

Depth contours in arrangements of halfplanes*

Sariel Har-Peled[†]

Micha Sharir[‡]

September 25, 2016

I'm not aware of too many things
I know what I know, if you know what I mean
Philosophy is the talk on a cereal box
Religion is the smile on a dog
I'm not aware of too many things
I know what I know, if you know what I mean
Choke me in the shallow waters
Before I get too deep.
– What I Am, Edie Brickell and Kenny Withrow.

Abstract

Let H be a set of n halfplanes in \mathbb{R}^2 in general position, and let $k < n$ be a given parameter. We show that the number of vertices of the arrangement of H that lie at depth exactly k (i.e., that are contained in the interiors of exactly k halfplanes of H) is $O(nk^{1/3} + n^{2/3}k^{4/3})$. The bound is tight when $k = \Theta(n)$. This generalizes the study of Dey [Dey98], concerning the complexity of a single level in an arrangement of lines, and coincides with it for $k = O(n^{1/3})$.

1 Introduction

Given a set of n points in the plane, the *k -set* problem is to provide a tight bound on the maximum number of lines that pass through a pair of points and contain exactly k points on one of their sides. The currently best known lower bound is the slightly superlinear bound $ne^{\Omega(\sqrt{\log k})}$, due to Tóth [Tót01] and Nivasch [Niv08], and the best known upper bound is $O(nk^{1/3})$, due to Dey [Dey98]. This problem has a long history, and closing this gap is one of the major open problems in discrete geometry.

In the dual setting, we are given a set L of n lines in the plane, and the problem is to bound the maximum number of vertices of the arrangement $\mathcal{A}(L)$ that have exactly k lines that pass strictly below them. This is known as the *k -level* problem.

In this paper, we study a natural variant of the k -level problem. We first note that an equivalent formulation of the k -level problem can be obtained by associating with each line of ℓ the (open) halfplane ℓ^+ bounded by ℓ and lying above it, and by observing that the vertices of $\mathcal{A}(L)$ at level k are exactly those vertices that are contained in exactly k of the resulting halfplanes. It is thus natural to ask what happens if we do not restrict the given halfplanes to lie above their bounding lines.

*Work by Sariel Har-Peled was partially supported by NSF AF awards CCF-1421231 and CCF-1217462. 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 Computer Science 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; <http://sarielhp.org/>.

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

In this generalized setting, we are thus given a set H of n open halfplanes, and we define the *depth* of a point $q \in \mathbb{R}^2$ to be the number of halfplanes of H that contain q . The *k -contour* problem is to bound the maximum possible number of vertices (of the arrangement of the bounding lines of the halfplanes of H) and edges at depth exactly k . See [Figure 2.2](#) for an example of these contours. To simplify the analysis, we require the halfplanes of H to be in *general position*, which means that none of their bounding lines is vertical, no pair of these lines are parallel, and no triple of lines are concurrent.

It is not hard to see that the k -contour can have quadratic complexity in the worst case, albeit only for $k = \Theta(n)$; see [Lemma 2.9](#), [Remark 2.10](#), and [Figure 2.1](#). This demonstrates that the k -level and the k -contour are fundamentally different structures, as the complexity of the k -level is at most $O(n^{4/3})$, for any k . Our main result, stated next, offers two surprises: (a) the quadratic lower bound can only arise when $k = \Theta(n)$, and (b) the complexity of the k -contour is asymptotically the same as that of the k -level for $k = O(n^{1/3})$.

Theorem 1.1. *Let H be a set of n halfplanes in \mathbb{R}^2 in general position, and let $k < n$ be a given parameter. The number of vertices of $\mathcal{A}(H)$ that lie at depth exactly k is $O(nk^{1/3} + n^{2/3}k^{4/3})$. The bound is worst-case tight when $k = \Theta(n)$.*

We actually derive a sharper form of this theorem, in which the bound is expressed also in terms of the number of *x -extreme* vertices at depth k . We give the refined version after introducing the necessary notations, in the following section.

We note that the k -contour problem, as well as its predecessors, the k -set and the k -level problems, are difficult because we focus on a single fixed level or depth. If we care about the first k levels, or the first k depths, the problems become much simpler, and tight bounds are known. Specifically, Alon and Györi have established the worst-case tight bound $n(k+1)$ for the number of vertices at level at most k in an arrangement of n lines [[AG86](#)], and, as we show below, the same bound also holds for the number of vertices at depth at most k in an arrangement of n halfplanes. Our analysis also provides a simple alternative proof of the bound in [[AG86](#)], for the case of levels.

It might be interesting to note that the k -contour problem can also be expressed in a dual, perhaps somewhat artificial, setting, in which we have a set of n points in the plane, some of which are red and some blue, and a parameter $k < n$, and we want to bound the number of lines that pass through a pair of points and are such that the number of red points below the line plus the number of blue points above the line is exactly k .

The analysis and organization of the paper. After introducing the terminology and deriving various basic properties of depth contours, in [Section 2](#), we follow the approach of Agarwal *et al.* [[AACS98](#)] and Dey [[Dey98](#)], and construct concave (and convex) chains that cover the vertices of the k -contour. We then bound the number of common tangencies between pairs of chains, which, as in [[Dey98](#)], allows us, using the Crossing Lemma in a dual setting, to bound the number of vertices of the k -contour. In our case, though, the chain decomposition is somewhat more involved, and requires care, since the structures we work with are not x -monotone in nature. In particular, ensuring the crucial property that the chains are pairwise non-overlapping, which is easy in the case of levels, is rather less trivial in our case. [Section 3](#) introduces the chains, and derives various properties thereof, including the property of being pairwise non-overlapping.

Bounding the number of common tangencies between chains takes place in [Section 4](#). Agarwal *et al.* [[AACS98](#)] state (without details) a general bound on the number of such tangencies. Applied in our case it yields a bound of $O(n^{4/3} + n^{2/3}c^{2/3})$, where c is the number of chains, which is weaker than our bound for $k = O(n^{1/3})$. To get the improved bound, we need a somewhat more refined analysis of the tangencies and their structure (in the k -contour case). This is taken care of in [Section 4](#).

A brief conclusion, including the potential application of our result to levels and depth in higher-dimensional arrangements, is given in the final [Section 5](#).

2 Preliminaries

Notations. Let H be a set of n (open) halfplanes in \mathbb{R}^2 in general position, as defined in the introduction. Let $k < n$ be a given parameter. A point $p \in \mathbb{R}^2$ is at **depth** k if p is contained in exactly k halfplanes of H . Let L denote the set of lines bounding the halfplanes of H . Consider the arrangement $\mathcal{A}(L)$ of L ; we will sometimes also refer to it as the arrangement of H , and denote it as $\mathcal{A}(H)$. Each vertex, edge, and face of $\mathcal{A}(L)$ has a fixed depth. Let E_k denote the collection of all the edges of $\mathcal{A}(L)$ at depth exactly k . Our goal is to derive an upper bound for $|E_k|$. Let C_k denote the union of (the closures of) the edges of E_k . We refer to C_k as the **depth- k contour** (or just **k -contour**). For each edge e in E_k , each endpoint of e is incident to exactly one other edge of E_k , as is easily checked (see [Figure 3.2](#)). Hence, C_k is a union of pairwise disjoint cycles and unbounded paths.

A vertex at depth k is **x -extreme** if it is a local x -minimum or x -maximum of C_k . A halfplane is an **upper halfplane** (resp., **lower halfplane**) if it lies above (resp., below) its bounding line. (By our general position assumption, no halfplane in H is bounded by a vertical line, so we ignore such halfplanes in what follows.)

A special case of this setup, discussed in the introduction, is when all the halfplanes in H are upper halfplanes (or all are lower halfplanes). Assuming that they are all upper halfplanes, C_k is the **k -level** of $\mathcal{A}(L)$. It is an x -monotone unbounded connected path, and we recall that, as shown by Dey [[Dey98](#)], the number of its vertices is $O(nk^{1/3})$.

In this paper we address the following general problem.

Problem 2.1. For numbers n, k and τ , what is the maximum possible number of vertices of the k -contour $C_k = C_k(H)$ in the arrangement of a set H of n open halfplanes in the plane (in general position), when $C_k(H)$ has at most τ x -extreme vertices (at depth k), where the maximum is taken over all such sets H .

The following refined version of [Theorem 1.1](#) gives a sharper upper bound for this quantity.

Theorem 2.2. *Let H be a set of n halfplanes in \mathbb{R}^2 in general position, let $k < n$ be a given parameter, and let τ denote the number of x -extreme vertices of $C_k(H)$. Then the overall number of vertices of $C_k(H)$ is $O(nk^{1/3} + n^{2/3}\tau^{2/3})$.*

As we will show below, we always have $\tau = O(k^2)$, so the new bound is indeed always dominated by the bound in [Theorem 1.1](#), and is smaller when $\tau = o(k^2)$ (and when the second term dominates). In particular, in the k -level problem, all the halfplanes are upper halfplanes, and then $C_k(H)$ does not contain any x -extreme vertex, so the bound in [Theorem 2.2](#) coincides asymptotically with Dey's bound $O(nk^{1/3})$, but the coincidence already takes place for $\tau = O(\sqrt{nk})$.

2.1 The number of vertices of depth at most k

The following simple (probably folklore) lemma is a crucial ingredient of our analysis.

Lemma 2.3. *Let λ be an arbitrary line in \mathbb{R}^2 . All the points of the form $\lambda \cap \ell$, over all lines $\ell \in L \setminus \{\lambda\}$, that are at depth at most k are contained in a contiguous interval I of intersection points along λ , all of whose elements are at depth at most $2k$. The overall number of intersection points in I is at most $2k + 2$.*

Proof: Let u and v be the extreme intersection points on λ at depth at most k . (One of u, v could lie at infinity; as the proof will shortly show, it is impossible that both lie at infinity when $k < n/2$.) We claim that the interval I of intersection points between u and v (including u, v) is the desired interval. Clearly, every intersection point along λ at depth at most k lies in I . Suppose to the contrary that I contained an intersection point q at depth greater than $2k$. Each (open) halfplane that contains q must contain either u or v , so at least one of these points has to be at depth larger than k , a contradiction that establishes the claim.

We next bound the number of intersection points in the interior of I . Each of these points is incident to at least one line of L , and all these lines are distinct. Each halfplane bounded by one of these lines must contain

either u or v . Thus, if the number of inner points were larger than $2k$, one of u, v would have to be at depth larger than k , a contradiction. ■

Remark 2.4. **Lemma 2.3** extends to any higher dimension (and also to $d = 1$), with essentially an identical proof. That is, for a set H of halfspaces in \mathbb{R}^d , and for any segment pq , whose endpoints are at depth at most k , the relative interior of pq intersects at most $2k$ boundary (hyper)planes of the halfspaces of H , all at depth at most $2k$. In particular, if the line supporting pq is a line of $\mathcal{A}(H)$ then its (relatively open) portion between p and q contains at most $2k$ vertices of $\mathcal{A}(H)$ (all at depth at most $2k$).

At most k -level. As an interesting application of **Lemma 2.3**, we obtain an immediate proof of the well known tight bound on the number of vertices in an arrangement of lines at level at most k (see [AG86, GP84]), and at the same time establish the same tight bound for the number of vertices at depth at most k in an arrangement of halfplanes. The proof is (arguably) slightly simpler than the proof of [AG86] (the proof of [GP84] produces a larger constant).

Lemma 2.5. *For a set H of n halfplanes in \mathbb{R}^2 , the number of vertices of $\mathcal{A}(H)$ at depth at most k is at most $n(k + 1)$. The bound is tight in the worst case.*

Proof: Simply observe that, by **Lemma 2.3**, there are at most $2k + 2$ such vertices along each line of L , and each vertex is counted exactly twice. The tightness of the bound follows from the tightness of the bound in [AG86] for the case of levels. ■

2.2 Setup

We somewhat simplify the setting by clipping the unbounded rays of C_k , if any, so as to make C_k (artificially) bounded. Let U be the set of all finite vertices of $\mathcal{A}(H)$ at depth at most k (this omits virtual vertices at “infinity” lying on rays that have depth at most k), and let P' be the convex hull of U . Clearly, P' is a bounded convex polygon. To simplify the analysis, slightly expand P' (e.g., by taking its Minkowski sum with a sufficiently small square), so that the resulting convex polygon P does not contain any vertex of $\mathcal{A}(H)$ on its boundary, nor does it fully contain any edge. By construction, only the unbounded rays of C_k , if any, reach the outside of P , and we clip each such ray to its portion within P ; the clipping occurs at exactly one point on each ray, as is easily seen. Thus each unbounded ray is replaced by a “prefix” subsegment, one of whose endpoints lies on ∂P and the other is the original endpoint of the ray.

Lemma 2.6. *The region P satisfies the following properties.*

- (i) *All the vertices of $\mathcal{A}(H)$ at depth at most k are contained in the interior of P .*
- (ii) *The maximum depth of any point on ∂P is $2k + 2$.*
- (iii) *The maximum depth of any point in P is $k^+ := 3k + 2$.*

Proof: Part (i) is immediate from the construction.

For part (ii), any point on $\partial P'$ is either in U , or on a segment connecting two points of U , and therefore, by (the proof of) **Lemma 2.3**, has depth at most $2k$. Expanding P' to get P increases the maximum depth of a point in P by at most 2 (assuming that the displacement of P' is sufficiently small).

For part (iii), any point $p \in P$ is on a segment connecting a point $q \in \partial P$ and a point $r \in U$. Arguing as in the proof of **Lemma 2.3**, p has depth at most $\text{depth}(q) + \text{depth}(r) \leq k + 2k + 2 = 3k + 2$. ■

The polygon P is the *domain* of interest, since all the vertices of $\mathcal{A}(H)$ of depth k are contained in its interior. It suffices to restrict our analysis to the portion of the arrangement lying in P .

Observation 2.7. *By **Lemma 2.3**, every line of L contains at most $O(k)$ vertices of $\mathcal{A}(H)$ at depth at most $3k + 2$, which, by **Lemma 2.6(iii)**, implies that P contains at most $O(nk)$ vertices of $\mathcal{A}(H)$.*

2.3 Counting paths and cycles

The modified (i.e., clipped) C_k consists of pairwise disjoint bounded paths and bounded cycles, where each bounded path terminates at points that lie on ∂P , and each bounded cycle remains unchanged from its version in the original C_k .

Each cycle of C_k can be charged to the at least two x -extreme vertices that it contains, so the number of cycles is at most $\tau/2$. Each path c of (the modified) C_k can be charged to its two endpoints, both lying on ∂P . The number of charged points is at most $2n$, because each line of L has at most two unbounded rays in the original E_k , and each of these rays has at most one intersection with ∂P . Consequently, the number of bounded paths in C_k is at most n .

We have the following well known upper bound on τ [Mat95].

Lemma 2.8. *The number τ of x -extreme vertices of C_k is $O(k^2)$.*

Proof: This is a consequence of the Clarkson-Shor analysis. Indeed, the number of such vertices (in fact, the number of these vertices over all depths $j = 0, 1, \dots, k$) is $O(k^2)$ times the number of x -extreme vertices of the 0-contour in the arrangement of a random sample of n/k halfplanes of H . However, the (unclipped) 0-contour, for any subset $H' \subseteq H$, consists of a single bounded convex cycle or unbounded convex path, which has at most two x -extreme vertices, and this establishes the claim. ■

We note that this bound is tight in the worst case. This is shown in the following simple construction. It is not new, and seems to be folklore [Mat95]. We include it here for the sake of completeness.

Lemma 2.9. *The number of x -extremal vertices of C_k can be $\Omega(k^2)$ in the worst case.*

Proof: The construction is depicted in [Figure 2.1](#). Specifically, we use the following $2k$ halfplanes (assuming k to be even).

$$\begin{array}{ll} x + y \geq j, & \text{for } j = 1, \dots, k/2, \\ x + y \leq j + \frac{1}{2}, & \text{for } j = 1, \dots, k/2, \\ -x + y \geq j, & \text{for } j = 1, \dots, k/2, \\ -x + y \leq j + \frac{1}{2}, & \text{for } j = 1, \dots, k/2, \end{array}$$

and ‘pad’ the construction with $n - 2k$ additional halfplanes, all very close to $y \geq k$, say. The first $2k$ boundary lines form a grid-like structure, with $\Theta(k^2)$ edges that lie at depth exactly k . These edges are arranged in a collection of pairwise disjoint 4-cycles, each enclosing a cell of the grid, and each such cycle has two x -extreme vertices. ■

Remark 2.10. Note that if we want the constructed grid to have $\Omega(n^2)$ vertices, its edges will be at depth $\Theta(n)$. That is, this construction is *not* quadratic in n when $k \ll n$. That a quadratic-size construction is possible only when $k = \Theta(n)$ follows from our main [Theorem 1.1](#). Note also that in the construction the overall number of vertices of C_k is $\Omega(k^2)$. This however is not known to be a tight bound, when $k \ll n$, when compared with the upper bound in [Theorem 1.1](#).

3 The complexity of the k -contour

We now adapt the analysis technique of Dey [Dey98], originally developed for *levels* in line arrangements, with several nontrivial modifications and adjustments.

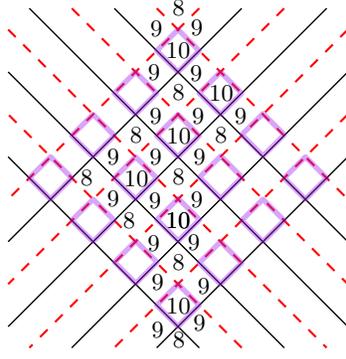


Figure 2.1: An example showing that C_k can contain $\Omega(k^2)$ x -extreme vertices and bounded cycles. Upper halfplanes are bounded by black lines, and lower halfplanes are bounded by red (dashed) lines. Traversing this arrangement, crossing (in the upward direction) a black line increases the depth by one, while crossing a red line decreases it by one. Here, the purple 4-sided cycles form C_9 . Observe that while there are faces at depth 10 in this arrangement, there are no edges with this depth. A number in a face indicates its depth.

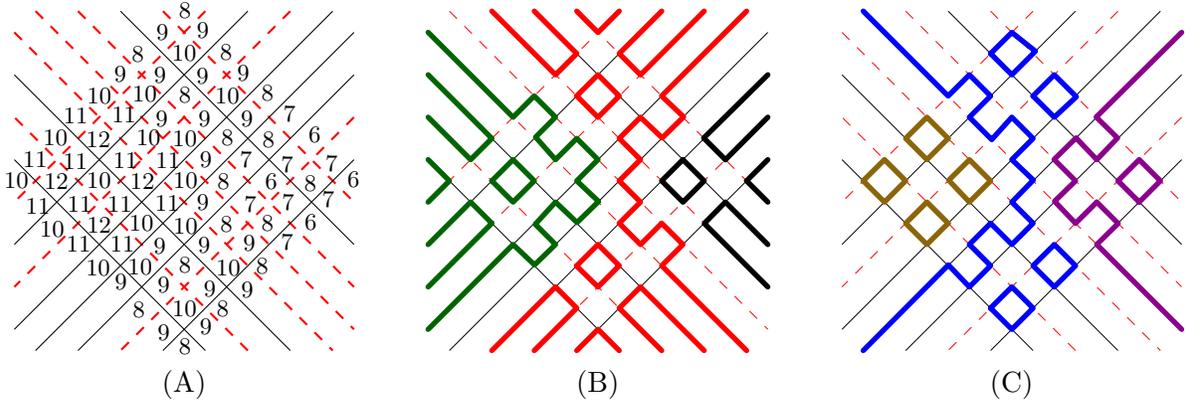


Figure 2.2: Another example of contours and their structure. (A) The lines bounding the corresponding halfplanes and the induced depth of the faces. As before black lines bound upper halfplanes, and red (dashed) lines bound lower halfplanes. (B) Contours 6, 8, and 10. (C) Contours 7, 9, and 11.

3.1 Decomposing C_k into chains

3.1.1 The different types of vertices

We classify each vertex w of C_k according to the structure of its local neighborhood as we sweep the plane from left to right. We get the following four types of vertices:

- (r) **right turn**: C_k makes a right turn at w and continues rightwards;
- (l) **left turn**: C_k makes a left turn at w and continues rightwards;
- (x) **extreme**: w is an x -extreme vertex, which means, as we recall, that w is a locally x -extremal (locally rightmost or leftmost) point of C_k . As above, we denote by τ the total number of such vertices;
- (b) **boundary**: an intersection point of an unbounded ray of (the original) C_k with ∂P .

We have already bounded the number τ of x -vertices by $O(k^2)$ (Lemma 2.8), and the number of b -vertices by $2n$. The real challenge is of course to bound the number of r -vertices and l -vertices. We explicitly bound the number of r -vertices, and a symmetric variant of the same argument yields the same bound on the number of l -vertices.

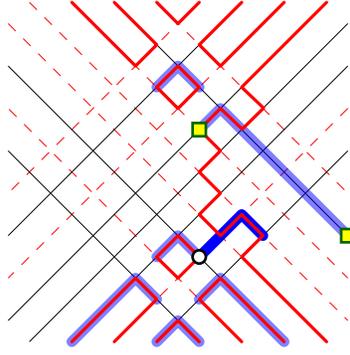


Figure 2.3: The concave chain decomposition of the 8-contour for the configuration of Figure 2.2. The circled marked vertex is an x -vertex that delimits two chains. The bolder chain reaches it along an edge that is not of depth k . The chain with the two square vertices starts on the left at an x -vertex, and ends on the right at a b -vertex.

3.1.2 The decomposition

The first step is to “represent” C_k as a collection of *concave chains*, or more precisely, cover the r -vertices of C_k by such chains. More precisely, we ensure that every r -vertex of C_k will be an inner vertex of (exactly) one of the chains. The l -vertices are not guaranteed to be covered by the chains, but we will also apply a symmetric construction of *convex chains*, which will have the l -vertices as inner vertices.

To this end, start at an arbitrary r -vertex v of C_k . Trace C_k from v to the right, along the line ℓ supporting the right edge of C_k incident to v , and follow ℓ until we reach the first r -, x -, or b -vertex w (in particular, the chain marches through l -vertices of C_k without changing direction).

If w is an r -vertex, we turn right at w , and continue the tracing of the chain in this manner, until we reach a “terminal” x - or b -vertex. If w is an x -vertex or a b -vertex on ∂P , it becomes the right endpoint of the chain. We then go back to the initial v , and apply a symmetric tracing to the left.

Each such “two-sided tracing” produces one concave chain. We now repeat the process, starting at another r -vertex of C_k that has not yet been visited, and trace a new chain from it. We keep doing this until all the r -vertices of C_k are encountered.

See Figure 2.3 for an example of the chain cover. This figure also demonstrates, that unlike b -vertices, each lying on a single chain, x -vertices might be shared by several (up to four—see below) chains.

Let \mathcal{C} denote the collection of these concave chains.

3.1.3 Properties

The chains of \mathcal{C} satisfy several useful properties. The first few properties are easy, and are given without a proof.

- (a) All the inner vertices of a chain are r -vertices of C_k (at which C_k makes right turns), and every r -vertex appears as an inner vertex of some chain.
- (b) Each constructed chain is not necessarily confined to C_k , and may go out of the contour and back in many times. (This happens when the chain goes through l -vertices.) Each chain however is well defined, given the “seed” vertex v .

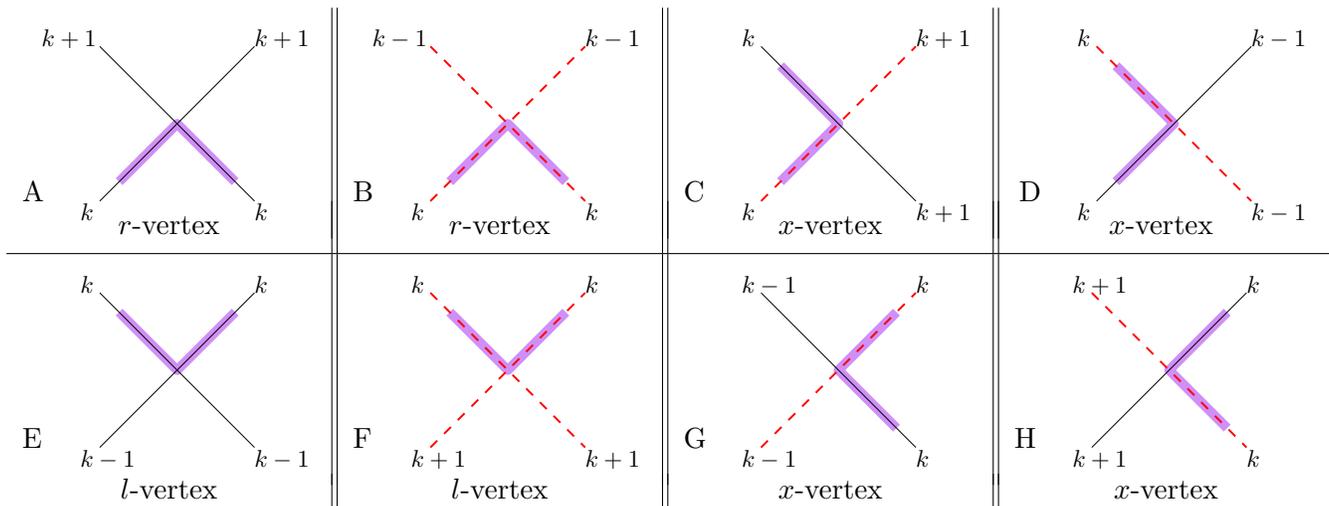


Figure 3.2: The k -contour near every possible type of crossing. Upper halfplanes are bounded by black lines, and lower halfplanes are bounded by red (dashed) lines. The labels are the depths of the associated edges. The captions are the classification of the vertices according to [Section 3.1.1](#).

(c) Not all vertices of C_k have necessarily been encountered by this process. Any unaccounted vertex is either an x -vertex or an l -vertex of C_k . See [Figure 3.1](#) for an illustration. Some of these vertices (only l -vertices) may become internal points of edges of the chains, while others may have not been visited at all. This is no cause for alarm, because, as already mentioned, we also apply a symmetric construction for the l -vertices, which will produce symmetrically defined *convex chains*, and the two phases together will account for all r - and l -vertices of C_k (some x - and b -vertices might still remain uncovered). In what follows, though, we will only focus on concave chains (and r -vertices).

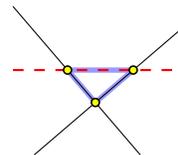


Figure 3.1: Two x -vertices and one l -vertex that are not covered in the concave chain decomposition.

- (d) Since C_k is contained in P , the overall number of vertices of $\mathcal{A}(L)$ that are contained in the chains (not necessarily chain vertices) is $O(nk)$; see [Observation 2.7](#). (This bound also includes the additional $2n$ “fake” b -vertices, which do not appear as original vertices in $\mathcal{A}(L)$.)
- (e) By the construction of P , all the chains of \mathcal{C} are bounded.

Remark 3.1. A chain might have one b -vertex and one x -vertex as endpoints; see [Figure 2.3](#).

Lemma 3.2. *Let c be a concave chain in \mathcal{C} . The halfplanes whose bounding lines support the edges of c are either all upper halfplanes or all lower halfplanes.*

Proof: Since all the inner vertices of c are r -vertices, it suffices to show that, for any r -vertex v , the two halfplanes whose bounding lines support the two edges of C_k incident to v are either both upper or both lower. This follows by a straightforward case analysis, depicted in [Figure 3.2](#), where all upper / lower combinations are considered, and where only the two with the same halfplane directions (with a local structure depicted in [Figure 3.2\(A,B\)](#)) yield r -vertices. ■

3.1.4 Chains do not overlap

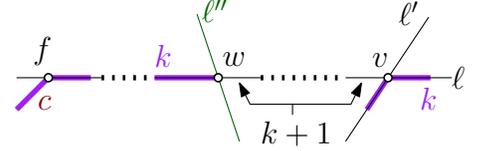
Two chains c_1, c_2 *overlap*, if there exists a line ℓ , such that $c_1 \cap c_2 \cap \ell$ is an interval of nonzero length. We next show that such overlaps cannot occur.

An r -vertex v has the (characterizing) property that, as we sweep from left to right through v , the k -contour reaches v from its incident bottom-left edge, and leaves it on its incident bottom-right edge; see [Figure 3.2\(A,B\)](#).

Lemma 3.3. *Let v be an r -vertex that lies on some chain $c \in \mathcal{C}$. Then v must be an inner vertex of c .*

Proof: We assume that c is defined by upper halfplanes (the other possibility is handled in a symmetric fashion—see a note at the end of the proof). Thus v is formed by two upper halfplanes, bounded from below by two respective lines ℓ, ℓ' ; see Figure 3.2(A).

Assume to the contrary that v is an interior point of some edge e of c , along the line ℓ . Let f and g denote the left and right endpoints of e , respectively. Assume further that the slope of ℓ' is larger than that of ℓ , as depicted in the figure (the complementary situation is handled symmetrically—see below).



First note that f cannot be a terminal vertex of c . Indeed, in this case f would have to be encountered as we trace c to the left (starting from or passing through g) through v , but then, by construction, c should have made a turn at v , contrary to assumption. Thus f is also an r -vertex, formed by two upper halfplanes (one of which is the one bounding ℓ), at depth exactly k . The argument just given also implies that v is encountered as we traverse c to the right (from or through f), because traversing it to the left would necessarily cause c to make a turn at v .

Hence, as we move from f towards v along ℓ , the depth is k just as we leave f , and is $k + 1$ just before reaching v . Hence, it must go up, from k to $k + 1$, in at least one in-between vertex w , at which we cross a line ℓ'' bounding some halfplane h'' of H . If h'' is a lower halfplane then w is an x -extreme vertex at depth k —see Figure 3.2(C) or (D), depending on the relative slope of ℓ'' with respect to ℓ . But then c would have terminated at w (or at some other interior point of e), a contradiction. Hence, h'' must be an upper halfplane, and ℓ'' must have smaller slope than that of ℓ (or else the depth would go down at w , rather than up—see Figure 3.2(A)), and w is an r -vertex. However, since the tracing of c is from f to the right, as we have argued above, we would have turned at w and never reach v at all, another contradiction that implies the claim.

If the slope of ℓ' is smaller than that of ℓ , we apply a symmetric argument, in which we consider (the only possible) right-to-left traversal of c , and replace f by the right endpoint g of e . Finally, if c is defined by lower halfplanes, we apply a symmetric argument, in which the depth has to go down between f and v (or between g and v), rather than up. ■

Lemma 3.4. *No pair of chains of \mathcal{C} overlap.*

Proof: Suppose to the contrary that some pair c_1, c_2 of distinct chains in \mathcal{C} have two respective overlapping edges e_1, e_2 , along the same line ℓ of L . Clearly, if $e_1 = e_2$ then necessarily $c_1 = c_2$ too. This is obvious if both chains are traced in the same direction through $e_1 = e_2$, but it also holds when they are traced in opposite directions, because of Lemma 3.3. Thus $e_1 \neq e_2$, in which case one of these edges must have an endpoint in the interior of the other edge, which again contradicts Lemma 3.3. ■

3.1.5 Number of chains and their crossings

Lemma 3.5. *There are at most $n + 2\tau$ chains, where τ is the number of x -vertices in C_k .*

Proof: We charge each chain to its left and right endpoints. Each such point is either an x -vertex or a b -vertex. As noted earlier, each b -vertex lies on a unique chain, and each x -vertex may be shared by at most four chains (it is incident to four edges of $\mathcal{A}(L)$, and each of them belongs to at most one chain, by Lemma 3.4). Since the number of b -vertices is at most $2n$, the claim follows. ■

Lemma 3.6. *The number of crossings between pairs of chains is $O(nk)$.*

Proof: Indeed, as stated in property (d), each crossing point of a pair of chains is a vertex of $\mathcal{A}(L)$ contained in P , and each such vertex can lie on at most two chains (by Lemma 3.4). By Observation 2.7, the number of vertices in P is $O(nk)$, and the claim follows. ■

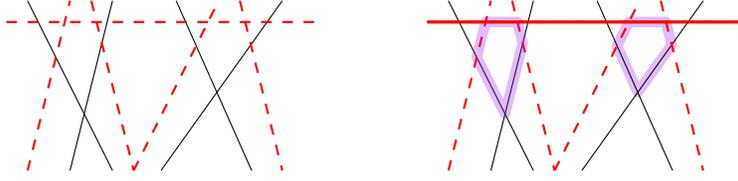


Figure 3.4: A configuration of halfplanes (with the same drawing conventions as in the previous figures), and a degenerate common tangency which supports two edges on the same line of two chains of the highlighted 5-contour.

3.2 Bounding the number of vertices of the chains

We are therefore in the following scenario. We have a collection \mathcal{C} of at most $n + 2\tau$ pairwise nonoverlapping concave chains, with $O(nk)$ crossings between them. Let V denote the set of their inner vertices (each of which is an r -vertex). We want to bound $|V|$. To do so, following Dey [Dey98], we introduce another construct, and denote by X the number of *common tangents* between pairs of chains. Formally, a common tangent between two chains c_1, c_2 is a segment uv , where u is a vertex of c_1 and v is a vertex of c_2 , u and v are inner vertices of their respective chains, and both c_1 and c_2 lie fully below the line supporting uv . See Figure 3.3 for an illustration.

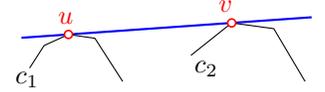


Figure 3.3: A common tangent between two chains.

We note that there may also exist *degenerate* common tangencies, where the tangent line contains an edge of each of the two chains (by the general position assumption, the tangent line cannot contain an edge of one chain and support the other chain at a single vertex); see Figure 3.4. In what follows we will only count in X the non-degenerate tangencies, and, with a few exceptions, will refer to them simply as tangencies.

We derive both a lower bound and an upper bound on X , and the comparison between these bounds will give us the desired upper bound on $|V|$.

3.2.1 Bounding the number of vertices via the number of tangencies

For a lower bound, we use the Crossing Lemma for simple graphs drawn in the plane (see, e.g., [PA95]), applied in a dual setting. That is, proceeding as in Dey [Dey98], we pass to the dual plane, in which lines are mapped to points and points to lines. We denote the point dual to a line ℓ as ℓ^* . We construct a graph G over the points dual to the lines of L , where two dual points ℓ_1^*, ℓ_2^* are connected by an edge, drawn as the straight segment $\ell_1^* \ell_2^*$, if and only if, in the primal plane, the vertex $\ell_1 \cap \ell_2$ belongs to V . The resulting graph G is *simple*, and has n vertices and $|V|$ edges.

Now, a pair of vertices $u, v \in V$ that belong to distinct respective chains c_1, c_2 , define a (non-degenerate) common tangent between c_1 and c_2 if and only if each of u, v lies in the interior of the (left-right) double wedge defined by the other vertex and its incident lines. Equivalently, an r -vertex u , incident to two lines ℓ_1, ℓ_2 , and an r -vertex v , incident to two lines ℓ_3, ℓ_4 , determine a (non-degenerate) common tangent between their respective chains if and only if the two dual segments $\ell_1^* \ell_2^*$ and $\ell_3^* \ell_4^*$ cross each other in the dual plane.¹ Namely, every non-degenerate common tangency corresponds to a crossing in (our concrete drawing of) G .

Since $|E(G)| = |V|$, the Crossing Lemma then implies that either $|V| \leq 4n$ or $X = \Omega(|V|^3/n^2)$. Together, these inequalities can be written as

$$|V| = O\left(n^{2/3}X^{1/3} + n\right), \quad (3.1)$$

so it remains to derive an upper bound on X .

¹A degenerate tangency is transformed in the dual plane to two segments with a common endpoint, which is not considered as a crossing in the Crossing Lemma.



Figure 4.1: (i) Crossing tangency; (ii) Disjoint tangency.

4 Bounding the number of common tangencies

Let c_1, c_2 be two distinct chains in \mathcal{C} , let u be an r -vertex of c_1 and v an r -vertex of c_2 , so that uv is a common tangent to c_1 and c_2 . We charge the tangent uv as follows. Assume, without loss of generality, that u lies to the left of v . Trace c_1 from u to the right, and trace c_2 from v to the left. The following complementary cases can arise:

Crossing tangency: We reach a crossing point q of c_1 and c_2 . See Figure 4.1(i).

In this case we charge uv to q , and note that the charging is unique. See Dey [Dey98] for a justification of this claim. In our setting, Dey’s argument has to be augmented by Lemma 3.4, which implies that q determines c_1 and c_2 uniquely; Dey’s argument then implies that u and v are also uniquely determined. By what has been argued, the number of such events is $O(nk)$.

Disjoint tangency: Both chains end before a crossing is reached. See Figure 4.1(ii).

In this case, the traced portions of the chains are disjoint.² Denote the right endpoint of c_1 by w_1 , and the left endpoint of c_2 by w_2 . Two cases can arise: Either w_1 lies to the left of w_2 , in which case the x -spans of the two chains are disjoint, or w_1 lies to the right of w_2 (as depicted in Figure 4.1(ii)), in which case the x -spans of the chains overlap. Since w_1 is an endpoint of a chain it is necessarily an x -vertex or a b -vertex (see Section 3.1.1). We then charge uv to the pair (c_1, c_2) , and claim that the charging is unique, in the sense that, given c_1 and c_2 , and the information that uv is a common tangent between them so that $u \in c_1$, $v \in c_2$, u lies to the left of v , and c_1 terminates to the right of u before reaching c_2 , then u and v are unique. This follows from the easy observation that there cannot exist another vertex u' of c_1 between u and w_1 at which a similar tangency with c_2 occurs. Then u is uniquely determined, and, consequently, v too.

Since crossing tangencies are already accounted for, it remains to bound the number of disjoint tangencies.

4.0.1 Counting disjoint tangencies with remote crossings

The **left headlight** of a chain $c \in \mathcal{C}$ is the ray emanating from the left endpoint of c along the line containing the leftmost edge of c , in the direction away from c . The **right headlight** of c is defined analogously for the right endpoint (and rightmost edge) of c . See Figure 4.2(A).

Lemma 4.1. *Either headlight of any chain can intersect at most $O(k)$ other chains.*

Proof: Any such intersection is a vertex of $\mathcal{A}(L)$ contained in P , and its depth is therefore at most $k^+ = 3k + 2$ (see Lemma 2.6(iii)). By Lemma 2.3, there are at most $O(k)$ such vertices along any single line, and thus along any single headlight. ■

We say that two chains c_1, c_2 have a **remote crossing** if a headlight of one of them crosses the other chain (being somewhat pedantic, we emphasize that an intersection of two headlights is not a remote crossing). See Figure 4.2(B).

Lemma 4.2. *There are at most $O(\tau k)$ disjoint tangencies between pairs of chains with remote crossings.*

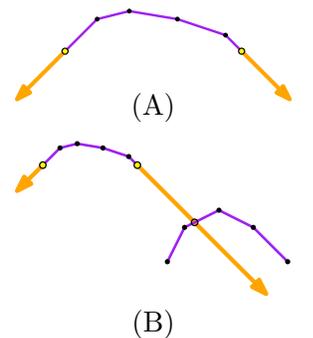


Figure 4.2: (A) Left and right headlights. (B) A remote crossing.

²Note that c_1 and c_2 might still cross one another to the left of u or to the right of v (but not both).

Proof: Recall that each such tangency is charged to the corresponding pair of chains (c_1, c_2) . For each chain c , its headlights cross $O(k)$ other chains, as follows from [Lemma 4.1](#). However, a headlight can cross another chain only if it emanates from an x -vertex (a headlight from a b -vertex lies fully outside P , and can never intersect a chain). We then charge this tangency to the relevant x -vertex, and it get charged at most $O(k)$ times. ■

Lemma 4.3. *Let c_1 and c_2 be two chains with a common disjoint tangency, so that c_1 and c_2 have overlapping x -spans. Then c_1 and c_2 have a remote crossing. Consequently, there are at most $O(\tau k)$ disjoint tangencies between such pairs of chains.*

Proof: Assume without loss of generality that the tangency point on c_1 is to the left of the tangency point on c_2 . It is easily checked that either (a) the right endpoint w_1 of c_1 lies vertically above c_2 , or (b) the left endpoint of c_2 lies vertically above c_1 . If neither of these holds, the chains either have disjoint x -spans or do not have a disjoint tangency of the sort under consideration. Assuming that (a) holds (as depicted in [Figure 4.3](#)), the right headlight of c_1 must cross c_2 , or else the disjoint tangency could not materialize (since then c_2 would have to lie entirely below the headlight). This completes the proof. ■

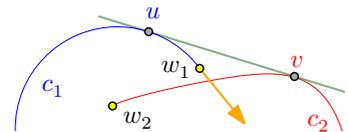


Figure 4.3: Chains with overlapping x -spans.

It therefore remains to bound the number of disjoint common tangencies between pairs of chains that have disjoint x -spans and have no remote crossing. We refer to pairs of chains satisfying both conditions as *separated*.

4.0.2 Counting disjoint tangencies between separated chains

An edge e of $\mathcal{A}(L)$ *dominates* a chain c , if the halfplane bounded by the line supporting e contains c in its interior. A chain c_1 *dominates* another chain c_2 , if c_1 has an edge that dominates c_2 .

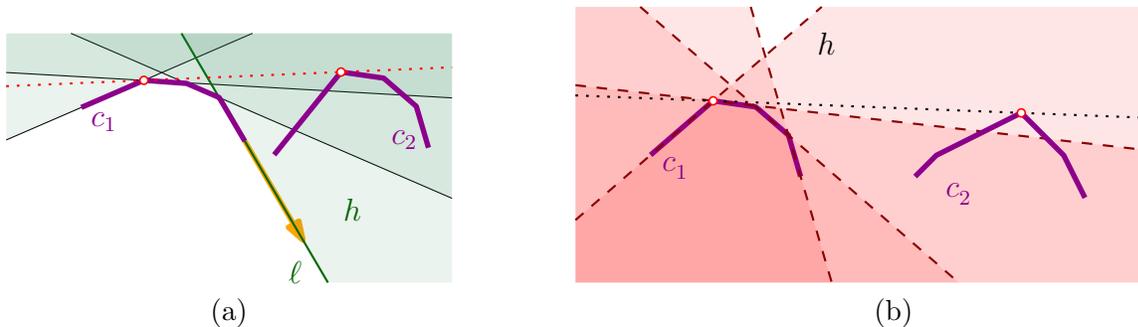


Figure 4.4: Proof of [Lemma 4.4](#): (a) The case of upper halfplanes. (b) The case of lower halfplanes.

Lemma 4.4. *Let c_1 and c_2 be two separated chains with a common disjoint tangency. Then c_1 dominates c_2 , and c_2 dominates c_1 . More precisely, either the leftmost or the rightmost edge of c_1 dominates c_2 , and either the leftmost or the rightmost edge of c_2 dominates c_1 .*

Proof: Assume without loss of generality that c_1 lies to the left of c_2 . If c_1 is defined by upper halfplanes (recall [Lemma 3.2](#)), then the halfplane h defining the rightmost edge of c_1 contains c_2 . Indeed, since c_1 and c_2 are separated, and in particular have no remote crossing, the line ℓ bounding h does not cross c_2 (only the portion of ℓ that is the right headlight of c_1 can reach the x -span of c_2 , and, by assumption, this headlight does not cross c_2). The claim then follows by noting that c_2 cannot lie fully below ℓ , or else there would be no disjoint tangency between the two chains, see [Figure 4.4\(a\)](#).

Otherwise, c_1 is defined by lower halfplanes, and then the halfplane h defining the leftmost edge of c_1 must contain c_2 . Indeed, if any vertex of c_2 lied above the line bounding h then no common tangency between c_1 and c_2 could exist, see [Figure 4.4\(b\)](#). ■

An *interior* chain has at least one of its endpoints in the interior of P (any such endpoint is an x -vertex).

Lemma 4.5. *The number of disjoint tangencies between separated pairs of interior chains is $O(\tau^2)$.*

Proof: Each such tangency is uniquely charged to the pair of its chains. Since there are only $O(\tau)$ interior chains, the claim is immediate. ■

Chains that are not interior are called *b -chains*. Both endpoints of a b -chain are b -vertices. Each of the leftmost and rightmost edges of a b -chain is called a *b -leg*.

Observation 4.6. *Each line $\ell \in L$ contains at most two b -legs.*

Indeed, a b -leg has one of its endpoints on ∂P , and is contained in P . Every line of L intersects ∂P in at most two points, and can therefore support at most two b -legs.

Lemma 4.7. *For each chain c (either interior or a b -chain), the number of b -chains that form with c a separated pair with a disjoint tangency is at most $2k$.*

Proof: Let c_1, \dots, c_t denote all the b -chains that form with c a separated pair with a disjoint tangency. By [Lemma 4.4](#), each chain c_i has a b -leg e_i that dominates c . By [Observation 4.6](#), each line of L contains at most two b -legs, which implies that c is fully contained in at least $t/2$ distinct halfplanes of H . Since the depth of any inner vertex of c is exactly k , we have $t/2 \leq k$, which implies the claim. ■

Lemma 4.8. (a) *There are at most $O(nk)$ disjoint tangencies between separated pairs of b -chains.*

(b) *There are at most $O(\tau k)$ disjoint tangencies between separated pairs of a b -chain and an interior chain.*

Proof: Both (a) and (b) follow as there are only $O(n)$ b -chains and $O(\tau)$ interior chains, and, by [Lemma 4.7](#), each of these chains can participate in at most $2k$ tangencies of the sort considered in the lemma. ■

By combining all the bounds collected so far, we get the following summary result.

Lemma 4.9. *There are at most $O(nk + \tau k + \tau^2)$ tangencies between the chains of \mathcal{C} .*

4.1 Putting it all together

Plugging the bound $X = O(nk + \tau k + \tau^2)$ on the number of tangencies between the chains of \mathcal{C} , as provided by [Lemma 4.9](#), into the bound in [Eq. \(3.1\)](#), we thus get

$$|V| = O\left(n^{2/3}X^{1/3} + n\right) = O\left(n^{2/3}(nk + \tau k + \tau^2)^{1/3} + n\right) = O\left(nk^{1/3} + n^{2/3}k^{1/3}\tau^{1/3} + n^{2/3}\tau^{2/3}\right).$$

It remains to observe that the middle term in the last bound is always dominated by the first and the last terms. Indeed, if $\tau = O(\sqrt{nk})$ then τ is also $O(n)$, and then the first term dominates the second term. If $\tau = \Omega(\sqrt{nk})$ then $\tau = \Omega(k)$ too, and then the third term dominates the second term. We thus have $|V| = O(nk^{1/3} + n^{2/3}\tau^{2/3})$.

This completes the proof of [Theorem 2.2](#). The proof of [Theorem 1.1](#) is then completed by plugging the bound $\tau = O(k^2)$, provided by [Lemma 2.8](#), into the bound just derived.

Remark 4.10. The bound in [Theorem 2.2](#) implies that, as long as $\tau = O(\sqrt{nk})$, the complexity of the k -contour is $O(nk^{1/3})$, which is asymptotically the same as the bound for the complexity of the k -level in an arrangement of lines. In particular, in view of [Lemma 2.8](#), this is always the case when $k = O(n^{1/3})$.

Note that when $k = \Theta(n)$, the upper bound is $O(n^2)$, which is tight in the worst case, due to the construction given in [Lemma 2.9](#).

5 Conclusions

The two natural open problems for further research are to extend our analysis to higher dimensions, and to derive a better lower bound on the complexity of the k -contour where $k \ll n^{1/3}$.

Concerning the former goal, our result offers a convenient tool that could be used for such an extension. Consider for example the k -depth (or k -level) problem in three dimensions. The data consists of n halfspaces bounded by n respective planes. We take each plane h , and intersect it with all the other $n - 1$ halfspaces. We get a planar subproblem involving $n - 1$ halfplanes, and bound the complexity of the k -contour using [Theorem 1.1](#) of [Theorem 2.2](#). Summing up these bounds yields a bound on the complexity of the k -level or k -contour in the original three-dimensional problem. Unfortunately, without any further enhancements, the resulting bound is too weak. For example, with a naive application of [Theorem 1.1](#), this approach only yields the bound $O(nk^2)$, as is easily checked, which is a (tight) bound for the complexity of the first k levels or depths. It is an interesting challenge to use this approach in a cleverer way to reconstruct, and perhaps even improve, the best known upper bound $O(nk^{3/2})$ for the k -level problem [[SST01](#)].

References

- [AAC98] [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.
- [AG86] [N. Alon](#) and [E. Györi](#). The number of small semispaces of a finite set of points in the plane. *J. Combin. Theory Ser. A*, 41:154–157, 1986.
- [Dey98] [T. K. Dey](#). Improved bounds for planar k -sets and related problems. *Discrete Comput. Geom.*, 19(3):373–382, 1998.
- [GP84] [J. E. Goodman](#) and [R. Pollack](#). On the number of k -subsets of a set of n points in the plane. *J. Combin. Theory Ser. A*, 36:101–104, 1984.
- [Mat95] [J. Matoušek](#). On geometric optimization with few violated constraints. *Discrete Comput. Geom.*, 14:365–384, 1995.
- [Niv08] [G. Nivasch](#). An improved, simple construction of many halving edges. In [J. E. Goodman](#), [J. Pach](#), and [R. Pollack](#), editors, *Surveys in Discrete and Computational Geometry: Twenty Years Later*, volume 453 of *Contemporary Mathematics*, pages 299–305. Amer. Math. Soc., Providence, RI, 2008.
- [PA95] [J. Pach](#) and [P. K. Agarwal](#). *Combinatorial Geometry*. John Wiley & Sons, New York, 1995.
- [SST01] [M. Sharir](#), [S. Smorodinsky](#), and [G. Tardos](#). An improved bound for k -sets in three dimensions. *Discrete Comput. Geom.*, 26(2):195–204, 2001.
- [Tót01] [G. Tóth](#). Point sets with many k -sets. *Discrete Comput. Geom.*, 26(2):187–194, 2001.