mid x \in D_g\}$ in \mathbb{R}^{d+1} is called a δ -terrain iff $g(x)$ is continuous over each connected component of D_g and the intersection of G^δ with an arbitrary line in direction x_{d+1} is either an interval or empty. For $A \subseteq \mathbb{R}^d$, the δ -neighborhood of A is $A^\delta := A \oplus \mathbf{B}_{d+1}^\delta$, where \mathbf{B}_{d+1}^δ is a closed ball of radius δ ; the shape of this ball is determined by the underlying metric ρ .

We proceed with some basic definitions. 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* $\vec{\delta}_H(A, B)$ is defined as $\vec{\delta}_H(A, B) := \max_{x \in A} \min_{y \in B} \rho(x, y)$. The *(undirected) Hausdorff distance* $\delta_H(A, B)$ is defined as $\delta_H(A, B) := \max\{\vec{\delta}_H(A, B), \vec{\delta}_H(B, A)\}$. The Hausdorff distance is a natural way to extend a metric ρ to the class of compact sets. Note that δ_H is indeed a metric, while $\vec{\delta}_H$ 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.

We start by considering the directed Hausdorff distance. Let a parameter $\delta > 0$, a polyhedral set F of complexity m (an arbitrary collection of m simplices; not necessarily a terrain), and a convex-domain terrain $G = \{(x, g(x)) \mid x \in D_g\}$ of complexity n in \mathbb{R}^{d+1} be given. We address the problem to decide whether there exists a $t \in \mathbb{R}^{d+1}$ such that $\vec{\delta}_H(F + t, G) \leq \delta$. This is equivalent to $F + t \subseteq G^\delta$, which in turn can be expressed as $\overline{G^\delta \oplus (-F)} \neq \emptyset$. We consequently define *the feasible region* I :

$$I := \overline{G^\delta \oplus (-F)} \quad (5)$$

The following lemma will be instrumental in the sequel.

Lemma 3.3. Let $G = \{(x, g(x)) \mid x \in D_g\}$ be a δ -terrain in \mathbb{R}^{d+1} for a projectable metric. Then there exist terrains G_u, G_l with domain D_g^δ , and interior-disjoint sets $U, L \subseteq \mathbb{R}^{d+1}$ contained in the vertical closed cylinder $D_g^\delta \times \mathbb{R}$ above D_g , such that $\overline{G^\delta} \cap (D_g^\delta \times \mathbb{R}) = U \cup L$, and the lower envelope of U is G_u , and the upper envelope of L is G_l .

Proof. Let G_u (resp., G_l) be the set of all upper (resp., lower) endpoints of all line segments which are obtained as the intersection of a vertical line (i.e., in direction x_{d+1}) with G^δ . Endpoints of degenerate segments that consist of one point only are assigned to G_u and G_l . Let U be the closed set of all points vertically above or on G_u , and let L be the closed set of all points vertically below or on G_l . By construction the claimed envelope properties hold. Also, U and L are interior-disjoint, and $U \cup L$ equals $\overline{G^\delta} \cap (D_g^\delta \times \mathbb{R})$. \square

3.2.1 Properties of δ -terrains

We now digress slightly from our consideration of matching problems to prove some structural results concerning δ -terrains.

Lemma 3.4. *Every convex-domain terrain in \mathbb{R}^{d+1} is a δ -terrain for any $\delta > 0$, with respect to any convex polyhedral metric or the Euclidean metric.*

Proof. Let G be a convex-domain terrain, and let $\delta > 0$. Let l be an arbitrary vertical line, and assume for the sake of contradiction that its intersection with G^δ consists of at least two intervals. Let the two lowermost intervals be $[a, b]$ and $[c, d]$. Since interval endpoints lie on the boundary of G^δ , there can be no point of G in $\mathring{\mathbf{B}}_{d+1}^\delta(a) \cup \mathring{\mathbf{B}}_{d+1}^\delta(b) \cup \mathring{\mathbf{B}}_{d+1}^\delta(c) \cup \mathring{\mathbf{B}}_{d+1}^\delta(d)$, where $\mathring{\mathbf{B}}_{d+1}^\delta(a)$ is the open δ -ball in \mathbb{R}^{d+1} centered at a . Now let $p, q \in G$ be such that $\mathbf{B}_{d+1}^\delta(p) \cap [a, b] \neq \emptyset$ and $\mathbf{B}_{d+1}^\delta(q) \cap [c, d] \neq \emptyset$; these points are guaranteed to exist. Both p and q are contained in the cylinder $l \oplus \mathbf{B}_{d+1}^\delta$.

Let p' (resp., q') be the projection of p (resp., q) onto the domain D_g of G . By the convexity of D_g the line segment s between p' and q' has to be contained in D_g . Since G is a continuous terrain there must exist a path π on G from p to q whose projection onto D_g is s . Thus π is also contained in $l \oplus \mathring{\mathbf{B}}_{d+1}^\delta$. However, observe that $\mathring{\mathbf{B}}_{d+1}^\delta(b)$ divides the cylinder $l \oplus \mathring{\mathbf{B}}_{d+1}^\delta$ into two parts, and that p lies below $\mathring{\mathbf{B}}_{d+1}^\delta(b)$ while q above $\mathring{\mathbf{B}}_{d+1}^\delta(b)$. Thus π , and consequently G , have to intersect $\mathring{\mathbf{B}}_{d+1}^\delta(b)$, which is a contradiction. \square

Since the domain of a terrain in two dimensions is one-dimensional we know that for $d = 1$ the result of Lemma 3.4 holds for any connected domain.

It is tempting to hope that the complexity of G^δ for a δ -terrain G is only linear, at least if G is defined over a well-behaved two-dimensional domain in \mathbb{R}^3 . However the following lemma shows that this complexity can be quadratic, even if the domain is convex.

Lemma 3.5. *Let G be a terrain of complexity n in \mathbb{R}^3 , and let $\delta > 0$. Then the complexity of G^δ is $O(n^{2+\varepsilon})$ for any $\varepsilon > 0$. A lower bound for this complexity is $\Omega(n^2)$, and it holds even if G is a convex-domain terrain.*

Proof. G consists of $O(n)$ pairwise interior-disjoint triangles, segments, and points. Agarwal and Sharir [4] have shown that the complexity and the computation time for the δ -neighborhood for such a collection of objects is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. For the lower bound example we construct a terrain over a rectangular domain that consists of $n/2$ rectangles and $n/2$ pyramidal spikes, see Figure 4. The rectangles are almost horizontal at height h , are adjacent to each other, have width $2\delta/(n+1)$ in the x_1 -direction and are $(\delta + \gamma)n/2$ long in the x_2 -direction, for some $0 \leq \gamma < \delta$. They build the upper half of a convex polyhedron. The spikes are located right next to the rectangles, one after the other at $2\delta + \gamma$ spacing in x_2 -direction. The spikes have x_3 -height $h + \delta$, and their square bases have side length γ . The holes between the spikes and rectangles are filled so that the terrain is defined over a convex domain. The spikes and the rectangles are located in a way that ensures that the δ -neighborhood of every spike intersects the δ -neighborhood of every rectangle in the overall δ -neighborhood G^δ , resulting in $\Omega(n^2)$ intersections overall. This construction equally applies to the L_2 -metric and to convex polyhedral metrics. \square




Figure 4: Example of a convex-domain terrain in \mathbb{R}^3 , whose δ -neighborhood has quadratic complexity.

3.2.2 The Structure of the Feasible Region

We assume that G is a convex-domain terrain. Lemma 3.3 states that over the domain D_g^δ , the complement $\overline{G^\delta}$ can be split into two interior-disjoint parts, $\overline{G^\delta} = U \cup L$, where U is the “upper” and L the “lower” part, with respect to x_{d+1} .

Since $I \subseteq G^\delta$, it suffices to restrict $\overline{G^\delta}$ to the cylinder above the domain D_g^δ . Thus we can substitute $\overline{G^\delta} = U \cup L$ into (5) and obtain

$$I = \overline{U \oplus (-F)} \cap \overline{L \oplus (-F)}. \quad (6)$$

Since G_u and G_l are both terrains, we know that over D_g^δ

$$\begin{aligned} U \oplus (-F) &= \text{lower envelope of } G_u \oplus (-F) =: \underline{\mathcal{E}} \\ \text{and } L \oplus (-F) &= \text{upper envelope of } G_l \oplus (-F) =: \overline{\mathcal{E}} \end{aligned}$$

Let $\underline{\mathcal{E}}_\downarrow$ (resp., $\overline{\mathcal{E}}_\uparrow$) be the $((d+1)$ -dimensional) region below $\underline{\mathcal{E}}$ (resp., above $\overline{\mathcal{E}}$). Plugging this into (6) yields

$$I = \underline{\mathcal{E}}_\downarrow \cap \overline{\mathcal{E}}_\uparrow, \quad (7)$$

which is the region sandwiched between the two envelopes $\underline{\mathcal{E}}$ and $\overline{\mathcal{E}}$.

Let $\gamma_d(n)$ denote the combinatorial complexity of G_u and G_l . We can construct I as follows (see Section 3.2.3 for details):

1. Compute G_u (resp., G_l) as the upper (resp., lower) envelope of the δ -neighborhoods of all simplices of G .
2. Compute the envelopes $\underline{\mathcal{E}}$ and $\overline{\mathcal{E}}$. G_u and G_l are composed of $\gamma_d(n)$ surface patches whose shape depends on the underlying metric. We compute $\underline{\mathcal{E}}$ and $\overline{\mathcal{E}}$ by computing for each such patch and each $a \in F$ the $O(m\gamma_d(n))$ individual Minkowski sums, and finally computing the envelope of these Minkowski sums.
3. Compute $\underline{\mathcal{E}}_\downarrow \cap \overline{\mathcal{E}}_\uparrow = I$.

3.2.3 Computing the Feasible Region

Let us now consider how fast the algorithm for computing I can be implemented for various metrics in various dimensions.

In $d + 1 = 2$ dimensions G consists of n pairwise non-intersecting line segments. For any convex metric, the δ -neighborhoods of single segments are pseudo-disks [26], i.e., the boundaries of two δ -neighborhoods intersect at most twice. Kedem et al. [26] have shown that the union of n pseudo-disks has complexity $O(n)$. Therefore G^δ and hence also G_u and G_l have complexity $O(n)$. For projectable metrics, the boundary of the δ -neighborhood of a line segment consists of line segments. The upper envelope G_u and the lower envelope G_l of these $O(n)$ segments can be computed in time which is a log-factor slower than the complexity, hence $O(n \log(n))$ [34]. G_l and G_u consist of line segments. The algorithm requires computing the Minkowski sum of each such line segment with each $a \in F$. Such sum of two line segments is simply the convex hull of the vector sums of pairs of endpoints, one from each line segment. The lower and upper envelopes $\underline{\mathcal{E}}$ and $\overline{\mathcal{E}}$ of the resulting $O(mn)$ line segments have complexity $O(mn\alpha(mn))$, and can be computed in time $O(mn \log(mn))$ [24]. Moreover, these envelopes intersect at most $O(mn\alpha(mn))$ times [34], and a simple sweep-based procedure suffices to check whether I is empty. For $d = 1$ and projectable metrics we therefore have a runtime of $O(mn \log(mn))$.



Figure 5: Combinatorially different Minkowski sums of a line segment and a circular arc.

For the Euclidean metric in $d + 1 = 2$ dimensions, G_l and G_u still have complexity $O(n)$, but they may also contain circular arcs. In the next step we also need to compute the Minkowski sum of each such circular arc with each $a \in F$. This computation can be carried out in constant time by copying the circular arc at both ends of the line segment and adding a copy of the segment between the two points on the arcs at which the segment is tangent; see Figure 5 for an illustration of two possible cases. Therefore $\underline{\mathcal{E}}$ and $\overline{\mathcal{E}}$ also consist of line segments and circular arcs, have complexity $O(mn2^{\alpha(mn)})$, and can be computed in time $O(mn2^{\alpha(mn)} \log(mn))$ [34]. The overall running time for $d = 1$ and the L_2 metric is therefore $O(mn2^{\alpha(mn)} \log(mn))$.

Let us now consider an arbitrary dimension $d \geq 2$, and a projectable convex polyhedral metric of constant description complexity. We compute the Minkowski sum of each simplex of G with the polyhedron \mathbf{B}_{d+1}^δ , and compute the upper envelope of all the obtained simplices (which yields G_u), as well as their lower envelope (which yields G_l). The Minkowski sum of two convex polyhedra of constant description complexity is the convex hull of the vector sums of pairs of extreme points, one from each of the two polyhedra. The individual Minkowski sums can thus be computed in constant time each. The envelope of $O(n)$ simplices can be computed in time $O(n^d \alpha(n))$ (Theorem 2.9). In fact, the construction algorithm of Theorem 2.9 provides a decomposition of the envelopes into simplices. Using this fact we now compute the Minkowski sums for each of the $O(n^d \alpha(n))$ simplices of G_u or G_l with each $a \in F$. Applying the results of Theorem 2.9, we compute the sandwiched region I between the two envelopes $\underline{\mathcal{E}}$ and $\overline{\mathcal{E}}$ in $O((mn^d \alpha(n))^d \alpha(mn^d \alpha(n))) = O(m^d n^{d^2} \alpha^{d+1}(m+n))$

time. The algorithm runs in overall time $O(m^d n^{d^2} \alpha^{d+1}(m+n))$, which implies the following theorem.

Theorem 3.6. *Let F be a polyhedral set and G be a convex-domain polyhedral terrain (or a polyhedral δ -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:*

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

3.2.4 Undirected Hausdorff Distance

We now consider the case of the undirected Hausdorff distance, and wish to derive a decision procedure corresponding to the one described above. Here we assume that both F and G are convex-domain terrains (or, in more generality, δ -terrains). We can apply virtually the same approach as above. We split ∂F^δ into F_u and F_l , and define U^F (resp., L^F) to be the region above F_u (resp., below F_l). The corresponding regions for G are denoted by U^G and L^G . Define the feasible region I as $I^F \cap I^G$, where

$$\begin{aligned} I^F &:= \overline{U^F \oplus (-G)} \cap \overline{L^F \oplus (-G)} \\ I^G &:= \overline{U^G \oplus (-F)} \cap \overline{L^G \oplus (-F)}. \end{aligned}$$

It is easy to see that the undirected Hausdorff distance between F and G is at most δ iff the feasible region is not empty. As in the case of the directed Hausdorff distance, we check this by computing the feasible region. This is equivalent to computing

$$\underline{\mathcal{E}}_G \downarrow \cap \underline{\mathcal{E}}_F \downarrow \cap \overline{\mathcal{E}}_G \uparrow \cap \overline{\mathcal{E}}_F \uparrow$$

where $\underline{\mathcal{E}}_G \downarrow$ is the region below the lower envelope of $G_u \oplus (-F)$, $\overline{\mathcal{E}}_G \uparrow$ is the region above the upper envelope of $G_l \oplus (-F)$, $\underline{\mathcal{E}}_F \downarrow$ is the region below the lower envelope of $F_u \oplus (-G)$, and $\overline{\mathcal{E}}_F \uparrow$ is the region above the upper envelope of $F_l \oplus (-G)$.

Since we have two lower envelopes and two upper envelopes we can simply compute $\underline{\mathcal{E}}_{F,G}$ as the lower envelope of $G_u \oplus (-F) \cup F_u \oplus (-G)$, and $\overline{\mathcal{E}}_{F,G}$ as the upper envelope of $G_l \oplus (-F) \cup F_l \oplus (-G)$, and use the fact that

$$\underline{\mathcal{E}}_{F,G} \downarrow \cap \overline{\mathcal{E}}_{F,G} \uparrow = \underline{\mathcal{E}}_G \downarrow \cap \underline{\mathcal{E}}_F \downarrow \cap \overline{\mathcal{E}}_G \uparrow \cap \overline{\mathcal{E}}_F \uparrow$$

The above algorithms carry over almost verbatim, and we have:

Corollary 3.7. *Let F and G be convex-domain polyhedral terrains (or polyhedral δ -terrains) in \mathbb{R}^{d+1} , with respective complexities m and n . Put $N = m + n$. We can test whether there exists a translation in \mathbb{R}^{d+1} that brings F into undirected Hausdorff distance at most δ of G in the following time:*

- When $d = 1$ and the underlying metric is projectable, the running time is $O(N^2 \log N)$.
When the metric is Euclidean, the running time is $O(N^2 2^{\alpha(N)} \log N)$.
- When $d \geq 2$ and the metric is projectable, the running time is $O(N^{d^2+d} \alpha^{d+1}(N))$.

3.2.5 Parametric Searching

Parametric searching, originally introduced by Megiddo [31], is a general technique that facilitates the design of optimization algorithms by essentially reducing this task to the design of a parallel algorithm for the corresponding decision problem. Since the technique has become a mainstream practice in geometric optimization we do not review it in detail and instead refer the reader to the excellent survey by Agarwal and Sharir [3].

In brief, in order to apply parametric searching in our setting we need to derive a parallel algorithm (in a rather weak model of parallel computation, see [3] for details) for deciding whether the Hausdorff distance of F and G , under translations, is less than, equal to, or greater than a given parameter δ . The (serial) decision algorithms presented above reduce this task to constructing the sandwich region I . The Hausdorff distance of F and G is less than δ (resp., equal to δ , greater than δ) if the interior of I is non-empty (resp., the closure of I is non-empty but its interior is empty, the closure of I is empty). However, these algorithms are inherently serial because of the subprocedures that are used to construct the sandwich region I —the construction algorithm of Theorem 2.9 proceeds by adding the objects one by one, an inherently serial approach.

To work around this difficulty we use a randomized divide-and-conquer algorithm for constructing the sandwich region $\mathcal{S}(\mathcal{F}, \mathcal{G})$ (using the same notation and definitions as in Theorem 2.9). The randomized divide-and-conquer approach to such problems has become widespread and an analogous algorithm for constructing $\mathcal{E}(\mathcal{F})$ is described by Sharir and Agarwal [34, Section 7.2.2]. The price we pay for the simplicity and ease of parallelization is in the slightly suboptimal running time of $O(n^{d+\varepsilon})$ for any $\varepsilon > 0$.

The algorithm proceeds by taking a random sample $R_{\mathcal{F}}$ of r surfaces of \mathcal{F} and a random sample $R_{\mathcal{G}}$ of r surfaces of \mathcal{G} , for a sufficiently large constant r , and constructing the simplicial decomposition of the sandwich region $\mathcal{S}(R_{\mathcal{F}}, R_{\mathcal{G}})$. By the theory of $(1/r)$ -nets [23], each simplex of this decomposition is intersected by at most n/r objects of \mathcal{F} and at most n/r objects of \mathcal{G} , which allows us to recurse independently in each of the simplices. Since the number of simplices is $O(r^d \alpha(r))$ [19] an easy calculation shows that the recursive process completes in time $O(n^{d+\varepsilon})$ for any $\varepsilon > 0$. Since the recursive computations in each simplex are independent of each other this algorithm is easily parallelizable, yielding a parallel decision procedure that runs in $O(\log n)$ steps on $O(n^{d+\varepsilon})$ processors. Skipping the standard details we conclude that this procedure can be plugged into the parametric searching technique to achieve the following main result of this subsection.

Theorem 3.8. *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. [17]. We would also like to thank Pankaj K. Agarwal and Micha Sharir for helpful discussions. The anonymous referees are acknowledged for the detailed and thoughtful

comments that helped us improve the presentation of this paper.

References

- [1] P. K. Agarwal, S. Har-Peled, M. Sharir, and Y. Wang. Hausdorff distance under translation for points, disks, and balls. In *Proc. 17th ACM Symposium on Computational Geometry*, 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. Efficient algorithms for geometric optimization. *ACM Comput. Surv.*, 30:412–458, 1998.
- [4] P. K. Agarwal and M. Sharir. Pipes, cigars, and kreplach: The union of Minkowski sums in three dimensions. In *Proc. 15th ACM Symposium on Computational Geometry*, pages 143–153, 1999.
- [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] P. K. Agarwal, M. Sharir, and S. Toledo. Applications of parametric searching in geometric optimization. *Journal of Algorithms*, 17:292–318, 1994.
- [7] H. Alt, O. Aichholzer, and G. Rote. Matching shapes with a reference point. *Internat. J. Comput. Geom. Appl.*, 7:349–363, 1997.
- [8] H. Alt, B. Behrends, and J. Blömer. Approximate matching of polygonal shapes. *Ann. Math. Artif. Intell.*, 13:251–266, 1995.
- [9] H. Alt, P. Braß, M. Godau, C. Knauer, and C. Wenk. Computing the Hausdorff distance of geometric patterns and shapes. *Discrete and Computational Geometry - The Goodman-Pollack-Festschrift*, 2002.
- [10] 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, Amsterdam, 2000.
- [11] 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.
- [12] J.-D. Boissonnat and M. Yvinec. *Algorithmic Geometry*. Cambridge University Press, UK, 1998.
- [13] B. Chazelle. Cutting hyperplanes for divide-and-conquer. *Discrete Comput. Geom.*, 9(2):145–158, 1993.
- [14] L. P. Chew, M. T. Goodrich, D. P. Huttenlocher, K. Kedem, J. M. Kleinberg, and D. Kravets. Geometric pattern matching under Euclidean motion. *Comput. Geom. Theory Appl.*, 7:113–124, 1997.

- [15] K. L. Clarkson. A randomized algorithm for closest-point queries. *SIAM Journal on Computing*, 17:830–847, 1988.
- [16] M. de Berg, O. Cheong, O. Devillers, M. van Kreveld, and M. Teillaud. Computing the maximum overlap of two convex polygons under translations. *Theory of Computing Systems*, 31:613–628, 1998.
- [17] M. de Berg, L. J. Guibas, and D. Halperin. Vertical decompositions for triangles in 3-space. *Discrete Comput. Geom.*, 15:35–61, 1996.
- [18] H. Edelsbrunner. The upper envelope of piecewise linear functions: Tight complexity bounds in higher dimensions. *Discrete and Computational Geometry*, 4:337–343, 1989.
- [19] H. Edelsbrunner, L. J. Guibas, and M. Sharir. The upper envelope of piecewise linear functions: algorithms and applications. *Discrete and Computational Geometry*, 4:311–336, 1989.
- [20] M. Godau. *On the complexity of measuring the similarity between geometric objects in higher dimensions*. Ph.D. thesis, Free University Berlin, Berlin, Germany, 1998.
- [21] M. T. Goodrich, J. S. B. Mitchell, and M. W. Orletsky. Approximate geometric pattern matching under rigid motions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(4):371–379, 1999.
- [22] D. Halperin and M. Sharir. New bounds for lower envelopes in three dimensions, with applications to visibility in terrains. *Discrete and Computational Geometry*, 12:313–326, 1994.
- [23] D. Haussler and E. Welzl. Epsilon-nets and simplex range queries. *Discrete Comput. Geom.*, 2:127–151, 1987.
- [24] J. Hershberger. Finding the upper envelope of n line segments in $O(n \log n)$ time. *Inform. Process. Lett.*, 33:169–174, 1989.
- [25] D. P. Huttenlocher and K. Kedem. Computing the minimum Hausdorff distance for point sets under translation. In *Proc. 6th Annu. ACM Sympos. Comput. Geom.*, pages 340–349, 1990.
- [26] K. Kedem, R. Livne, J. Pach, and M. Sharir. On the union of jordan regions and collision-free translational motion amidst polygonal obstacles. *Discrete Comput. Geom.*, 1, 1986.
- [27] V. Koltun and M. Sharir. The partition technique for overlays of envelopes. *SIAM Journal on Computing*, to appear.
- [28] V. Koltun and C. Wenk. Matching polyhedral terrains using overlays of envelopes. In *Proc. 9th Scandinavian Workshop on Algorithm Theory*, 2004.
- [29] C. Lee. Subdivisions and triangulations of polytopes. In J. E. Goodman and J. O’Rourke, editors, *Discrete and Computational Geometry*, pages 271–290. CRC Press, 1997.
- [30] J. Matoušek. Range searching with efficient hierarchical cuttings. *Discrete Comput. Geom.*, 10(2):157–182, 1993.

- [31] N. Megiddo. Applying parallel computation algorithms in the design of serial algorithms. *J. ACM*, 30(4):852–865, 1983.
- [32] 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.
- [33] M. Sharir. Almost tight upper bounds for lower envelopes in higher dimensions. *Discrete and Computational Geometry*, 12:327–345, 1994.
- [34] M. Sharir and P. K. Agarwal. *Davenport-Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, New York, 1995.
- [35] B. Tagansky. A new technique for analyzing substructures in arrangements of piecewise linear surfaces. *Discrete and Computational Geometry*, 16:455–479, 1996.
- [36] C. Wenk. *Shape Matching in Higher Dimensions*. Ph.D. thesis, Free University Berlin, Berlin, Germany, 2002.
- [37] A. Wiernik and M. Sharir. Planar realizations of nonlinear Davenport-Schinzel sequences by segments. *Discrete Comput. Geom.*, 3:15–47, 1988.