

Approximating the k -Level in Three-Dimensional Plane Arrangements*

Sariel Har-Peled[†]

Haim Kaplan[‡]

Micha Sharir[§]

October 13, 2015

Abstract

Let H be a set of n non-vertical planes in three dimensions, and let $r < n$ be a parameter. We give a simple alternative proof of the existence of a $O(1/r)$ -cutting of the first n/r levels of $\mathcal{A}(H)$, which consists of $O(r)$ semi-unbounded vertical triangular prisms. The same construction yields an approximation of the (n/r) -level by a terrain consisting of $O(r/\varepsilon^3)$ triangular faces, which lies entirely between the levels $(1 \pm \varepsilon)n/r$. The proof does not use sampling, and exploits techniques based on planar separators and various structural properties of levels in three-dimensional arrangements and of planar maps. The proof is constructive, and leads to a simple randomized algorithm, that computes the terrain in $O(n + r^2\varepsilon^{-6} \log^3 r)$ expected time. An application of this technique allows us to mimic Matoušek’s construction of cuttings in the plane [36], to obtain a similar construction of “layered” $(1/r)$ -cutting of the entire arrangement $\mathcal{A}(H)$, of optimal size $O(r^3)$. Another application is a simplified optimal approximate range counting algorithm in three dimensions, competing with that of Afshani and Chan [1].

1. Introduction

Cuttings. Let H be a set of n (non-vertical) hyperplanes in \mathbb{R}^d , and let $r < n$ be a parameter. A $(1/r)$ -

cutting of the arrangement $\mathcal{A}(H)$ is a collection of pairwise openly disjoint simplices (or other regions of constant complexity) whose union covers \mathbb{R}^d , such that each simplex is crossed (meets in its interior) at most n/r hyperplanes of H . Cuttings have proved to be a powerful tool for a variety of problems in discrete and computational geometry, because they provide an effective divide-and-conquer mechanism for tackling such problems, see Agarwal [7] for a survey. Applications include a variety of range searching techniques [11], partition trees [37], incidence problems involving points and lines, curves, and surfaces [24], and many more.

The first (albeit suboptimal) construction of cuttings is due to Clarkson [23]. This concept was formalized later on by Chazelle and Friedman [22], who gave a sampling-based construction of optimal-size cuttings (see below). An optimal deterministic construction algorithm was provided by Chazelle [20]. Matoušek [39] studied the number of cells in a $(1/r)$ -cutting in the plane (see also [28]). See Agarwal and Erickson [11] for a comprehensive review of this topic.

To be effective, it is imperative that the number of simplices in the cutting be asymptotically as small as possible. Chazelle and Friedman [22] were the first to show the existence of a $(1/r)$ -cutting of the entire arrangement of n hyperplanes in \mathbb{R}^d , consisting of $O(r^d)$ simplices, which is asymptotically the best possible bound. (We note in passing that cuttings of optimal size are not known for arrangements of (say, constant-degree algebraic) surfaces in \mathbb{R}^d , except for $d = 2$, where the known bound, $O(r^2)$, is tight, and for $d = 3, 4$, where nearly tight bounds are known [21, 33, 34].)

For additional work related to cuttings and their applications, see [2, 3, 4, 5, 6, 8, 9].

Shallow cuttings. The *level* of a point p in the arrangement $\mathcal{A}(H)$ of H is the number of hyperplanes lying vertically below it (that is, in the $(-x_d)$ -direction). The k -*level*, denoted as L_k , is the closure of all the points that lie on some hyperplane of H and are at level exactly k , and the $(\leq k)$ -level, denoted as $L_{\leq k}$, is the union of all the j -levels, for $j = 0, \dots, k$. A collection of pairwise openly disjoint simplices whose union covers $L_{\leq k}$, such that each simplex is crossed at most n/r hyperplanes

*Work by Sariel Har-Peled was partially supported by NSF AF awards CCF-1421231 and CCF-1217462. Work by Haim Kaplan was partially supported by grant 1161/2011 from the German-Israeli Science Foundation, by grant 822/10 from the Israel Science Foundation, and by the Israeli Centers for Research Excellence (I-CORE) program (center no. 4/11). Work by Micha Sharir has been supported by Grant 2012/229 from the U.S.-Israel Binational Science Foundation, by Grant 892/13 from the Israel Science Foundation, by the Israeli Centers for Research Excellence (I-CORE) program (center no. 4/11), and by the Hermann Minkowski-MINERVA Center for Geometry at Tel Aviv University.

[†]Department of Computer Science, University of Illinois, 201 N. Goodwin Avenue, Urbana, IL, 61801, USA. E-mail: sariel@illinois.edu; url: <http://sarielhp.org/>

[‡]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel. E-mail: haimk@post.tau.ac.il

[§]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel. E-mail: michas@post.tau.ac.il

of H , is called a k -shallow $(1/r)$ -cutting. Naturally, the parameters k and r can vary independently, but the interesting case, which is the one that often arises in many applications, is the case where $k = \Theta(n/r)$. In fact, shallow cuttings for any value of k can be reduced to this case—see Chan and Tsakalidis [19, Section 5].

In his seminal paper on reporting points in half-spaces [38], Matoušek has proved the existence of small-size shallow cuttings in arrangements of hyperplanes in any dimension, showing that the bound on the size of the cutting can be significantly improved for shallow cuttings. Specifically, he has shown the existence of a k -shallow $(1/r)$ -cutting, for n hyperplanes in \mathbb{R}^d , whose size is $O(q^{\lceil d/2 \rceil} r^{\lfloor d/2 \rfloor})$, where $q = k(r/n) + 1$. For the interesting special case where $k = \Theta(n/r)$, we have $q = O(1)$ and the size of the cutting is $O(r^{\lfloor d/2 \rfloor})$, a significant improvement over the general bound $O(r^d)$. (For example, in three dimensions, we get $O(r)$ simplices, instead of $O(r^3)$ simplices for the whole arrangement.) This has led to improved solutions of many range searching and related problems.

In his paper, Matoušek presented a deterministic algorithm that can construct such a shallow cutting in polynomial time; the running time improves to $O(n \log r)$ but only when r is small, i.e., $r < n^\delta$ for a sufficiently small constant δ (that depends on the dimension d). Later, Ramos [45] presented a (rather complicated) randomized algorithm for $d = 2, 3$, that constructs a hierarchy of shallow cuttings for a geometric sequence of $O(\log n)$ values of r , in $O(n \log n)$ expected time. Recently, Chan and Tsakalidis [19] provided a deterministic $O(n \log r)$ -time algorithm for computing $O(n/r)$ -shallow $(1/r)$ -cutting. Their algorithm can also construct a hierarchy of shallow cuttings for a geometric sequence of $O(\log n)$ values of r in $O(n \log n)$ deterministic time. Interestingly, they use, as a black box in the analysis of their algorithm, Matoušek’s theorem on the existence of an $O(n/r)$ -shallow $(1/r)$ -cutting of size $O(r)$.

Each simplex Δ in the cutting has a *conflict list* associated with it, which is the set of hyperplanes intersecting Δ . The algorithms mentioned above for computing cuttings also compute the *conflict lists* associated with the simplices of the cutting. Alternatively give the cutting one can produce the conflict lists in $O(n \log r)$ time using a result of Chan [16], as we outline in [Section 3.2](#).

Matoušek’s proof of the existence of small-size shallow cuttings, as well as subsequent studies of this technique, are fairly complicated. They rely on random sampling, combined with a clever variant of the so-called exponential decay lemma of [22], and with several additional (and rather intricate) techniques.

Approximating a level. An early study of Matoušek [36] gives a construction of a $(1/r)$ -cutting of small (optimal) size in arrangements of lines in the plane. The construction chooses a sequence of r levels, n/r apart from one another, and approximates each of them by a coarser polygonal line, by choosing every $n/(2r)$ th vertex of the level, and by connecting them by an x -monotone polygonal path. Each approximate level does not deviate much from its original level, so they remain disjoint from one another. Then, partitioning the region between every pair of consecutive approximate levels into vertical trapezoids produces a total of $O(r^2)$ such trapezoids, each crossed by at most $O(n/r)$ lines.

It is thus natural to ask whether one can approximate, in a similar fashion, a k -level of an arrangement of planes in 3-space. This is significantly more challenging, as the k -level is now a polyhedral terrain, and while it is reasonably easy to find a good set of vertices that “represent” this level (in an appropriate sense, detailed below), it is less clear how to triangulate them effectively to form an xy -monotone terrain, such that (i) none of its triangles is crossed by too many planes of H , and (ii) it remains close to the original level. To be more precise, given k and $\varepsilon > 0$, we want to find a polyhedral terrain with a small number of faces, which lies entirely between the levels k and $(1 + \varepsilon)k$ of $\mathcal{A}(H)$. A simple tweaking of Matoušek’s technique produces such an approximation in the planar case, but it is considerably more involved to do it in 3-space.

Algorithms for terrain approximation, such as in [10], do not apply in this case, as they have a quadratic blowup in the output size, compared to the optimal approximation. Also, they are not geared at all to handle our measure of approximation (in terms of lying close to a specified level, in the sense that no point on the approximation is separated by too many planes from the level).

Such an approximation to the k -level, whose size is optimal up to polylogarithmic factors, can be obtained by using a *relative-approximation* sample of the planes, and by extracting the appropriate level in the sample [29]. A more natural approach, of using the triangular faces of an optimal-size shallow cutting to form an approximate k -level, seems to fail in this case, as the shallow cutting is in general just a collection of simplices, stacked on top of one another, with no clearly defined xy -monotonicity. Such a monotonicity is obtained in Chan [17], but the resulting cuttings do not lead to sharp approximations of the level, of the sort we seek.

In short, a simple, effective, and optimal technique for approximating a level in three dimensions (let alone in higher dimensions) does not follow easily from exist-

ing techniques.

An additional advantage of such an approximation is that it immediately yields a simply-shaped shallow cutting of the first k levels of $\mathcal{A}(H)$, by replacing each triangle Δ of the approximate level by the semi-unbounded triangular prism Δ^* having Δ as its top face. Such a cutting (by prisms) has already been constructed by Chan [17], but it does not yield (that is, come from) an approximation to the level. Such a shallow cutting, by triangular semi-unbounded prisms, was a central tool in Chan’s algorithm for dynamic convex hulls in three dimensions [18].

Thus, resolving the question of approximating the k -level by an xy -monotone terrain of small, optimal size is not a mere technical issue, but rather a tool that will shed more light on the geometry of arrangements of planes, and that has applications to a variety of problems. For example, it yields an efficient algorithm for approximating the level of a point in an arrangement of planes in \mathbb{R}^3 —see Section 4.2 for details. (Afshani and Chan [1] present a similar approach which is slightly more involved, as they do not have the desired terrain property.)

1.1. Our results In this paper we give an alternative, simpler and constructive proof of the existence of optimal-size shallow cuttings by vertical semi-unbounded triangular prisms in three dimensions. With a bit more care, the construction yields an optimal-size approximate level, as discussed above. Specifically, given r and ε , one can approximate the (n/r) -level in an arrangement of n non-vertical planes in \mathbb{R}^3 , by a polyhedral terrain of complexity $O(r/\varepsilon^3)$, that lies entirely between the levels n/r and $(1 + \varepsilon)n/r$. The same construction works for any values of the level k and the parameter r , with a somewhat more involved bound on the complexity of the approximation (which we do not spell out in this version).

The construction does not use sampling, nor does it use the exponential decay lemma of [22, 38]. It is based on the planar separator theorem of Lipton and Tarjan [35], or, more precisely, on recent separator-based decomposition techniques of planar maps, as in Klein *et al.* [32] (see also Frederickson [27]), and on several insights into the structure and properties of levels in three dimensions and of planar maps, which we believe to be of independent interest.

As what we believe to be an interesting application of our technique, we extend Matoušek’s construction [36] of cuttings in planar arrangements to three dimensions. That is, we construct a “layered” $(1/r)$ -cutting of the entire arrangement $\mathcal{A}(H)$ of a set H of n non-vertical planes in \mathbb{R}^3 , of optimal size $O(r^3)$, by

approximating each level in a suitable sequence of levels, and then by triangulating each layer between consecutive levels in the sequence. The analysis becomes considerably more involved in three dimensions, and requires several known but interesting and fairly advanced properties of plane arrangements.

Another application of our technique is to approximate range counting. Specifically, we show how to preprocess a set H of n non-vertical planes in \mathbb{R}^3 , and a prescribed error parameter $\varepsilon > 0$, in near-linear time (in n), into a data structure of size $O(n/\varepsilon^{8/3})$, so that, given a query point $q \in \mathbb{R}^3$, we can compute the number of planes of H lying below q , up to a factor $1 \pm \varepsilon$, in $O(\log(n/(\varepsilon k)))$ expected time. As noted, this competes with Afshani and Chan’s technique [1]. The general approach is similar in both solutions, but our solution is somewhat simpler, due to the availability of approximating terrains, and the dependence on ε in our solution is explicit and reasonable (this dependence is not explicit in [1]).

The thrust of this paper is to show, via alternative, simpler, and more geometric methods, the existence of cuttings and approximate levels of optimal size. The proofs are constructive, but naive implementations thereof would be rather inefficient. Nevertheless, using standard random sampling techniques, we can obtain simple randomized algorithms that perform (suitable variants of) these constructions efficiently. Specifically, they run in near-linear expected time (which becomes linear when r is not too large).

Sketch of our technique. The k -level in a plane arrangement in three dimensions is an xy -monotone polyhedral terrain. After triangulating each of its faces, its xy -projection forms a (straight-edge) triangulated biconnected planar graph. Since the average complexity of the first k levels is $O(nk^2)$ (see, e.g., [25]), we may assume, by moving from a specified level to a nearby one, that the complexity of our level is $O(nk)$. The decomposition techniques of planar graphs mentioned above (as in [32]) allow us to partition the level into $O(n/k)$ clusters, where each cluster has $O(k^2)$ vertices and at most ck boundary vertices (that also belong to other clusters), for some sufficiently small prescribed fraction c . (In the terminology of [32], this is a k^2 -division of the graph.) Each such cluster, projected to the xy -plane, is a polygon with $O(k)$ boundary edges (and with $O(k^2)$ interior diagonals). We show that, replacing each such projected polygon by its convex hull results in a collection of $O(n/k)$ convex *pseudo-disks*, namely, each region is simply connected, and the boundaries of any pair of regions intersect at most twice. Moreover, the decomposition has the property that, for each triangle Δ that is a fully contained in

such a pseudo-disk, lifting its vertices back to the k -level yields a triple of points that span a triangle Δ' with a small number of planes crossing it, so it lies close to the k -level.

An old result of Bambah and Rogers [15], proving a statement due to L. Fejes-Tóth, and reviewed in [44, Lemma 3.9], shows that a union of m convex pseudo-disks that covers the plane induces a triangulation of the plane by $O(m)$ triangles, such that each triangle is fully contained inside one of the pseudo-disks. (As a matter of fact, it shows that the pseudo-disks can be replaced by smaller pairwise openly disjoint convex polygons, with the same union, so that the total number of edges of the polygons is at most $6m$; the desired triangulation is obtained by simply triangulating each of these polygons.) Lifting (the vertices of) this triangulation to the k -level, with a corresponding lifting of its triangular faces, results in the desired terrain approximating the level. A significant technical contribution of this paper is to provide an alternative proof of this result. The original proof in [15] appears to be fairly involved, although its presentation in [44] is simplified. Still, it does not seem to lead to a sufficiently efficient construction. Our proof in contrast does lead to such a construction, as described in Section 2.

A shallow cutting of the first k levels is obtained by simply replacing each triangle Δ in the approximate level by the semi-unbounded vertical prism of points lying below Δ .

Confined triangulations. The idea of decomposing the union of objects (pseudo-disks here) into pairwise openly disjoint simply-shaped fragments, each fully contained in some original object, is implicit in algorithms for efficiently computing the union of objects; see the work of Ezra *et al.* [26], which was in turn inspired by Mulmuley’s work on hidden surface removal [42]. Mustafa *et al.* [43] use a more elaborate version of such a decomposition, for situations where the objects are weighted. While these decompositions are useful for a variety of applications, they still suffer from the problem that the complexity of a single region in the decomposition might be arbitrarily large. In contrast, the triangulation scheme that we use (following [15]) is simpler, optimal, and independent of the complexity of the relevant pseudo-disks. We are pleased that this nice property of pseudo-disks is (effectively) applicable to the problems studied here, and expect it to have many additional potential applications.

In particular, we extend our analysis, and show that such a decomposition exists for arbitrary convex shapes, with the number of pieces being proportional to the union complexity, and with each region being a triangle or a cap (i.e., the intersection of an input shape with

a halfplane). This provides a representation of “most” of the union by triangles, where the more complicated caps are only used to fill in the “fringe” of the union. We believe that this triangulation should be useful in practice, in situations where deciding if a point belongs to an input shape is significantly more expensive than deciding if a point is inside a triangle.

Paper organization. We start by presenting the construction of the confined triangulation in Section 2. We then describe the construction of approximate levels, and the construction of shallow cuttings that it leads to, in Section 3. We then present applications of our results in Section 4. Specifically, in Section 4.1 we show how to build a layered cuttings of the whole arrangement, and in Section 4.2 we show how to answer approximate range counting queries for halfspaces.

2. Triangulating the union of convex shapes

In this section we show that, given a finite collection of m convex pseudo-disks covering the plane, one can construct a triangulation of the plane, consisting of $O(m)$ triangles, such that each triangle is contained in a single original pseudo-disk—see Theorem 2.1 below for details. As a matter of fact, our result can be extended to situations where the union of the pseudo-disks is not the entire plane; see below. This claim is a key ingredient in our construction of approximate k -levels, detailed in Section 3, but it is not new, as it is an immediate consequence of an old result of Bambah and Rogers [15] (proving a statement by L. Fejes-Tóth), which can also be found in Pach and Agarwal [44, Lemma 3.9]. We present here a constructive proof that is different from those in [15, 44], which leads to an $O(m \log m)$ -time algorithm for constructing the triangulation for a set of m pseudo-disks, in a suitable model of computation. (As an aside, we also think that such a nice property deserves more than one proof.) We also establish an extension of this result for more general convex shapes.

2.1. Preliminaries

The notion of a triangulation that we use here is slightly non-standard, as it is a triangulation of the entire plane, and not just of the convex hull of some input set of points. As such, it contains unbounded triangles, where the boundary of each such triangle consists of one bounded segment and two unbounded rays.

Given a convex shape D , a *cap* of D is the region formed by the intersection of D with a halfplane. A *crescent* is a portion of a cap obtained by removing





Figure 2.1: A crescent.

from it a convex polygon that has the base chord of the cap as an edge, see [Figure 2.1](#).

DEFINITION 2.1. *Given a collection \mathcal{D} of convex shapes in the plane, a decomposition \mathcal{T} of their union into pairwise openly disjoint regions is a **confined triangulation**, if (i) every region in \mathcal{T} is either a triangle or a cap, and (ii) every such region is fully contained in one of the original input shapes.*

See [Figure 2.2](#) for an example of a confined triangulation, and [Theorem 2.2](#) for a formal statement of the result.

2.2. Construction We are given a collection \mathcal{D} of m convex pseudo-disks, and our purpose is to construct a confined triangulation for \mathcal{D} , as described above. In what follows we consider both the case where the union of \mathcal{D} covers the plane, and the case where it does not.

2.2.1. Painting the union from front to back.

The combinatorial complexity of the boundary of the union $\mathcal{U} := \mathcal{U}(\mathcal{D})$ of \mathcal{D} is at most $6m - 12$, where we ignore the complexity of individual members of \mathcal{D} , and just count the number of intersection points of pairs of boundaries of members of \mathcal{D} that lie on $\partial\mathcal{U}$; see [31]. For convenience, we also include the leftmost and rightmost points of each $D \in \mathcal{D}$ in the set of intersection points (if they lie on the union boundary), thus increasing the complexity of the union by at most $2m$, and assume general position of the pseudo-disks. In general, an intersection point v of a pair of boundaries is at *level* k (of the arrangement $\mathcal{A}(\mathcal{D})$ of \mathcal{D}) if it is contained in the interiors of exactly k members of \mathcal{D} . The boundary intersections are thus at level 0, and a simple application of the Clarkson–Shor technique [25] implies that the number of boundary intersection points that lie at level 1 is also $O(m)$. Hence there exists at least one pseudo-disk $D \in \mathcal{D}$ that contains at most c intersection points at levels 0 or 1 (including leftmost and rightmost points of disks), for some suitable absolute constant c . Clearly, these considerations also apply to any subset of \mathcal{D} .

This allows us to order the members of \mathcal{D} as D_1, \dots, D_m , so that the following property holds. Set $\mathcal{D}_i := \{D_1, \dots, D_i\}$, for $i = 1, \dots, m$. Then D_i contains at most c intersection points at levels 0 and 1 of $\mathcal{A}(\mathcal{D}_i)$. Equivalently, for each i , the boundary of $D_i^0 := D_i \setminus \mathcal{U}(\mathcal{D}_{i-1})$ contains at most c intersection

points.

To prepare for the algorithmic implementation of the construction in this proof, which will be presented later, we note that this ordering is not easy to obtain efficiently in a deterministic manner. Nevertheless, a random insertion order (almost) satisfies the above property: The expected sum of the complexities of the regions D_i^0 , for a random insertion order, is $O(m)$. See later for more details.

We thus have $\mathcal{U}(\mathcal{D}_j) = \bigcup_{i \leq j} D_i^0$ (as an openly disjoint union), for each j ; for the convenience of presentation (and for the algorithm to follow), we interpret this ordering as an incremental process, where the pseudo-disks of \mathcal{D} are inserted, one after the other, in the order D_1, \dots, D_m , and we maintain the partial unions $\mathcal{U}(\mathcal{D}_j)$, after each insertion, by the formula $\mathcal{U}(\mathcal{D}_j) = \mathcal{U}(\mathcal{D}_{j-1}) \cup D_j^0$.

2.2.2. Decomposing the union into vertical trapezoids.

Since the boundary of $D_i^0 = D_i \setminus \mathcal{U}(\mathcal{D}_{i-1})$ contains at most c intersection points, we can decompose D_i^0 into $O(1)$ *vertical pseudo-trapezoids*, using the standard vertical decomposition technique; see, e.g., [47]. Let \mathcal{T}_j be the collection of pseudo-trapezoids in the decomposition of $\mathcal{U}(\mathcal{D}_j)$, collected from the decompositions of the regions D_i^0 , for $i = 1, \dots, j$, and let V_j be the set of vertices of these pseudo-trapezoids, each of which is either an intersection point (boundary intersection or x -extreme point) of $\mathcal{A}(\mathcal{D}_j)$, or an intersection between some ∂D_i and a vertical segment erected from an intersection point of $\mathcal{A}(\mathcal{D}_j)$.

Each of the pseudo-trapezoids in \mathcal{T}_j is bounded by (at most) two vertical segments, a portion of the boundary of a single pseudo-disk as its top edge, and a portion of the boundary of (another) single pseudo-disk as its bottom edge; see [Figure 2.3](#). We have $D_1^0 = D_1$, which we regard as a single pseudo-trapezoid, in which the vertical sides degenerate to the leftmost and rightmost points of ∂D_1 ; see [Figure 2.3\(1\)](#). Note that, in the vertical decomposition of D_i^0 we split it by vertical segments through the intersection points on its boundary, but not through vertices of V_{i-1} on ∂D_i^0 which are not intersection points of $\mathcal{A}(\mathcal{D})$. (Informally, these vertices are “internal” to $\mathcal{U}(\mathcal{D}_{i-1})$, and are not “visible” from the outside.) See, e.g., [Figure 2.3\(4\)](#). The set V_i is obtained by adding to V_{i-1} the vertices of the pseudo-trapezoids in the decomposition of D_i^0 .

If D_i^0 is bounded then each pseudo-trapezoid τ in its decomposition has a top boundary and a bottom boundary, but one or both of the vertical sides may be missing (see, e.g., [Figure 2.3\(1\)](#) for the single pseudo-trapezoid $D_1^0 = D_1$ and [Figure 2.3\(3\)](#) for the left pseudo-trapezoid of 3). From the point of view of τ , each of the top and bottom boundaries of τ may be





Figure 2.2: A union of three disks, and its decomposition into triangles and caps. Note, that the decomposition computed by our algorithm is somewhat different for this case.

either convex (if it is a subarc of ∂D_i on ∂D_i^0), or concave (if it is part of the boundary of some previously inserted pseudo-disk); If D_i^0 is not bounded then some of the vertical pseudo-trapezoids covering D_i^0 will also be unbounded and missing some of their boundaries. Note that D_i^0 is not necessarily connected; in case it is not connected we separately decompose each of its connected components into vertical pseudo-trapezoids in the above manner, see Figure 2.3(4).

At the end of the incremental process, after inserting all the pseudo-disks in \mathcal{D} , the pseudo-trapezoids in $\mathcal{T} := \mathcal{T}_m$ cover $\mathcal{U}(\mathcal{D})$, which may or may not be the entire plane, and they are pairwise openly disjoint. By construction, each pseudo-trapezoid in \mathcal{T} is contained in a single pseudo-disk of \mathcal{D} . Moreover, since the complexity of each D_i^0 is $O(1)$, the total number of pseudo-trapezoids in \mathcal{T} is $O(m)$. So \mathcal{T} possesses some of the properties that we want, but it is not a triangulation.




2.2.3. Polygonalizing the pseudo-trapezoids. To get a triangulation, we associate a polygonal vertical pseudo-trapezoid τ^* with each pseudo-trapezoid $\tau \in \mathcal{T}$. We obtain τ^* from τ by replacing the bottom boundary τ_b and the top boundary τ_t of τ by respective polygonal chains τ_b^* and τ_t^* , that are defined as follows.¹ Let D_i be the pseudo-disk during whose insertion τ was created; in particular, $\tau \subseteq D_i^0$. Let u and v denote the endpoints of τ_b . Consider the region R_{τ_b} between τ_b and the straight


¹The term “polygonal” is somewhat misleading, as some of the boundaries of the pseudo-disks of \mathcal{D} may also be polygonal. To avoid confusion think of the boundaries of the pseudo-disks of \mathcal{D} as smooth convex arcs (as drawn in the figures) even though they might be polygonal.

segment uv ; clearly, by the convexity of D_i , R_{τ_b} is fully contained in D_i . See figure on the right.

If R_{τ_b} contains no vertices of V_i , other than u and v (this will always be the case when $R_{\tau_b} \subseteq \tau$), we replace τ_b by $\tau_b^* = uv$. Otherwise, we replace τ_b by the chain τ_b^* of edges of the convex hull of $V_i \cap R_{\tau_b}$, other than the edge uv . We define τ_t^* analogously, and take τ^* to be the polygonal vertical pseudo-trapezoid that has the same vertical edges as τ , and its top (resp., bottom) part is τ_t^* (resp., τ_b^*). See figure on the right.



Note that, by construction, τ_b^* is a convex polygonal chain. From the point of view of τ , it is convex (resp., concave) if and only if τ_b is convex (resp., concave). An analogous property holds for τ_t^* and τ_t . We denote the crescent-like region bounded by τ_b and τ_b^* by \bar{R}_{τ_b} ; \bar{R}_{τ_t} is defined analogously. (Formally, $\bar{R}_{\tau_b} = R_{\tau_b} \setminus CH(V_i \cap R_{\tau_b})$ and $\bar{R}_{\tau_t} = R_{\tau_t} \setminus CH(V_i \cap R_{\tau_t})$.) Let \mathcal{T}_i^* be the set of polygonal vertical pseudo-trapezoids associated in this manner with the pseudo-trapezoids in \mathcal{T}_i .



Note that R_{τ_b} and R_{τ_t} need not be disjoint. Nevertheless, τ_b^* and τ_t^* cannot cross one another, as follows from Invariant (I2) that we establish below (see Lemma 2.1). This implies that τ^* is well defined. If τ_b^* and τ_t^* are not disjoint then they may only be pinched together at common vertices (in the extreme case they may be identical). This pinching, if it occurs, causes the interior of τ^* to be disconnected (into at most two pieces; it may also be empty, as is the case for D_1^0 , illustrated in Figure 2.3(1)).

2.2.4. Filling the cavities. The insertion of D_i may in general split some arcs of $\partial \mathcal{U}(\mathcal{D}_{i-1})$ into subarcs, whose new endpoints are either points of contact between ∂D_i and $\partial \mathcal{U}(\mathcal{D}_{i-1})$, or endpoints of vertical seg-

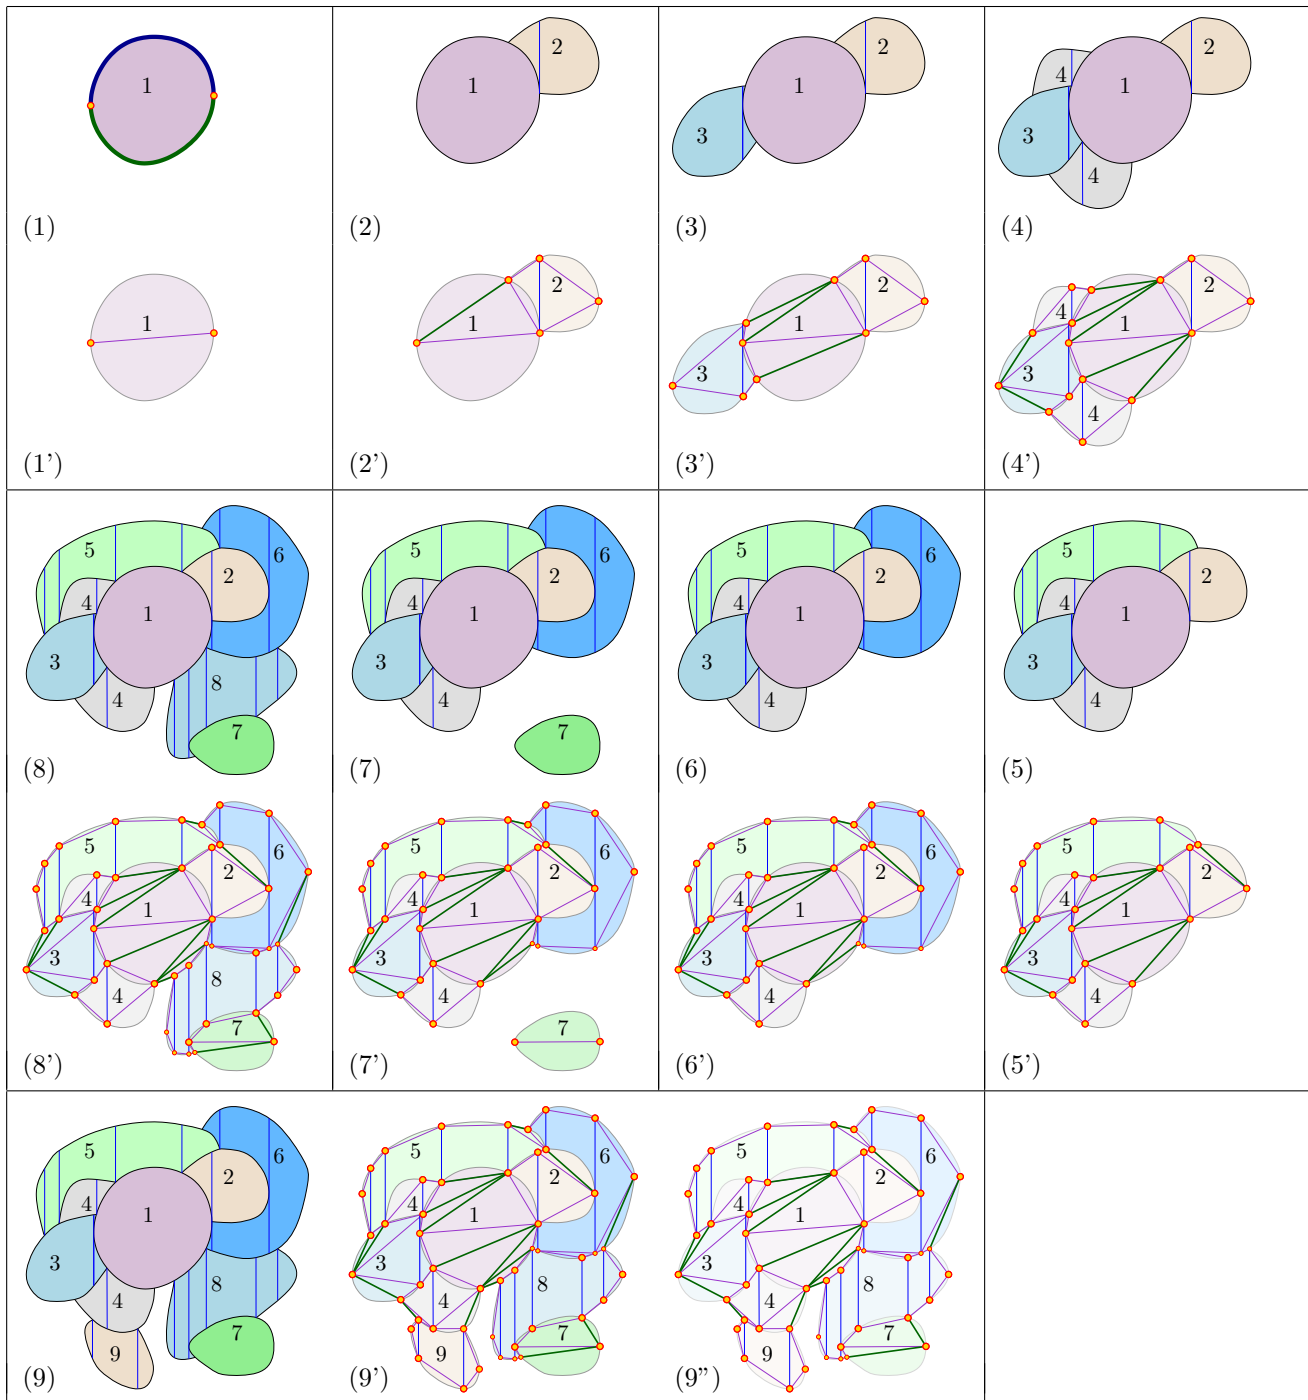


Figure 2.3: A step-by-step illustration of the decomposition \mathcal{T} into pseudo-trapezoids and of the polygonalization of the union. See Section 2.2.4. An animation of this figure is available online at <http://sarielhp.org/blog/?p=8920>; see also Figure 4.1.


ments erected from other vertices of D_i^0 . This can be seen all over [Figure 2.3](#). For example, see the subdivision of the top arc of D_7 caused by the insertion of D_8 in [Figure 2.3\(8'\)](#). Some of these subarcs are boundaries of the new pseudo-trapezoids of D_i^0 and thus do not belong to $\partial\mathcal{U}(\mathcal{D}_i)$, and some remain subarcs of $\partial\mathcal{U}(\mathcal{D}_i)$. We refer to subarcs of the former kind as *hidden*, and to those of the latter kind as *exposed*.

We take each new exposed arc γ , with endpoints u, v , and apply to it the same polygonalization that we applied above to τ_b and τ_t . That is, we take the region R_γ enclosed between γ and the segment uv , and define γ^* to be either uv , if R_γ does not contain any vertex of V_i , or else the boundary of $\text{CH}(R_\gamma \cap V_i)$, except for uv . We note that γ^* is a convex polygonal chain that shares its endpoints with γ , and denote the region enclosed between γ and γ^* as \bar{R}_γ .

Let E_i denote the collection of all straight edges in the polygonal boundaries of the pseudo-trapezoids in \mathcal{T}_i^* and in the polygonal chains γ^* corresponding to new exposed subarcs γ of $\partial\mathcal{U}(\mathcal{D}_{j-1})$, $1 \leq j \leq i$, which were created and polygonalized when adding the corresponding pseudo-disk D_j .

2.2.5. Putting it all together

When the pseudo-disks cover the plane. When the polygonalization process terminates, there are no more regions R_γ , for boundary arcs γ of the union (because there is no boundary), so we are left with a straight-edge planar map M with E_m as its set of edges. (Invariant (1) in [Lemma 2.1](#) below asserts that the edges in E_m do not cross each other.) By Euler's formula, the complexity of M is $O(m)$. We then triangulate each face of M , and, as the analysis in the next subsection will show, obtain the desired triangulation.



The general case. In general, the construction decomposes the union into (pairwise openly disjoint) triangles and crescent regions. To complete the construction, we decompose each crescent region into triangles and caps. A crescent region with $t \geq 2$ vertices on its concave boundary can be decomposed into $t - 2$ triangles and at most $t - 1$ caps. The case $t = 2$ is vacuous, as the crescent is then a cap, so assume that $t \geq 3$. To get such a decomposition, take an extreme edge of the concave polygonal chain, and extend it till it intersects the convex boundary of the crescent, at some point w , thereby chopping off a cap from the crescent. We then create the triangles that w spans with all the concave edges that it sees, and then recurse on the remaining crescent; see

figure on the right. It is easily seen that this results in $t - 2$ triangles and at most $t - 1$ caps, as claimed. After this fix-up, we get a decomposition of the union into triangles and caps.

2.3. Analysis The correctness of the construction is established in the following lemma.

LEMMA 2.1. *The pseudo-trapezoids in \mathcal{T}_i^* and the edges of E_i satisfy the following invariants:*

- (I1) *The segments in E_i do not cross one another.*
- (I2) *Each subarc γ of $\partial\mathcal{U}(\mathcal{D}_i)$ with endpoints u and v has an associated convex polygonal arc $\gamma^* \subseteq E_i$ between u and v . The chains γ^* are pairwise openly disjoint, and their union is the boundary of a polygonal region $\mathcal{U}_i^* \subseteq \mathcal{U}(\mathcal{D}_i)$.*
- (I3) *The pseudo-trapezoids in \mathcal{T}_i^* are pairwise openly disjoint, and each of them is fully contained in some pseudo-disk of \mathcal{D}_i .*
- (I4) $\mathcal{U}(\mathcal{D}_i) \setminus \bigcup_{\tau^* \in \mathcal{T}_i^*} \tau^*$ *consists of a collection of pairwise openly disjoint holes. Each hole is a region between two x -monotone convex chains or between two x -monotone concave chains, with common endpoints, where either both chains are polygonal, or one is polygonal and the other is a portion of the boundary of a single pseudo-disk that lies on $\partial\mathcal{U}(\mathcal{D}_i)$ (i.e., a crescent). The union of the crescents is $\mathcal{U}(\mathcal{D}_i) \setminus \mathcal{U}_i^*$. Each hole, of either kind, is fully contained in some pseudo-disk D_j , $j \leq i$.*

We refer to holes of the former (resp., latter) kind in (I4) of the lemma as *internal polygonal holes* (resp., *external half-polygonal holes*).

Proof: We prove that these invariants hold by induction on i . The invariants clearly hold for \mathcal{T}_1^* after starting the process with $D_1^0 = D_1$. Concretely, \mathcal{T}_1^* consists of the single degenerate pseudo-trapezoid uv , where u and v are the leftmost and rightmost points of \mathcal{D}_1 , respectively, and $E_1 = \{uv\}$. The (external half-polygonal) holes are the portions of D_1 lying above and below uv . It is obvious that (I1)–(I4) hold in this case.

Suppose the invariants hold for \mathcal{T}_{i-1}^* and E_{i-1} . We first prove (I1) for E_i . By construction, the new edges in $E_i \setminus E_{i-1}$ form a collection of convex or concave polygonal chains, where each chain γ^* starts and ends at vertices u, v of either ∂D_i^0 or $\partial\mathcal{U}(\mathcal{D}_{i-1})$. Moreover, by construction, u and v are connected to one another by a single arc γ of the respective boundary ∂D_i^0 or $\partial\mathcal{U}(\mathcal{D}_{i-1})$ (γ is either an exposed or a hidden subarc of $\partial\mathcal{U}(\mathcal{D}_{i-1})$, or a subarc of ∂D_i along ∂D_i^0), and the

region \overline{R}_γ between γ and γ^* does not contain any vertex of V_i in its interior.

Clearly, the edges in a single chain γ^* do not cross one another. Suppose to the contrary that an edge e of some (new) chain γ^* is crossed by an edge e' of some other (new or old) chain. Then either e' has an endpoint inside \overline{R}_γ , contradicting the construction, or e' crosses γ too, to exit from \overline{R}_γ , which again is impossible by construction, since no edge crosses ∂D_i^0 or $\partial \mathcal{U}(\mathcal{D}_{i-1})$. This establishes (I1).

(I2) follows easily from the construction and from the preceding discussion. Note that, for each polygonal chain γ^* , each of its endpoints is also an endpoint of exactly one neighboring arc $\hat{\gamma}^*$, so the union of these arcs consists of closed polygonal cycles, which bounds some polygonal region, which we call \mathcal{U}_i^* , as claimed.

By construction, the vertical boundaries of the new polygonal pseudo-trapezoids of D_i^0 are contained in D_i^0 and do not cross any boundaries of other polygonal pseudo-trapezoids. This, together with (I1), imply that the new pseudo-trapezoids are pairwise openly disjoint, and are also openly disjoint from the polygonal pseudo-trapezoids in \mathcal{T}_{i-1}^* . It is also clear from the construction that each new pseudo-trapezoid $\sigma^* \in \mathcal{T}_i^* \setminus \mathcal{T}_{i-1}^*$ is contained in D_i . So (I3) follows.

Finally consider (I4). Each new hole that is created when adding D_i^0 is of one of the following kinds:

(a) The hole is a region of the form \overline{R}_{τ_b} or \overline{R}_{τ_t} , for some $\tau \in \mathcal{T}_i \setminus \mathcal{T}_{i-1}$, such that \overline{R}_{τ_b} or \overline{R}_{τ_t} is contained in τ (if it lies outside τ , it becomes part of τ^*). Such a hole is contained in D_i , and is bounded by two concave or two convex chains, one of which is polygonal, and the other is part of ∂D_i^0 .

(b) The hole is a region of the form \overline{R}_γ , for an exposed subarc γ of an arc of $\partial \mathcal{U}(\mathcal{D}_{i-1})$, that got delimited by a new vertex (an endpoint of some arc of ∂D_i). These holes are similar to those of type (a).

(c) The hole was part of a hole of type (a) or (b) in $\mathcal{U}(\mathcal{D}_{i-1})$, bounded by an arc γ of $\partial \mathcal{U}(\mathcal{D}_{i-1})$ and its associated polygonal chain γ^* , so that γ has been split into several subarcs (some hidden and some exposed) when adding D_i . For each of these subarcs ζ , we construct an associated polygonal chain ζ^* , either as a top or bottom side of some polygonal pseudo-trapezoid τ^* (constructed from a pseudo-trapezoid τ that has ζ as its top or bottom side), or as the polygonalization of an exposed subarc. The concatenation of the chains ζ^* results in a convex polygonal chain that is contained in \overline{R}_γ and connects the endpoints of γ . The region enclosed between γ^* and ζ^* is an internal polygonal hole. Again, holes of type (c) can be seen all over Figure 2.3; for example, see the top part of D_1 in Figure 2.3(2').

Holes of type (a) and (b) are *boundary half-*

polygonal holes, whereas holes of type (c) are *internal polygonal holes*. Using the induction hypothesis that (I4) holds for $\mathcal{U}(\mathcal{D}_{i-1})$, we get that the union of the new holes of type (a) and (b), together with the old holes of type (a) and (b) corresponding to subarcs of $\partial \mathcal{U}(\mathcal{D}_i) \cap \partial \mathcal{U}(\mathcal{D}_{i-1})$, is $\mathcal{U}(\mathcal{D}_i) \setminus \mathcal{U}_i^*$. This completes the proofs of (I1)–(I4). ■

LEMMA 2.2. (a) Let \mathcal{D} be a collection of $m \geq 3$ planar convex pseudo-disks, whose union covers the plane. Then there exists a set V of $O(m)$ points and a triangulation T of V , such that each triangle $\Delta \in T$ is fully contained in some member of \mathcal{D} .

(b) If $\mathcal{U}(\mathcal{D})$ is not the entire plane, it can be partitioned into $O(m)$ pairwise openly disjoint triangles and caps, such that each triangle or cap is fully contained in some member of \mathcal{D} .

Proof: Since the number of vertices of M is $O(m)$, Euler's formula implies that $|E_m| = O(m)$ too. It is easily seen from the construction and from the invariants of Lemma 2.1, that each face of M is fully contained in some original pseudo-disk, so the same holds for each triangle. This establishes (a). Part (b) follows in a similar manner from the construction. ■

2.4. Efficient construction of the triangulation

With some care, the proof of Lemma 2.2 can be turned into an efficient algorithm for constructing the required triangulation. This is a major advantage of the new proof over the older one. The algorithm is composed of building blocks that are variants of well-known tools, so we only give a somewhat sketchy description thereof

2.4.1. Construction of the original pseudo-trapezoids.

(A similar approach is mentioned in Matoušek *et al.* [40].) The construction proceeds by inserting the pseudo-disks of \mathcal{D} in a *random* order, which, for simplicity, we denote as D_1, \dots, D_m . (Unlike the deterministic construction given above, here we do not guarantee that each D_i^0 has constant complexity. Nevertheless, as argued below, the random nature of the insertion order guarantees that this property holds on average.) As before, we put $\mathcal{D}_i = \{D_1, \dots, D_i\}$ for each i , and we maintain $\mathcal{U}(\mathcal{D}_i)$ after each insertion of a pseudo-disk. To do so efficiently, we maintain a vertical decomposition K_i of the complement \mathcal{U}_i^c of the union $\mathcal{U}(\mathcal{D}_i)$ into vertical pseudo-trapezoids, and maintain, for each $\tau \in K_i$, a *conflict list*, consisting of all the pseudo-disks D_j that have not yet been inserted (i.e., with $j > i$), and that intersect τ .

Since the number of pseudo-trapezoids in the decomposition of the complement of the union of any k pseudo-disks is $O(k)$ (an easy consequence of the linear

bound on the union complexity [31]), a simple application of the Clarkson-Shor technique (similar to those used to analyze many other randomized incremental algorithms) shows that the expected overall number of these “complementary” pseudo-trapezoids that arise during the construction is $O(m)$, and that the expected overall size of their conflict lists is $O(m \log m)$.

When we insert a pseudo-disk D_i , we retrieve all the pseudo-trapezoids of K_{i-1} that intersect D_i . The union $\bigcup_{\tau \in K_{i-1}} (D_i \cap \tau)$ is precisely D_i^0 . For each $\tau \in K_{i-1}$, the intersection $D_i \cap \tau$ decomposes τ into $O(1)$ sub-trapezoids (this follows from the property that each of the four sides of τ crosses ∂D_i at most twice), some of which lie inside D_i (and, as just noted, form D_i^0), and some lie outside D_i , and form part of the new complement of the union \mathcal{U}_i^c .

Typically, the new pseudo-trapezoidal pieces are not necessarily real pseudo-trapezoids, as they may contain one or two “fake” vertical sides, because the feature that created such a side got “chopped off” by the insertion of D_i , and is no longer on the pseudo-trapezoid boundary. In this case, we “glue” these pieces together, across common fake vertical sides, to form the new real pseudo-trapezoids. We do it both for pseudo-trapezoids that are interior to D_i , and for those that are exterior. (This gluing step is a standard theme in randomized incremental constructions; see, e.g., [46].) This will produce (a) the desired vertical decomposition of D_i^0 , and (b) the vertical decomposition K_i of the new union complement \mathcal{U}_i^c . The conflict lists of the new exterior pseudo-trapezoids (interior ones do not require conflict lists) are assembled from the conflict lists of the pseudo-trapezoids that have been destroyed during the insertion of D_i , again, in a fully standard manner.

To recap, this procedure constructs the vertical decompositions of all the regions D_i^0 , so that the overall expected number of these pseudo-trapezoids is $O(m)$, and the total expected cost of the construction is $O(m \log m)$.

2.4.2. Construction of the polygonal chains and the triangulation. By (I2) of Lemma 2.1, before D_i was inserted, each arc γ of $\partial \mathcal{U}(\mathcal{D}_{i-1})$ has an associated convex polygonal arc γ^* with the same endpoints. The union of the arcs γ^* forms a (possibly disconnected) polygonal curve within $\mathcal{U}(\mathcal{D}_{i-1})$, which partitions it into two subsets, the *interior*, \mathcal{U}_i^* , which is disjoint from $\partial \mathcal{U}(\mathcal{D}_{i-1})$ (except at the endpoints of the arcs γ^*), and the *exterior*, which is simply the (pairwise openly disjoint) union of the regions R_γ .

To construct the triangulation, we maintain, for each polygonal chain γ^* of the boundary between the interior and the exterior, a list of its segments, sorted in left-to-right order of their x projections, in a binary

search tree (since the leftmost and rightmost points of each pseudo-disk are vertices in the construction, each chain γ^* is x -monotone). We also maintain a triangulation of the interior. When we add D_i we update the lists representing the arcs γ and extend the triangulation of the interior to cover the “newly annexed” interior, as follows.

When D_i is inserted, some of the arcs γ of $\partial \mathcal{U}(\mathcal{D}_{i-1})$ are split into several subarcs, some of which still appear on $\partial \mathcal{U}(\mathcal{D}_i)$ (call them, as above, *exposed arcs*), while others are now contained in D_i (call them *hidden*). Each endpoint of any new subarc is either an intersection point of ∂D_i with $\partial \mathcal{U}(\mathcal{D}_{i-1})$, or an endpoint of a vertical segment erected from some other vertex of D_i^0 . (This also includes the case where an arc of $\partial \mathcal{U}(\mathcal{D}_{i-1})$ is fully “swallowed” by D_i and becomes hidden in its entirety.) In addition, $\partial \mathcal{U}(\mathcal{D}_i)$ contains *fresh arcs*, which are subarcs of ∂D_i along ∂D_i^0 . The fresh subarcs and the hidden subarcs form the top and bottom sides of the new pseudo-trapezoids in the decomposition of D_i^0 (where each top or bottom side may be either fresh or hidden). To obtain the top or bottom sides of some new pseudo-trapezoids we may have to concatenate several previously exposed subarcs of $\partial \mathcal{U}(\mathcal{D}_{i-1})$. These subarcs are connected at vertices of $\partial \mathcal{U}(\mathcal{D}_{i-1})$ which are not intersection points of the arrangement but intersections of vertical sides of pseudo-trapezoids which we already generated within $\mathcal{U}(\mathcal{D}_{i-1})$.




Figure 2.4

The algorithm needs to construct, for each new exposed, hidden, and fresh arc γ , its associated polygonal curve γ^* . It does so in two stages, first handling exposed and hidden arcs, and then the fresh ones. Let γ be an exposed or hidden subarc, let δ denote the arc of $\partial \mathcal{U}(\mathcal{D}_{i-1})$, or the concatenation of several such arcs, containing γ , and let δ^* be its associated polygonal chain, or, in case of concatenation, the concatenation of the corresponding polygonal chains. As already noted, since the x -extreme points of each pseudo-disk boundary are vertices in the construction, δ and δ^* are both x -monotone.

If $\gamma = \delta$, we do nothing, as $\gamma^* = \delta^*$. Otherwise, let u and v be the respective left and right endpoints

of γ . If uv does not intersect δ^* then γ^* is just the segment uv . Otherwise, γ^* is obtained from a portion of δ^* , delimited on the left by the point u' of contact of the right tangent from u to δ^* , and on the right by the point v' of contact of the left tangent from v to δ^* , to which we append the segments uu' on the left and $v'v$ on the right. See Figure 2.4 for an illustration.

After applying this procedure to every subarc of the old δ , the endpoints of δ (and of δ^*) are now connected by a new convex polygonal chain $\hat{\delta}^*$, which visits each of the new vertices along δ and lies in between δ and δ^* . The region between $\hat{\delta}^*$ and δ^* is a new interior polygonal hole, and we triangulate it, e.g., into vertical trapezoids, by a straightforward left-to-right scan.

Recall that some arcs τ_b and τ_t of new trapezoids τ may be concatenations of several hidden subarcs γ_i (connected at vertices which are not vertices of new trapezoids as explained above). For each such arc, say τ_b , we obtain τ_b^* by concatenating the polygonal chains γ_i^* in x -monotone order.

We next handle the fresh arcs. Each such arc is the top or bottom side of some new pseudo-trapezoid τ , say it is the bottom side τ_b . If τ_t is also fresh, then τ is a convex pseudo-trapezoid, and we replace each of τ_b , τ_t by the straight segment connecting its endpoints. If τ_t is hidden, we take its associated chain τ_t^* , which we have constructed in the preceding stage, and form τ_b^* from it using the same procedure as above: Letting u and v denote the endpoints of τ_b , we check whether uv intersects τ_t^* . If not, τ_b^* is the segment uv . Otherwise, we compute the tangents from u and v to τ_t^* , and form τ_b^* from the tangent segments and the portion of τ_t^* between their contact points. We triangulate each polygonal pseudo-trapezoid τ once we have computed τ_b^* and τ_t^* .

2.4.3. Further implementation details. The actual implementation of the construction of the polygonal chains γ^* proceeds as follows. Given a new arc γ , which is a subarc of an old arc δ , we construct γ^* from δ^* as follows. Let u and v be the endpoints of γ . We (binary) search the list of edges of δ^* for the edge e_u whose x -projection contains the x -projection of u and for the edge e_v whose x -projection contains the x -projection of v . We then walk along the list representing δ^* from e_u towards e_v until we find the point u' of contact of the right tangent from u to δ^* . We perform a similar search from e_v towards e_u to find v' . (If we have traversed the entire portion of δ^* between e_u and e_v without encountering a tangent, we conclude that uv does not intersect δ^* , and set $\gamma^* := uv$.) We extract the sublist between u' and v' from δ^* by splitting δ^* at u' and v' and we insert the segments uu' and $v'v$ at the endpoints of this sublist to obtain

γ^* . We create the polygonalization of fresh arcs from their hidden counterparts in an analogous manner. Note that we destroy the representation of δ^* to produce the representation of γ^* . So in case the arc δ is split into several new subarcs, γ_i , some care has to be taken to maintain a representation of the remaining part of δ^* after producing each γ_j^* from which we can produce the representation of the remaining subarcs γ_i .

For the analysis, we note that to produce γ^* we perform two binary searches to find e_u and e_v , each of which takes $O(\log m)$ time, and then perform linear scans to locate u' and v' . Each edge e traversed by these linear scans (except for $O(1)$ edges) drops off the boundary of the interior so we can charge this step to e and the total number of such charges is linear in the size of the triangulation.

2.5. The result

The computation model. In the preceding description, we implicitly assume a convenient model of computation, in which each primitive geometric operation that is needed by the algorithm, and that involves only a constant number of pseudo-disks (e.g., deciding whether two pseudo-disks or certain subarcs thereof intersect, computing these intersection points, and sorting them along a pseudo-disk boundary) takes constant time. In our application, described in the next section, the pseudo-disks are convex polygons, each having $O(k)$ edges. In this case, each primitive operation can be implemented in $O(\log k)$ time, so the running time should be multiplied by this factor.

The preceding analysis implies the following theorem.

THEOREM 2.1. *We can construct a triangulation of the union of m pseudo-disks covering the plane, with $O(m)$ triangles, such that each triangle is contained in a single pseudo-disk, in $O(m \log m)$ randomized expected time, in a suitable model of computation where every primitive operation takes $O(1)$ time. If the union does not cover the plane, it can be decomposed into $O(m)$ triangles and caps, with similar properties and at the same asymptotic cost.*

COROLLARY 2.1. *Given m convex polygons that are pseudo-disks, that cover that plane, each with at most k edges, one can compute a confined triangulation of the plane, in $O(m \log m \log k)$ expected time.*

2.6. Extension to general convex shapes **Theorem 2.1** uses only peripherally the property that the input shapes are pseudo-disks, and a simple modification (of the analysis, not of the construction itself) allows

us to extend it to general convex shapes. Specifically, let \mathcal{D} be a collection of m simply-shaped convex regions in the plane, such that the union complexity of any i of them is at most $u(i)$, where the complexity is measured, as before, by the number of boundary intersection points on the union boundary, and where $u(\cdot)$ is a monotone increasing function satisfying $u(i) = \Omega(i)$. We assume that the regions in \mathcal{D} are simple enough so that the boundaries of any pair of them intersect only a constant number of times, and so that each primitive operation on them can be performed in reasonable time (which we take to be $O(1)$ in the statement below). The interesting cases are those in which $u(i)$ is small (that is, near-linear). They include, e.g., the case of fat triangles, or a low-density collection of convex regions; see [13] and references therein.

Deploying the algorithm of [Theorem 2.1](#) results in the desired confined triangulation of $\mathcal{U}(\mathcal{D})$. Extending the analysis to this general setup, we obtain the following theorem.

THEOREM 2.2. *Let \mathcal{D} be a collection of n convex shapes in the plane, such that the union complexity of any i of them is at most $u(i)$, where $u(i)$ is a monotone increasing function with $u(i) = \Omega(i)$. Then one can compute, in $O(u(m)\log m)$ expected time, a confined triangulation of $\mathcal{U}(\mathcal{D})$ with $O(u(m))$ triangles and caps (or just triangles if the union covers the entire plane), under the assumption that every primitive geometric operation takes $O(1)$ time.*

3. Construction of shallow cuttings and approximate levels

We begin by presenting a high-level description of the technique, filling in the technical details in subsequent subsections. The high-level part does not pay too much attention to the efficiency of the construction; this is taken care of later in this section.

3.1. Sketch of the construction Assume that, for a given parameter r , we want to approximate level $k = n/r$ of $\mathcal{A}(H)$. Note that when r is too close to n , that is, when k is a constant, we can simply compute the k -level explicitly and use it as its own approximation. The complexity of such a level is $O(n)$, and it can be computed in near-linear time [18, 12]. We therefore assume in the remainder of this section that $r \ll n$.

Put $k_1 := (1 + c)n/r$ and $k_2 := (1 + 2c)n/r$, for a suitable sufficiently small (but otherwise arbitrary) constant fraction c . The analysis of Clarkson and Shor [25] implies that the overall complexity of $L_{\leq k_2}$ (the first k_2 levels of $\mathcal{A}(H)$) is $O(nk^2)$. This in turn implies that there exists an index $k_1 \leq \xi \leq k_2$ for

which the complexity $|L_\xi|$ of L_ξ is $O(nk^2/(cn/r)) = O(nk/c) = O(n^2/(cr))$. We fix such a level ξ , and continue the construction with respect to L_ξ (slightly deviating from the originally prescribed value of k). However, to simplify the notation for the current part of the analysis, we use k to denote the nearby level ξ , and will only later return to the original value of k .

The next step is to decompose the xy -projection of the k -level L_k into a small number of connected polygons, from which the approximate level will be constructed. We first review the existing machinery, already mentioned in the introduction, for this step.

Decomposing a level into a small number of polygons. Let H , k , and L_k be as above. It is convenient to assume that the faces of L_k are triangles; this can be achieved by triangulating each face, without affecting the asymptotic complexity of L_k . In particular, the k -level (or, rather, its xy -projection) can then be interpreted as a planar, triangulated and biconnected graph (a graph is *biconnected*, if any pair of vertices are connected by at least two vertex-disjoint paths).

As has been discovered over the years, planar graphs can be efficiently decomposed into smaller pieces that are well behaved. This goes back to the planar separator theorem of Lipton and Tarjan [35], Miller’s cycle separator theorem [41], and Frederickson divisions [27], and has eventually culminated in the fast κ -division algorithm of Klein *et al.* [32]. Specifically, for a (specific drawing of a) planar triangulated and biconnected graph G with N vertices, and for a parameter $\kappa < n$, a κ -division of G is a decomposition of G into several connected subgraphs G_1, \dots, G_m , such that (i) $m = O(N/\kappa)$; (ii) each G_i has at most κ vertices; (iii) each G_i has at most $O(\sqrt{\kappa})$ boundary vertices, namely, vertices that belong to at least one additional subgraph; and (iv) each G_i has at most $O(1)$ holes, namely, faces of the induced drawing of G_i that are not faces of G (as they contain additional edges and vertices of G). Such a division can be computed in $O(N)$ time [32].²

The construction, continued. We set

$$t := (cn/r - 6)/12;$$

since $r \ll n$ we have $t > 1$. We apply the planar subdivision algorithm of [32, 27], as just reviewed, and construct a t^2 -division of the xy -projection L'_k of L_k , that is, of the planar graph consisting of the vertices and edges of L'_k . This subdivision produces

$$m := O(|L_k|/t^2) = O\left(\frac{n^2/(cr)}{c^2n^2/r^2}\right) = O(r/c^3)$$

²The algorithm of [32] in fact constructs k -divisions for a geometrically increasing sequence of values of the parameter k , in overall $O(N)$ time.

connected polygons, P_1, \dots, P_m , with pairwise disjoint interiors, whose union covers the entire xy -plane, and whose edges are projections of (some) edges of L_k .

By construction, each P_i is connected and has at most $O(t)$ edges (and also contains $O(t^2)$ edges and vertices of the k -level in its interior). Let C_i denote the convex hull of P_i , for $i = 1, \dots, m$. As we show in [Corollary 3.1](#) in [Section 3.3](#) below, $\mathcal{C} := \{C_1, \dots, C_m\}$ is a collection of m (possibly unbounded) convex *pseudo-disks* whose union is the entire plane.

We then apply [Lemma 2.2](#) to \mathcal{C} and obtain a set S of $O(m)$ points in the xy -plane, and a triangulation T of S , such that each triangle $\Delta \in T$ is fully contained in some hull C_i in \mathcal{C} . For a point p in the xy -plane, we denote by $\uparrow_k(p)$ the *lifting* of p to the k -level, i.e., the point on the level that is co-vertical with p . Let T' denote the corresponding collection of triangles in \mathbb{R}^3 , given by $T' = \{\uparrow_k(\Delta) \mid \Delta \in T\}$. For a bounded triangle Δ , $\uparrow_k(\Delta)$ is the triangle spanned by the lifted images of the three vertices of Δ . We lift an unbounded triangle Δ with vertices p and q by lifting pq to $\uparrow_k(p)\uparrow_k(q)$ as before, and lifting each of its rays, say $[p, \infty)$, to a ray $\uparrow([p, \infty))$ emanating from $\uparrow_k(p)$ in a direction parallel to the plane which is vertically above $[p, \infty)$ at infinity. If the liftings $\uparrow([p, \infty))$, and $\uparrow([q, \infty))$, and the edge $\uparrow_k(p)\uparrow_k(q)$ are not on the same plane, we add another ray, say r , emanating from p parallel to $[q, \infty)$. We add to T' the unbounded triangle spanned by $\uparrow([q, \infty))$, $\uparrow_k(p)\uparrow_k(q)$, and r , and the unbounded wedge spanner by r and $\uparrow([p, \infty))$.

Note that the triangles of T' are in general not contained in L_k . However, for each triangle $\Delta' \in T'$, its vertices lie on L_k , and, as we show in [Lemma 3.4](#) below, at most $12t + 6$ planes of H can cross Δ' . This implies, returning now to the original value of k , that Δ' fully lies between the levels

$$\xi \pm (12t + 6) = \xi \pm cn/r$$

of $\mathcal{A}(H)$. In particular, Δ' lies fully above the level

$$\xi - cn/r \geq k_1 - cn/r = n/r = k,$$

and fully below the level

$$\xi + cn/r \leq k_2 + cn/r = (1 + 3c)n/r = (1 + 3c)k.$$

The lifted triangulation T' forms a polyhedral terrain that consists of $O(r/c^3)$ triangles and is contained between the levels $k = n/r$ and $(1 + 3c)k$. That is, for a given $\varepsilon > 0$, choosing $c = \varepsilon/3$ makes T' an ε -approximation of L_k , and we obtain the following result.

THEOREM 3.1. *Let H be a set of n non-vertical planes in \mathbb{R}^3 in general position, and let $r \leq n$, $\varepsilon > 0$ be*

given parameters. Then there exists a polyhedral terrain consisting of $O(r/\varepsilon^3)$ triangles, that is fully contained between the levels n/r and $(1 + \varepsilon)n/r$ of $\mathcal{A}(H)$.

To turn this approximate level into a shallow cutting, replace each $\Delta' \in T'$ (including each of the unbounded triangles just constructed) by the semi-unbounded vertical prism Δ^* consisting of all the points that lie vertically below Δ' . This yields a collection Ξ of prisms, with pairwise disjoint interiors, whose union covers $L_{\leq n/r}$, so that, for each prism τ of Ξ , we have (a) each vertex of τ lies at level (at least k and) at most $(1 + \frac{2}{3}\varepsilon)k$, and (b) the top triangle of τ is crossed by at most $\frac{1}{3}\varepsilon k$ planes of H . Hence, as is easily seen, each prism of Ξ is crossed by at most $(1 + \varepsilon)n/r$ planes, so Ξ is the desired shallow cutting. That is, we have the following result.

THEOREM 3.2. *Let H be a set of n non-vertical planes in \mathbb{R}^3 in general position, let $k < n$ and $\varepsilon > 0$ be given parameters, and put $r = n/k$. Then there exists a k -shallow $((1 + \varepsilon)/r)$ -cutting of $\mathcal{A}(H)$, consisting of $O(r/\varepsilon^3)$ vertical prisms (unbounded from below). The top of each prism is a triangle that is fully contained between the levels k and $(1 + \varepsilon)k$ of $\mathcal{A}(H)$, and these triangles form a polyhedral terrain (which thus approximates the k -level L_k up to a relative error of ε).*

3.2. Efficient implementation We next turn our constructive proof into an efficient algorithm, and show:

THEOREM 3.3. *One can construct the k -shallow $((1 + \varepsilon)/r)$ -cutting given in [Theorem 3.2](#), or, equivalently, the ε -approximating terrain of the k -level in [Theorem 3.1](#), in $O(n + r^2\varepsilon^{-6} \log^3 r)$ randomized expected time. Computing the conflict lists of the vertical prisms takes an additional $O(n \log r)$ time.*

Proof: Recall that $k = n/r$. Let (H, \mathcal{R}) denote the range space where each range in \mathcal{R} corresponds to some vertical segment or ray e , and is equal to the subset of the planes of H that are crossed by e . Clearly, (H, \mathcal{R}) has finite VC dimension (see, e.g., [47]). We draw a random sample S of $n' = \frac{br}{\varepsilon^2} \log r$ planes from H , where b is a suitable constant. For b sufficiently large, such a sample is a *relative* $(\frac{1}{r}, \varepsilon)$ -approximation for (H, \mathcal{R}) , with probability $\geq 1 - 1/r^{O(1)}$; see [29] for the definition and properties of relative approximations. This means that each vertical segment that intersects $x \geq n/r$ planes of H intersects between $(1 + \varepsilon)\frac{n'}{n}x$ and $(1 - \varepsilon)\frac{n'}{n}x$ planes of S , and each vertical segment that intersects $x < n/r$ planes of H intersects at most $\frac{n'}{n}x + \varepsilon\frac{n'}{n}r$ planes of S . (This holds, with probability $\geq 1 - 1/r^{O(1)}$, for all vertical segments.)

The strategy is to use (the smaller) S instead of H in the construction, as summarized in [Theorem 3.2](#), and then argue that a suitable approximate level in $\mathcal{A}(S)$ is also an approximation to level k in $\mathcal{A}(H)$. Set

$$k' = \frac{b(1+\varepsilon)}{\varepsilon^2} \log r, \quad \text{and} \quad t' = \frac{b(1+\varepsilon)}{\varepsilon} \log r = \varepsilon k'.$$

We apply [Theorem 3.2](#) so that it approximates level k' in $\mathcal{A}(S)$. For this, we choose a random index ξ in the range $[k', k'+t']$, construct the ξ -level of $\mathcal{A}(S)$, and then apply the construction to that level.

Before doing that, we first show that the ξ -level is a good approximation to level k in $\mathcal{A}(H)$. Consider a point p on level k of $\mathcal{A}(H)$. By the property specified above of a relative relative $(\frac{1}{r}, \varepsilon)$ -approximation, it follows that the level of p in $\mathcal{A}(S)$ is at most $(1+\varepsilon)(n'/n)(n/r) = k'$. Similarly, let p be a point at level larger than say $(1+4\varepsilon)k$ of $\mathcal{A}(H)$. Then the level of p in $\mathcal{A}(S)$ is at least $(1-\varepsilon)(n'/n)(1+4\varepsilon)(n/r) \geq k'+t'$ for $\varepsilon < 1/5$. Since this hold with high probability, for every point p , we conclude that the entire ξ -level is between level k and $(1+4\varepsilon)k$ of $\mathcal{A}(H)$.

We can now apply the machinery in [Theorem 3.2](#). The first step in this analysis is to construct the ξ -level in $\mathcal{A}(S)$. Rather than just constructing that level, we compute all the first $k'+t'$ levels, using an algorithm of Chan [16], which takes

$$\begin{aligned} O(n' \log n' + n'(k')^2) &= O\left(\frac{r \log r}{\varepsilon^2} \left(\log \frac{r}{\varepsilon} + \frac{\log^2 r}{\varepsilon^4}\right)\right) \\ &= O(\varepsilon^{-6} r \log^3 r) \end{aligned}$$

expected time. We can then easily extract the desired (random) level ξ . In expectation (over the random choice of ξ), the complexity of the ξ -level is

$$n_1 = O(n'(k')^2/(\varepsilon k')) = O(n'k'/\varepsilon) = O\left(\frac{r}{\varepsilon^5} \log^2 r\right),$$

and we assume in what follows that this is indeed the case.

We now implement the construction in a straightforward way. We already have the random ξ -level. We project it onto the xy -plane, and construct a $(t')^2$ -division of the projection, in $O(n_1)$ time. It consists of

$$m := O(n_1/(t')^2) = O\left(\frac{n'}{\varepsilon^3 k'}\right) = O(r/\varepsilon^3)$$

pieces, each with $O(t') = O(\frac{1}{\varepsilon} \log r)$ edges. We compute their convex hulls and construct the corresponding confined triangulation, in overall time

$$O(m \log m \log t') = O\left(\frac{r}{\varepsilon^3} \log \frac{r}{\varepsilon} \log O\left(\frac{1}{\varepsilon} \log r\right)\right).$$

Finally, we need to lift the vertices of the resulting triangles to the ξ -level of $\mathcal{A}(S)$. This can be done, in brute force, in $O(r|S|) = O\left(\frac{r^2}{\varepsilon^2} \log r\right)$ time. We obtain a terrain T' , with the claimed number of triangles, which is an ε -approximation of the k' -level of $\mathcal{A}(S)$, and which lies above that level. That is the level in $\mathcal{A}(S)$ of each point on T' is between k' and $(1+\varepsilon)(k'+t') = (1+\varepsilon)^2 k' < (1+3\varepsilon)k'$ (for $\varepsilon < 1/2$). We now repeat the preceding analysis, with 3ε replacing ε , and conclude that T' lies fully between level k and level $(1+12\varepsilon)k$ of $\mathcal{A}(H)$. A suitable scaling of ε gives us the desired approximation in $\mathcal{A}(H)$.

Combining this with the preceding analysis, we conclude that (with high probability) T' is also a (2ε) -approximation for the k -level in $\mathcal{A}(H)$. This at last completes the construction. Its overall expected cost is $O(n + \varepsilon^{-6} r^2 \log^3 r)$.

The final stage is to compute for every vertical prism of T' the planes of H that intersect it (i.e., the prism's conflict list). To this end, we put the vertices of T' into the range reporting data-structure of Chan [16] – specifically, after preprocessing, in $O(r \log r)$ expected time, given a query half-space h^+ , one can report the points $h^+ \cap T'$ in $O(\log r + |h \cap T'|)$ expected time.³ We query this data structure with every plane $h \in H$, and for each vertex x of T' which is reported we add h to the conflict lists of the prisms incident to x . This take $O(n \log r)$ since the total size of the conflict lists is linear. ■

We now proceed to fill in the details of the various steps of the construction, and of the corresponding algorithm.

3.3. The convex hulls of pairwise openly disjoint polygons are pseudo-disks

LEMMA 3.1. *Let P and P' be two connected polygons in the plane with disjoint interiors, and let C and C' denote their respective convex hulls. Then ∂C and $\partial C'$ cross each other at most twice.*

Proof: For simplicity of exposition, we assume that P and P' are in general position, in a sense that will become more concrete from the proof. It is easily argued that this can be made without loss of generality.

Assume, for the sake of contradiction, that ∂C and $\partial C'$ cross more than twice (in general position, the boundaries do not overlap). This implies that each of $\partial C \setminus C'$, $\partial C' \setminus C$ is disconnected, and thus there exist four vertices u, w, v , and z of the boundary of $C^* = \text{CH}(C \cup C')$, that appear along ∂C^* in this

³The paper of Chan [16] does not use shallow cuttings.




Figure 3.1: Illustration of the proof of Lemma 3.1.

circular order, so that $u, v \in \partial C \setminus C'$ and $w, z \in \partial C' \setminus C$. See Figure 3.1. Clearly, u and v are also vertices of P , and w and z are vertices of P' .

We show that this scenario leads to an impossible planar drawing of K_5 . For this, let o be an arbitrary point outside C^* . Connect o to each of u, v, w, z by noncrossing arcs that lie outside C^* , and connect u, w, v , and z by the four respective portions of ∂C^* between them. Finally, connect u to v by a path contained in P , and connect w to z by a path contained in P' . The resulting ten edges are pairwise noncrossing, where, for the last pair of edges, the property follows from the disjointness of (the interiors of) P and P' . The contradiction resulting from this impossible planar drawing of K_5 establishes the claim. ■

(Note that the above proof does not require the polygons to be simply connected.)

COROLLARY 3.1. *Let $\mathcal{P} = \{P_1, \dots, P_m\}$ be a set of m pairwise openly disjoint connected polygons in the plane, and let C_i denote the convex hull of P_i , for $i = 1, \dots, m$. Then $\mathcal{C} := \{C_1, \dots, C_m\}$ is a collection of m convex pseudo-disks.*

3.4. Crossing properties of the planar subdivision e set $t := (cn/r - 6)/12$; since $r \ll n$ we have $t > 1$. We apply the planar subdivision algorithm of [32, 27], as just reviewed, and construct a t^2 -division of the xy -projection L'_k of L_k , that is, of the planar graph consisting of the vertices and edges of


Recall that our construction computes a t^2 -division of the xy -projection L'_k of L_k where $t := (cn/r - 6)/12$ (recall that $k = n/r$, $r \ll n$). Our goal in the rest of this section is to show that the lifting $\uparrow(\Delta)$ of any triangle Δ contained in the convex hull C of a subgraph (“piece”) P of this decomposition intersects at most ck planes of H . We prove this for bounded triangles, the proof for unbounded triangles is similar.

Recall that for a point p in the xy -plane, we denote by $\uparrow_k(p)$ the (unique) point that lies on L_k and is co-vertical with p . The *crossing distance* $\mathbf{cr}(p, q)$ between any pair of points $p, q \in \mathbb{R}^3$, with respect to H , is the number of planes of H that intersect the closed segment pq . The crossing distance is a quasi-metric, in that it is symmetric and satisfies the triangle inequality. For a connected set $X \subseteq \mathbb{R}^3$, the *crossing number* $\mathbf{cr}(X)$ of X is the number of planes of H intersecting X (thus $\mathbf{cr}(p, q)$ is the crossing number of the segment pq).

LEMMA 3.2. *Let p, q, r be three collinear points in the xy -plane, such that $q \in pr$, and let $p' = \uparrow_k(p)$, $q' = \uparrow_k(q)$, and $r' = \uparrow_k(r)$; these points, that lie on the k -level, are in general not collinear. Let q'' be the intersection of the vertical line through q with the segment $p'r'$. Then we have $\mathbf{cr}(q'', q') \leq \mathbf{cr}(p', r')$.*


Proof: We have

$$\begin{aligned} \mathbf{cr}(q'', q') &= |\text{level}(q') - \text{level}(q'')| = |k - \text{level}(q'')| \\ &= |\text{level}(p') - \text{level}(q'')| \leq \mathbf{cr}(p', q'') \\ &\leq \mathbf{cr}(p', r'). \end{aligned}$$



In what follows, we consider polygonal regions contained in L_k , where each such region R is a connected union of some of the faces of L_k . The xy -projection of R is a connected polygon in the xy -plane, and, for simplicity, we refer to R itself also as a polygon.

LEMMA 3.3. *Let H be a set of n non-vertical planes in \mathbb{R}^3 in general position. Let P' be a bounded connected polygon with t edges that lies on the k -level L_k of $\mathcal{A}(H)$, such that all the boundary edges of P' are edges of L_k . Let p' be a vertex of the external boundary of P' , and let q be any point in the convex hull C of the xy -projection P of P' . Then the crossing distance between p' and $q' = \uparrow_k(q)$ is at most $4t + 2$.*



Proof: Since q lies in C , we can find two points u, v on the external boundary of P such that $q \in uv$. Put $q' = \uparrow_k(q)$, $u' = \uparrow_k(u)$, and $v' = \uparrow_k(v)$, and denote by q'' the point that lies on the segment $u'v'$ and is co-vertical with q . We have

$$\begin{aligned} \mathbf{cr}(p', q') &\leq \mathbf{cr}(p', u') + \mathbf{cr}(u', q'') + \mathbf{cr}(q'', q') \\ &\leq \mathbf{cr}(p', u') + \mathbf{cr}(u', v') + \mathbf{cr}(q'', q'). \end{aligned}$$

Let π_1 and π_2 be the two portions of the external boundary that connect p' and u' , and u' and v' , respectively. We may assume, without loss of generality, that these portions do not overlap. Now, by [Lemma 3.2](#), we have $\mathbf{cr}(q'', q') \leq \mathbf{cr}(u', v')$, so we get

$$\begin{aligned} \mathbf{cr}(p', q') &\leq \mathbf{cr}(p', u') + 2\mathbf{cr}(u', v') \\ &\leq \mathbf{cr}(\pi_1) + 2\mathbf{cr}(\pi_2) \leq 2\mathbf{cr}(\partial P'), \end{aligned}$$

where $\partial P'$ denotes the external boundary of P' .

To bound the number of planes of H that intersects $\partial P'$, consider its vertices p_1, p_2, \dots, p_t . Observe that p_1 is contained in three planes. For each i , p_i lies on at most two planes that do not contain p_{i-1} (there are two such planes when $p_{i-1}p_i$ is a diagonal of an original face of the untriangulated level L_k). Furthermore, the open segment $p_{i-1}p_i$ does not intersect any plane that does not intersect one of its endpoints. Therefore, the number $\mathbf{cr}(\partial P')$ of planes of H that intersect $\partial P'$ satisfies $\mathbf{cr}(\partial P') \leq 3 + 2(t-1) = 2t + 1$. (Note that this analysis is somewhat conservative—for example, if the polygon P' uses only original edges of the k -level, the number drops to roughly t .) ■

LEMMA 3.4. *Let H be a set of n non-vertical planes in \mathbb{R}^3 in general position, and let P' be a connected polygon with t edges, such that P' lies on the k -level L_k of $\mathcal{A}(H)$, and such that all the boundary edges of P' are edges of L_k . Then, for any triangle $\Delta = \Delta pqr$ that is fully contained in the convex hull of the xy -projection of P' , the number $\mathbf{cr}(\Delta')$ of planes of H that cross the triangle $\Delta' = \Delta' p'q'r'$, where $p' = \uparrow_k(p)$, $q' = \uparrow_k(q)$, $r' = \uparrow_k(r)$, is at most $12t + 6$.*

Proof: Let p, q, r be the vertices of Δ , and put $p' = \uparrow_k(p)$, $q' = \uparrow_k(q)$, $r' = \uparrow_k(r)$; that is, $\Delta' = \Delta p'q'r'$. Let w be any vertex of the external boundary of P' . Any plane that crosses Δ' must also cross two of its sides. Moreover, by [Lemma 3.3](#) and the triangle inequality,

$$\mathbf{cr}(p', q') \leq \mathbf{cr}(w, p') + \mathbf{cr}(w, q') \leq 2(4t + 2),$$

and similarly for $\mathbf{cr}(p', r')$ and $\mathbf{cr}(q', r')$. Adding up these bounds and dividing by 2, implies the claim. ■

4. Applications

4.1. Constructing layered cuttings of the whole arrangement

4.1.1. Preliminaries. To construct such a cutting, we need the following technical tools.

LEMMA 4.1. *Let H be a set of n non-vertical planes in \mathbb{R}^3 in general position. The number of pairs of edges (e, e') of $\mathcal{A}(H)$ such that the xy -projections of e and of e' cross each other, and the unique vertical segment connecting e and e' crosses no other plane of H , is $O(n^3)$.*

Proof: The number of such pairs of edges is at most $\sum_{c \in \mathcal{A}(H)} |c|^2$, where the sum ranges over all three-dimensional cells c of $\mathcal{A}(H)$, and where $|c|$ denotes the overall complexity of c . This latter sum is known to be $O(n^3)$ —it is an easy consequence of the Zone Theorem in three dimensions; see Aronov *et al.* [14]. ■

LEMMA 4.2. *Let H be a set of n non-vertical planes in \mathbb{R}^3 in general position, and let q be a parameter. The number of pairs of edges (e, e') of $\mathcal{A}(H)$ such that the xy -projections of e and of e' cross each other, and the unique vertical segment connecting e and e' crosses at most q planes of H , is $O(n^3q)$.*

Proof: This follows by a standard application of the Clarkson-Shor technique [25] to the bound stated in [Lemma 4.1](#). ■

4.1.2. Constructing a layered cutting of $\mathcal{A}(H)$ Let H be a set of n non-vertical planes in \mathbb{R}^3 in general position, and let $r < n$ be a parameter. Our goal is to construct a $(1/r)$ -cutting of the entire $\mathcal{A}(H)$, of optimal size $O(r^3)$. To do so, consider some fixed sequence of $2r$ levels

$$k_1^- < k_2^- < k_1^+ < k_3^- < k_2^+ < \dots < k_r^- < k_{r-1}^+ < k_r^+,$$

where each pair of consecutive indices in this sequence are at distance at least $n/(4r)$. That is, we form a sequence of overlapping intervals $[k_1^-, k_1^+], \dots, [k_r^-, k_r^+]$, so that each interval starts after the preceding one starts and before it ends, and no three intervals have a common index. We choose such a sequence in the following random manner. Fix the intervals

$$I_1^- = [1, n/(4r)]$$

$$I_i^- = [(i-3/2)n/r + 1, (i-5/4)n/r], \quad \text{for } i = 2, \dots, r$$

$$I_i^+ = [in/r + 1, (i+1/4)n/r], \quad \text{for } i = 1, \dots, r-1$$

$$I_r^+ = [n - n/(2r) + 1, n - n/(4r)].$$

Then choose k_i^- (resp., k_i^+) uniformly at random from I_i^- (resp., I_i^+), for $i = 1, \dots, r$.

The strategy goes as follows. For each index $i = 1, \dots, r$, consider the pair of levels $L_{k_i^-}, L_{k_i^+}$, which

we denote shortly and respectively as L_i^- , L_i^+ , and approximate both of them simultaneously, using the following refinement of the algorithm summarized in [Theorem 3.1](#), with the same parameter $t = cn/r$ for all pairs, where $c \ll 1/4$ is a sufficiently small constant. Project L_i^- and L_i^+ onto the xy -plane, and overlap the resulting planar maps M_i^- , M_i^+ into a single map M_i^* . Each vertex of M_i^* is either the projection of a vertex of one of the levels L_i^- , L_i^+ , or a crossing point of a pair of projected edges, one from each level.

We now apply the preceding analysis to M_i^* , and get a triangulation T_i of the xy -plane, whose combinatorial complexity is $m_i = O(|M_i^*|/t^2)$. We lift its vertices up to both levels L_i^- , L_i^+ , resulting in a corresponding pair of triangulated terrains T_i^- , T_i^+ , with identical xy -projections. We claim that T_i^- approximates L_i^- and T_i^+ approximates L_i^+ . Indeed, each boundary edge of a piece of the t^2 -division of M_i^* is a portion of a projected edge of either L_i^- or L_i^+ . Hence, traversing any portion of the boundary of a piece, we encounter at most $O(t)$ edges of each of the levels L_i^- , L_i^+ , and the arguments used above imply that, when lifted to either of the two levels, the crossing number of the corresponding path is at most $O(t)$, from which the claim follows. Each triangle Δ of T_i is lifted to a pair of triangles $\Delta^- \in T_i^-$, $\Delta^+ \in T_i^+$, and we connect them by a vertical triangular prism Δ^* that has them as its bases. These prisms are pairwise openly disjoint, and their union is the layer Λ_i between T_i^- and T_i^+ . Let Ξ_i denote the collection of these prisms. For $i = 1$ (resp., $i = r$), we extend each prism Δ^* in Ξ_1 (resp., in Ξ_r) to a semi-unbounded prism that contains all the points vertically below (resp., above) the original Δ^* .

Repeating this process for each $i = 1, \dots, r$ results in a collection $\Xi = \bigcup_i \Xi_i$ of vertical prisms whose union is the entire 3-space. These prisms are not pairwise openly disjoint, but each point in \mathbb{R}^3 is contained in the interiors of at most two prisms. Informally, the layers $\Lambda_1, \dots, \Lambda_r$ overlap in pairs (but no three layers overlap), so that each layer is fully triangulated by pairwise openly disjoint vertical prisms.

LEMMA 4.3. *The expected size of Ξ is $O(r^3)$.*

Proof: The overall number of prisms in Ξ is, by [Theorem 3.1](#),

$$(4.1) \quad |\Xi| = O\left(\frac{1}{t^2} \sum_{i=1}^r |M_i^*|\right).$$

We have $|M_i^*| = O(|L_{k_i^-}| + |L_{k_i^+}| + |X_i|)$, where X_i is the set of pairs (e, e') of edges, where e is an edge of $L_{k_i^-}$, e' is an edge of $L_{k_i^+}$, and the xy -projections of e and e' cross each other.

Estimating $\sum_i (|L_{k_i^-}| + |L_{k_i^+}|)$ is easy. Each level of $\mathcal{A}(H)$ appears in this sum with probability at most $4r/n$ (note that some levels will never be chosen), so the expected value is at most proportional to $4r/n$ times the complexity of $\mathcal{A}(H)$, namely, $O((r/n) \cdot n^3) = O(n^2r)$.⁴

To estimate the expected value of $\sum_i |X_i|$, we note that each pair (e, e') that is counted in this sum belongs to the set X_0 of pairs that are accounted for in the bound in [Lemma 4.2](#), with $q = 7n/(4r)$, but our pairs constitute only a small subset of X_0 . Specifically, by our choice of random levels, the probability of a pair $(e, e') \in X_0$ to appear in one of the sets X_i is at most proportional to

$$(4r/n)^2 \cdot |X_0| = O((r/n)^2 \cdot n^3 \cdot (n/r)) = O(n^2r).$$

Substituting the separate bounds obtained so far in [Eq. \(4.1\)](#), we get that the expected size of Ξ satisfies

$$|\Xi| = O\left(\frac{1}{t^2} \sum_{i=1}^r |M_i^*|\right) = O\left(\frac{r^2}{n^2} \cdot n^2r\right) = O(r^3). \quad \blacksquare$$

We therefore obtain the following result.

THEOREM 4.1. *For a set H of n planes in \mathbb{R}^3 , and a parameter $r < n$, one can construct a layered $(1/r)$ -cutting of $\mathcal{A}(H)$ of size $O(r^3)$. Specifically, we cover space by a set Ξ of $O(r^3)$ vertical triangular prisms, such that each point is covered at most twice, and each prism is crossed by $\leq n/r$ planes of H .*

4.2. Approximate halfspace range counting In its dual setting, the problem is: Let H be a set of n nonvertical planes in \mathbb{R}^3 in general position, and let $\varepsilon > 0$ be an error parameter. We wish to preprocess H into a data structure that supports queries of the form: For a query point q , count the number of planes lying below q , up to a multiplicative factor of $1 \pm \varepsilon$. That is, if q lies at level k , the answer should be between $(1 - \varepsilon)k$ and $(1 + \varepsilon)k$.

Let $m = O(1/\varepsilon^{4/3})$. We construct and store the first m levels of $\mathcal{A}(H)$ explicitly, each level as its own terrain. Formally, we set $k_i = i$, for $i = 0, \dots, m$. Next, for deeper levels, we use the approximate level construction. Take the sequence of levels $k_{m+i} := m(1 + \varepsilon)^i$, for $i = 1, \dots, m'$, where $m' = \lceil \log_{1+\varepsilon} \frac{n}{m} \rceil \approx \frac{1}{\varepsilon} \log n$. For each $i = m+1, \dots, m+m'$, approximate the level L_{k_i} up to an additive error of εk_i , let T_i denote the underlying triangulation in the xy -plane of the projection of the approximation, and let T'_i denote the approximating

⁴Each vertex of $\mathcal{A}(H)$ appears in three consecutive levels, and each edge appears in two, so features of $\mathcal{A}(H)$ may be drawn more than once, but at most three times.

terrain, namely, the appropriate lifting of T_i . By construction, it is easily checked that the terrains T'_i do not intersect one another, and are therefore stacked on top of one another. To answer an approximate (dual) halfspace range counting query with a point q , we simply need to find two consecutive terrains T'_i, T'_{i+1} so that q lies between them, and return $m(1 + \varepsilon)^{i-m}$, say, as the approximate count, when $i > m$, or i itself otherwise.

By [Theorem 3.1](#), for $i = 1, \dots, m'$, the complexity of T_{m+i} (and of T'_{m+i}) is

$$|T_{m+i}| = O\left(\frac{n}{\varepsilon^3 k_{m+i}}\right) = O\left(\frac{n}{\varepsilon^3 m(1 + \varepsilon)^i}\right).$$

Summing these bounds over i , we get

$$\sum_{i=1}^{m'} |T_{m+i}| = O\left(\frac{n}{\varepsilon^3 m}\right) \sum_{i=1}^{m'} \frac{1}{(1 + \varepsilon)^i} = O\left(\frac{n}{\varepsilon^4 m}\right).$$

Storing the first m levels of $\mathcal{A}(G)$ requires $O(nm^2)$ space (this bounds their overall complexity), so both bounds are $O(n/\varepsilon^{8/3})$, for $m = O(1/\varepsilon^{4/3})$. This bounds the storage used by our data structure. The preprocessing time is

To answer a query with some point q , we run a binary search over the terrains T'_i , and locate the xy -projection of q in the relevant planar maps T_i . The total cost is therefore

$$O\left(\log\left(\frac{1}{\varepsilon} \log n\right) \cdot \log \frac{n}{\varepsilon}\right).$$

Afshani and Chan [1] showed how to avoid the binary search for finding the right level, using a data structure of Kaplan *et al.* [30]. Afshani and Chan use this latter structure to find a rough approximation to the level. Specifically, they find an estimate $\hat{\ell}$ to the level which is off by a factor of b with probability $1/b$. Then, instead of doing a binary search, they linearly search for the right level, starting from the level in the hierarchy closest to $\hat{\ell}$. The expected number of searches that they perform is then $O(1)$ and these searches take $O(\log(n/(\varepsilon k)))$ time. We can apply the exact same technique using our approximate levels instead of the more complicated refined shallow cuttings used in [1], and then get the following.

THEOREM 4.2. *There exists a data structure of size $O(n/\varepsilon^{8/3})$, that can be constructed in near-linear expected time, and that can answer approximate level queries in an arrangement of n planes in \mathbb{R}^3 , up to a relative error of ε , in $O(\log(n/(\varepsilon k)))$ expected time, where k is the exact level of the query.*

Acknowledgments. We thank János Pach for pointing out that a variant of [Lemma 2.2](#) is already known.

References

- [1] P. Afshani and T. M. Chan. On approximate range counting and depth. *Discrete Comput. Geom.*, 42(1):3–21, 2009.
- [2] P. Afshani and T. M. Chan. Optimal halfspace range reporting in three dimensions. In *Proc. 20th ACM-SIAM Sympos. Discrete Algs. (SODA)*, pages 180–186, Philadelphia, PA, USA, 2009. SIAM.
- [3] P. Afshani, T. M. Chan, and K. Tsakalidis. Deterministic rectangle enclosure and offline dominance reporting on the RAM. In *Proc. 41st Internat. Colloq. Automata Lang. Prog. (ICALP)*, volume 8572 of *Lect. Notes in Comp. Sci.*, pages 77–88. Springer, 2014.
- [4] P. Afshani and K. Tsakalidis. Optimal deterministic shallow cuttings for 3d dominance ranges. In *Proc. 25th ACM-SIAM Sympos. Discrete Algs. (SODA)*, pages 1389–1398, 2014.
- [5] P. K. Agarwal. Partitioning arrangements of lines I: an efficient deterministic algorithm. *Discrete Comput. Geom.*, 5:449–483, 1990.
- [6] P. K. Agarwal. Partitioning arrangements of lines: II. Applications. *Discrete Comput. Geom.*, 5:533–573, 1990.
- [7] P. K. Agarwal. Geometric partitioning and its applications. In J. E. Goodman, R. Pollack, and W. Steiger, editors, *Computational Geometry: Papers from the DIMACS Special Year*, pages 1–37. Amer. Math. Soc., 1991.
- [8] P. K. Agarwal. *Intersection and Decomposition Algorithms for Planar Arrangements*. Cambridge University Press, 1991.
- [9] P. K. Agarwal, B. Aronov, T. M. Chan, and M. Sharir. On levels in arrangements of lines, segments, planes, and triangles. *Discrete Comput. Geom.*, 19:315–331, 1998.
- [10] P. K. Agarwal and P. K. Desikan. An efficient algorithm for terrain simplification. In *Proc. 8th ACM-SIAM Sympos. Discrete Algs. (SODA)*, pages 139–147, 1997.
- [11] P. K. Agarwal and J. Erickson. Geometric range searching and its relatives. In B. Chazelle, J. E. Goodman, and R. Pollack, editors, *Advances in Discrete and Computational Geometry*, pages 1–56. Amer. Math. Soc., 1999.
- [12] P. K. Agarwal and J. Matoušek. Dynamic half-space range reporting and its applications. *Algorithmica*, 13:325–345, 1995.
- [13] B. Aronov, M. de Berg, E. Ezra, and M. Sharir. Improved bounds for the union of locally fat objects in the plane. *SIAM Journal on Computing*, 43(2):543–572, 2014.
- [14] B. Aronov, J. Matoušek, and M. Sharir. On the sum of squares of cell complexities in hyperplane arrangements. *J. Combin. Theory Ser. A*, 65:311–321, 1994.
- [15] R. P. Bambah and C. A. Rogers. Covering the

- plane with convex sets. *J. London Math. Soc.*, s1-27(3):304–314, 1952.
- [16] **T. M. Chan**. Random sampling, halfspace range reporting, and construction of ($\leq k$)-levels in three dimensions. *SIAM Journal on Computing*, 30(2):561–575, 2000.
- [17] **T. M. Chan**. Low-dimensional linear programming with violations. *SIAM Journal on Computing*, pages 879–893, 2005.
- [18] **T. M. Chan**. A dynamic data structure for 3-d convex hulls and 2-d nearest neighbor queries. *J. Assoc. Comput. Mach.*, 57(3), 2010.
- [19] **T. M. Chan** and **K. Tsakalidis**. Optimal deterministic algorithms for 2-d and 3-d shallow cuttings. In *Proc. 31st Int. Annu. Sympos. Comput. Geom. (SoCG)*, pages 719–732, 2015.
- [20] **B. Chazelle**. Cutting hyperplanes for divide-and-conquer. *Discrete Comput. Geom.*, 9(2):145–158, 1993.
- [21] **B. Chazelle**, **H. Edelsbrunner**, **L. J. Guibas**, and **M. Sharir**. A singly-exponential stratification scheme for real semi-algebraic varieties and its applications. *Theo. Comp. Sci.*, 84:77–105, 1991.
- [22] **B. Chazelle** and **J. Friedman**. A deterministic view of random sampling and its use in geometry. *Combinatorica*, 10(3):229–249, 1990.
- [23] **K. L. Clarkson**. New applications of random sampling in computational geometry. *Discrete Comput. Geom.*, 2:195–222, 1987.
- [24] **K. L. Clarkson**, **H. Edelsbrunner**, **L. J. Guibas**, **M. Sharir**, and **E. Welzl**. Combinatorial complexity bounds for arrangements of curves and spheres. *Discrete Comput. Geom.*, 5:99–160, 1990.
- [25] **K. L. Clarkson** and **P. W. Shor**. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.
- [26] **E. Ezra**, **D. Halperin**, and **M. Sharir**. Speeding up the incremental construction of the union of geometric objects in practice. *Comput. Geom. Theory Appl.*, 27(1):63–85, 2004.
- [27] **G. N. Frederickson**. Fast algorithms for shortest paths in planar graphs, with applications. *SIAM Journal on Computing*, 16(6):1004–1022, 1987.
- [28] **S. Har-Peled**. Constructing planar cuttings in theory and practice. *SIAM Journal on Computing*, 29(6):2016–2039, 2000.
- [29] **S. Har-Peled** and **M. Sharir**. Relative (p, ε) -approximations in geometry. *Discrete Comput. Geom.*, 45(3):462–496, 2011.
- [30] **H. Kaplan**, **E. Ramos**, and **M. Sharir**. Range minima queries with respect to a random permutation, and approximate range counting. *Discrete Comput. Geom.*, 45(1):3–33, 2011.
- [31] **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.*, 1:59–71, 1986.
- [32] **P. N. Klein**, **S. Mozes**, and **C. Sommer**. Structured recursive separator decompositions for planar graphs in linear time. In *Proc. 45th Annu. ACM Sympos. Theory Comput. (STOC)*, pages 505–514, 2013.
- [33] **V. Koltun**. Almost tight upper bounds for vertical decompositions in four dimensions. *J. Assoc. Comput. Mach.*, 51(5):699–730, 2004.
- [34] **V. Koltun** and **M. Sharir**. Curve-sensitive cuttings. *SIAM Journal on Computing*, 34(4):863–878, 2005.
- [35] **R. J. Lipton** and **R. E. Tarjan**. A separator theorem for planar graphs. *SIAM J. Appl. Math.*, 36:177–189, 1979.
- [36] **J. Matoušek**. Construction of ε -nets. *Discrete Comput. Geom.*, 5:427–448, 1990.
- [37] **J. Matoušek**. Efficient partition trees. *Discrete Comput. Geom.*, 8:315–334, 1992.
- [38] **J. Matoušek**. Reporting points in halfspaces. *Comput. Geom. Theory Appl.*, 2(3):169–186, 1992.
- [39] **J. Matoušek**. On constants for cuttings in the plane. *Discrete Comput. Geom.*, 20:427–448, 1998.
- [40] **J. Matoušek**, **N. Miller**, **J. Pach**, **M. Sharir**, **S. Sifrony**, and **E. Welzl**. Fat triangles determine linearly many holes. In *Proc. 32nd Annu. IEEE Sympos. Found. Comput. Sci. (FOCS)*, pages 49–58, 1991.
- [41] **G. L. Miller**. Finding small simple cycle separators for 2-connected planar graphs. *Journal of Computer and System Sciences*, 32(3):265–279, 1986.
- [42] **K. Mulmuley**. An efficient algorithm for hidden surface removal, II. *J. Comp. Sys. Sci.*, 49:427–453, 1994.
- [43] **N. H. Mustafa**, **R. Raman**, and **S. Ray**. Settling the APX-hardness status for geometric set cover. In *Proc. 55th Annu. IEEE Sympos. Found. Comput. Sci. (FOCS)*, pages 541–550, 2014.
- [44] **J. Pach** and **P. K. Agarwal**. *Combinatorial Geometry*. John Wiley & Sons, 1995.
- [45] **E. A. Ramos**. On range reporting, ray shooting and k -level construction. In *Proc. 15th Annu. Sympos. Comput. Geom. (SoCG)*, pages 390–399, 1999.
- [46] **R. Seidel**. A simple and fast incremental randomized algorithm for computing trapezoidal decompositions and for triangulating polygons. *Comput. Geom. Theory Appl.*, 1:51–64, 1991.
- [47] **M. Sharir** and **P. K. Agarwal**. *Davenport-Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, New York, 1995.

Figure 4.1: Animation of algorithm – you would need Acrobat reader to see the animation - click the figure to make it start.