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

Sariel Har-Peled† Haim Kaplan‡ Micha Sharir§

April 23, 2021

Abstract

Let H be a set of n non-vertical planes in three dimensions, and let $r < n$ be a parameter. We give a construction that approximates the (n/r)-level of the arrangement $\mathcal{A}(H)$ of H by a terrain consisting of $O(r/\varepsilon^3)$ triangular faces, which lies entirely between the levels n/r and $(1+\varepsilon)n/r$. 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. This leads to conceptually cleaner constructions of shallow cuttings in three dimensions.

On the way, we get two other results that are of independent interest: (a) We revisit an old result of Bambah and Rogers [BR52] about triangulating a union of convex pseudo-disks, and provide an alternative proof that yields an efficient algorithmic implementation. (b) We provide a new construction of cuttings in two dimensions.

1 Introduction

A tribute to Jirka Matoušek. We were very fortunate to have Jirka as a friend and colleague. He has entered our community in the late 1980's, and has been a giant lighthouse

---

*A preliminary version of this paper appeared in *Proc. 27th Annu. ACM-SIAM Sympos. Discrete Algs.* (SODA), 2016, 1193–1212 [HKS16a]. 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), by the Blavatnik Research Fund at Tel Aviv University, 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`

ever since, showing us the way into new discoveries, solving mysteries for us, and providing us with new tools, ideas, and techniques, that have made our work much more interesting and productive. He has been everywhere, making seminal contributions to so many topics in computational and discrete geometry (and to other fields too). We have been avid readers of his many books, most notably *Lectures on Discrete Geometry*, and have been admiring his clear yet precise style of exposition and presentation. We have also learned to appreciate his personality, his dry but touching sense of humor, his love for nature, his infinite devotion to science on one hand, and to his family and friends on the other hand. His departure has been painful to us, and we will miss him badly. We thank you, Jirka, for all the gifts you gave us, and may your soul be blessed.

This paper is about a topic that Jirka has worked on, rather extensively, during the early 1990s, concerning *cuttings* and related techniques for decompositions of arrangements or of point sets, and their applications to range searching and other algorithmic and combinatorial problems in geometry. In particular, in 1992 he has written a seminal paper on "Reporting points in halfspaces" [Mat92c], where he introduced and analyzed *shallow cuttings*, a technique that had many applications during the following decades.

In a later paper, following his earlier work [Mat90] (probably his first entry into computational geometry), Jirka [Mat98] presented a construction of $(1/r)$-cuttings, for a set of lines in the plane, with $\leq 8r^2+6r+4$ cells. This construction uses, as a basic building block, a strikingly simple procedure for approximating a level in a line arrangement: Since a specific level is an x-monotone polygonal chain, one can pick every qth vertex, for $q \approx n/r$, and connect these vertices consecutively to form an approximate level, which is at line-crossing distance at most $q/2$ from the original level. As is well known, this construction is asymptotically optimal for any arrangement of lines in general position. This elegant level approximation algorithm, in two dimensions, raises the natural question of whether one can approximate a level in three dimensions for a given set of planes, by an xy-monotone polyhedral terrain constructed directly, in an analogous manner, from the original level.

This paper provides an affirmative answer to this question, thereby pushing Jirka's work further, for the special case of three-dimensional arrangements of planes. It refines the shallow cuttings technique of [Mat92c], and applies the new technique to obtain cleaner and more efficient solutions for several related problems. Our new scheme for approximating a level by a terrain, while significantly more involved than Jirka's two-dimensional construction, still echoes and generalizes his basic idea of "shortcutting" the original level by a coarser triangular mesh (instead of a simplified polygonal chain in the plane) spanned by selected vertices of the level.

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) such that the closure of their union covers $\mathbb{R}^d$, and each region is crossed (intersected in its interior) by 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 [Aga91a] for an early survey. Applications include a variety of range searching techniques [AE99], partition trees [Mat92a], incidence problems involving points and lines, curves, and surfaces [CEG+90], and many more.

The first (albeit suboptimal) construction of cuttings is due to Clarkson [Cla87]. This concept was formalized later on by Chazelle and Friedman [CF90], who gave a sampling-based construction of optimal-size cuttings (see below). An optimal deterministic construction algorithm was provided by Chazelle [Cha93]. Matoušek [Mat98] studied the number of cells in a $(1/r)$-cutting in the plane (see also [Har00]). See Agarwal and Erickson [AE99] and Chazelle [Cha04] for comprehensive reviews 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 [CF90] 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, i.e., nearly cubic and quartic in r, respectively, are known [CEGS91, Kol04, KS05].)

For additional works related to cuttings and their applications, see [Aga90a, Aga90b, Aga91b, AACS98, AC09b, ACT14, AHZ10, AT14, CT15, Har00, Mat92a, Mat92b, Ram99].

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). For a given parameter $0 \le k \le n-1$, 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 $(\le k)$-level, denoted as $L_{\le k}$, is the union of all the j-levels, for $j = 0, \ldots, k$. A collection of pairwise openly disjoint simplices such that the closure of their union covers $L_{\le k}$, and such that each simplex is crossed by at most n/r hyperplanes of H, is 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 $k = \Theta(n/r)$. Furthermore, shallow cuttings for any value of k can be reduced to this case—see Chan and Tsakalidis [CT15, Section 5].

In his seminal paper on reporting points in halfspaces [Mat92c], 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 $O\left(q^{\lceil d/2\rceil} r^{\lfloor d/2\rfloor}\right)$, 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 $O\left(r^{\lfloor d/2\rfloor}\right)$, 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 lead 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 [Ram99] 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, where for each r the corresponding cutting is a $(1/r)$-cutting of the first $\Theta(n/r)$ levels of $\mathcal{A}(H)$. Ramos's algorithm runs in $O(n \log n)$ total expected time. Recently, Chan and Tsakalidis [CT15] provided a deterministic $O(n \log r)$-time algorithm for computing an $O(n/r)$-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, as above, in $O(n \log n)$ deterministic time. Interestingly, they use Matoušek's theorem on the existence of an $O(n/r)$-shallow $(1/r)$-cutting of size $O(r)$ in the analysis of their algorithm.

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, given the cutting, one can produce the conflict lists in $O(n \log r)$ time using a result of Chan [Cha00], as we outline in Section 4.2.

Matoušek's proof of the existence of small-size shallow cuttings, as well as subsequent studies of this technique, are technically involved. They rely on random sampling, combined with a clever variant of the so-called exponential decay lemma of [CF90], and with several additional (and rather intricate) techniques.

Approximating a level. An early study of Matoušek [Mat90] 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 a set H of n 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 (suitably small) 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 [AD97, AS98], do not apply immediately in this case, as they produce a suboptimal output, of size larger than the optimal by a logarithmic factor. More importantly, they are not geared 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). Nevertheless, we note that the algorithm in [AS98] can be modified to provide a logarithmic approximation in our sense, but the resulting running time (at least $\Omega(n^8)$, and probably much worse) is quite large. Perhaps more significantly, without the results in the present work, it is not even clear that such an optimal-size approximation exists at all.

Such an approximation to the k-level, whose size is almost optimal up to a polylogarithmic factor, can be obtained by using a *relative-approximation* sample of H, and by extracting the appropriate level in the sample [HS11]. 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 [Cha05], by replacing a standard shallow cutting by the upper convex hull of its simplices. However, the resulting cuttings do not lead to a sharp approximation of the level, of the sort we seek.

In short, an effective and optimal technique for approximating a level in three dimensions as a terrain (let alone in higher dimensions) does not follow easily from existing 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 vertical semi-unbounded triangular prism Δ^* having Δ as its top face, and consisting of all points that lie on or vertically below Δ. Such a cutting (by prisms) has already been constructed by Chan [Cha05], but it does not yield (that is, come from) a $(1+\varepsilon)$-approximation to the level. Such a shallow cutting, by vertical semi-unbounded triangular prisms, was a central tool in Chan's algorithm for dynamic convex hulls in three dimensions [Cha10].

This discussion suggests that 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 in three dimensions.

1.1 Our results

In this paper we give an alternative and constructive proof of the existence of optimal-size shallow cuttings in a three-dimensional plane arrangement, by vertical semi-unbounded triangular prisms. 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 with $O(r/\varepsilon^3)$ triangular faces, 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 $r \le n/k$, with a somewhat more involved bound on the complexity of the approximation.

The construction does not use sampling, nor does it use the exponential decay lemma of [CF90, Mat92c]. It is based on the planar separator theorem of Lipton and Tarjan [LT79], or, more precisely, on recent separator-based decomposition techniques of planar maps, as in Klein *et al.* [KMS13] (see also Frederickson [Fre87]), 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.

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 map. Since the overall complexity of the first k levels is $O(nk^2)$ (see, e.g., [CS89]), we may assume, by moving from a specified level to a nearby one, that the complexity of our level is near the average value $O(nk)$. The decomposition techniques of planar graphs mentioned above (as in [Fre87]) allow us to partition the level into $O(n/k)$ clusters, where each cluster has $O(k^2)$ vertices and $O(k)$ boundary vertices (vertices that also belong to other clusters). In the terminology of [Fre87], 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 projected edges of the original level). 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 hull is (trivially) simply connected, and the boundaries of any pair of hulls cross at most twice. Moreover, the decomposition has the property that, for each triangle Δ that is 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 [BR52], proving a statement due to L. Fejes-Tóth, and reviewed in [PA95, Lemma 3.9] (and also briefly below), 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 each pseudo-disk can be shrunk into a convex polygon so that these polygons are pairwise openly disjoint, with the same union, and the total number of edges of the polygons is at most $6m$; the desired triangulation is obtained by simply triangulating, arbitrarily, 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 shallow cutting of the first k levels is obtained by simply replacing each triangle Δ of the approximate level by the semi-unbounded vertical prism of points lying below Δ.

Planar cuttings. Interestingly, a simplified version of the algorithm for approximating the k-level in 3-space can also be applied to arrangement of lines in the plane, yielding a new construction of cuttings in the plane, which is different from previous approaches. We present this warm-up exercise in Section 3, and believe it to be of independent interest.

Confined triangulations. One of the main contributions of this work is providing an alternative proof of the aforementioned result of Bambah and Rogers [BR52]. The original proof in [BR52], and its simplified presentation in [PA95], do not seem to lead to a sufficiently efficient construction. In contrast, the new proof leads to an algorithm with near linear running time that constructs a triangulation with the desired properties; see Section 2.

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.* [EHS04], which was in turn inspired by Mulmuley's work on hidden surface removal [Mul94]. Mustafa *et al.* [MRR14] 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 [BR52]) is simpler, optimal, and independent of the complexity of the relevant pseudo-disks. We are pleased that this nice property of convex 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 (and are absent when the union covers the entire plane, as in [BR52]). We believe that this triangulation could be useful in practice, in situations where, given a query point q, one wants to decide whether q is inside the union, and if so, provide a witness shape that contains q. For this, we simply locate the triangle in our triangulation that contains q, from which the desired witness shape is immediately available. This is significant in situations where deciding whether a point belongs to an input shape is considerably more expensive than deciding whether it lies inside a triangle.

Additional applications. Two additional applications of our construction, that are described in the arxiv version of this paper [HKS16b], are the following:

(a) We extend Matoušek's construction [Mat90] 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.

(b) 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. This competes with Afshani and Chan's technique [AC09a]. 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 given explicitly in [AC09a]).

(c) We present yet another construction of cuttings in two dimensions that uses packing arguments together with the new techniques of this paper to construct cuttings in the plane.

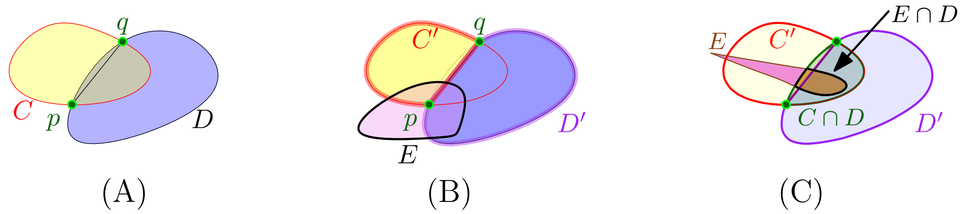

Figure 2.1: The proof of Bambah and Rogers.

Paper organization. We start by presenting the construction of the confined triangulation in Section 2. As a warm-up exercise, we use this result in Section 3 to present a new algorithm for constructing cuttings of arrangements of lines in the plane. We then describe the construction of approximate levels, and the construction of shallow cuttings that it leads to, in Section 4.

2 Triangulating the union of convex pseudo-disks and other 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 2.4 below for details. 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 4, but it is not new, as it is an immediate consequence of an old result of Bambah and Rogers [BR52] (proving a statement by L. Fejes-Tóth), whose proof is sketched below. Our analysis provides an alternative constructive proof.

We use this result (i.e., Theorem 2.4) as a black box later on in the paper, and the impatient reader might want to skip this (somewhat tedious) section for later reading, and go directly to Section 3.

Bambah and Rogers' proof. For the sake of completeness, we briefly sketch the proof of Bambah and Rogers (as presented in Pach and Agarwal [PA95, Lemma 3.9]). Let $\mathcal{K}$ be a collection of m convex pseudo-disks in the plane, and assume, for simplicity, that their union is a triangle T (extending this simpler scenario to the more general case is straightforward). We may also assume that no pseudo-disk of $\mathcal{K}$ is contained in the union of the other regions of $\mathcal{K}$, as one can simply throw away any such redundant pseudo-disk. Finally, since the construction will create regions with overlapping boundaries, we use the more general definition of pseudo-disks, requiring, for each pair $C, D \in \mathcal{K}$, that $C \setminus D$ and $D \setminus C$ be both connected. See Figure 2.1 (A).

Let C and D be two pseudo-disks of $\mathcal{K}$, such that the intersection $\mathrm{int}(C) \cap \mathrm{int}(D)$ of their interiors is nonempty and minimal in terms of containment (that is, it does not contain any other such intersection). Let p and q be the two intersection points of ∂C and ∂D (since

$C \cap D$ has a nonempty interior, ∂C and ∂D cannot overlap, so p and q are well defined). Cut C and D along the segment pq, and let $C' \subseteq C$ and $D' \subseteq D$ be the two resulting pieces whose union is $C \cup D$, see Figure 2.1 (B). Let $\mathcal{K}' = (\mathcal{K} \setminus \{C, D\}) \cup \{C', D'\}$. We claim that $\mathcal{K}'$ is a collection of m pseudo-disks covering T.

Indeed, consider a pseudo-disk $E \in \mathcal{K}'$ other than C', D'. We need to show that $E \setminus C'$ and $C' \setminus E$ are both connected, and similarly for E and D'. The pseudo-disk property is immediate if $E \cap pq$ is empty. If E contains p (resp., q), then it is easy to verify, by convexity, that E and C' are pseudo-disks, and similarly for E and D'. Assume then E does not contain p or q, but still intersects the segment pq. By assumption, $E \setminus (C \cup D)$ is not empty, so we may assume, without loss of generality, that E intersects the boundary of $C \setminus D$. But then $E \cap D \subseteq C \cap D$, as otherwise E would intersect the boundary of C in four points, which is impossible. This in turn contradicts the minimality of $C \cap D$, see Figure 2.1 (C).

We thus replace $\mathcal{K}$ by $\mathcal{K}'$, and repeat this process till all the pseudo-disks in the resulting collection are pairwise interior disjoint. At this point, $\mathcal{K}$ is a pairwise openly disjoint cover of the triangle T, by m convex polygons (each contained inside its original pseudo-disk). By Euler's formula, these polygons have a total of $O(m)$ edges, and can thus be triangulated into $O(m)$ triangles with the desired property.

This elegant proof is significantly simpler than what follows, but it does *not* seem to lead to an efficient algorithm for constructing the desired triangulation in near-linear running time. We present here a different alternative (efficiently) constructive proof, 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 extensions of this result to the case where the union of the pseudo-disks does not cover the plane, and for more general convex shapes, not necessarily pseudo-disks.

2.1 Preliminaries

The notion of a triangulation that we use here is slightly non-standard, as it might be 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 (where the segment might degenerate into a single point, in which case the triangle becomes a wedge).

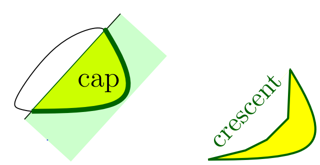

Figure 2.2

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 from it a convex polygon that has the base chord of the cap as an edge, but is otherwise contained in the interior of the cap. See Figure 2.2.

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

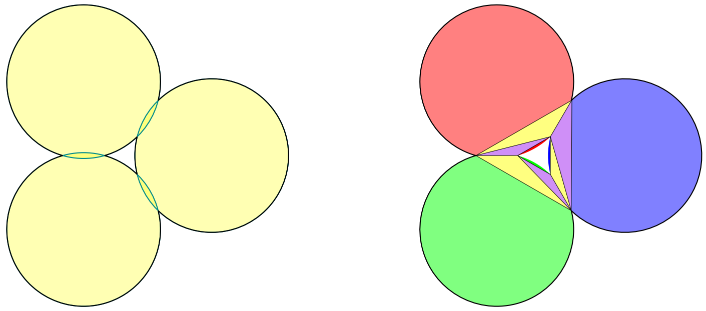

Figure 2.3: A union of three disks, and its decomposition into triangles and caps. Note that the decomposition computed by our algorithm is somewhat different for this case.

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 2.3.

2.2 Construction

We are given a collection $\mathcal{D}$ of m convex pseudo-disks, and our goal is to construct a confined triangulation for $\mathcal{D}$, as described above, with $O(m)$ pieces. 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

A basic property of a collection $\mathcal{D}$ of m pseudo-disks is that the combinatorial complexity of the boundary of the union $\mathcal{U} := \mathcal{U}(\mathcal{D}) = \bigcup_{C \in \mathcal{D}} C$ 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 [KLPS86]. For convenience, we also (i) 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 (ii) assume general position of the pseudo-disks.

An intersection point v of a pair of boundaries is at *depth* 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 intersection points are thus at depth 0, and a simple application of the Clarkson–Shor technique [CS89] implies that the number of boundary intersection points that lie at depth 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 depths 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 $D_1, \ldots, D_m$, so that the following property holds. Set $\mathcal{D}_i := \{D_1, \ldots, D_i\}$, for $i = 1, \ldots, m$. Then D_i contains at most c intersection points at depths 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

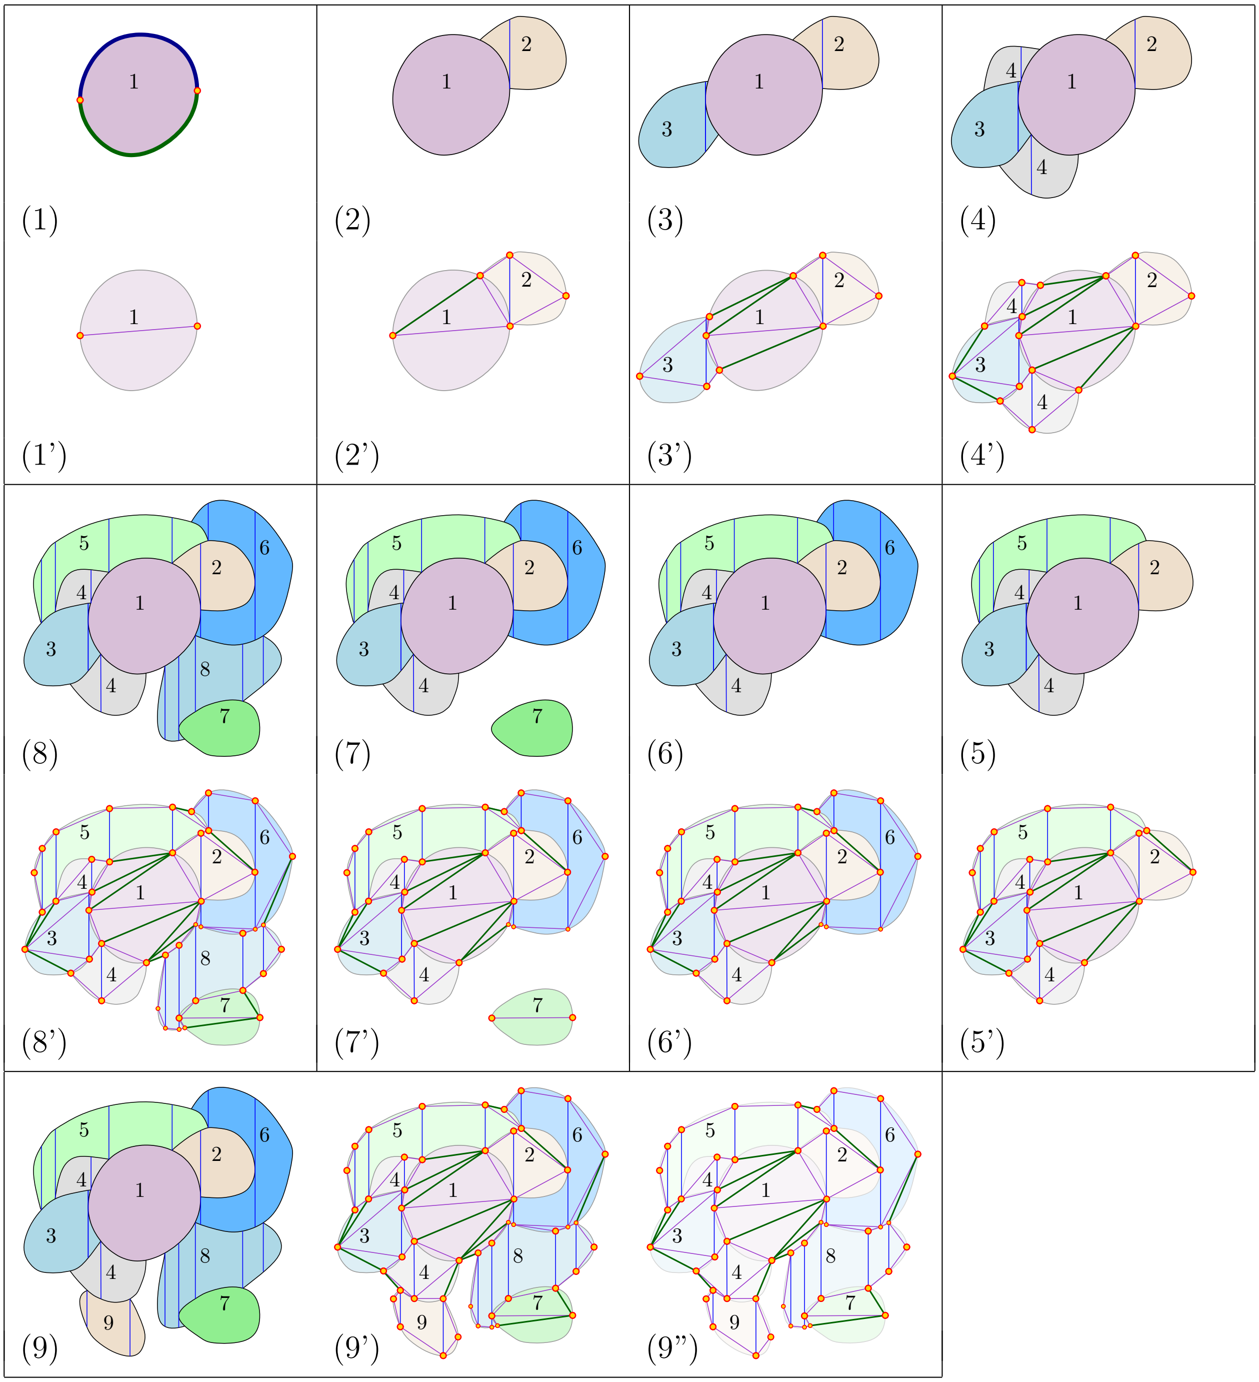

Figure 2.4: A step-by-step illustration of the decomposition $\mathcal{T}$ into pseudo-trapezoids and of the polygonalization of the union. See Section 2.2.4. An animation of the steps depicted in this figure is available online at `http://sarielhp.org/blog/?p=8920`, see also the animated figure in the arxiv version of the paper [HKS16b].

deterministic manner. Nevertheless, a random insertion order (almost) satisfies the above property: As we will show, the expected sum of the complexities of the regions D_i^0, for a random insertion order, is $O(m)$, which is the property that our analysis really needs. See later for more details.

We thus have $\mathcal{U}(\mathcal{D}_j) = \bigcup_{i \le j} D_i^0$ (as a pairwise 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 $D_1, \ldots, D_m$, 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., [SA95]. 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 $i = 1, \ldots, j$, and let V_j be the set of vertices of these pseudo-trapezoids, each of which is either an intersection point (more precisely, a boundary intersection or an x-extreme point) of $\mathcal{A}(\mathcal{D}_j)$, or an intersection between some $\partial\mathcal{D}_i$ 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 the top parts of the subfigures in Figure 2.4. 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 2.4(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 $\mathcal{U}(\mathcal{D}_{i-1})$ that 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 2.4(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 2.4(1) for the single pseudo-trapezoid $D_1^0 = D_1$ and Figure 2.4(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 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 2.4(4).

At the end of the incremental process, after inserting all the m 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

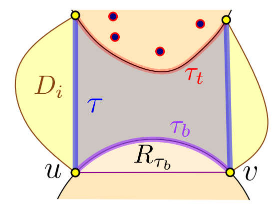

To get a triangulation, we associate a polygonal vertical pseudo-trapezoid τ^* with each pseudo-trapezoid $\tau \in \mathcal{T}$. We obtain τ^* from τ by replacing the bottom boundary τ_b and the top boundary τ_t of τ by respective polygonal chains τ_b^* and τ_t^*, that are defined as follows.[1] Let 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 segment uv; clearly, by the convexity of D_i, R_{τ_b} is fully contained in D_i. See figure on the right.

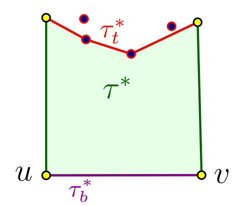

If R_{τ_b} contains no vertices of V_i, other than u and v (this will always be the case when $R_{\tau_b} \subseteq \tau$), we replace τ_b by $\tau_b^* = uv$. Otherwise, we replace τ_b by the chain τ_b^* of edges of the convex hull of $V_i \cap R_{\tau_b}$, other than the edge uv. We define τ_t^* analogously, 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.

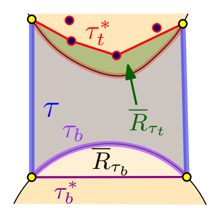

Note that, by construction, τ_b^* is a convex polygonal chain. From the point of view of τ, it is convex (resp., concave) if and only if τ_b is convex (resp., concave). (These statements become somewhat redundant when τ_b^* is the straight segnment uv.) An analogous property holds for τ_t^* and τ_t.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

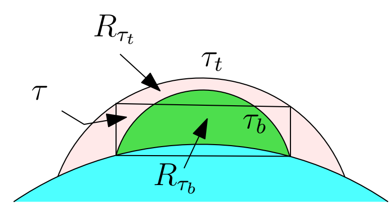

Note that R_{τ_b} and R_{τ_t} need not be disjoint, as illustrated in the figure on the right. Nevertheless, τ_b^* and τ_t^* cannot cross one another, as follows from Invariant (I2) that we establish below (in Lemma 2.2). This implies that τ^* is well defined. If τ_b^* and τ_t^* are not disjoint then they may only be pinched together at common vertices, or overlap in a single common connected portion (in the extreme case they may be identical).

[1] The 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.

This pinching or overlap, if it occurs, causes the interior of τ^* to be disconnected (into at most two pieces, as depicted in the figure to the right; it may also be empty, as is the case for D_1^0, illustrated in Figure 2.4(1)).

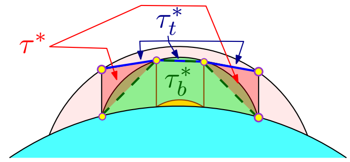

2.2.4 Filling the cavities

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 segments erected from other vertices of D_i^0. This can be seen all over Figure 2.4. For example, see the subdivision of the top arc of D_7 caused by the insertion of D_8 in Figure 2.4(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*. Note that, among the subarcs into which an arc of $\partial\mathcal{U}(\mathcal{D}_{i-1})$ is split, only the leftmost and rightmost extreme subarcs can be exposed (this follows from the pseudo-disk property of the objects of $\mathcal{D}$).

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 $\mathrm{CH}(R_\gamma \cap V_i)$, 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 $\overline{R}_\gamma$.

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 $\partial\mathcal{U}(\mathcal{D}_{j-1})$, $1 \le j \le i$, which were created and polygonalized when adding the corresponding pseudo-disk D_j. See figure on the right.

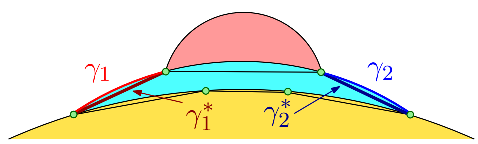

2.2.5 Putting it all together

When the pseudo-disks cover the plane. When the polygonalization process terminates, there are no more regions $\overline{R}_\gamma$, 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 (I1) in Lemma 2.2 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 $t \ge 2$ 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

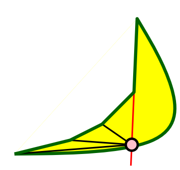

is then a cap, so assume that $t \geq 3$. To get such a decomposition, take an extreme edge of the concave polygonal chain, and extend it till it intersects the convex boundary of the crescent, at some point w, thereby chopping off a cap from the crescent. We then create the triangles that w spans with 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. Here too, by Euler's formula, the complexity of M is $O(m)$.

2.3 Analysis

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

Lemma 2.2. *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 forms 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 $\partial\mathcal{U}(\mathcal{D}_i)$. (Each of the latter holes is a crescent-like region of the form $\overline{R}_{\tau_b}$, $\overline{R}_{\tau_t}$, for some trapezoid τ, or $\overline{R}_\gamma$, for some exposed arc γ, as defined above.) The union of the holes of the latter kind (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 $\mathcal{T}_1^*$ and E_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 $\overline{R}_\gamma$ 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 $\overline{R}_\gamma$, contradicting the construction, or e' crosses γ too, to exit from $\overline{R}_\gamma$, 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 bound 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:

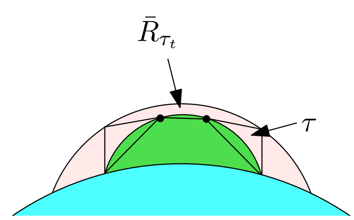

(a) The hole is a region of the form $\overline{R}_{\tau_b}$ or $\overline{R}_{\tau_t}$, for some $\tau \in \mathcal{T}_i \setminus \mathcal{T}_{i-1}$, such that $\overline{R}_{\tau_b}$ or $\overline{R}_{\tau_t}$ is contained in τ (if it lies outside τ, it becomes part of τ^*). Such a hole is contained in D_i, and is bounded by two concave or two convex chains, one of which, call it ζ^*, is polygonal, and the other, ζ, is part of ∂D_i^0. Moreover, ζ^*, if different from the chord e connecting the endpoints of ζ, passes through inner vertices of $\partial\mathcal{U}(\mathcal{D}_{i-1})$ that "stick into" the corresponding portion R_{τ_b} or R_{τ_t} of τ; see figure on the right.

(b) The hole is a region of the form $\overline{R}_\gamma$, 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 $\overline{R}_\gamma$ 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 2.4; for example, see the top part of D_1 in Figure 2.4(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). ∎

Theorem 2.3. *(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 that covers the plane, 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 and 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 2.2, 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 Theorem 2.3 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.* [MMP+91].) The construction proceeds by inserting the pseudo-disks of $\mathcal{D}$ in a *random* order, which, for simplicity, we denote as $D_1, \ldots, D_m$. (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 $\mathcal{D}_i = \{D_1, \ldots, D_i\}$ 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 (as depicted in the figure on the right), 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 (as depicted in Figure 2.5) is $O(k)$ (an easy consequence of the linear bound on the union complexity [KLPS86]), 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 $O(m \log m)$.

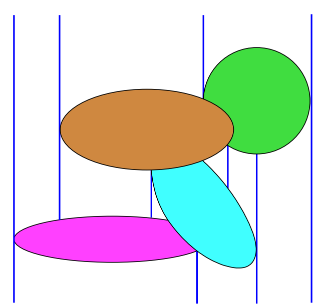

Figure 2.5

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 $D_i \cap \tau$ (if nonempty) 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., [Sei91].) 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 (dominated by the cost of handling the conflict lists) is $O(m \log m)$.

2.4.2 Construction of the polygonal chains and the triangulation

By (I2) of Lemma 2.2, 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 (polygonal) *interior*, $\mathcal{U}_{i-1}^*$, which is disjoint from $\partial\mathcal{U}(\mathcal{D}_{i-1})$ (except at the endpoints of the arcs γ^*), and the (half-polygonal) *exterior*, which is simply the (pairwise openly disjoint) union of the corresponding regions $\overline{R}_\gamma$.

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 separate binary search tree (since the leftmost and rightmost points of each pseudo-disk are vertices in the construction, each chain γ^* is indeed 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. At most two of these arcs still appear on $\partial\mathcal{U}(\mathcal{D}_i)$, and each of them is an extreme subarc of γ (we call them, as above, *exposed* arcs). All the others are now contained in D_i (we 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 "inner" 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 $\mathcal{U}(\mathcal{D}_{i-1})$.)

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 2.6 for an illustration.

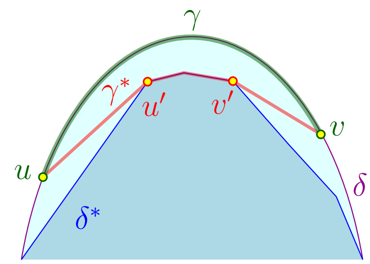

Figure 2.6

Note that the old arc δ may contain several new exposed or hidden arcs γ, so we apply the above procedure to each such arc γ. After doing this, 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 (the endpoints of the new arcs γ) and lies in between δ and δ^*. The region between $\hat{\delta}^*$ and δ^* is a new interior polygonal hole, and we partition it into simple cells, 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 inner 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.

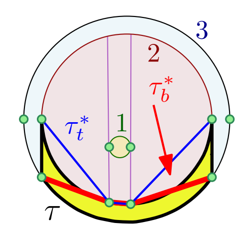

We next handle the fresh arcs. Each such arc is the top or bottom side of some new pseudo-trapezoid τ, say it is the bottom side τ_b. If τ_t is also fresh, then τ is a convex pseudo-trapezoid, and we replace each of τ_b, τ_t by the straight segment connecting its endpoints. If τ_t is hidden, we take 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. See figure on the right. 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 vv' 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 in the standard (say, real RAM) model, so the running time should be multiplied by this factor.

The preceding analysis implies the following theorem.

Theorem 2.4. *A triangulation of the union of m convex pseudo-disks covering the plane, with $O(m)$ triangles, such that each triangle is contained in a single pseudo-disk, can be computed, 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 2.5. *Given m convex polygons that are pseudo-disks, that cover the 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. A statement analogous to the second part of Theorem 2.4 holds in this case too.*

2.6 Extension to general convex shapes

Theorem 2.4 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 [AdBES14] and references therein.

Deploying the algorithm of Theorem 2.4 results in the desired confined triangulation of $\mathcal{U}(\mathcal{D})$. Extending the analysis to this general setup (and omitting the straightforward technical details), we obtain the following theorem.

Theorem 2.6. *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, 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), can be computed, in $O(u(m) \log m)$ expected time, under the assumption that every primitive geometric operation takes $O(1)$ time.*

3 Warm-up exercise: Constructing cuttings in the plane

In this section we apply the machinery developed in the previous section to obtain a new construction of $(1/r)$-cuttings in arrangements of lines in the planes.

Let L be a set of n lines in the plane in general position, and let $0 < r \le n$ be a parameter. In the planar setup, a $(1/r)$-cutting for L is a partition of the plane into $O(r^2)$ pairwise openly disjoint triangles, such that (the interior of) each triangle is crossed by at most n/r lines of L.

The construction of cuttings in the plane that we present here is similar in spirit to the more involved scheme for approximating the level in arrangements of planes in three dimensions, as presented in the following Section 4.

3.1 Tools

3.1.1 Divisions

We begin by reviewing the construct of a ***κ-division*** of a planar graph, which is a decomposition of such a graph into subgraphs, and a refined and stronger variant of the planar separator theorem of Lipton and Tarjan [LT79] and Miller's cycle separator theorem [Mil86]. It goes back to Frederickson's 30-years-old work [Fre87], and has eventually culminated in the fast κ-division algorithm of Klein *et al.* [KMS13]. We remind the reader that a graph is *biconnected* if any pair of vertices are connected by at least two vertex-disjoint paths.

Definition 3.1 (Fredman [Fre87]). Given a non-crossing plane drawing of a planar triangulated and biconnected graph G with N vertices, and a parameter $\kappa < N$, a ***κ-division*** of G is a decomposition of G into m connected subgraphs $G_1, \ldots, G_m$, such that

(i) $m = O(N/\kappa)$,
(ii) each G_i has at most κ vertices,
(iii) each G_i has at most $\beta\sqrt{\kappa}$ *boundary vertices*, for some absolute constant β, 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).

As shown in Klein *et al.* [KMS13], a κ-division of a planar triangulated and biconnected graph with N vertices can be computed in $O(N)$ time.[2]

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

Another tool that we need is the following folklore result, whose proof is included for the sake of completeness.

Lemma 3.2. *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. The analysis easily extends to the more general case too.

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 the convex hull $C^* = \mathrm{CH}(C \cup C')$, that appear along ∂C^* in this 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'.

[2] The algorithm of [KMS13] constructs κ-divisions for a geometrically increasing sequence of values of the parameter κ, in overall $O(N)$ time.

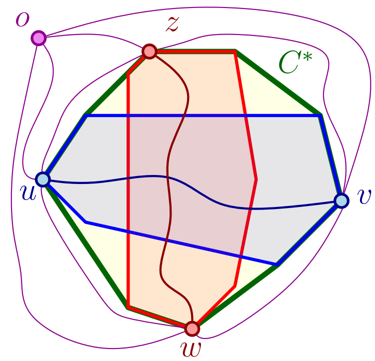

Figure 3.1

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 3.3. *Let $\mathcal{P} = \{P_1, \ldots, P_m\}$ 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 $i = 1, \ldots, m$. Then $\mathcal{C} := \{C_1, \ldots, C_m\}$ is a collection of m convex (polygonal) pseudo-disks.*

3.2 Construction of cuttings in two dimensions

Combining the tools from the previous subsection, we obtain the following new construction of cuttings in the plane.

Theorem 3.4. *Given a set L of n lines in the plane, in general position, and a parameter $0 < r \le n$, one can decompose the plane into $O(r^2)$ pairwise openly disjoint triangles, such that (the interior of) each triangle is crossed by at most n/r lines of L.*

Proof: Consider the arrangement $\mathcal{A}(L)$, add a "fake" vertex at infinity, which serves as a common endpoint of all the unbounded rays in $\mathcal{A}(L)$, and triangulate every (bounded or unbounded) face of $\mathcal{A}(L)$ with more than three boundary edges, by adding diagonals. Let G be the resulting planar graph, whose vertices are the vertices of $\mathcal{A}(L)$, and each of whose edges is either an original (bounded or unbounded) edge of $\mathcal{A}(L)$, or one of the added diagonals. Clearly, G is planar, triangulated and, as is easily checked, also biconnected. It has $N := 1 + \binom{n}{2}$ vertices.

Construct a κ-division of G, for $\kappa = \left(\frac{n}{\beta r}\right)^2$, where β is the constant from the construction of κ-divisions. We get a partition of the plane into $m = O(N/\kappa) = O(n^2/\kappa) = O(r^2)$

subgraphs $G_1, \ldots, G_m$, and we turn each subgraph G_i into a (not necessarily simple) polygon P_i by forming the union of all the faces of G_i that are also faces of G. By the properties of κ-divisions, each P_i has at most $\beta\sqrt{\kappa} \le n/r$ vertices (and edges) of $\mathcal{A}(L)$ on its boundary.

We clean up the construction, as follows. If one of the polygons B in this collection has a hole, we remove the hole from B (i.e., add its area to B), and remove all the polygons contained inside the hole from the collection. We repeat this process until all the polygons are simple, and obtain a partition of the plane into $m = O(r^2)$ simple pairwise openly disjoint polygons $B_1, \ldots, B_m$, such that each polygon has at most $t := n/r$ vertices of $\mathcal{A}(L)$ on its boundary.

Since every line of L intersects ∂B at a vertex, and every line that intersects the interior of B must cross its boundary at least twice, it follows that the interior of B is crossed by at most t lines of L.

Now form the convex hulls $C_i = \mathrm{CH}(B_i)$, for $i = 1, \ldots m$. By Corollary 3.3, the set $\{C_1, \ldots, C_m\}$ is a set of m convex pseudo-disks. Hence, by Theorem 2.6, one can compute a triangulation of the plane into $O(m)$ triangles, such that each triangle is fully contained in one of these hulls. Since a line intersects the interior of C_i if and only if it intersects the interior of B_i, it follows that at most t lines of L can intersect the interior of C_i, for $i = 1, \ldots, m$, and therefore every triangle in the triangulation is crossed by at most $t = n/r$ lines. This shows that $\mathcal{T}$ is a $(1/r)$-cutting of $\mathcal{A}(L)$ of size $O(r^2)$, as desired. ∎

We remark that a disadvantage of this construction is that it takes $O(N) = O(n^2)$ time to perform. This also holds, by the way, for a naive implementation of Matoušek's deterministic construction of planar cuttings [Mat90].

4 Construction of shallow cuttings and approximate levels in three dimensions

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.

4.1 Sketch of the construction

Let H be a set of n planes in three dimensions in general position. Assume that, for a given parameter $0 < r \le n$, 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 $O(n \log n)$ time [Cha10, AM95] (better than what is stated in Theorem 4.3 below for such a large value of r). 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 [CS89] 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, using the κ-division technique reviewed in Section 3.1.1. Specifically, we set

$$\kappa := \left(\frac{cn - 43.5r}{9\beta r}\right)^2,$$

where β is the constant from property (iii) of κ-divisions (see Section 3.1.1). Notice that since $r \ll n$ we have $\kappa > 1$. Let L'_k denote the projection of L_k onto the xy-plane. We turn L'_k into a triangulated and biconnected planar graph G'_k, similarly to the way in which we handled planar arrangements of lines in Section 3.2. That is, we add a new vertex v_∞ at infinity, replace each ray $[p, \infty)$ of L'_k by the edge (p, v_∞), and triangulate each bounded or unbounded face, if needed, by adding diagonals. The resulting graph is planar and triangulated, and, as is easily checked, is also biconnected. We can therefore apply to G'_k the planar κ-division algorithm of Klein *et al.* [KMS13], as reviewed in Section 3.1.1, with the value of κ given above. The resulting κ-division of G'_k consists of

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

connected, possibly unbounded, polygons, $P_1, \ldots, P_m$, with pairwise disjoint interiors. The union of $P_1, \ldots, P_m$ covers the entire xy-plane, and the edges of these polygon are projectioins of (some) edges of L_k (including of diagonals drawn to triangulate the original faces of L_k).

By construction, each P_i is connected and has at most $\beta\sqrt{\kappa} = (cn - 43.5r)/(9r)$ edges (and also contains $O(\kappa)$ edges and vertices of the projected k-level in its interior). Let C_i denote the convex hull of P_i, for $i = 1, \ldots, m$. As shows in Corollary 3.3, $\mathcal{C} := \{C_1, \ldots, C_m\}$ is a collection of m (possibly unbounded) convex pseudo-disks whose union is the entire plane.

We then apply Theorem 2.3 to $\mathcal{C}$ and obtain a set S of $O(m)$ points in the xy-plane, and a triangulation T of S that covers the plane, 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 unique point on the level that is co-vertical with p. For a bounded triangle Δ of T, $\uparrow_k(\Delta)$ is defined as the triangle spanned by the lifted images of the three vertices of Δ. We lift an unbounded triangle Δ with vertices p, q, and v_∞ by lifting pq to $\uparrow_k(p)\uparrow_k(q)$, as before, and lifting each of its rays, say $[p, \infty)$, as follows. If $[p, \infty)$ is the projection of an original ray of L_k, we simply lift it to that ray. Otherwise, we lift $[p, \infty)$ to a ray $\uparrow([p, \infty))$ that emanates from $\uparrow_k(p)$ in a direction parallel to the plane which lies 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 coplanar, we add another ray r emanating from $\uparrow_k(p)$ parallel to $\uparrow([q,\infty))$. We add to T' the unbounded triangle spanned by $\uparrow([q,\infty))$, $\uparrow_k(p)\uparrow_k(q)$, and r, and the unbounded wedge spanned by r and $\uparrow([p,\infty))$. Let T' denote the corresponding collection of lifted (bounded and unbounded) triangles and wedges in $\mathbb{R}^3$, given by $T' = \{\uparrow_k(\Delta) \mid \Delta \in T\}$.

Note that the triangles of T' are in general not contained in L_k. However, for each triangle $\Delta' \in T'$, its (finite) vertices lie on L_k, and, as we show in Lemma 4.6 below, at most $9\beta\sqrt{\kappa} + 43.5 = cn/r$ planes of H can cross Δ'. This implies, returning now to the original value of k, that Δ' fully lies between the levels $\xi \pm cn/r$ 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 4.1. *Let H be a set of n non-vertical planes in $\mathbb{R}^3$ in general position, and let $0 < r \leq n$, $\varepsilon > 0$ 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 $\mathcal{A}(H)$, which can be computed in polynomial time.*

(The last assertion in the theorem is a consequence of the constructive nature of our analysis. Efficient implementation of this construction is described later in this section.)

To turn this approximate level into a shallow cutting, replace each $\Delta' \in T'$ (including each of the unbounded triangles and wedges, as constructed above) 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, as will be established in Lemma 4.6 below, (b) the top triangle of τ is crossed by at most $\frac{1}{3}\varepsilon k$ planes of H (in the preceding analysis, we wrote this bound as $\frac{cn}{r}$; this is the same value, recalling that $\varepsilon = 3c$ and $k = n/r$). 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 4.2. *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 or a wedge that is fully contained between the levels k and $(1+\varepsilon)k$ of $\mathcal{A}(H)$, and these triangles form a polyhedral terrain (we say that such a terrain approximates the k-level L_k up to a relative error of ε).*

4.2 Efficient implementation

We next turn our constructive proof into an efficient algorithm, and show:

Theorem 4.3. *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 we have:*

(a) *One can construct the k-shallow $((1+\varepsilon)/r)$-cutting of $\mathcal{A}(H)$ given in Theorem 4.2, or, equivalently, the ε-approximating terrain of the k-level in Theorem 4.1, in $O(n + r\varepsilon^{-6}\log^3 r)$ expected time.*

(b) *Computing the conflict lists of the vertical prisms takes an additional $O(n(\varepsilon^{-3} + \log \frac{r}{\varepsilon}))$ expected time.*

(c) *The algorithm, not including the construction of the conflict lists, computes a correct ε-approximating terrain with probability at least $1 - 1/r^{O(1)}$.*

(d) *If we also compute the conflict lists then we can verify, in $O(n/\varepsilon^3)$ time, that the cutting is indeed correct and thereby make the algorithm always succeed, at the cost of increasing its expected running time by a constant factor.*

Proof: **(a)** Let $(H, \mathcal{R})$ denote the range space in which 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 cross e. Clearly, $(H, \mathcal{R})$ has finite VC-dimension (see, e.g., [SA95]). We draw a random sample S of $n' = \dfrac{br}{\varepsilon^2}\log r$ planes from H, where b is a suitable constant. For b sufficiently large, such a sample is a *relative* $\left(\frac{1}{r}, \varepsilon\right)$*-approximation* for $(H, \mathcal{R})$, with probability $\geq 1 - 1/r^{O(1)}$; see [HS11] for full details concerning the definition and properties of relative approximations. In our context, this means (assuming that the sample is indeed a relative approximation) that each vertical segment or ray 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 or ray that intersects $x < n/r$ planes of H intersects at most $\frac{n'}{n}x + \varepsilon\frac{n'}{r}$ planes of S. (This holds, with probability $\geq 1 - 1/r^{O(1)}$, for all vertical segments and rays.)

The strategy is to use (the smaller) S instead of H in the construction, as summarized in Theorem 4.2, and then argue that a suitable approximate level in $\mathcal{A}(S)$ is also an approximation to level k in $\mathcal{A}(H)$ with the desired properties. 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 choose a random index ξ in the range $[k', k'+t']$, construct the ξ-level of $\mathcal{A}(S)$, and then apply the construction of the proof of Theorem 4.2 to this level, as will be detailed below.

Before doing this, we first show that the ξ-level in $\mathcal{A}(S)$ 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 $\left(\frac{1}{r}, \varepsilon\right)$-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 $(1-\varepsilon)(n'/n)(1+4\varepsilon)(n/r) \geq (1+\varepsilon)k' = k' + t'$, for $\varepsilon \leq 1/2$. Since this holds with probability $\geq 1 - 1/r^{O(1)}$, for every point p, we conclude that the entire ξ-level of $\mathcal{A}(S)$ is between levels k and $(1+4\varepsilon)k$ of $\mathcal{A}(H)$, with probability $\geq 1 - 1/r^{O(1)}$.

We can now apply the machinery in Theorem 4.2. 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 a randomized algorithm of Chan [Cha00],[3] which takes

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

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

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

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

We now continue the implementation of the construction in a straightforward manner. 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 $O(t') = O\left(\frac{1}{\varepsilon}\log r\right)$ edges. We compute their convex hulls in $O(mt') = O(\frac{r}{\varepsilon^4}\log r)$ time, and then construct the corresponding confined triangulation, using Corollary 2.5, in overall time

$$O(mt') + O(m \log m \log t') = O\left(\frac{r}{\varepsilon^4}\log r + \frac{r}{\varepsilon^3}\log\frac{r}{\varepsilon}\log\left(\frac{1}{\varepsilon}\log r\right)\right).$$

Finally, we need to lift the vertices of the resulting triangles to the ξ-level of $\mathcal{A}(S)$. This can be done, using a point location data structure over the xy-projection of this level, in $O(n_1 \log n_1 + m \log n_1) = O(\frac{r}{\varepsilon^5}\log^2 r \log\frac{r}{\varepsilon})$ time, where the $n_1 \log n_1$ is the preprocessing cost of the point-location data-structure. 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; the last two properties follow from Theorem 4.2, applied to $\mathcal{A}(S)$ with suitable parameters. 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 $\varepsilon < 1$). 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 $\mathcal{A}(H)$.

This at last completes the construction (excluding the construction of the conflict lists). Its overall expected cost is $O(n + r\varepsilon^{-6}\log^3 r)$. This establishes (a).

[3]The paper of Chan [Cha00] does not use shallow cuttings, so we are not using "circular reasoning" in applying his algorithm.

(b) To complete the construction, we next turn to its final stage which is to compute, for every semi-unbounded vertical prism Δ^* stretching below a triangle $\Delta' \in T'$, the set of planes of H that intersect it (i.e., the conflict list of the prism). To this end, we put the vertices of T' into the range reporting data structure of Chan [Cha00]. In this structure, after preprocessing, in $O(\frac{r}{\varepsilon^3}\log\frac{r}{\varepsilon})$ expected time, one can report, for any given query half-space h^+, the points in $h^+ \cap T'$, in $O(\log\frac{r}{\varepsilon} + |h^+ \cap T'|)$ expected time (we recall again that this data range reporting structure of Chan is simple and does not use shallow cuttings). We query this data structure with the set of halfspaces h^+, bounded from below by the respective planes $h \in H$, and, for each vertex x of T' that we report, we add h to the conflict lists of the prisms incident to x. This takes $O(n\log\frac{r}{\varepsilon} + \frac{n}{\varepsilon^3})$ expected time, since the total size of the conflict lists is $O(\frac{r}{\varepsilon^3}\cdot\frac{n}{r}) = O(\frac{n}{\varepsilon^3})$ (in expectation and with probability $\geq 1 - 1/r^{O(1)}$).

(c) and **(d)** (c) follows because the probability that the sample S fails to be a relative $\left(\frac{1}{r},\varepsilon\right)$-approximation for $(H,\mathcal{R})$ is at most $1/r^{O(1)}$. When the sample does indeed fail, T' may fail to be the desired k-shallow $((1+\varepsilon)/r)$-cutting. Such a failure happens if and only if there exists a vertex of T' whose conflict list is of size smaller than k or larger than $(1+12\varepsilon)k$. When we detect such a conflict list, we repeat the entire computation. Since the failure probability is small the expected number of times we will repeat the computation is (a small) constant. ■

We now proceed to fill in the details of the various steps of the construction.

4.3 Crossing properties of the planar subdivision

Recall that our construction computes a t^2-division of the xy-projection L'_k of L_k where $t := (cn/r - 43.5r)/9\beta r$ (and recall that $k = n/r$, $r \ll n$, and β is a constant). Our goal in the rest of this section is to show that the lifting $\uparrow_k(\Delta)$ of any triangle Δ contained in the convex hull C of a polygon P of this decomposition intersects at most ck planes of H. We prove this explicitly for bounded triangles, and the proof for unbounded triangles (or wedges) 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 closed segment pq).

Lemma 4.4. *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, which 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 $\mathbf{cr}(q'',q') \leq \frac{1}{2}\mathbf{cr}(p',r') + 8.5$.*

Proof: For a point u, we denote by level(u) the number of planes lying vertically strictly below u. Put $k'' = \text{level}(q'')$. The point p' lies at level k, which is the closure of all points of level k. Thus the number of planes lying vertically strictly below q is k if p' is in the relative interior of a face of level k, at least $k-1$ if p' is in the relative interior of an edge of level k, and at least $k-2$ if p' is a vertex of level k. In either case, we have $\text{level}(p') \ge k-2$, and similarly for r', and thus

$$\mathbf{cr}(p', q'') \ge |\text{level}(p') - \text{level}(q'')| \ge |k - k''| - 2 \ ,$$

and

$$\mathbf{cr}(q'', r') \ge |\text{level}(p') - \text{level}(q'')| \ge |k - k''| - 2 \ .$$

On the other hand we have

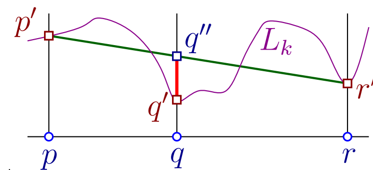

$$\mathbf{cr}(q', q'') \le |k'' - \text{level}(q')| + 3 \le |k - k''| + 5 \ .$$

(Indeed, if q'' lies above q' then $|k'' - \text{level}(q')| \le |k'' - (k-2)| \le |k'' - k| + 2$, and if q' lies above q'' then $|k'' - \text{level}(q')| \le |k'' - k|$. In addition, the difference in the levels of q' and q'' does not count the at most three planes that intersect $q'q''$ at q'', if q'' is above q', or at q', otherwise; this accounts for the terms 3 and 5 in the preceding inequality.) Hence,

$$\mathbf{cr}(q', q'') \le \frac{1}{2}\Big(\mathbf{cr}(p', q'') + \mathbf{cr}(q'', r') + 4\Big) + 5 \le \frac{1}{2}\big(\mathbf{cr}(p', r') + 3\big) + 7 = \frac{1}{2}\mathbf{cr}(p', r') + 8.5,$$

where the term 3 in the next-to-last expression is due to the potential double counting of the (up to) three planes passing through q'', in both terms $\mathbf{cr}(p', q'')$ and $\mathbf{cr}(q'', r')$. ∎

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.

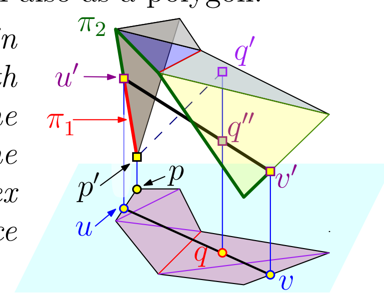

Lemma 4.5. *Let H be a set of n non-vertical planes in $\mathbb{R}^3$ in general position. Let P' be a bounded connected polygon with t edges that lies on the k-level L_k of $\mathcal{A}(H)$, such that all the boundary edges of P' are edges of L_k. Let p' be a vertex of the external boundary of P', and 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 $3t + 14.5$.*

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

$$\mathbf{cr}(p', q') \le \mathbf{cr}(p', u') + \mathbf{cr}(u', q'') + \mathbf{cr}(q'', q') \le \mathbf{cr}(p', u') + \mathbf{cr}(u', v') + \mathbf{cr}(q'', q').$$

Let π_1 and π_2 be the two portions of the external boundary that connect p' and u', and u' and v', respectively, and that do not overlap. Now, by Lemma 4.4, we have $\mathbf{cr}(q'', q') \le \frac{1}{2}\mathbf{cr}(u', v') + 8.5$, so we get

$$\mathbf{cr}(p', q') \le \mathbf{cr}(p', u') + \frac{3}{2}\mathbf{cr}(u', v') + 8.5 \le \mathbf{cr}(\pi_1) + \frac{3}{2}\mathbf{cr}(\pi_2) + 8.5 \le \frac{3}{2}\mathbf{cr}(\partial P') + 13,$$

where $\partial P'$ denotes the external boundary of P', and where the last inequality follows because $\frac{3}{2}\mathbf{cr}(\pi_1) + \frac{3}{2}\mathbf{cr}(\pi_2)$ double counts the planes that pass through u', adding at most $\frac{3}{2} \cdot 3 = 4.5$ to the bound.

To bound the number of planes of H that intersect $\partial P'$, consider its vertices $p_1, p_2, \ldots, p_t$ (the actual number of vertices might be smaller since P' may not be simply connected). 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 cross any plane, and each plane that contains it contains both its endpoints. Therefore, the number $\mathbf{cr}(\partial P')$ of planes of H that intersect $\partial P'$ satisfies $\mathbf{cr}(\partial P') \leq 3 + 2(t-1) = 2t + 1$, from which the lemma follows. (Note that this analysis is somewhat conservative—for example, if the polygon P' uses only original edges of the k-level, the bound drops to $t + 2$.) ■

Lemma 4.6. *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 $9t + 43.5$.*

Proof: 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 4.5 and the triangle inequality,

$$\mathbf{cr}(p', q') \leq \mathbf{cr}(w, p') + \mathbf{cr}(w, q') \leq 2(3t + 14.5),$$

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

Acknowledgments. We thank János Pach for pointing out that a variant of Theorem 2.3 is already known. We also thank the anonymous referees for their useful feedback.

References

[AACS98] 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.

[AC09a] P. Afshani and T. M. Chan. On approximate range counting and depth. *Discrete Comput. Geom.*, 42(1):3–21, 2009.

[AC09b] 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, 2009.

[ACT14] 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.

[AD97] 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.

[AdBES14] B. Aronov, M. de Berg, E. Ezra, and M. Sharir. Improved bounds for the union of locally fat objects in the plane. *SIAM J. Comput.*, 43(2):543–572, 2014.

[AE99] 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., Providence, RI, 1999.

[Aga90a] P. K. Agarwal. Partitioning arrangements of lines I: an efficient deterministic algorithm. *Discrete Comput. Geom.*, 5:449–483, 1990.

[Aga90b] P. K. Agarwal. Partitioning arrangements of lines: II. Applications. *Discrete Comput. Geom.*, 5:533–573, 1990.

[Aga91a] 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., Providence, RI, 1991.

[Aga91b] P. K. Agarwal. *Intersection and Decomposition Algorithms for Planar Arrangements.* Cambridge University Press, New York, 1991.

[AHZ10] P. Afshani, C. H. Hamilton, and N. Zeh. A general approach for cache-oblivious range reporting and approximate range counting. *Comput. Geom. Theory Appl.*, 43(8):700–712, 2010.

[AM95] P. K. Agarwal and J. Matoušek. Dynamic half-space range reporting and its applications. *Algorithmica*, 13:325–345, 1995.

[AS98] P. K. Agarwal and S. Suri. Surface approximation and geometric partitions. *SIAM J. Comput.*, 27(4):1016–1035, 1998.

[AT14] 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.

[BR52] R. P. Bambah and C. A. Rogers. Covering the plane with convex sets. *J. London Math. Soc.*, s1-27(3):304–314, 1952.

[CEG+90] 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.

[CEGS91] B. Chazelle, H. Edelsbrunner, L. J. Guibas, and M. Sharir. A singly-exponential stratification scheme for real semi-algebraic varieties and its applications. *Theoret. Comput. Sci.*, 84:77–105, 1991. Also in *Proc. 16th Int. Colloq. on Automata, Languages and Programming*, pages 179–193, 179–193.

[CF90] B. Chazelle and J. Friedman. A deterministic view of random sampling and its use in geometry. *Combinatorica*, 10(3):229–249, 1990.

[Cha93] B. Chazelle. Cutting hyperplanes for divide-and-conquer. *Discrete Comput. Geom.*, 9(2):145–158, 1993.

[Cha00] T. M. Chan. Random sampling, halfspace range reporting, and construction of ($\leq k$)-levels in three dimensions. *SIAM J. Comput.*, 30(2):561–575, 2000.

[Cha04] B. Chazelle. Cuttings. In D. P. Mehta and S. Sahni, editors, *Handbook of Data Structures and Applications*, chapter 25. Chapman and Hall/CRC, Boca Raton, Florida, 2004.

[Cha05] T. M. Chan. Low-dimensional linear programming with violations. *SIAM J. Comput.*, 34(4):879–893, 2005.

[Cha10] T. M. Chan. A dynamic data structure for 3-d convex hulls and 2-d nearest neighbor queries. *J. Assoc. Comput. Mach.*, 57(3):1–15, 2010. Art. 16.

[Cla87] K. L. Clarkson. New applications of random sampling in computational geometry. *Discrete Comput. Geom.*, 2:195–222, 1987.

[CS89] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.

[CT15] 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.

[EHS04] 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.

[Fre87] G. N. Frederickson. Fast algorithms for shortest paths in planar graphs, with applications. *SIAM J. Comput.*, 16(6):1004–1022, 1987.

[Har00] S. Har-Peled. Constructing planar cuttings in theory and practice. *SIAM J. Comput.*, 29(6):2016–2039, 2000.

[HKS16a] S. Har-Peled, H. Kaplan, and M. Sharir. Approximating the k-level in three-dimensional plane arrangements. In *Proc. 27th ACM-SIAM Sympos. Discrete Algs.* (SODA), pages 1193–1212, 2016.

[HKS16b] S. Har-Peled, H. Kaplan, and M. Sharir. Approximating the k-level in three-dimensional plane arrangements. *CoRR*, abs/1601.04755, 2016.

[HS11] S. Har-Peled and M. Sharir. Relative (p, ε)-approximations in geometry. *Discrete Comput. Geom.*, 45(3):462–496, 2011.

[KLPS86] 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.

[KMS13] 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.

[Kol04] V. Koltun. Almost tight upper bounds for vertical decompositions in four dimensions. *J. Assoc. Comput. Mach.*, 51(5):699–730, 2004.

[KS05] V. Koltun and M. Sharir. Curve-sensitive cuttings. *SIAM J. Comput.*, 34(4):863–878, 2005.

[LT79] R. J. Lipton and R. E. Tarjan. A separator theorem for planar graphs. *SIAM J. Appl. Math.*, 36:177–189, 1979.

[Mat90] J. Matoušek. Construction of ε-nets. *Discrete Comput. Geom.*, 5:427–448, 1990.

[Mat92a] J. Matoušek. Efficient partition trees. *Discrete Comput. Geom.*, 8:315–334, 1992.

[Mat92b] J. Matoušek. Range searching with efficient hierarchical cutting. *Discrete Comput. Geom.*, 10:157–182, 1992.

[Mat92c] J. Matoušek. Reporting points in halfspaces. *Comput. Geom. Theory Appl.*, 2(3):169–186, 1992.

[Mat98] J. Matoušek. On constants for cuttings in the plane. *Discrete Comput. Geom.*, 20:427–448, 1998.

[Mil86] G. L. Miller. Finding small simple cycle separators for 2-connected planar graphs. *J. Comput. Syst. Sci.*, 32(3):265–279, 1986.

[MMP+91] 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.

[MRR14] 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.

[Mul94] K. Mulmuley. An efficient algorithm for hidden surface removal, II. *J. Comput. Syst. Sci.*, 49:427–453, 1994.

[PA95] J. Pach and P. K. Agarwal. *Combinatorial Geometry.* John Wiley & Sons, New York, 1995.

[Ram99] E. A. Ramos. On range reporting, ray shooting and k-level construction. In *Proc. 15th Annu. Sympos. Comput. Geom.* (SoCG), pages 390–399. ACM, 1999.

[SA95] M. Sharir and P. K. Agarwal. *Davenport-Schinzel Sequences and Their Geometric Applications.* Cambridge University Press, New York, 1995.

[Sei91] 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.

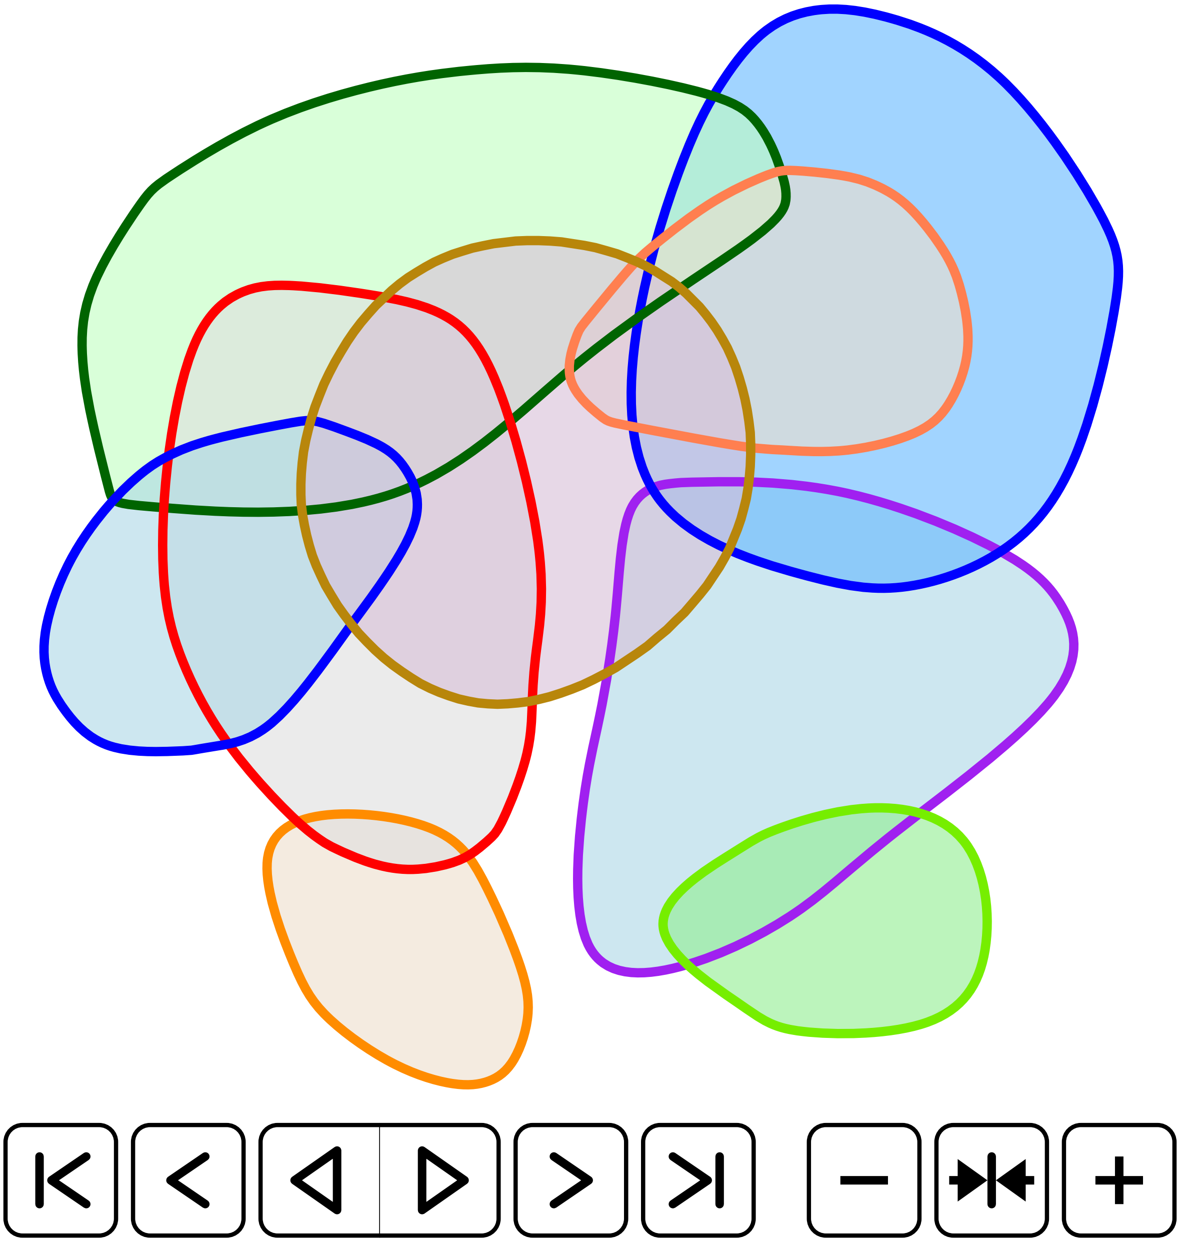

Figure 4.1: Animation of the algorithm in Section 2; one needs `Acrobat reader` to see the animation (as of April 2021, `okular` also works). Click on the figure to make it start.