Lower Envelopes of Surface Patches in 3-Space

Pankaj K. Agarwal
Department of Computer Science, Duke University, Durham, NC, USA

Esther Ezra
Department of Computer Science, Bar-Ilan University, Ramat-Gan, Israel

Micha Sharir
School of Computer Science, Tel Aviv University, Tel Aviv, Israel

Abstract

Let Σ be a collection of n surface patches, each being the graph of a partially defined semi-algebraic function of constant description complexity, and assume that any triple of them intersect in at most $s = 2$ points. We show that the complexity of the lower envelope of the surfaces in Σ is $O(n^2 \log^{6+\varepsilon} n)$, for any $\varepsilon > 0$. This almost settles a long-standing open problem posed by Halperin and Sharir [26], thirty years ago, who showed the nearly-optimal albeit weaker bound of $O(n^2 \cdot 2^{c\sqrt{\log n}})$ on the complexity of the lower envelope, where $c > 0$ is some constant. Our approach is fairly simple and is based on *hierarchical cuttings* and *gradations*, as well as a simple *charging scheme*. We extend our analysis to the case $s > 2$, under a "favorable cross section" assumption, in which case we show that the bound on the complexity of the lower envelope is $O(n^2 \log^{11+\varepsilon} n)$, for any $\varepsilon > 0$. Incorporating these bounds with the randomized incremental construction algorithms of Boissonnat and Dobrindt [16], we obtain efficient constructions of lower envelopes of surface patches with the above properties, whose overall expected running time is $O(n^2 \text{polylog} n)$, as well as efficient data structures that support point location queries in their minimization diagrams in $O(\log^2 n)$ expected time.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Hierarchical cuttings, surface patches in 3-space, lower envelopes, charging scheme, gradation

Digital Object Identifier 10.4230/LIPIcs.ESA.2024.6

Funding *Pankaj K. Agarwal*: Partially supported by NSF grants IIS-18-14493, CCF-20-07556 and CCF-22-23870 and by a US-Israel Binational Science Foundation Grant 2022131.
Esther Ezra: Partially supported by Israel Science Foundation Grant 800/22 and US-Israel Binational Science Foundation Grant 2022131.
Micha Sharir: Partially supported by Israel Science Foundation Grant 495/23.

1 Introduction

Let $\Sigma = \{\sigma_1, \ldots, \sigma_n\}$ be a collection of n surface patches in $\mathbb{R}^3$ of constant complexity. Specifically, we assume that each surface patch $\sigma_i \in \Sigma$ is the graph of a partially defined semi-algebraic function of constant complexity.[1]

The lower envelope of Σ, denoted by $E(\Sigma)$, is defined as the graph of the partially defined function

$$E(\Sigma)(x, y) = \min_{i=1,\ldots,n} f_i(x, y),$$

[1] Roughly speaking, a semi-algebraic set in $\mathbb{R}^d$ is the set of points in $\mathbb{R}^d$ that satisfy a Boolean formula over a set of polynomial inequalities; the complexity of a semi-algebraic set is the number of polynomials defining the set and their maximum degree.

where $f_i(x,y)$ is set to $+\infty$ if f_i is undefined at (x,y). The projection of $E(\Sigma)$ onto the xy-plane is called the *minimization diagram* of $E(\Sigma)$, and is denoted by $M(\Sigma)$. This is a planar subdivision, where over each face of $M(\Sigma)$ there is at most a single surface patch of Σ that attains $E(\Sigma)$. Each edge of $M(\Sigma)$ either lies in the projection of the intersection curve of a pair of patches in Σ, or is a portion of the projection of the boundary of a single patch. Each vertex v of $M(\Sigma)$ is either an *inner vertex*, i.e., the projection of an intersection point of the relative interior of three surface patches in Σ, or an *outer vertex*, which is the projection of a vertex of $E(\Sigma)$, either the projection of a single patch, or obtained by a pair of patch boundaries whose xy-projections intersect (at v), or formed by a triple of patches where the intersection curve of two of them passes above the boundary of the third one. The combinatorial complexity of the envelope $E(\Sigma)$ is defined as the overall number of vertices, edges, and faces of $M(\Sigma)$. We refer to [26, 38, 39] for further details concerning these definitions and properties.

The main result of this paper is a tighter upper bound for the combinatorial complexity of $E(\Sigma)$, for the case where the relative interiors of any triple of the surface patches in Σ intersect at most twice in their relative interiors.

Related work. Consider first the two-dimensional setting, where the input consists of n arcs in the plane and each pair of them intersect in at most s points, for some constant parameter s. In this case the complexity of the lower envelope of these arcs is at most $\lambda_{s+2}(n)$, the maximum length of a Davenport–Schinzel sequence of order $s+2$ on n symbols (see [39]).

In three dimensions, Halperin and Sharir [26] showed that the combinatorial complexity of $E(\Sigma)$ (where Σ is defined as above) is $O(n^2 \cdot 2^{c\sqrt{\log n}})$, for some constant c, given, as we assume, that the relative interiors of any triple of the surface patches in Σ intersect in at most $s = 2$ points. A follow-up work by Sharir [38] studied the general case, where any triple of surfaces intersect $O(1)$ times, and established the bound $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$ (here, and everywhere where bounds of these form are mentioned, the constant of proportionality depends on ε). In fact, Sharir's work presented an extension to the d-dimensional case, where it was shown that the lower envelope complexity of n surface patches in $\mathbb{R}^d$ (satisfying some mild assumptions) is $O(n^{d-1+\varepsilon})$, for any $\varepsilon > 0$.

In spite of these significant results, which are nearly tight in the worst case, the prevailing conjecture is that the bound should be of the form $O(n^{d-2}\lambda_q(n))$ for some constant $q > 0$ that depends on the shape and degree of the surfaces. So far such a bound was shown for only a few special cases, such as the case of $(d-1)$-simplices in $\mathbb{R}^d$, where the bound is $O(n^{d-2}\lambda_3(n)) = O(n^{d-1}\alpha(n))$ (where $\alpha(\cdot)$ is the (extremely slowly growing) near-constant inverse Ackermann function) [34, 39, 40]. This bound is asymptotically tight in the worst case [41], and thus, in particular, it also follows that the bound in [38] for surface patches is nearly tight. If the surfaces in Σ are graphs of totally defined bivariate functions of constant complexity, Schwartz and Sharir [37] showed that the conjectured bound holds under some further assumptions on the input surfaces. Specifically, for the case $s = 2$, if the intersection curve of each pair of surfaces is connected and the functions are totally defined then the complexity of $E(\Sigma)$ is $O(n^2)$. For $s > 2$ and totally defined functions, if the intersection curve of each pair of surfaces intersects every plane $x = const$ in at most two points (in a sense, an x-monotone curve), then the complexity of $E(\Sigma)$ is $O(n\lambda_{s+2}(n))$. Still, the cases of surface patches and of more general intersection curves have remained elusive, and we are yet not aware of any improvement of the bound $O(n^{2+\varepsilon})$, shown in [26, 38].

Our results. Our main result is a tighter upper bound for the complexity of the lower envelope of Σ, under the assumption that the relative interiors of any triple of elements in Σ intersect in at most two points (i.e., the case $s = 2$). Specifically, we show in Section 2:

▶ **Theorem 1.** *Let Σ be a collection of n surface patches in $\mathbb{R}^3$, each being the graph of a partially defined semi-algebraic function of constant complexity. If the relative interiors of any triple of surfaces of Σ intersect in at most two points, then the complexity of the lower envelope of Σ is $O(n^2 \log^{6+\varepsilon} n)$, for any $\varepsilon > 0$.*

As a concrete example, we apply Theorem 1 to the setting of spherical caps of constant complexity (where such a cap is a connected portion of the sphere that is bounded by a semi-algebraic curve of constant complexity):

▶ **Corollary 2.** *The complexity of the lower envelope of n spherical caps of constant complexity in $\mathbb{R}^3$ is $O(n^2 \log^{6+\varepsilon} n)$, for any $\varepsilon > 0$.*

We also present an extension to the case where the relative interior of any triple of surface patches in Σ intersect in $s > 2$ points, under some restrictions that follow from the work in [37]. Specifically, we assume that the surface patches in Σ are portions of graphs of continuously defined bivariate functions $f(x, y)$ of constant maximum degree. We further assume that these bivariate functions $f(x, y)$ have "favorable" cross sections, that is, we assume that the intersection curve of any pair of graphs of such functions intersect each plane of the form $x = x_0$ (where x_0 is a fixed parameter) in at most two points. Although the favorable cross section assumption seems restrictive, it covers the case where the functions $f(x, y)$ have a relatively simple form. In particular, such a property is useful for bounding the number of combinatorial changes in geometric structures of moving objects (such objects can also continuously deform). For instance, Schwartz and Sharir [37] showed that the overall number of combinatorial changes in the convex hull of n moving points in the plane is $O(n\lambda_{s+2}(n))$, where $s > 0$ is a constant that depends on the kind of the trajectories along which the points are moving. The main observation in [37] was that the set of combinatorial changes corresponds to the vertices of the lower envelope of bivariate functions with favorable cross sections. As a dual problem, consider the problem of bounding the total number of combinatorial changes in the lower envelope of continuously moving lines in the plane. This problem was studied by Alexandron et al. [12], who, similarly to the work in [37], showed a reduction to lower envelopes. Note that the convex hull or the lower envelope deform continuously as the points or lines move, but their combinatorial description changes at discrete time instances (referred to as *events*); see e.g. [23, 37, 39].

Using a fairly simple approach developed in this paper, we obtain the following (details are deferred to the full version of this paper):

▶ **Theorem 3.** *Let Σ be a collection of n surface patches, each of which is a portion of a graph of a continuously defined bivariate function $f(x, y)$ of constant degree. If the intersection curve of any pair of such graphs intersect each plane parallel to the zy-plane at most twice, then the complexity of the lower envelope of Σ is $O(n^2 \log^{11+\varepsilon} n)$, for any $\varepsilon > 0$.*

As an application of Theorem 3, consider the question of how many times the lower envelope of a set S of segments in the plane changes as the segments move continuously. Each segment is specified by the position of its two endpoints, where each endpoint is moving along an "algebraic path" (such a model was considered, e.g., in [7]). In this case, the y-coordinate of any moving segment can be represented by a (partially defined) bivariate function $f(x, t)$, where the variables x and t represent the x-coordinate of the segment and time, respectively.

Note that each pair of graphs of such functions intersect the plane $t = t_0$ in at most a single point, which is the point of intersection of the corresponding segments at position t_0. We also observe that a combinatorial change in the lower envelope of the segments in S, as they move along the time axis, corresponds to a vertex of the lower envelope of the graphs of the functions $f(x,t)$ – this is a fairly standard property and was reviewed, e.g., in [12] where the lower envelope of moving lines in the plane was considered. Therefore, the resulting setting satisfies the properties stated in Theorem 3, from which we conclude that the overall number of combinatorial changes that the lower envelope undergoes is $O(n^2 \log^{11+\varepsilon} n)$.

On the algorithmic front, we can construct the lower envelope of Σ, using the randomized incremental construction of Boissonnat and Dobrindt [16] in overall expected time that is within one logarithmic factor from its combinatorial complexity bound. Specifically, we conclude:

▶ **Corollary 4.** *Let Σ be as in Theorem 1 (resp., Theorem 3). Then the lower envelope of Σ can be constructed by a randomized algorithm in $O(n^2 \log^{7+\varepsilon} n)$ (resp., $O(n^2 \log^{12+\varepsilon} n)$) expected time, for any $\varepsilon > 0$. Moreover, the minimization diagram $M(\Sigma)$ can be preprocessed within the same expected time bound in order to support point location queries that take $O(\log^2 n)$ expected time each.*

The main contribution of this work is to bypass the intricate charging scheme and recursive machinery developed in earlier works on lower envelopes [26, 38], which result in an artifact of an $O(n^\varepsilon)$ factor in the bound on the lower envelope complexity, and replace them with a simpler machinery that produces a tighter bound. At a very high level, we employ a hierarchical space decomposition scheme, which allows us to eliminate the "boundary effect" of the elements in Σ. Specifically, we use *hierarchical cuttings* for the boundary projections of the patches in Σ, based on the work of Chazelle [18] (see also [6, 11]). This forms a hierarchical partition of $\mathbb{R}^3$ into vertical prisms τ, where we classify each $\sigma \in \Sigma$ that intersects τ, as being either *wide* in τ (if its boundary does not intersect τ) or *narrow* otherwise. As a result, the inner vertices v of the lower envelope of Σ (formed by triple intersections) are classified as being either of type WWW (that is, all three surface patches forming v are wide in τ), WWN, WNN, or NNN (with similar meanings; see Section 2 for these details). Whereas NNN-vertices are handled by a recursive scheme exploiting our hierarchical cuttings, and WWW- and WWN-vertices are handled using a simple charging scheme (see Section 3), we are still left to bound the number of WNN-vertices (this is the hardest part of the analysis). This analysis is presented in Section 4. A main technical contribution of this step is to form a hierarchical decomposition for lower envelopes based on "gradations" (see Section 4.2 for the details of this construction). In a typical approach exploiting "standard" cuttings in a hierarchical manner, the constant of proportionality of the bound on the size of the cutting is amplified over the steps in the hierarchy, and, as a result, the overall storage and preprocessing time incur an additional factor of $O(n^\varepsilon)$, for any $\varepsilon > 0$ (see Section 2.2 for these details). In contrast, our approach exploits gradations, which enables to construct such a hierarchy of cuttings, but in a more controlled manner, where instead of the naïve approach we use a global measure to estimate the actual size of the cutting.

2 Overview of the Technique

In this section we present some preliminaries, and our general recursive frameworks that exploit standard hierarchical cuttings and gradations, and indicate how such a decomposition implies Theorem 1. The proof of the theorem relies on various combinatorial bounds that are established in Sections 3 and 4.

2.1 Preliminaries

Let Σ be a collection of n surface patches as above. We assume that the given patches are in general position, meaning that the coefficients of the polynomials defining the surfaces and their boundaries are algebraically independent over the rationals. This, e.g., implies that no four surfaces meet at a point, the boundary of one surface does not meet a curve of intersection of two other surfaces, two boundary curves of distinct surfaces do not meet at a point, etc. See [26, 38] for such an assumption and the property that it does not involve any loss of generality. In what follows, we assume that the relative interior of any triple of surfaces from Σ intersect in at most two points.

As already said, we measure the complexity of $E(\Sigma)$ by the number of vertices of its minimization diagram $M(\Sigma)$ (see once again the discussion in Section 1). Recalling the discussion in the introduction, we call such a vertex *inner* if it is the xy-projection of an intersection of the relative interiors of three surface patches of Σ, and *outer* if it is the xy-projection of some point on the boundary of some surface patch. There are four types of outer vertices: (i) the xy-projection of one of the original $O(n)$ vertices of the given patches, (ii) the xy-projection of one of the $O(n^2)$ points of intersection between the boundary of one surface patch and the relative interior of another patch, (iii) the intersection of the xy-projections of two such boundaries, or (iv) the xy-projection of an intersection curve of a pair of surface patches with the projected boundary of a third one, where the curve passes (locally) above the third surface.

The number of outer vertices of types (i)–(iii) is clearly $O(n^2)$ (it is only $O(n)$ for type (i) vertices). Concerning vertices of type (iv), we have the following straightforward lemma:

▶ **Lemma 5.** *The number of outer vertices of type (iv) is $O(n\lambda_q(n))$, for some constant q that depends on the complexity of the individual surface patches of Σ.*

Proof. For each surface $\sigma \in \Sigma$ let $\delta_\sigma := \partial\sigma$ denote its boundary, and let V_σ denote the upward vertical curtain erected from δ_σ, namely the union of the upward-directed vertical rays emanating from the points of δ_σ. Each other surface patch of Σ crosses V_σ in a (possibly empty, possibly disconnected) curve, and the vertices we are after are the xy-projections of the breakpoints of the lower envelope of these curves, where two of the curves meet. Each of the curves is semi-algebraic of constant complexity, so the number of breakpoints of their lower envelope is $O(\lambda_q(n))$, for a suitable constant parameter q; see, e.g., [39]. Repeating this argument for each boundary δ_σ, the asserted bound follows. (As a matter of fact, vertices of types (ii) and (iii) are also degenerate breakpoints of the envelope.) ◀

The main effort of the analysis is to bound the number of inner vertices. Our analysis is based on a recursive mechanism exploiting hierarchical cuttings, tailored for this setting, which we review next.

2.2 Hierarchical cuttings

We present a primary and a secondary space decompositions, based on hierarchical cuttings, where the first exploits the machinery of Chazelle [18], and the latter exploits the machinery of Ramos [35] using gradation. The primary decomposition is fairly simple and presented in this section, and the secondary one is more involved and presented in Section 4.

Cuttings, hierarchy and gradation. Given a parameter $1 \le r \le n$ and a region $L \subseteq \mathbb{R}^3$, a $(1/r)$-*cutting* for Σ that covers L is a set Ξ of interior-disjoint constant-complexity "elementary cells" (e.g., vertical pseudo-prisms, simplices, etc.), such that (i) the interior of every cell $\tau \in \Xi$ intersects at most n/r surfaces of Σ, and (ii) the union of the cells in Ξ covers L.

The *conflict list* Σ_τ of a cell $\tau \in \Xi$ is the set of the (at most n/r) surfaces of Σ that intersect its interior. The *defining set* of τ is a minimal-size subset $D_\tau \subseteq \Sigma$ (of interest are only cases where $|D_\tau|$ is bounded by a constant, which is the typical case, and is also the case in our setting), such that τ appears as an elementary cell in the decomposition of the arrangement $\mathcal{A}(D_\tau)$. We have $\Sigma_\tau \cap D_\tau = \emptyset$. The *size* of Ξ is the number of its cells.

In the standard notion of three-dimensional cuttings we have $L = \mathbb{R}^3$. In this work we construct efficient (i.e., small-size) cuttings for the case where L covers the lower envelope of Σ (see Section 4 for the specific construction).

In general, we can efficiently construct $(1/r)$-cuttings when r is a constant. To obtain a cutting with a larger non-constant value, one uses a recursive scheme, with some constant value r_0 of r, where at each step cells of the previous cutting are split into subcells, using a $(1/r_0)$-cutting of the conflict list of the parent cell. The standard way of doing this is called a *hierarchical cutting*, and we discuss it below. A more refined scheme is a *gradation*, which better controls the number of cells that arise in the construction. Here the approach is to form a nested sequence of random samples of the form $\emptyset = \Sigma_{-1} \subseteq \Sigma_0 \subseteq \cdots \subseteq \Sigma_{l-1} \subseteq \Sigma_l = \Sigma$, where the index l is the height of the gradation. For each index $0 \le i \le l$ (for $i = -1$ no action is required), we consider the elementary cells appearing in the decomposition of the portion of the arrangement $\mathcal{A}(\Sigma_i)$ that covers L. These properties are discussed in detail in Section 4.2 and are inspired by the work of Ramos [35].

The construction. We project the boundary curves of the surfaces in Σ onto the xy-plane, and denote by Γ be the resulting set of planar curves.

We next construct a hierarchy of cuttings for Γ based on Chazelle's construction [18], as follows.[2] Fix some sufficiently large constant parameter $r_0 > 0$. We proceed through at most $\lceil \log_{r_0} n \rceil$ steps, where at each step i we construct a $(1/r_0^i)$-cutting of Γ, denoted by Ξ_i. We have $\Xi_0 = \mathbb{R}^2$, and at any further step $i > 0$, Ξ_i is a refinement of Ξ_{i-1}. Specifically, we refine each cell $\tau_0 \in \Xi_{i-1}$ by constructing a $(1/r_0)$-cutting within τ_0 for the curves from Γ in the conflict list of τ_0. This forms a collection of (sub)cells τ, which, taken over all cells τ_0, comprises Ξ_i. Each resulting cell is an *open* pseudo-trapezoid, formed by the *vertical decomposition* of some planar arrangement of a suitable sample of the curves [10]. By the analysis in [18] (see also [6, 11]), it follows that: (i) each cell $\tau \in \Xi_i$ intersects at most n/r_0^i curves of Γ, and (ii) the size of Ξ_i is bounded by $O(r_0^{2i} \log^2 r_0)$. We emphasize that property (ii) follows from a global argument, which was a key idea in the analysis presented in [18], in contrary to a naive approach where the bound incurs an artifact of an $O(n^\varepsilon)$ factor.

We next form a three-dimensional hierarchical cutting, by lifting every resulting pseudo-trapezoidal cell τ from all the "layers" Ξ_i in the z-direction into a prism $\tau \times \mathbb{R}$. This results in a set of unbounded vertical prisms, which comprises, over all levels i of the hierarchy, our (primary) hierarchical construction. Observe that a cell $\tau \in \Xi_i$ (in the planar cutting) intersects a curve $\gamma \in \Gamma$ if and only if the lifted cell of τ intersects the boundary of the corresponding surface patch (which projects to γ). In what follows, and with a slight abuse of notation, we keep denoting by Ξ_i the ith layer of the three-dimensional (lifted) cutting, and its cells by τ.

Let τ be a cell in Ξ_i. Each surface $\sigma \in \Sigma$ that intersects τ is either *wide* in τ, if σ fully covers the xy-projection of τ, or *narrow* in τ otherwise. In particular, it is easy to verify that if σ is wide in τ then σ intersects τ but $\partial\sigma \cap \tau = \emptyset$, and if σ is narrow in τ then $\partial\sigma \cap \tau \neq \emptyset$. Let $\Sigma_\tau^W \subseteq \Sigma$ (resp., $\Sigma_\tau^N \subseteq \Sigma$) be the set of the wide (resp., narrow) surfaces in τ. By the above discussion, we conclude (recall that r_0 is a constant):

[2] The construction of Chazelle is described for the setting of hyperplanes in any dimension [18], but the same technique also applies to semi-algebraic arcs in the plane, as observed in [6, 11].

▶ **Lemma 6.** *For any step $0 \le i \le \lceil \log_{r_0} n \rceil$, the size of Ξ_i is $O(r_0^{2i} \log^2 r_0)$, and any cell $\tau \in \Xi_i$ intersects at most n/r_0^i surfaces from Σ that are narrow in τ. In particular, at the last step we obtain a decomposition into $O(n^2)$ prism cells, each intersecting $O(1)$ locally narrow surfaces.*

We remark that the lemma provides no control on the number of wide surfaces at a prism. Generally speaking, our strategy is to construct the (primary) hierarchical cutting for the narrow surfaces, as outlined above. As space is progressively partitioned into smaller vertical prisms, narrow surfaces at some cell τ are more likely to become wide in many (in general, most) of the subcells of τ that they cross. In several non-trivial steps in our analysis (see the following sections for details) we show how to dispose of all the wide surfaces in a cell, in a direct nonrecursive manner, so that only narrow surfaces are processed recursively.

Specifically, let τ be a cell of the hierarchical cutting at some level i. We classify the (inner) vertices v of the lower envelope $E(\Sigma)$ within τ as being of type WWW, WWN, WNN, or NNN, so that, as already defined, the first type indicates that v is formed by the intersection of three wide surfaces in τ, the second type indicates that it is formed by the intersection of two wide surfaces and one narrow surface in τ, and the other two types are defined analogously. In a main step of the analysis, we show (in what follows, we use the notation $\beta(n) = \lambda_q(n)/n$, where $q > 0$ is a constant that depends on the complexity of the surfaces in Σ):

▶ **Lemma 7.** *Let τ be a cell of the hierarchical cutting, and let n_τ be the overall number of narrow and wide surfaces of Σ that intersect τ. Then the following bounds hold: (i) The number of WWW-vertices of $E(\Sigma)$ that lie in τ is $O(n_\tau^2)$. (ii) The number of WWN-vertices, as above, is $O(n_\tau^2 \beta(n_\tau))$. (iii) The number of WNN-vertices, as above, is $O(n_\tau^2 \log^{5+\varepsilon} n_\tau)$, for any $\varepsilon > 0$.*

We prove parts (i)–(ii) of Lemma 7 in Section 3, and the more technically challenging part (iii) in Section 4. NNN-vertices are processed recursively.

2.3 The recursive framework

Given Lemmas 6 and 7, we proceed as follows. Assume, without loss of generality, that at step $i = 0$ all surface patches in Σ are narrow, implying that all inner vertices of $E(\Sigma)$ are of type NNN. [3] We next apply the following recursive scheme in order to bound their overall number. At step $i = 1$ we construct Ξ_1, which, according to Lemma 6, results in $O(r_0^2 \log^2 r_0)$ cells (prisms) τ, each of which intersects at most n/r_0 surfaces of Σ that are narrow in τ. The number of wide surfaces in τ is clearly at most n. As a result, an NNN-vertex v from the previous step either remains an NNN-vertex (if the three surfaces forming v all remain narrow in τ), or becomes a WWW-, WWN- or WNN-vertex (in the respective cases where all three, only two, or only one of them become wide in τ). Let n_τ denote the number of surfaces (wide and narrow) that cross τ. Then, using Lemma 7, we bound the number of the non-NNN vertices in each cell τ by $O(n_\tau^2 \log^{5+\varepsilon} n_\tau)$, for a total bound of $O\left(\sum_\tau n_\tau^2 \log^{5+\varepsilon} n_\tau\right) = O(r_0^2 n^2 \log^{5+\varepsilon} n \log^2 r_0)$ vertices, over all cells. Since all that remains is to bound the number of NNN-vertices that "survived", we can dispose of all the (locally) wide surfaces in each cell τ, and continue to process recursively only the still narrow surfaces. In general, this implies that at each step $i \ge 1$ of the construction, the cutting Ξ_i at this

[3] If some of these surfaces are graphs of totally defined functions, we can dispose of them immediately using Lemma 7 – see below for details.

level is such that each cell $\tau \in \Xi_i$ intersects at most n/r_0^i locally narrow surfaces and at most n/r_0^{i-1} locally wide surfaces of Σ, where the latter bound follows from the inductive bound on the number of those narrow surfaces in the parent cell τ_0 of τ that just became wide in τ (because all former wide surfaces have already been discarded). This implies that at each step $0 \le i \le \lceil \log_{r_0} n \rceil$, except for the last one, the overall number of (inner) vertices that we count, excluding those passed down the hierarchy, is (as above, we use the bound in Lemma 7(iii), which dominates the other two bounds):

$$O\left(r_0^{2i} \log^2 r_0 \left(\frac{n}{r_0^{i-1}}\right)^2 \log^{5+\varepsilon}\left(n/r_0^{i-1}\right)\right) = O\left(n^2 r_0^2 \log^{5+\varepsilon} n \log^2 r_0\right). \tag{1}$$

At the last step, we bound the number of NNN-vertices (as well as those of all the other three types) in a brute-force manner, which results in a total bound of $O(n^2)$, which is the number of final prisms, according to Lemma 6, as the number of surfaces at each prism is $O(1)$. Summing over the at most $\lceil \log_{r_0} n \rceil$ levels of the hierarchical cutting, and recalling that r_0 is a constant, we obtain that the total number of inner vertices of $E(\Sigma)$ is

$$O\left(\sum_{i=0}^{\lceil \log_{r_0} n \rceil} n^2 r_0^2 \log^{5+\varepsilon} n \log^2 r_0\right) = O(n^2 \log^{6+\varepsilon} n), \tag{2}$$

for any $\varepsilon > 0$, as asserted in Theorem 1.

3 Bounding the number of WWW- and WWN-Vertices

In this section we analyze the number of WWW- and WWN-vertices in a prism cell τ anywhere along the hierarchy, and establish parts (i) and (ii) of Lemma 7. Our solution is based on a charging scheme, which bears some resemblance to charging schemes used in earlier works on lower envelopes [26, 38], but it is different in flavor and simpler in nature.

Let $\mathcal{W} := \mathcal{W}_\tau$, $\mathcal{N} := \mathcal{N}_\tau$ be the respective sets of wide and narrow surfaces from Σ in τ. Following the notation of Lemma 7, $|\mathcal{W}| \le n_\tau$ and $|\mathcal{N}| \le n_\tau$. We denote by $E(\mathcal{W} \cup \mathcal{N})$ the lower envelope of the surfaces in $\mathcal{W} \cup \mathcal{N}$, where each surface is clipped to its portion within τ.

Since the relative interiors of any triple of the surface patches of Σ intersect at most twice, we can classify each inner vertex v as either a *left vertex* or a *right vertex*, depending on whether v is the leftmost (resp., rightmost) of the two intersection points of the three surface patches incident to v; if this triple intersect only once, it does not matter how we classify v, and we assume general position to ensure that no two vertices have the same x-coordinate (or else just classify them arbitrarily as left and right).

The analysis uses repeatedly the following argument and notation. Without loss of generality, it suffices to bound the number of right vertices. Let v be such a vertex, incident to three surfaces a, b, c. The vertex v is incident to three intersection curves $\gamma_{ab} = a \cap b$, $\gamma_{ac} = a \cap c$, $\gamma_{bc} = b \cap c$. Each of these curves is split at v into a *lower portion* and an *upper portion*, where the lower portion is the part that appears on $E(\mathcal{W} \cup \mathcal{N})$ locally near v (below the third surface incident to v), and the upper portion is the complementary part, which passes above $E(\mathcal{W} \cup \mathcal{N})$ locally near v, hidden from the envelope by the third surface. The lower portions of the three curves γ_{ab}, γ_{ac}, γ_{bc} emanate from v in three directions that span the circle of orientations, meaning that their xy-projections are not all contained (locally) in a halfplane bounded by a line through the projection of v (this property is easy to verify under the general position assumption, and is also stated in [26, 38]). The same holds for the upper portions of the three curves. In particular there is at least one curve (and at most two curves) γ whose upper portion emanates from v to the right. See Figure 1 for an illustration.

Our general strategy is to apply a charging scheme, in which we trace the upper portion of γ to the right, and charge v to some feature that we encounter along γ. That is, suppose that we trace a curve γ from v away from (i.e., above) the envelope, to the right. If we reach along γ, before it shows up on the envelope again, a singularity or a locally x-extremal point, we charge v to that point. The charging is essentially unique, meaning that each such point is charged only $O(1)$ times (assuming general position), and there are only $O(1)$ such points along each curve, where the constant of proportionality depends on the description complexity of the surfaces in $\mathcal{W} \cup \mathcal{N}$. Therefore the total number of these charges is $O(n_\tau^2)$, implying that the number of vertices v with this property is $O(n_\tau^2)$. We may thus assume in what follows that none of these events occur along the upper portions of the three curves incident to v.

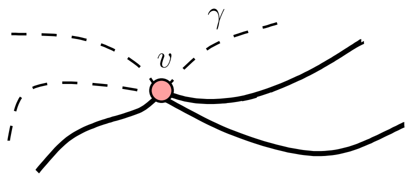

Figure 1 The local structure of the incident edges at a vertex v of the envelope. The solid lines represent the portion of the edges on the envelope, whereas the dashed lines are the edge portions above the lower envelope. In particular, the upper portion γ goes to the right.

We present the remaining technical details of our charging scheme in the following lemma:

▶ **Lemma 8.** *(i) The complexity of the lower envelope $E(\mathcal{W})$ of the wide surfaces at a prism τ is $O(n_\tau^2)$, which implies that the number of WWW-vertices of $E(\mathcal{W} \cup \mathcal{N})$ is $O(n_\tau^2)$. (ii) The number of WWN-vertices on $E(\mathcal{W} \cup \mathcal{N})$ is $O(n_\tau^2 \beta(n_\tau))$, where $\beta(\cdot)$ is as defined above.*

Proof. We first prove part (i) of the lemma. Without loss of generality, it suffices to bound the number of right WWW-vertices (as defined earlier). We follow the notation used above, denoting the three surfaces incident to v as a, b, c, and their corresponding intersection curves as γ_{ab}, γ_{ac} and γ_{bc}. Suppose that γ_{ab} is a curve whose upper portion proceeds from v to the right; that is, as we trace γ_{ab} from v to the right, it is hidden from the envelope by the third surface c. Since c is wide and v is a right vertex, c does not intersect this portion of γ_{ab} again, so it does not show up on the envelope. We can then charge v, in a unique manner, to the right endpoint of this upper portion, which is an intersection point of γ_{ab} with $\partial\tau$ or lies at infinity when the xy-projection of τ is unbounded (we recall that we excluded the cases of singularity or locally x-extremal endpoints). It is easy to verify that the number of such endpoints, for each curve γ_{ab}, is $O(1)$ (the constant of proportionality here depends on the complexity of τ, which is a constant too), for an overall bound of $O(n_\tau^2)$. Hence the number of WWW-vertices on the envelope is $O(n_\tau^2)$, and the proof is complete.

WWN-vertices. We next prove part (ii) of the lemma. We say that an intersection curve of two surfaces is a WW-curve (resp., WN-curve) if the two intersecting surfaces are both wide (resp., one is wide and the other is narrow); for curves incident to a WWN-vertex, these are the only two possibilities.

Let v be an inner vertex of type WWN, and let a, b and c be the three surfaces incident to v, where a and b are wide at τ and c is narrow. Similarly to what has been defined earlier, we consider the three intersection curves γ_{ab}, γ_{ac}, γ_{bc} which pass through v, and each curve is split at v into its lower portion and its upper portion, as defined above.

Suppose, without loss of generality, that v is a right vertex. As already argued, there is at least one curve (and at most two curves) γ whose upper portion emanates from v to the right.

Let v be a right WWN-vertex, and let γ be a curve incident to v whose upper portion proceeds from v to the right. If γ is a WN-curve then, as reasoned earlier, γ cannot appear on the envelope again, because it is hidden from it by the third, wide surface incident to v, which it cannot cross as it proceeds to the right, since v is a right vertex. In this case we charge v to the corresponding right endpoint of γ. Any such curve γ can be charged in this manner only $O(1)$ times, which implies that the number of such vertices v is $O(n_\tau^2)$. We may therefore assume that there is a single curve γ whose upper portion emanates from v to the right, and that $\gamma = \gamma_{ab}$ is the incident WW-curve (we once again recall that we do not consider singular or locally x-extremal endpoints).

The only case we need to consider is that γ appears on the envelope again, somewhere to the right of v. As the tracing of γ begins, the third, narrow incident surface c lies below γ. Hence γ must pass above ∂c in order to show up on $E(\mathcal{W} \cup \mathcal{N})$ again.

If, before this happens, γ meets another wide surface, we encounter a WWW-vertex along γ. Let w be the first such vertex. Then w is a vertex of the lower envelope of the wide surfaces at τ. We then charge v to w, and conclude that the number of such vertices v is $O(n_\tau^2)$.

Hence we may assume that when γ reaches the upward vertical curtain H_c above ∂c, no wide surface passes below γ on H_c. Hence $\gamma \cap H_c$ is a breakpoint of the lower envelope of the (cross sections of the) wide surfaces within H_c, and the number of such breakpoints is $O(n_\tau \beta(n_\tau))$ for each c, for a total of $O(n_\tau^2 \beta(n_\tau))$ such breakpoints, where $\beta(\cdot)$ is as defined above. We thus conclude that the number of vertices v with this property is $O(n_\tau^2 \beta(n_\tau))$. We emphasize that H_c may also contain cross sections of narrow surfaces, but we simply ignore them in the above reasoning. In conclusion, this shows that the number of WWN-vertices on the lower envelope is $O(n_\tau^2 \beta(n_\tau))$, as asserted. ◀

This completes the proof of parts (i) and (ii) of Lemma 7.

4 The Analysis of WNN-Vertices

In this section we describe a recursive scheme for bounding the overall number of WNN-vertices in a prism cell τ of the primary decomposition, described in Section 2. This scheme is based on a secondary decomposition, applied to the narrow and wide surfaces in τ. This decomposition is different from the one described in Section 2, as it is based on both narrow and wide surfaces (whereas the primary decomposition is based only on the narrow surfaces), and is therefore more involved. The general idea is to progressively partition space into prism cells Δ (of a somewhat different kind than those in the preliminary decomposition). Once a narrow surface σ becomes wide in a cell Δ, we count all the *new* WWW- and WWN-vertices that it is involved in (that is, these vertices were of type WNN in the parent cell of Δ). We do it collectively for all such newly wide surfaces σ, and then dispose of all these surfaces, σ, and continue to process recursively (i.e., hierarchically) only the leftover narrow surfaces in Δ, and the associated WNN-vertices that they form. We emphasize that we distinguish between the originally given wide surfaces, which define the underlying set of WNN-vertices, and the newly wide surfaces that were narrow in a parent cell. We do not keep processing this latter type of wide surfaces, but dispose of them immediately, in contrary, the original wide surfaces remain intact and are processed by a recursive mechanism – see below.

We first review a useful structural property concerning the overlay of the lower envelope of the wide surfaces with the vertical walls erected from the boundaries of the narrow surfaces, and then describe the actual construction and its analysis.

4.1 Overlaying the lower envelope of wide surfaces with vertical walls

Let τ be a prism cell from the primary decomposition, and denote by τ^* its xy-projection. As in Section 3, we denote by $\mathcal{W} := \mathcal{W}_\tau$ (resp., $\mathcal{N} := \mathcal{N}_\tau$) the sets of surfaces from Σ that intersect τ and are wide (resp., narrow) in τ. We assume that $\mathcal{W} \neq \emptyset$, for otherwise there are no WNN-vertices in τ. For simplicity of presentation, we denote the sizes of both $\mathcal{W}$ and $\mathcal{N}$ by m (then m replaces the notation n_τ used earlier).

Let $\mathcal{F} \subseteq \mathcal{W}$, $\mathcal{G} \subseteq \mathcal{N}$ be two subsets, each of at most $k \leq m$ elements; these will be appropriate random samples, whose precise definition will be given shortly. We form the minimization diagram $M(\mathcal{F})$ of the surfaces of $\mathcal{F}$ (confined to τ^*). Arguing as above, we claim that the overall number of intersections of edges of $M(\mathcal{F})$ with $\partial\tau^*$ is $O(\lambda_q(k))$, for an appropriate constant q. Indeed, any such intersection is obtained as (the xy-projection of) a breakpoint of the lower envelope of the cross sections of the surfaces in $\mathcal{F}$ within H_{τ^*}, where H_{τ^*} is the vertical wall erected from $\partial\tau^*$. Since τ^* has constant complexity, the number of such breakpoints is indeed bounded by the bound stated above, with a suitable value of q. By combining this property with Lemma 7(i), it follows that the complexity of $M(\mathcal{F})$ clipped to τ^* is $O(k^2)$.

Next, we project the boundary curves of the surfaces in $\mathcal{G}$ onto the xy-plane, and let Γ be the resulting collection of planar curves, all clipped to τ^*. A curve in Γ may be split into several connected portions by this clipping, but since τ^* has constant complexity, it follows that $|\Gamma| = O(k)$.

▶ **Lemma 9.** *The complexity of the overlay $\mathcal{O}(\mathcal{F}, \Gamma)$ of $M(\mathcal{F})$ and Γ (confined to τ^*) is $O(k^2\beta(k))$.*

Proof. Observe that any vertex of the overlay (lying in τ^*) is either an original vertex of $M(\mathcal{F})$, or an intersection between a pair of curves in Γ, or an intersection between a curve $\gamma \in \Gamma$ with an edge of $M(\mathcal{F})$. The number of vertices of the first kind is $O(k^2)$ by part (i) of Lemma 7. The number of vertices of the second kind is also $O(k^2)$, as each such vertex is a vertex in the arrangement $\mathcal{A}(\Gamma)$. The number of vertices of the third kind is bounded as follows.

Let γ be a curve in Γ, and consider the vertical wall V_γ consisting of the union of all z-vertical lines passing through the points of γ. Consider the collection $\mathcal{F}_\gamma$ of all the intersection curves of V_γ with all the surfaces in $\mathcal{F}$. Then a vertex of the overlay (of the third kind) corresponds to a vertex of the lower envelope of the curves of $\mathcal{F}_\gamma$ (within the vertical wall V_γ). Therefore its complexity is[4] $O(\lambda_q(|\mathcal{F}_\gamma|)) = O(\lambda_q(k))$, for some absolute constant q that depends on the complexity of the surfaces in Σ. As before, we write this bound as $O(k\beta(k))$. Summing this bound over all curves in Γ, we obtain an overall complexity of $O(k^2\beta(k))$, as asserted. ◀

Let $H(\Gamma)$ be the set of the vertical walls erected from the curves in Γ. Consider the three-dimensional arrangement $\mathcal{A}(\mathcal{F} \cup H(\Gamma))$, clipped to τ, and its portion that lies below the lower envelope $E(\mathcal{F})$ (note that the only effect of $H(\Gamma)$ is to form curves on $E(\mathcal{F})$ where

[4] The various constant parameters q may differ from one another; for uniformity, we replace each of them by their maximum value.

its walls cross the envelope, but it otherwise does not hide any portion of the envelope). By Lemma 9, we obtain a decomposition of this portion into $O(k^2\beta(k))$ semi-unbounded prisms. Indeed, since the complexity of $\mathcal{O}(\mathcal{F},\Gamma)$ is $O(k^2\beta(k))$, this is also an asymptotic bound on its corresponding (planar) vertical decomposition [10]. By clipping the cells of this decomposition (pseudo-trapezoids, or trapezoids for short) to τ^* and lifting them in the z-direction, until they hit the lower envelope $E(\mathcal{F})$, we obtain a three-dimensional decomposition of the portion of $\mathcal{A}(\mathcal{F} \cup H(\Gamma))$ that lies below $E(\mathcal{F})$. Each resulting cell is a semi-unbounded *open* prism Δ of constant complexity, with a unique (wide) ceiling surface. It is fairly standard to show that Δ is defined by at most five surfaces, four defining its base (that is, its xy-projection) and one forming its ceiling; see [10] for further details. A structure of this kind will be applied in our construction of the secondary decomposition, as described next.

4.2 Hierarchical construction via gradation

Our construction is based on a framework introduced by Ramos [35], and proceeds as follows. We set a parameter $r = \log^c m$, for some constant $c > 1$,[5] and put $l := \log_r m < \log m/\log\log m$. We say that $R \subseteq Y$ is a p-sample of a ground set Y if it is formed by drawing each element $y \in Y$ independently with probability p. We construct a *gradation*

$$\emptyset = \mathcal{W}_{-1} \subseteq \mathcal{W}_0 \subseteq \cdots \subseteq \mathcal{W}_{l-1} \subseteq \mathcal{W}_l = \mathcal{W}$$

for the wide surfaces, where $\mathcal{W}_i$ is a $1/r$-sample of $\mathcal{W}_{i+1}$, for $i = 0,\ldots,l-1$. It therefore follows that $\mathcal{W}_i = \mathcal{W}_{l-(l-i)}$ is a (r^i/m)-sample of $\mathcal{W} = \mathcal{W}_l$, as is easily verified (for example, $\mathcal{W}_{l-1}$ is a $(r^{l-1}/m = 1/r)$-sample, as per its definition). Put $p_i := r^i/m$. We define the gradation

$$\emptyset = \mathcal{N}_{-1} \subseteq \mathcal{N}_0 \subseteq \cdots \subseteq \mathcal{N}_{l-1} \subseteq \mathcal{N}_l = \mathcal{N}$$

for the narrow surfaces in an analogous manner. As done in [35], we consider a "forward view" of this sampling, that is, we write $\mathcal{W}_i$ as $\mathcal{W}_{i-1} \cup R_i^W$, where $R_i^W \subseteq \mathcal{W}_i$ is the subset of elements not chosen into $\mathcal{W}_{i-1}$, so it is itself a random $(1-1/r)$-sample of $\mathcal{W}_i$, which we can also interpret as a q_i-sample from $\mathcal{W} \setminus \mathcal{W}_{i-1}$ with q_i satisfying $p_i = q_i(1-p_{i-1}) + p_{i-1}$, or $q_i = (p_i - p_{i-1})/(1-p_{i-1})$. It is easy to verify that $q_i < 2p_i$, for any $r \geq 2$. Similarly, $\mathcal{N}_i$ is written as $\mathcal{N}_{i-1} \cup R_i^N$, where R_i^N is a q_i-sample from $\mathcal{N} \setminus \mathcal{N}_{i-1}$.

The construction. We proceed over the iterations i of the gradation, and incrementally construct, for each i, a structure Ξ_i, defined as follows. Initially, $i = -1$, and Ξ_i consists of the single cell τ from the primary structure, its conflict list consists of the wide surfaces in $\mathcal{W}$ and the narrow surfaces in $\mathcal{N}$. At any further iteration $i \geq 0$, we refine Ξ_{i-1}, in order to obtain Ξ_i, using the following procedure.

Denote by $H(\mathcal{N}_i)$ the set of vertical walls erected from the boundaries of the surfaces in $\mathcal{N}_i$ (clipped to τ). We next form the portion of the arrangement $\mathcal{A}(\mathcal{W}_i \cup H(\mathcal{N}_i))$ lying below the lower envelope of $\mathcal{W}_i$, as discussed in Section 4.1 above. Denote this structure by $E(\mathcal{W}_i, H(\mathcal{N}_i))$.[6] We clip $E(\mathcal{W}_i, H(\mathcal{N}_i))$ to each cell $\Delta_0 \in \Xi_{i-1}$. Specifically, we construct the lower envelope of $\mathcal{W}_i$, and overlay it with the vertical walls in $H(\mathcal{N}_i)$. We then take

[5] This choice of r is determined by our analysis – see Lemma 10 and its proof.
[6] The structure $E(\mathcal{W}_i, H(\mathcal{N}_i))$ is in fact the overlay of the lower envelope $E(\mathcal{W}_i)$ of $\mathcal{W}_i$ and the vertical walls in $H(\mathcal{N}_i)$.

the resulting structure and clip it, for each parent cell $\Delta_0 \in \Xi_{i-1}$, to the unbounded prism formed by the vertical walls of Δ_0. Observe that this step also forms a clipping to the actual cell Δ_0, since the (wide) surface defining the ceiling of Δ_0 is already included in $\mathcal{W}_i$ (by the gradation property).

We construct the vertical decomposition of $E(\mathcal{W}_i, H(\mathcal{N}_i))$, clipped to the vertical prism spanned by Δ_0, by projecting its vertices and edges onto the xy-plane, by forming the two-dimensional vertical decomposition of the resulting planar decomposition, and then by lifting each of its cells (trapezoids) in the z-direction until it hits the lower envelope of $\mathcal{W}_i$, and by extending them all the way down in the negative z-direction. (Note that Δ_0 is contained in the primal "master" cell τ, so this construction is in particular confined to the vertical prism τ. The unboundedness in the negative z-direction is not affected by this constraint.) This resulting set of semi-unbounded prisms Δ comprises the refinement of the parent cell Δ_0. We repeat this process for each $\Delta_0 \in \Xi_{i-1}$, and collect all newly constructed prisms Δ to obtain Ξ_i.

Given a cell $\Delta \in \Xi_i$, we denote by $\mathcal{W}_\Delta \subseteq \mathcal{W}$ (resp., $\mathcal{N}_\Delta \subseteq \mathcal{N}$) the set of the wide (resp., narrow) surfaces that intersect Δ, where a surface σ intersecting Δ is wide in Δ if the xy-projection of σ fully covers the xy-projection of Δ, and narrow otherwise (as defined as in Section 2). Note that some surface patches in $\mathcal{N}$ may become wide in Δ. In this case, we dispose of such surfaces immediately and, in particular, they are not used in any further step of the construction. The rationale for this removal is that we are after WNN-vertices, and when a surface becomes newly wide, any WNN-vertex it was involved in becomes a WWW- or a WWN-vertex, and we know how to bound the number of these vertices; see below for more details. We emphasize once again that we distinguish between the originally given wide surfaces, which we keep processing, and the newly wide surfaces from which we dispose, as above. We also comment that, by definition, if σ is wide in $\Delta_0 \in \Xi_{i-1}$, then it is also wide in any prism cell Δ of Ξ_i, whose xy-projection is contained in the xy-projection of Δ_0, given that $\sigma \cap \Delta \neq \emptyset$. Therefore the "wideness" property is preserved when we refine cells along the gradation. We next show:

▶ **Lemma 10.** *For any $0 \le i \le l-1$ we have:*

(1) Small conflict lists. *With high probability,*[7] *each cell $\Delta \in \Xi_i$ satisfies:*

$$|\mathcal{W}_\Delta| = O((m/r^i)\log r^i) \qquad \text{and} \qquad |\mathcal{N}_\Delta| = O((m/r^i)\log r^i).$$

(2) Small size. *The expected number of cells in Ξ_i is $O(r^{2i}\beta(r^i))$, for some absolute constant of proportionality, independent of r.*

In particular, properties (i)-(ii) hold, for all $0 \le i \le l-1$, with at least some constant probability.

Proof. Part (1) is straightforward by well-known properties of ε-nets [28], where we choose $\varepsilon = O\left(\frac{\log(r^i)}{r^i}\right)$. This implies that the expected size of $\mathcal{W}_i$ is indeed $O\left(\frac{1}{\varepsilon}\log\frac{1}{\varepsilon}\right)$. This suffices to guarantee that, with high probability, the size of each conflict list is at most εm, which is the asserted bound, for a suitable choice of the constant of proportionality. The asserted bound for the narrow surfaces follows by similar considerations. See [28] for details.

Regarding part (2), we bound the expected number of cells $\Delta \in \Xi_i$ by bounding the expected number of edges in the undecomposed structure $E(\mathcal{W}_i, H(\mathcal{N}_i))$, clipped to each of the vertical prisms spanned by the cells $\Delta_0 \in \Xi_{i-1}$. By standard properties of planar

[7] Specifically, this probability is at least $1 - 1/\mathrm{poly}(r^i)$.

vertical decompositions, it follows that the total complexity of Ξ_i is proportional to this measure. An edge of this kind either (i) is incident to a vertex of $E(\mathcal{W}_i, H(\mathcal{N}_i))$ that lies inside Δ_0, or (ii) is not incident to such a vertex, but is part of $E(\mathcal{W}_i, H(\mathcal{N}_i))$, that is, such an edge is a closed or an unbounded curve with no vertices, or (iii) crosses the boundary of (the closure of) the parent cell $\Delta_0 \in \Xi_{i-1}$ and is not incident to any vertex of $E(\mathcal{W}_i, H(\mathcal{N}_i))$ that lies inside Δ_0. Edges of types (i)–(ii) are charged to the overall expected complexity of $E(\mathcal{W}_i, H(\mathcal{N}_i))$, which, according to Lemma 9, is proportional to

$$\mathbf{Exp}\Big[(|\mathcal{W}_i| + |\mathcal{N}_i|)^2\beta(|\mathcal{W}_i| + |\mathcal{N}_i|)\Big] = O(r^{2i}\beta(r^i)),$$

by the fact that $\mathcal{W}_i$ (resp., $\mathcal{N}_i$) is a (r^i/m)-sample of $\mathcal{W}$ (resp., $\mathcal{N}$), and by the fact that this is a Bernoulli random variable, strongly concentrated around its mean.

We next analyze the expected number of edges of type (iii). Recall that $\mathcal{W}_i = \mathcal{W}_{i-1} \cup R_i^W$ and $\mathcal{N}_i = \mathcal{N}_{i-1} \cup R_i^N$. Let $\mathcal{W}_{\Delta_0}^{(i)} \subseteq R_i^W$ be the subset of the wide surfaces in R_i^W that intersect Δ_0, and similarly define $\mathcal{N}_{\Delta_0}^{(i)} \subseteq R_i^N$. Recall that R_i^W is a q_i-sample of $\mathcal{W} \setminus \mathcal{W}_{i-1}$, where $q_i < 2r^i/m$. Arguing as in part (1), with $i-1$ instead of i (because Δ_0 is obtained at step $i-1$ of the gradation), it follows that, with high probability, Δ_0 intersects at most $O((m/r^{i-1})\log r^{i-1})$ surfaces of $\mathcal{W}$. Therefore, by standard considerations from probability theory it follows that the expected size of $\mathcal{W}_{\Delta_0}^{(i)}$ is at most q_i times the number of these surfaces, namely it is at most

$$O\left(\frac{r^i}{m} \cdot \frac{m}{r^{i-1}} \log r^{i-1}\right) = O(r \log r^{i-1}).$$

Similar considerations imply that the expected size of $\mathcal{N}_{\Delta_0}^{(i)}$ is $O(r \log r^{i-1})$ as well.

A crossing edge e has to hit a vertical wall H_0 of Δ_0, and since e lies on $E(\mathcal{W}_i, H(\mathcal{N}_i))$, its crossing point must be (i) a breakpoint of the (univariate) lower envelope, within H_0, of a collection of $O(|\mathcal{W}_{\Delta_0}^{(i)}|)$ constant-complexity curves, which are formed by the intersections of H_0 with the surfaces in $\mathcal{W}_{\Delta_0}^{(i)}$, or (ii) an intersection point of this lower envelope with a vertical line formed by the intersection of a vertical wall in $H(\mathcal{N}_{\Delta_0}^{(i)})$ with H_0. By the above bounds on the expected sizes of $\mathcal{W}_{\Delta_0}^{(i)}$ and $\mathcal{N}_{\Delta_0}^{(i)}$, the expected number of such points of types (i) and (ii) is $O(\lambda_q(r \log r^{i-1}))$, for some absolute constant q, a bound that we write, as before, as $O(ri\beta(ri)\log r)$ (there is no asymptotic difference between $\beta(ri \log r)$ and $\beta(ri)$). Thus the expected number of edges of type (iii), summing over the at most four vertical walls of Δ_0, is $O(ri\beta(ri)\log r)$. Summing over all parent prisms $\Delta_0 \in \Xi_{i-1}$, and recalling that, by induction on i, $|\Xi_{i-1}| \le A \cdot (r^{2i-2}\beta(r^{i-1}))$, for a sufficiently large constant $A > 0$ (that does not depend on r), we conclude that the overall expected number of edges of type (iii) is at most

$$A \cdot (r^{2i-1} i\beta(r^{i-1})\beta(ri)\log r). \tag{3}$$

We note that this bound is negligible with respect to the other terms contributed by edges of type (i) and (ii). Indeed, recall that $r = \log^c m$ (where $c > 1$) and $i \le l < \log m/\log\log m$, from which it follows that the term $i\beta(ri)\log r$ is at most $O(r^{1/c}\beta(r))$. Therefore the bound in (3) is at most $A \cdot (r^{2i-1+1/c}\beta(r^{i-1})\beta(r))$, and is thus asymptotically significantly smaller than $A \cdot (r^{2i}\beta(r^i))$, as is easily verified.

Using standard properties from probability theory (e.g., Chernoff's bound), one can show that properties (1)-(2) are satisfied with constant probability, for all $0 \le i \le l$. We omit the straightforward details. ◀

Bounding the number of WNN-vertices. Let $\mathcal{W}$, $\mathcal{N}$, τ, and m be as above. Our goal is to bound the number of WNN-vertices in τ. We proceed over the iterations i of the gradation. At iteration $i = -1$ we do not take any action. For any further iteration $i \geq 0$, let Δ be a cell of the current structure Ξ_i, and let $\Delta_0 \in \Xi_{i-1}$ be the parent cell of Δ in the structure at iteration $i-1$. At the first iteration $i = 0$ we have $O(1)$ cells in Ξ_0. Note that some of the narrow surfaces in $\mathcal{N}_{\Delta_0}$ may become wide in Δ. For any such newly wide surface σ, each of the WNN-vertices that it formed in Δ_0 must become either a WWW- or a WWN-vertex in Δ. By Lemma 7(i,ii), the number of these vertices is at most $O(m^2\beta(m))$, and this is also the asymptotic bound, over all cells Δ of Ξ_0. We dispose of the newly wide surfaces in each cell Δ, and proceed to the next iteration.

Assume that $0 < i < l$. By Lemma 10(1), $|\mathcal{W}_{\Delta_0}| = O((m/r^{i-1})\log r^{i-1})$, $|\mathcal{N}_{\Delta_0}| = O((m/r^{i-1})\log r^{i-1})$, and $|\mathcal{W}_\Delta| = O((m/r^i)\log r^i)$, $|\mathcal{N}_\Delta| = O((m/r^i)\log r^i)$. As above, some of the narrow surfaces σ in $\mathcal{N}_{\Delta_0}$ may become newly wide in Δ, and these are the only newly wide surfaces that we encounter, since we have disposed of all the previous newly wide surfaces (which originally were narrow surfaces from $\mathcal{N}$). Hence the number of these newly wide surfaces does not exceed $O((m/r^{i-1})\log r^{i-1})$. As above, we need to consider, for each newly wide surface σ, the set of WWW- and WWN-vertices in Δ that it forms. By Lemma 7(i,ii), the number of these vertices is at most

$$O\left(\left(\frac{m\log r^{i-1}}{r^{i-1}}\right)^2 \beta\left(\frac{m\log r^{i-1}}{r^{i-1}}\right)\right) = O\left(\left(\frac{m\log r^{i-1}}{r^{i-1}}\right)^2 \beta(m)\right).$$

By Lemma 10(2), there are $O(r^{2i}\beta(r^i))$ cells $\Delta \in \Xi_i$. Therefore the total number of new WWW- and WWN-vertices, over all these cells, is

$$O(r^2m^2\beta(m)\beta(r^i)\log^2 r^{i-1}) = O(r^2m^2i^2\beta^2(m)\log^2 r).$$

Recall that $r = \log^c m$ and $i < \log m/\log\log m$, so the above bound is $O(m^2\beta^2(m)\log^{2+2c} m)$.

Consider now the last iteration $i = l$. For each cell $\Delta \in \Xi_l$, the sets $\mathcal{W}_\Delta$, $\mathcal{N}_\Delta$ become empty, however, Δ may intersect new wide surfaces, which were narrow at the parent cell Δ_0. In this case, a WNN-vertex in Δ is obtained by the intersection of a pair of new wide surfaces and the ceiling of Δ (which is a portion of a surface $\sigma \in \mathcal{W}$). We bound the number of such vertices by brute-force. Observe that the number of new wide surfaces in Δ is $O(\log^{c+1} m)$, and therefore the brute-force bound is $O(\log^{2c+2} m)$, for a total of $O(m^2\beta(m)\log^{2+2c} m)$, over all cells $\Delta \in \Xi_l$, which is subsumed by the aforementioned bound $O(m^2\beta^2(m)\log^{2+2c} m)$.

Summing these bounds over the at most $l + 1 < \log m/\log\log m + 1$ iterations, we obtain that the total number of WNN-vertices in τ is $O(m^2\beta^2(m)\log^{3+2c} m/\log\log m)$, which we can write as $O(m^2\log^{5+\varepsilon} m)$, for any $\varepsilon > 0$, as asserted in Lemma 7(iii).

It is easy to verify that any WNN-vertex v in τ is counted by this procedure. Indeed, either v becomes a WWW- or a WWN-vertex during the iterative process, and is counted when that happens, or v is counted at some "leaf cell" where $i = l$.

This at last completes the proof of Lemma 7(iii). ◀

This finally concludes the proof of Theorem 1.

Concluding remarks

In this paper we have presented a bound of $O(n^2\text{polylog}n)$ on the combinatorial complexity of the lower envelope of graphs of a partially defined semi-algebraic function of constant description complexity, where the relative interior of any triple of them intersect in at most two points. Whereas this improves the previous bound $O(n^2 \cdot 2^{c\sqrt{\log n}})$ (where $c > 0$ is a constant) shown by Halperin and Sharir [26], a further improvement, and, in particular, showing the conjectured bound $O(n^2\beta(n))$, has still remained a challenging open problem.

Another intriguing open problem is with regard to graphs of continuously defined bivariate functions, where the relative interior of triples of them intersect in $s \geq 3$ points. We conjecture that the bound should resemble the bound $O(n\lambda_{s+2}(n))$ shown in [37] under the favorable cross section assumption. However, showing such a bound without this assumption, even just for the case $s = 3$, has still remained elusive.

References

1 P. Afshani and K. Tsakalidis, Optimal deterministic shallow cuttings for 3-d dominance ranges. *Algorithmica*, 80(11): 3192–3206 (2018).

2 P. Afshani, T. M. Chan, K. Tsakalidis, Deterministic rectangle enclosure and offline dominance reporting on the RAM. In *Proc. Forty-First International Colloq. Automata, Languages Programming*, 8572(77–88), (2014).

3 P. K. Agarwal, B. Aronov, E. Ezra, and J. Zahl. An efficient algorithm for generalized polynomial partitioning and its applications. *SIAM J. Comput.*, 50:760–787, (2021).

4 P. K. Agarwal, B. Aronov, and M. Sharir, Computing envelopes in four dimensions with applications. *SIAM J. Comput.* 26(6): 1714–1732 (1997).

5 P. K. Agarwal, A. Efrat, and M. Sharir, Vertical decomposition of shallow levels in 3-dimensional arrangements and its applications. *SIAM J. Comput.* 29(3): 912–953 (1999).

6 P. K. Agarwal, B. Aronov, and M. Sharir, On the complexity of many faces in arrangements of pseudo-segments and of circles, in *Discrete and Computational Geometry: The Goodman-Pollack Festschrift* (B. Aronov, S. Basu, J. Pach, and M. Sharir, eds.), Springer Verlag, Berlin, pp. 1–24, (2003).

7 P. K. Agarwal, J. Erickson, and L. J. Guibas, Kinetic binary space partitions for intersecting segments and disjoint triangles, in *Proc. 9th Annual ACM-SIAM Sympos. Discrete Algorithms*, pp. 107–116, (1998).

8 P. K. Agarwal, J. Matoušek, and O. Schwarzkopf, Computing many faces in arrangements of lines and segments. *SIAM J. Comput.* 27(2): 491–505 (1998).

9 P. Agarwal and M. Sharir, Efficient randomized algorithms for some geometric optimization problems. *Discrete Comput. Geom.*, 16:317–337 (1996).

10 P. K. Agarwal and M. Sharir, Arrangements and their applications, In *Handbook of Computational Geometry*, (J. Sack and J. Urrutia, eds.), Elsevier, Amsterdam, pages 973–1027, 2000.

11 P. K. Agarwal and M. Sharir, Pseudoline arrangements: Duality, algorithms and applications, *SIAM J. Comput.* 34:526–552, (2005).

12 G. Alexandron, H. Kaplan, and M. Sharir, Kinetic and dynamic data structures for convex hulls and upper envelopes. *Comput. Geom.*, 36(2):144–158 (2007).

13 N. Alon, and J. H. Spencer, The Probabilistic Method, Third Edition. Wiley-Interscience series in discrete mathematics and optimization, Wiley 2008.

14 M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars, *Computational Geometry: Algorithms and Applications*, 3rd Ed., SpringerVerlag, Berlin-Heidelberg, 2008.

15 M. de Berg, and O. Schwarzkopf, Cuttings and applications. *Int. J. Comput. Geom. Appl.*, 5(4):343–355 (1995).

16 J. D. Boissonnat and K. Dobrindt, On-line construction of the upper evelope of triangles and surface patches in three dimensions, *Comput. Geom.* 5:303–320 (1995)

17 Timothy M. Chan and Konstantinos Tsakalidis, Optimal deterministic algorithms for 2-d and 3-d shallow cuttings. *Discrete Comput. Geom.*, 56(4): 866-881 (2016).

18 B. Chazelle, Cutting hyperplanes for divide-and-conquer. *Discrete Comput. Geom.*, 9(1), 145–158 (1993).

19 B. Chazelle, H. Edelsbrunner, L. Guibas and M. Sharir, A singly exponential stratification scheme for real semi–algebraic varieties and its applications, *Theoret. Comput. Sci.* 84:77–105,

(1991). Also in *Proc. 16th Int. Colloq. on Automata, Languages and Programming*, 179–193, 1989.

20 B. Chazelle and J. Friedman, A deterministic view of random sampling and its use in geometry, *Combinatorica* 10:229–249 (1990).

21 K. L. Clarkson, New applications of random sampling in computational geometry, *Discrete Comput. Geom.*, 2:195–222 (1987).

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

23 L. J. Guibas, Kinetic data structures: A state of the art report, In *Robotics: The Algorithmic Perspective, Proceedings of the 3rd Workshop on the Algorithmic Foundations of Robotics* (P. K. Agarwal, L. E. Kavraki, and M. T. Mason, eds.), 1998, A K Peters/CRC Press, pp. 191–209.

24 L. Guth. Polynomial partitioning for a set of varieties. *Math. Proc. Camb. Phil. Soc.*, 159:459–469 (2015).

25 L. Guth and N. H. Katz, On the Erdős distinct distances problem in the plane, *Annals Math.* 181:155–190 (2015).

26 D. Halperin and M. Sharir, New bounds for lower envelopes in three dimensions, with applications to visibility in terrains. *Discrete Comput. Geom.*, 12: 313-326 (1994).

27 S. Har-Peled, H. Kaplan, and M. Sharir, Approximating the k-level in three dimensional plane arrangements. In M. Loebl, J. Nešetřil, and R. Thomas, editors, *Journey through Discrete Mathematics: A Tribute to Jiri Matoušek*, pages 467–504. Springer-Verlag, Berlin Heidelberg, 2017.

28 D. Haussler and E. Welzl, Epsilon-nets and simplex range queries, *Discrete Computat. Geom.*, 2:127–151 (1987).

29 V. Koltun, Almost tight upper bounds for vertical decompositions in four dimensions, *J. ACM* 51(5):699–730 (2004).

30 V. Koltun, Sharp bounds for vertical decompositions of linear arrangements in four dimensions, *Discrete Comput. Geom.* 31(3):435–460 (2004).

31 V. Koltun and M. Sharir, The partition technique for the overlay of envelopes. *SIAM J. Comput.* 32:841–863 (2003).

32 J. Matoušek, Efficient partition trees. *Discrete Comput. Geom.* 8:315–334 (1992).

33 J. Matoušek, Reporting points in halfspaces. *Comput. Geom. Theory Appl.*, 2:169–186 (1992).

34 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).

35 E. A. Ramos, On range reporting, ray shooting and k-level construction. In *Proc. Fifteenth Ann. Sympos. Comput. Geom.*, 390–399 (1999).

36 J. T. Schwartz and M. Sharir, On the Piano Movers' problem: II. General techniques for computing topological properties of real algebraic manifolds, *Advances in Appl. Math.*, 4:298–351 (1983).

37 J. T. Schwartz and M. Sharir, On the two-dimensional Davenport-Schinzel problem, *J. Symbolic Comput.*, 10:371–393 (1990).

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

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

40 B. Tagansky, A new technique for analyzing substructures in arrangements of piecewise linear surfaces, *Discrete Comput. Geom.*, 16:455–479 (1996).

41 A. Wiernik and M. Sharir, Planar realizations of nonlinear Davenport-Schinzel sequences by segments, *Discrete Comput. Geom.*, 3:15–47 (1988).