On Regular Vertices of the Union of Planar Convex Objects*

Esther Ezra† János Pach‡ Micha Sharir§

Abstract

Let $\mathcal{C}$ be a collection of n compact convex sets in the plane, such that the boundaries of any pair of sets in $\mathcal{C}$ intersect in at most s points, for some constant $s \geq 4$. We show that the maximum number of *regular* vertices (intersection points of two boundaries that intersect twice) on the boundary of the union U of $\mathcal{C}$ is[1] $O^*(n^{4/3})$, which improves earlier bounds due to Aronov *et al.* [5]. The bound is nearly tight in the worst case.

*Work by János Pach and Micha Sharir was supported by NSF Grant CCF-05-14079, and by a grant from the U.S.-Israeli Binational Science Foundation. Work by Esther Ezra and Micha Sharir was supported by grant 155/05 from the Israel Science Fund and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University. Work on this paper by the first author has also been supported by an IBM Doctoral Fellowship. A preliminary version of this paper has been presented in *Proc.* 23nd *Annu. ACM Sympos. Comput. Geom.*, 2007, pp. 220–226.

†School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel. E-mail: `estere@post.tau.ac.il` Current address: Department of Computer Science, Duke University, Durham, NC 27708-0129, USA. E-mail: `esther@cs.duke.edu`

‡Department of Computer Science, City College, CUNY, Convent Avenue at 138th Street, New York, NY 10031, USA; Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA; and Renyi Institute, P.O.Box 127, H-1364 Budapest, Hungary. E-mail: `pach@cims.nyu.edu`

§School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA. E-mail: `michas@post.tau.ac.il`

[1]In this paper, a bound of the form $O^*(f(n))$ means that the actual bound is $C_\varepsilon f(n) \cdot n^\varepsilon$, for any $\varepsilon > 0$, where C_ε is a constant that depends on ε (and generally tends to ∞ as ε decreases to 0).

1 Introduction

Let $\mathcal{C}$ be a collection of n compact convex sets in the plane, as in the abstract. Let U denote the union of $\mathcal{C}$. If the boundaries of a pair of sets in $\mathcal{C}$ intersect exactly twice, we refer to their two intersection points as *regular* intersections; all other boundary intersections are called *irregular*. Several recent papers have considered the problem of obtaining sharp bounds on the number of regular intersection points that can appear on the boundary of the union U. In the simplest instance of this problem, we assume that the boundaries of any pair of sets in $\mathcal{C}$ intersect at most twice, but make no other assumption on the shape of these sets; we then refer to $\mathcal{C}$ as a collection of *pseudo-disks*. In an early paper [11], Kedem *et al.* showed that in this case the boundary of the union contains at most $6n - 12$ intersection points, and that this bound is tight in the worst case. Pach and Sharir [14] have considered the general case, where the sets are not necessarily pseudo-disks, and have shown that, in the restricted case when $\mathcal{C}$ consists of convex sets, one has $R \leq 2I + 6n - 12$, where R (resp., I) denotes the number of regular (resp., irregular) points on ∂U, thus generalizing the result of Kedem *et al.*, in which $I = 0$.

The bound of Pach and Sharir is tight in the worst case, but since I can be large, it does not provide a good "absolute" upper bound (a bound that depends only on n) on R. In fact, I can be $\Omega(n^2)$ even when each pair of boundary sets are allowed to intersect in at most $s \geq 4$ points, where s is constant; see Figure 1(a). Moreover, there exist constructions in which both I and R are $\Theta(n^2)$ (see Figure 1(b)). However, in these constructions, some pairs of the boundaries of the sets in $\mathcal{C}$ intersect in an arbitrarily large number of points (that is, the assumption in the abstract does not hold). It is therefore interesting to seek bounds on R that are independent of I and depend only on n, in the case when any pair of sets share at most a bounded number of boundary points. This has been done by Aronov *et al.* [5]. Assuming, in addition to the above assumption, that the sets are convex, they obtained the upper bound $R = O^*(n^{3/2})$. For the more general case, where the sets in C are not necessarily convex, they show the existence of a positive constant δ, which depends only on s, so that $R = O(n^{2-\delta})$.

Our result. In this paper we consider the case where $\mathcal{C}$ satisfies the assumptions in the abstract, and derive an improved bound on R. Specifically, we show that $R = O^*(n^{4/3})$. This improves the first bound of [5]. Moreover, this bound is nearly tight in the worst case, already for $s = 4$ (the smallest interesting case), since one can easily construct n rectangles and disks that generate $\Theta(n^{4/3})$ regular vertices on the boundary of their union; see [14] and below for details.

A key step in our proof uses the fact that the sets in $\mathcal{C}$ are convex. We conjecture that the bound also holds for more general sets. Additional results in more general scenarios have been obtained in [8, Ch. 7]. These bounds, albeit being weaker than the bound presented in this paper, are much sharper than the general bound in [5]. See the end of this paper for a discussion.

2 Preliminaries and Overview

Let $\mathcal{C}$ be a collection of n compact convex sets in the plane, such that any pair of them share at most s boundary points, for some constant s.

Recall that two sets $C, C' \in \mathcal{C}$ are said to intersect *regularly* if $|\partial C \cap \partial C'| = 2$. Both of these intersection points are called *regular* vertices of the arrangement $\mathcal{A}(\mathcal{C})$ of (the boundary curves of

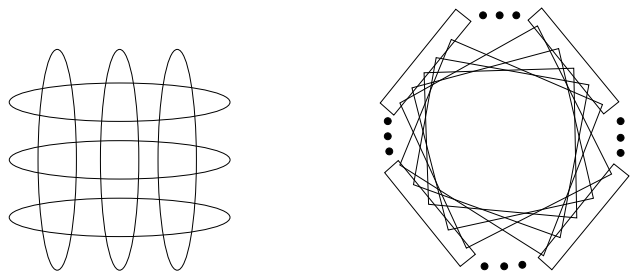

(a) (b)

Figure 1: (a) A construction with $I = \Theta(n^2)$ irregular vertices on ∂U. (b) A construction with $R = \Theta(n^2)$ regular vertices on ∂U.

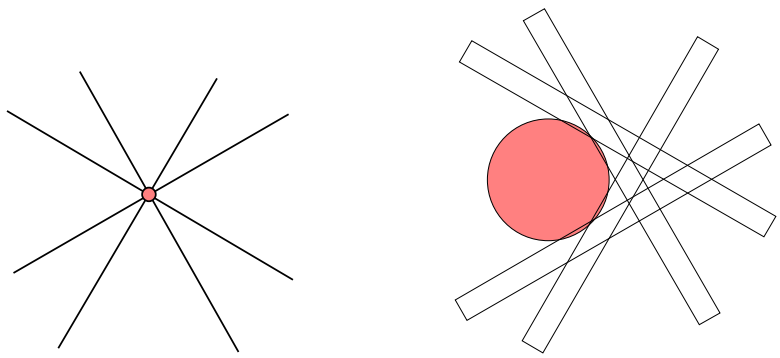

(a) (b)

Figure 2: From incidences to regular vertices: The lower bound.

the sets in) $\mathcal{C}$. All other vertices are *irregular*. If two boundary sets are externally tangent to each other at a single point, we regard this point as a pair of coincident regular vertices.

We establish an upper bound on the maximum number of regular vertices on the *boundary* of the union U of $\mathcal{C}$, which improves the earlier bound $O^*(n^{3/2})$, due to Aronov *et al.* [5].

Theorem 2.1. *Let $\mathcal{C}$ be a set of n compact convex sets as above. Then the number of regular vertices on the boundary of the union of $\mathcal{C}$ is at most $O^*(n^{4/3})$, where the constant of proportionality depends on s. This bound is nearly worst-case tight, as there are constructions that yield $\Omega\left(n^{4/3}\right)$ regular vertices that appear on the boundary of the union, already for $s = 4$.*

Lower bound. We use a construction given in [14], which we present for the sake of completeness. We start with a system of n lines and n points with $\Theta(n^{4/3})$ incidences between them (see, e.g., [15]). Then we map each line to a long and thin rectangle, and each point to a small disk, in such a way that, for each pair of a point p incident to a line ℓ, the disk into which p is mapped slightly penetrates the rectangle into which ℓ is mapped, and all the intersections between the boundaries of the disks and the rectangles are regular, and lie on the boundary of their union. See Figure 2. Clearly, $s = 4$ in this construction. Hence, we obtain a collection of $2n$ convex regions, each pair of whose boundaries intersect in at most four points, which have $\Theta(n^{4/3})$ regular vertices on the boundary of their union.

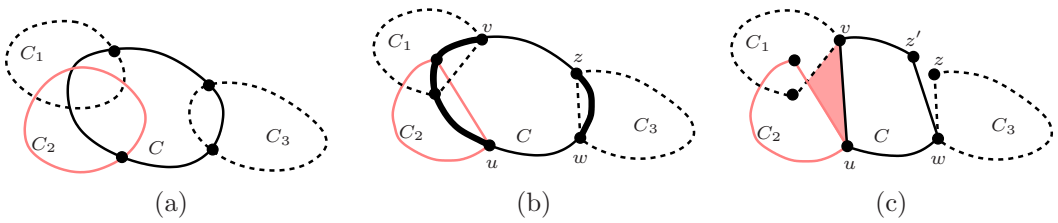

Figure 3: Demonstration of the transformation rule. (a) The boundary of C creates with the boundary of each of the three sets C_1, C_2, C_3 regular vertices that appear on the boundary of the union. (b) Each of C_1, C_2, C_3 is truncated by the chords connecting its intersections with ∂C. (c) Truncating C by similar shortcuts; wz is replaced by a nearby chord wz' to make C and C_3 touch at a single point. The shaded region is a newly created connected component of the complement of the union (a new "hole").

Upper bound: Overview of the proof. We apply the truncating transformation of Aronov *et al.* [5] to the input collection $\mathcal{C}$, as described in [5, Lemma 1] (see also the beginning of Section 3 for more details), such that at the end of this process, the transformed sets satisfy the following properties (see Figure 3 for an illustration of this process): (i) They are convex. (ii) Any two boundaries intersect at most s times. (iii) Any two sets $C, C' \in \mathcal{C}$ that intersected regularly before the transformation either become disjoint (if this intersection did not create vertices on ∂U), or *touch* each other at a single point (if at least one point of their boundary intersection appeared on ∂U). We continue to denote by $\mathcal{C}$ the collection of the truncated sets in order to simplify the notation.

For each $C \in \mathcal{C}$, the segment connecting the leftmost and the rightmost points of C is called the *spine* of C, and is denoted by σ_C. We can assume (without loss of generality, rotating the coordinate frame if necessary) that σ_C is uniquely determined. Since C is convex, we have $\sigma_C \subseteq C$. Note that, after the truncating transformation described above, any regular vertex on the boundary of the union must be formed by a pair of sets whose spines do not cross.

As a preliminary step, we dispose immediately of all regular vertices that appear as either a leftmost or a rightmost extreme point of one of their incident truncated sets. There are at most $2n$ such points. We can thus consider from now on only those regular vertices on the boundary of the union that are formed by a pair of sets whose spines are disjoint, and lie in the relative interior of the upper or lower hulls of their incident sets.

We then apply a decomposition scheme that consists of two phases. The first phase represents all pairs of sets of $\mathcal{C}$ with disjoint spines, so that one of these spines lies below the other, as the edge-disjoint union of complete bipartite graphs $A \times B$, whose overall complexity is sufficiently small, in a sense to be made precise below.

We then fix one such complete bipartite subgraph $A \times B$, in which the spines of the sets in A all lie below those of the sets of B, and analyze the number of regular vertices that it contributes to the union boundary. A crucial property of such a graph, established in Claim 3.2 below, is that each of these regular vertices must lie either on the upper envelope of the top boundaries of the sets in A, or on the lower envelope of the bottom boundaries of the sets in B. We then form, say, the upper envelope E_A^+ of the top boundaries of the sets in A, and decompose it into maximal connected arcs, each contained in the boundary of a single set, and having openly disjoint *x-spans* (i.e., projections onto the x-axis).

The fact that regular vertices are formed by touching pairs, suggests a second decomposition

phase, in which we decompose $A \times B$ into a union of complete bipartite subgraphs such that each such subgraph $A' \times B'$ is associated with some vertical strip Σ, and each spine σ of a set in B' lies, within the strip Σ, above every arc δ whose corresponding set belongs to A'. It is then fairly easy to show that the number of regular vertices of the union, induced by pairs of sets in $A' \times B'$, is only nearly linear in $|A'| + |B'|$.

The second decomposition phase is somewhat involved. It consists of decomposition steps that alternate between the primal and dual planes, where each step is based on a cutting of a certain line arrangement (in the primal) or pseudo-line arrangement (in the dual). While the dual decomposition is more "conventional" (in our problem it involves points and convex x-monotone curves), the primal one is trickier, because it involves the (lines containing the) spines and the arcs δ of the envelope. Since both families consist of curves, it is harder to control the efficiency of the process, and we need careful analysis of the way in which the arcs δ interact with the spines from the other set. Finally, we collect all these bounds, put them together, and obtain the bound asserted in the theorem.

3 The Number of Regular Vertices on the Boundary of the Union

Transforming the sets. We begin by applying to $\mathcal{C}$ the truncating transformation of Aronov *et al.* [5] to the input collection $\mathcal{C}$; see [5, Lemma 1]. For the sake of completeness, we briefly review it now. The transformation process is iterative. We arbitrarily order the regions in $\mathcal{C}$ as $(C_1, \ldots, C_n)$. Let $\mathcal{C}_i$ denote the collection after the first i steps of the transformation, with $\mathcal{C}_0 = \mathcal{C}$. The i-th step constructs $\mathcal{C}_i$ from $\mathcal{C}_{i-1}$ by considering the intersections of C_i with the other sets. Specifically, it shortcuts each set C which meets C_i regularly by the chord connecting the two intersections of ∂C and ∂C_i, and then shortcuts C similarly. See Figure 3 for an illustration of this process. At the end of the entire process, the transformed sets satisfy the following properties (see once again [5, Lemma 1] and Figure 3 for further details): (i) They are convex. (ii) Any two boundaries intersect at most s times. (iii) Any two sets $C, C' \in \mathcal{C}$ that intersected regularly before the transformation either become disjoint or *touch* each other at a single point. More precisely, if C, C' intersected regularly with at least one point of intersection of their boundaries on ∂U, the transformed sets are openly disjoint and touch each other at a point on the new union boundary. If they intersected regularly without creating vertices on ∂U, they become disjoint. To simplify the notation, from now on we let $\mathcal{C}$ denote the set of the transformed regions, and we refer to the tangency points (that correspond to regular vertices of the boundary of the union before the transformation) as regular vertices. We assume, without loss of generality, that the transformed sets in $\mathcal{C}$ are in general position, in the sense that no point is incident to more than two boundaries, and when two boundaries meet at a point, they cross each other there, except for the points that represent previous regular vertices, at which the respective sets now touch each other.

Let C and C' be two members of $\mathcal{C}$ that touch each other at a point that lies on ∂U. As already noted, we may assume that their spines σ_C and $\sigma_{C'}$ are disjoint, and one of them, say, σ_C, *lies below* the other, which means that (i) their x-spans have nonempty intersection J; (ii) σ_C lies below $\sigma_{C'}$ at each $x \in J$.

The first bi-clique decomposition. We collect all pairs of spines so that one of them lies below the other, as the edge-disjoint union of complete bipartite graphs (*bi-cliques*), so that the overall size of their vertex sets is $O^*(n^{4/3})$. More precisely, the following stronger property holds.

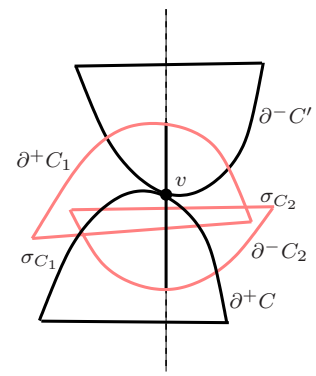

Figure 4: If the vertex v does not lie on either E_A^+ or E_B^-, then it is "hidden" from E_A^+ by $\partial^+ C_1$, and from E_B^- by $\partial^- C_2$, for some $C_1 \in A$, $C_2 \in B$. However, then v is contained in (the interior of) $C_1 \cup C_2$, contrary to the assumption that v is a boundary vertex of the union.

Lemma 3.1. *Given a collection $\mathcal{C}$ as above, let G be the graph whose vertices are the members of $\mathcal{C}$, and whose edges connect pairs of members (C, C') such that σ_C lies below $\sigma_{C'}$. Then there exists a decomposition $G = \bigcup_i A_i \times B_i$ into pairwise edge-disjoint bi-cliques such that*

$$\sum_i \Big(|A_i|^{2/3} |B_i|^{2/3} + |A_i| + |B_i| \Big) = O^*(n^{4/3}). \tag{1}$$

This is a fairly standard consequence of the techniques for "batched" range searching, which are based on multi-level data structures, as reviewed in [3, 13]. Results of a similar nature, yielding decompositions of other relations (although not the specific one in the lemma) into an edge-disjoint union of complete bipartite graphs, have appeared in the literature; see [10] for the case of points inside unit disks; and [2] for the case of intersecting segments. In fact, a decomposition of this kind has also been used in the previous paper [5] that has studied this problem, where in each bi-clique $A \times B$ the spine of each $C \in B$ lies fully above each of the sets in A, and the spine of each $C' \in A$ lies fully below each of the sets in B. We present a detailed proof of the lemma in the appendix in order to (a) aid readers unfamiliar with the range searching machinery, and (b) obtain an explicit derivation for the specific relation that we need to handle.

Handling a single bi-clique. Fix one of the resulting graphs $A \times B := A_i \times B_i$. All the spines of the sets in A lie below all the spines of the sets in B. Put $n_A = |A|$ and $n_B = |B|$.

Recall that our goal is to bound the overall number of tangencies between the boundaries of the sets in A and B. Let v be a regular vertex of the union lying on the top boundary $\partial^+ C$ and on the bottom boundary $\partial^- C'$, for two sets $C \in A$, $C' \in B$; this is the only possible situation, since v is an external tangency and is not the leftmost or rightmost point of either of the two incident sets. We have:

Claim 3.2. *The vertex v lies either on the upper envelope E_A^+ of the top boundaries of the sets in A, or on the lower envelope E_B^- of the bottom boundaries of the sets in B, or both.*

Proof: Suppose for the sake of contradiction that this were not the case. Then v would lie below some top boundary $\partial^+ C_1$, for $C_1 \in A$, and above some bottom boundary $\partial^- C_2$, for $C_2 \in B$; see

Figure 4. By construction, σ_{C_1} lies below σ_{C_2} at the x-coordinate x_v of v, which implies that the entire vertical segment connecting $\partial^+ C_1$ and $\partial^- C_2$ at x_v is fully contained in $C_1 \cup C_2$, so v cannot lie on the boundary of the union, a contradiction that establishes the claim. $\square$

Without loss of generality, we consider only the case where v lies on the upper envelope E_A^+ of the top boundaries of the sets in A. Since any pair of these boundaries intersect in at most s points, the number $m = m_A$ of connected portions of top boundaries that constitute E_A^+ satisfies $m \le \lambda_{s+2}(n_A)$, where $\lambda_t(q)$ is the maximal length of Davenport-Schinzel sequences of order t on q symbols (see [16]). Enumerate these arcs from left to right as $\delta_1, \ldots, \delta_m$, and let A^* denote the set of these arcs.

Let H_0 denote the subgraph of $A^* \times B$ consisting of all the pairs (δ, C), such that C forms with the set of A containing δ a regular vertex on ∂U (where the two sets touch each other), so that (a) the touching point lies on δ and on $\partial^- C$, and (b) δ lies fully below σ_C (i.e., the x-span of σ_C contains that of δ). If (b) does not hold, then an endpoint of σ_C lies above δ, and there can be at most two such arcs δ, for any fixed C, so the number of excluded pairs is at most $2n_B$. Hence, the number of regular "bichromatic" vertices formed by $A \cup B$, lying on E_A^+, and not counted in H_0, is only $O(n_B)$.

The second bi-clique decomposition. Our next step is to construct a collection of complete bipartite graphs $\{A_i^* \times B_i\}_i$, such that, for each i, $A_i^* \subset A^*$, $B_i \subset B$, and these graphs are edge-disjoint and cover H_0. In addition: (i) The sum $\sum_i (|A_i^*| + |B_i|)$ will be small, in a sense to be made precise below. (ii) For each i, there is an x-interval I_i such that, for each $\delta \in A_i^*$ and $C \in B_i$, the line ℓ_C containing the spine σ_C of C passes fully above δ over I_i. (We require the x-span of δ to contain I_i, but allow σ_C to end within I_i.) (iii) For each pair $(\delta, C) \in H_0$, there exists i such that $\delta \in \mathcal{A}_i^*$, $C \in B_i$, and the x-coordinate of the touching point $\delta \cap \partial^- C$ lies in I_i.

Suppose we have such a collection at hand. Fix one of the graphs $A_i^* \times B_i$. We now claim:

Claim 3.3. *For any $\delta \in A_i^*$, $C \in B_i$, such that $(\delta, C) \in H_0$, the relevant touching vertex v of $\delta \cap \partial^- C$ lies on the lower envelope $E_{B_i}^-$ of the bottom boundaries of the sets in B_i.*

Proof: The proof follows using the same arguments as in Claim 3.2 (see also [5]). That is, suppose to the contrary that v lies above the bottom boundary $\partial^- C'$ of another set $C' \in B_i$. By assumption, $\sigma_{C'}$ lies above v (because $v \in \delta$) and thus v lies in the interior of C', contradicting the assumption that v is a vertex of the union. See Figure 5(a). Note that it is crucial that the x-coordinate of v lies in the x-interval I_i as above, because the property that δ lies below $\sigma_{C'}$ is only enforced within I_i; see Figure 5(b). $\square$

In other words, each vertex v of this kind is an intersection point of $E_{B_i}^-$ and the concatenation of the arcs in A_i^*. Hence, by merging, in the x-order, the breakpoints of $E_{B_i}^-$ and the endpoints of the arcs in A_i^*, it easily follows that the number of such vertices is $O(\lambda_{s+2}(|B_i|) + |A_i^*|)$. Summing this bound over all subgraphs $A_i^* \times B_i$ yields an overall bound for the number of pairs $(\delta, C) \in H_0$ (to which we add the linear number of pairs that are not counted in H_0, as above). See below for the precise bound.

To obtain the desired cover of H_0, we proceed as follows. Let L denote the set of the lines supporting the spines of the sets in B. Fix a sufficiently large constant parameter r, and construct a $(1/r)$-cutting Ξ of the arrangement $\mathcal{A}(L)$, as in [6]. It consists of $O(r^2)$ vertical trapezoids, each crossed by at most n_B/r lines of L (and thus by at most n_B/r spines of the sets in B). We assume

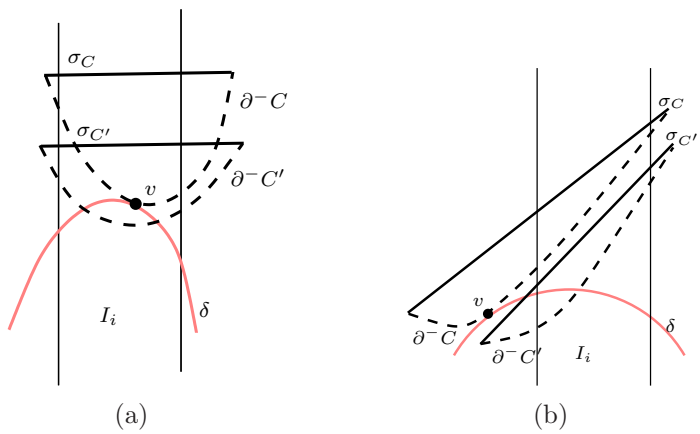

Figure 5: (a) If the vertex v does not lie on $E^-_{B_i}$, then it is hidden from $E^-_{B_i}$ by $\partial^- C'$, for some $C' \in B_i$. But then v is contained in (the interior of) C', contrary to the assumption that v appears on the boundary of the union. (b) The tangency point v, formed by the lower boundary of C and δ, and lying outside the interval I_i, can be "hidden" from E^-_B by the lower boundary of another $C' \in B$.

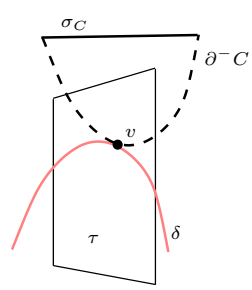

Figure 6: A touching occurs within a cell τ of the cutting.

in what follows that $m_A, n_B \geq r^2$; otherwise, we reach the bottom of the recursion, and different handling is applied—see below.

Consider a pair $(\delta, C) \in H_0$, where the touching between δ and $\partial^- C$ occurs in some cell τ of Ξ. In this case δ crosses τ (or has an endpoint inside τ), and σ_C either intersects τ (ending inside it or fully crossing it), or lies above τ (i.e., within the common x-span of C and τ, σ_C lies fully above τ). See Figure 6.

For technical reasons, we classify the arcs δ that cross τ as being either *short*, if either δ has an endpoint inside τ, or δ does not intersect the top edge of τ, or *tall*, if δ intersects the top edge and has no endpoint in τ. Let A^s_τ be the set of short arcs in τ, and A^t_τ the set of tall arcs in τ.

The next lemma shows that the overall number of short arcs, over all cells τ, is small.

Lemma 3.4. $\sum_\tau |A^s_\tau| = O(r^2 \log r + |A^*| \log r) = O(|A^*| \log r)$.

Proof: Note first that there are at most $2|A^*|$ pairs (δ, τ), such that δ ends inside τ. We may therefore ignore these short arcs. Construct a segment tree $\mathcal{T}$ on the x-projections of the cells of Ξ. Consider a node v of the tree, let Ξ_v denote the set of cells stored at v, and let I_v denote the x-span of v (the x-spans of the stored cells contain I_v but do not contain $I_{\text{parent}(v)}$). The cells in

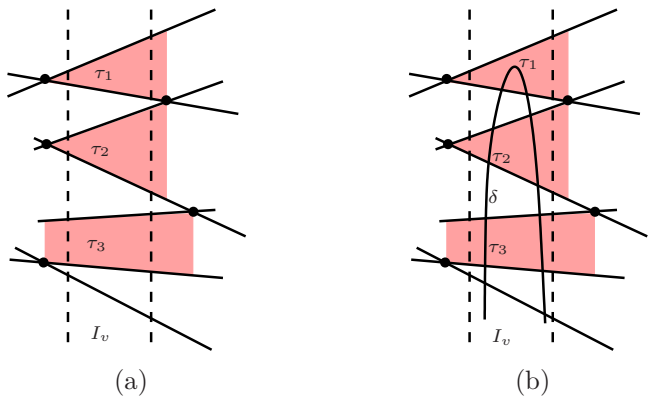

Figure 7: The cells τ_1, τ_2, τ_3 cross the vertical strip built on the x-span I_v of v from left to right. (a) Any vertical line in I_v crosses all these cells in the same order. (b) The arc δ crosses τ_1, τ_2, τ_3, where τ_1 is the unique cell whose top boundary edge is not crossed by δ; δ is short there and tall in τ_2, τ_3.

Ξ_v are linearly ordered in the y-direction, in the sense that for each $x_0 \in I_v$ the vertical line $x = x_0$ crosses all of them in a fixed order; see Figure 7(a).

In each cell τ, there are at most two (either tall or short) arcs, whose x-spans *overlap*, but are not contained in, I_v (the first intersects the vertical line through the left endpoint of I_v, and the second intersects the vertical line through its right endpoint). Since each cell τ is stored in $O(\log r)$ nodes of $\mathcal{T}$, there are $O(r^2 \log r)$ such arcs in total.

We thus continue the analysis for those (short) arcs δ of A^*, whose x-span is *contained* in I_v. There is at most one cell $\tau \in \Xi_v$ such that $\delta \in A^s_\tau$, namely the highest that δ meets; see Figure 7(b). The number of nodes v at which δ has this property is $O(\log r)$, because I_v contains the x-coordinate of an endpoint (actually, both endpoints) of δ. Hence, the contribution of arcs δ as above to $\sum_\tau |A^s_\tau|$ is $O(|A^*| \log r)$. Combining this with the previous bound, and recalling that we assume $|A^*| = m_A \geq r^2$, complete the proof of the lemma. $\square$

Remarks: (1) The fact that the arcs δ have pairwise openly disjoint x-projections is crucial for the bound that we obtain in Lemma 3.4. The decomposition of (a cover of) H_0 that we construct is a variant of the decomposition obtained in [5]; however, the analysis in [5] does not exploit the special structure of the arcs δ, and results in a suboptimal bound.

(2) An individual arc δ may cross $\Omega(r)$ cells τ, each of whose top boundary is disjoint from δ. However, Lemma 3.4 shows that the overall number of these crossings, summed over all arcs δ, is relatively small.

Each cell τ for which $|A^s_\tau| > \frac{|A^*| \log r}{r^2}$ is next split, by vertical lines, into subcells, such that each subcell τ' satisfies $|A^s_{\tau'}| \leq \frac{|A^*| \log r}{r^2}$. The number of cells is still $O(r^2)$, which follows from Lemma 3.4.

Fix a (new) cell τ, and form the complete bipartite graph $A^s_\tau \times \mathcal{C}^*_\tau$, where $\mathcal{C}^*_\tau$ consists of all sets $C \in B$ such that ℓ_C passes above τ. We associate the interval I_τ (the x-span of τ) with this graph. Since r is a constant, we have

$$\sum_\tau (|A^s_\tau| + |\mathcal{C}^*_\tau|) = O(m_A + n_B)$$

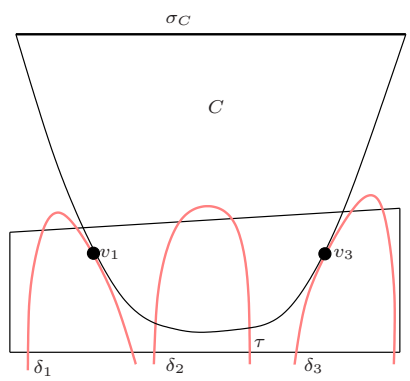

Figure 8: $\partial^- C$ touches δ_1 and δ_3 inside τ at v_1 and v_3, respectively, and cannot touch the intermediate tall arc δ_2.

(where the constant of proportionality depends on r), and the overall number of boundary touchings that they involve, over all cells τ, is $O(m_A + \lambda_{s+2}(n_B))$.

We next show:

Lemma 3.5. *The overall number of boundary touchings on the boundary of the union, occurring within a cell τ and involving a tall arc in τ, summed over all cells τ, is only linear in n_B.*

Proof: Let τ be a cell of Ξ, and let ℓ_C be the line containing the spine σ_C of a set C, so that ℓ_C intersects τ or passes fully above τ. We claim that there are at most two tall arcs in τ that touch $\partial^- C$ within τ, at a point that lies on ∂U. Indeed, suppose to the contrary that there are three such arcs $\delta_1, \delta_2, \delta_3$, which appear on E_A^+ in that order (from left to right). Consider the two respective boundary touchings that $\partial^- C$ forms with δ_1, δ_3, at two respective points v_1, v_3 inside τ. Then, due to the convexity of C, its portion γ between v_1 and v_3 lies below the top edge of τ, and δ_2 lies fully below (and touches) that portion, so it cannot be tall in τ (note that the x-span of δ is contained in that of γ), a contradiction that establishes the claim; see Figure 8. Thus the overall number of regular vertices of the above kind is $O(r^2 n_B) = O(n_B)$, as asserted. $\square$

Remark: Lemma 3.5 implies that we do not need to provide a compact representation for pairs of the above kind, and can simply add them as singletons to the output cover of H_0.

We thus conclude:

Corollary 3.6. *The overall number of boundary touchings involving both short and tall arcs, in all subcases considered so far, is $O(m_A + \lambda_{s+2}(n_B))$.*

We continue the construction recursively, within each cell τ, with A_τ^s and the subset $\mathcal{C}_\tau$ of those $C \in B$ whose line ℓ_C crosses τ. We have $|A_\tau^s| \le \frac{|A^*| \log r}{r^2} = \frac{m_A \log r}{r^2}$, and $|\mathcal{C}_\tau| \le \frac{n_B}{r}$. However, the next stage of the recursion is performed in the *dual* plane, and proceeds as follows. For each resulting cell τ, map A_τ^s and $\mathcal{C}_\tau$ to the dual plane. For each $C \in \mathcal{C}_\tau$, we map ℓ_C to a dual point ℓ_C^*, and each arc δ in A_τ^s is mapped to a convex x-monotone curve δ^*, which is the locus of all points dual to lines that are tangent to δ (possibly at one of its endpoints) and pass above δ (see [5] and [7] for further details). Thus a line ℓ_C lies above an arc δ if and only if the dual point ℓ_C^* lies above δ^*. Each pair of dual arcs δ_1^*, δ_2^* intersect each other exactly once, since any such intersection point is the dual of a common tangent to δ_1, δ_2 that passes above both of them, and since δ_1, δ_2 are two

convex curves that have openly disjoint x-spans, there is exactly one such common tangent. We now construct (for each cell τ obtained at the preceding step) a $(1/r)$-cutting of the arrangement of the dual arcs δ^*, obtaining $O(r^2)$ subcells, each of which is crossed by at most $\frac{|A^s_\tau|}{r} \le \frac{m_A \log r}{r^3}$ dual arcs δ^*, and contains at most $\frac{|\mathcal{C}_\tau|}{r^2} \le \frac{n_B}{r^3}$ dual points ℓ^*_C; see [4] for details.

As above, we construct, for each subcell τ' of this cutting, a complete bipartite graph that connects the dual points in τ' to the dual arcs that pass fully below that subcell. Again, since r is a constant, the sum of the sizes of the vertex sets of these graphs is $O(m_A + n_B)$. We are thus left with $O(r^4)$ subproblems, each involving at most $\frac{m_A \log r}{r^3}$ arcs of A^*, and at most $\frac{n_B}{r^3}$ sets in B. We now process each subproblem recursively, going back to the primal plane, and keep alternating in this manner, until we reach subproblems in which either $m_A^2 < n_B$, or $n_B^2 < m_A$. In the former (resp., latter) case, we continue the recursive construction *only* in the dual (resp., primal) plane, and stop as soon as one of m_A, n_B becomes smaller than r^2, in which case we output the complete bipartite graph $A^s_\tau \times \mathcal{C}_\tau$ involving the input sets to the subproblem.

Note that at the bottom of the recurrence, the boundary touchings do not necessarily lie on the lower envelope of the boundaries of the sets in $\mathcal{C}_\tau$, so we use the naive bound $|A^s_\tau| \cdot |\mathcal{C}_\tau|$ on their number, which however is $O(|A^s_\tau| + |\mathcal{C}_\tau|)$, where the constant of proportionality depends on r. We add to the output cover of H_0 only those pairs that result in boundary touchings, as singletons.

The preceding arguments imply that the union of all the bi-cliques constructed by this procedure, including the interactions with tall arcs and other "leftover" pairs detected by the decomposition, covers H_0. Indeed, for each such pair $(\delta, C) \in H_0$, the line ℓ_C containing the spine σ_C of C lies fully above δ. Our procedure detects all such pairs (δ, C) either (i) at the bottom of the recurrence, in which case all these pairs are reported in a brute force manner, or (ii) at a recursive step, performed in the primal plane and involving a cell τ in which the boundary touching appears, such that ℓ_C lies above τ and δ is short in τ, (iii) at a recursive step, performed in the dual plane and involving a cell τ', such that ℓ^*_C lies inside τ' and δ^* passes fully below it, or (iv) at some step where δ is tall at some primal cell τ.

Let $R(m_A, n_B)$ denote the maximum number of boundary touchings on the boundary of the union, that arise at a recursive step involving m_A arcs δ and n_B sets C, as above, and which are formed between one of the arcs δ and the bottom boundary of one of the sets C. As argued above, the number of such bichromatic touchings, that arise for any of the complete bipartite graphs generated at this stage, is nearly-linear in the sizes of the vertex sets of that graph. Hence R satisfies the following recurrence (where in the first three cases, $\min\{m_A, n_B\} \ge r^2$):

$$R(m_A, n_B) \le \begin{cases} O\left(m_A + \lambda_{s+2}(n_B)\right) + O(r^4) R\left(\frac{m_A \log r}{r^3}, \frac{n_B}{r^3}\right), & \text{if } m_A^2 \ge n_B \ge \sqrt{m_A}, \\ O\left(m_A + \lambda_{s+2}(n_B)\right) + O(r^2) R\left(\frac{m_A}{r}, \frac{n_B}{r^2}\right), & \text{if } n_B > m_A^2, \\ O\left(m_A + \lambda_{s+2}(n_B)\right) + O(r^2) R\left(\frac{m_A \log r}{r^2}, \frac{n_B}{r}\right), & \text{if } m_A > n_B^2, \\ O\left(m_A + n_B\right), & \text{if } \min\{m_A, n_B\} < r^2. \end{cases}$$

It is then easy to see, using induction on m_A and n_B, that the solution of this recurrence is

$$R(m_A, n_B) = O^*(m_A^{2/3} n_B^{2/3} + m_A + n_B) = O^*(n_A^{2/3} n_B^{2/3} + n_A + n_B). \tag{2}$$

Summing these bounds over all bi-cliques $A \times B$ of the first decomposition phase, and using the bound in (1), the upper bound of Theorem 2.1 follows.

This completes the proof of Theorem 2.1. $\square$

Open problems. A major open problem is to extend the bound to the case where the sets in $\mathcal{C}$ are not convex. A natural case to study is where the sets in $\mathcal{C}$ are x-monotone (i.e., each of the lower and upper portions of ∂C is an x-monotone curve), for each set $C \in \mathcal{C}$, the spine σ_C is contained in C, and each pair of boundaries intersect in at most some constant number, s, of points. In an expanded version of the paper (see [8, Ch. 7]), we obtained the upper bound $O^*(n^{(3s+1)/(2s+1)})$ for this case, which, albeit weaker than the bound derived above, is still considerably sharper than the general bound of [5]. We conjecture that the new bound $O^*(n^{4/3})$ also holds in this extended scenario, provided that each set in $\mathcal{C}$ has *constant description complexity* (that is, each set is defined as a Boolean combination of a constant number of polynomial equalities and inequalities of constant maximum degree).

Acknowledgments. The authors wish to thank Boris Aronov, Alon Efrat and an anonymous referee for their useful comments.

References

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

[2] P. K. Agarwal, M. de Berg, D. Halperin and M. Sharir, Efficient generation of k-directional assembly sequences, *Proc. 7th ACM-SIAM Sympos. Discrete Algo.* (1996), 122–131.

[3] P. K. Agarwal and J. Erickson, Geometric range searching and its relatives, in: *Advances in Discrete and Computational Geometry* (B. Chazelle, J. E. Goodman and R. Pollack, eds.), *Contemporary Mathematics*, vol. 223, Amer. Math. Soc. Press, Providence, RI, 1999, 1–56.

[4] P. K. Agarwal and M. Sharir, Pseudo-line arrangements: Duality, algorithms, and applications, *SIAM J. Comput.* 34(3) (2005), 526–552.

[5] B. Aronov, A. Efrat, D. Halperin and M. Sharir, On the number of regular vertices of the union of Jordan regions, *Discrete Comput. Geom.* 25 (2001), 203–220.

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

[7] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg, 1987.

[8] E. Ezra, *Geometric Arrangements: Substructures and Algorithms*, Ph.D. thesis, Computer Science Department, Tel-Aviv University, Tel-Aviv, August 2007. `http://www.math.tau.ac.il/~estere`

[9] E. Ezra and M. Sharir, Counting and representing intersections among triangles in three dimensions, *Comput. Geom. Theory Appl.* 32(3) (2005), 196–215.

[10] M. J. Katz and M. Sharir, An expander-based approach to geometric optimization, *SIAM J. Comput.* 26 (1997), 1384–1408.

[11] K. Kedem, R. Livne, J. Pach and M. Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, *Discrete Comput. Geom.* 1 (1986), 59–71.

[12] J. Matoušek, Cutting hyperplane arrangements, *Discrete Comput. Geom.* 6 (1991), 385–406.

[13] J. Matoušek, Geometric range searching, *ACM Computing Surveys* 26(4) (1994), 421–461.

[14] J. Pach and M. Sharir, On the boundary of the union of planar convex sets, *Discrete Comput. Geom.* 21 (1999), 321–328.

[15] J. Pach and M. Sharir, Geometric incidences, in *Towards a Theory of Geometric Graphs* (J. Pach, ed.), *Contemporary Mathematics, Amer. Math. Soc.* 342 (2004), 185–223.

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

Appendix - The First Bi-Clique Decomposition

Proof of Lemma 3.1: Let $\sigma = pq$, $\sigma' = p'q'$ be a pair of spines such that σ lies below σ'. This relationship can be expressed as the disjunction of a constant number of conjunctions of above/below relationships, over the possible x-orders of p, q, p' and q', where each atomic relationship asserts that an endpoint of one spine lies above or below the line containing the other spine. For example, if the x-order of the endpoints is p, p', q, q', then we require that p' lie above the line ℓ containing σ and that q lie below the line ℓ' containing σ'; see Figure 9(a)–(c). For simplicity of exposition, we describe the construction only for the subgraph of G that consists of pairs of sets with the above specific order of the endpoints of their spines; all other subcases are handled in a fully symmetric manner.

We apply a multi-level decomposition scheme (see [3]), where each level produces a decomposition into bi-cliques that satisfy some of the constraints, and each of them is passed to the next level to enforce additional constraints. At the two top levels, we produce a collection of pairwise edge-disjoint bi-cliques such that, for each of these graphs $A_1 \times B_1$, for each pair of spines $\sigma = pq \in A_1$ and $\sigma' = p'q' \in B_1$, the x-order of the endpoints is p, p', q, q', and such that the union of these graphs gives all such pairs of spines. This is easily done using a 2-dimensional range tree construction [1, 9]. The sum of the vertex sets of the resulting subgraphs is $O(n \log^2 n)$. Moreover, (a slightly sharper variant of) (1) is easily seen to hold for the decomposition thus far.

The next level enforces, for each resulting subgraph of $A_1 \times B_1$, the condition that p' lie above the line ℓ containing σ, for $\sigma \in A_1$ and p' the left endpoint of a spine $\sigma' \in B_1$. Put $m_1 = |A_1|$ and $n_1 = |B_1|$. To accomplish the task at hand, we choose a sufficiently large constant parameter r, and construct a $(1/r)$-cutting [12] of the arrangement of the lines that contain the spines of A_1. We obtain $O(r^2)$ cells, each of which is crossed by at most m_1/r lines and contains at most n_1/r^2 left endpoints of spines of B_1. (The latter property can be enforced by further splitting some cells of the cutting; also, assuming general position, we can construct the cutting so that no endpoint of any spine lies on the boundary of any of the cutting cells.) For each cell Δ, we form the bi-clique $A_2'(\Delta) \times B_2(\Delta)$, where $B_2(\Delta)$ consists of all spines whose left endpoints are in Δ, and where $A_2'(\Delta)$ consists of all spines whose supporting lines pass completely below Δ. These graphs are passed to

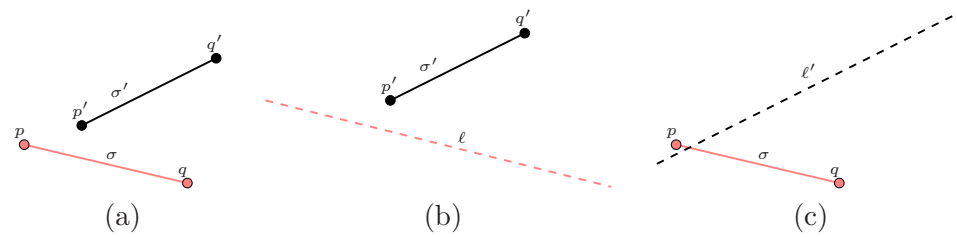

Figure 9: (a) The two spines σ, σ' are disjoint and σ lies below σ'. (b) The left endpoint p' of σ' lies above the line ℓ containing σ, and (c) the right endpoint q of σ lies below the line ℓ' containing σ'.

the next level of the structure. We then consider, for each cell Δ, the set $A_2(\Delta)$ of spines of A_1 that cross Δ, and the set $B_2(\Delta)$ as defined above. We pass to the dual plane, where the lines of the spines in $A_2(\Delta)$ are mapped to points and the left endpoints of spines in $B_2(\Delta)$ are mapped to lines. We construct a $(1/r)$-cutting of the arrangement of these dual lines, obtaining $O(r^2)$ cells, each of which is crossed by at most $|B_2(\Delta)|/r \le n_1/r^3$ lines and contains at most $|A_2(\Delta)|/r^2 \le m_1/r^3$ points. As above, we construct, for each cell of the cutting, a bi-clique from the dual points in the cell and the lines that pass fully above the cell, and pass all these graphs to the next level. We are left with $O(r^4)$ subproblems, each involving at most m_1/r^3 spines of A_1 and at most n_1/r^3 spines of B_1, which we process recursively. We continue to process each subproblem as above, going back to the primal plane, and keep alternating in this manner, until we reach subproblems in which either $m_1^2 < n_1$, or $n_1^2 < m_1$. In the former (resp., latter) case, we continue the recursive construction *only* in the primal (resp., dual) plane, and stop as soon as one of m_1, n_1 becomes smaller than r, in which case we produce a collection of singleton bi-cliques.

Suppose first that $\sqrt{m_1} \le n_1 \le m_1^2$. We show below that, for any fixed initial subgraph $A_1 \times B_1$, the resulting bi-clique decomposition $\{A_2'(\Delta) \times B_2(\Delta)\}_\Delta$, over all cells Δ of all the cuttings, satisfies

$$\sum_\Delta \left(|A_2'(\Delta)|^{2/3} |B_2(\Delta)|^{2/3} + |A_2'(\Delta)| + |B_2(\Delta)| \right) = O^* \left(|A_1|^{2/3} |B_1|^{2/3} + |A_1| + |B_1| \right), \qquad (3)$$

and the same holds for the corresponding decompositions in the dual spaces.

Indeed, let us consider only the primal decompositions, since the dual ones are handled in exactly the same manner. Since r is taken to be a constant, the sum in (3), over the graphs produced at the top level of the recursion, is at most $C(r)\left(|A_1|^{2/3}|B_1|^{2/3} + |A_1| + |B_1|\right)$, where $C(r)$ is a constant that depends on r. In the next level, we have at most $C'r^4$ subproblems, for some absolute constant $C' > 0$, each involving at most $|A_1|/r^3$ spines of A_1 and at most $|B_1|/r^3$ spines of B_1. The overall contribution to the sum in (3) by the bi-cliques produced at this level is at most

$$C'r^4 \cdot C(r) \left(\left(\frac{|A_1|}{r^3} \right)^{2/3} \left(\frac{|B_1|}{r^3} \right)^{2/3} + \frac{|A_1|}{r^3} + \frac{|B_1|}{r^3} \right) =$$

$$C'C(r) \left(|A_1|^{2/3} |B_1|^{2/3} + |A_1| r + |B_1| r \right).$$

Continuing in this manner, the contribution to the sum in (3) at the j-th level of the recursion is at most

$$(C')^j r^{4j} \cdot C(r) \left(\frac{|A_1|^{2/3}}{r^{2j}} \frac{|B_1|^{2/3}}{r^{2j}} + \frac{|A_1|}{r^{3j}} + \frac{|B_1|}{r^{3j}} \right) =$$

$$(C')^j C(r)\left(|A_1|^{2/3}|B_1|^{2/3} + |A_1|r^j + |B_1|r^j\right),$$

where, at the last level, $r^j = \min\left\{\frac{|A_1|^{2/3}}{|B_1|^{1/3}}, \frac{|B_1|^{2/3}}{|A_1|^{1/3}}\right\}$. Summing over the logarithmically many levels of the recursion, we obtain the overall bound $O^*(|A_1|^{2/3}|B_1|^{2/3})$, assuming r is sufficiently large.

It remains to consider the cases $|A_1|^2 < |B_1|$, $|B_1|^2 < |A_1|$. It suffices to consider only the first case. Here, after j levels of recursion (only in the primal plane), the contribution to (3) is at most

$$(C'')^j r^{2j} \cdot C(r)\left(\left(\frac{|A_1|}{r^j}\right)^{2/3}\left(\frac{|B_1|}{r^{2j}}\right)^{2/3} + \frac{|A_1|}{r^j} + \frac{|B_1|}{r^{2j}}\right) =$$

$$(C'')^j C(r)\left(|A_1|^{2/3}|B_1|^{2/3} + |A_1|r^j + |B_1|\right),$$

where C'' is another absolute constant, and where the last j satisfies $r^j = O(|A_1|)$. Substituting this value, summing over all j, and using the inequality $|A_1|^2 \le |B_1|$, we get the overall bound $O^*(|B_1|)$. Similarly, when $|B_1|^2 < |A_1|$, we get the overall bound $O^*(|A_1|)$.

Note that, for the case $\sqrt{|A_1|} \le |B_1| \le |A_1|^2$, when the recursion bottoms out, we have sets A', B' that satisfy $|A'_1|^2 \le |B'_1|$ or $|B'_1|^2 \le |A'_1|$, so the same analysis adds to (3) the terms $O^*(|A_1| + |B_1|)$, which thus completes the proof of the claim.

The final level of the structure enforces, for each resulting subgraph of $A_2 \times B_2$, the condition that q lie below the line ℓ' containing σ', for $\sigma' \in B_2$ and q the right endpoint of a spine $\sigma \in A_2$. This is done in a fully analogous manner to the preceding step. It is easily checked that, in complete analogy to the preceding analysis, the bi-clique decomposition $\{A_{2,\alpha} \times B_{2,\alpha}\}_\alpha$ that results from the fixed bi-clique $A_2 \times B_2$, over all cells of all the cuttings, satisfies

$$\sum_\alpha \left(|A_{2,\alpha}|^{2/3}|B_{2,\alpha}|^{2/3} + |A_{2,\alpha}| + |B_{2,\alpha}|\right) =$$

$$O^*\left(|A_2|^{2/3}|B_2|^{2/3} + |A_2| + |B_2|\right).$$

Combining this with (3), summing over the entire collection of these last-stage decompositions, and using the fact that (1) holds for the initial-level decomposition, we conclude that (1) holds for the overall final decomposition, thus completing the proof of the lemma. □

Remark: Inspecting the proof of the lemma, we see that it can be adapted to handle any relation which is a conjunction of several sub-relations, each involving points and curves in the plane, such that the given curves have two degrees of freedom (and constant description complexity), so that the problem can be dualized to another plane, where the curves become points and the points become curves. For example, the two instances mentioned above, of points inside unit disks and of intersecting segments, can be handled by the same proof technique, as can be many other related instances.