

On the Number of Regular Vertices of the Union of Jordan Regions*

B. Aronov,¹ A. Efrat,² D. Halperin,³ and M. Sharir^{3,4}

¹Department of Computer and Information Science,
Polytechnic University,
Brooklyn, NY 11201-3840, USA
aronov@ziggy.poly.edu

²Computer Science Department, Stanford University,
Stanford, CA 94305, USA
alon@cs.stanford.edu

³School of Mathematical Sciences, Tel Aviv University,
Tel-Aviv 69978, Israel
{halperin,sharir}@math.tau.ac.il

⁴Courant Institute of Mathematical Sciences,
New York University,
New York, NY 10012, USA

Abstract. Let \mathcal{C} be a collection of n Jordan regions in the plane in general position, such that each pair of their boundaries intersect in at most s points, where s is a constant. If the boundaries of two sets in \mathcal{C} cross exactly twice, then their intersection points are called *regular vertices* of the arrangement $\mathcal{A}(\mathcal{C})$. Let $R(\mathcal{C})$ denote the set of regular vertices on the boundary of the union of \mathcal{C} . We present several bounds on $|R(\mathcal{C})|$, depending on the type of

* The first and the fourth authors have been supported by a grant from the U.S.–Israeli Binational Science Foundation. Boris Aronov has also been supported by a Sloan Research Fellowship and by NSF Grant CCR–99-72568. Micha Sharir has also been supported by NSF Grants CCR-94-24398 and CCR-93-11127, and by a grant from the G.I.F., the German–Israeli Foundation for Scientific Research and Development. Dan Halperin has been supported by an Alon Fellowship, by the U.S.–Israel Binational Science Foundation, and by the Israel Science Foundation founded by the Israel Academy of Sciences and Humanities. Dan Halperin and Micha Sharir have also been supported by ESPRIT IV LTR Project No. 21957 (CGAL), and by the Hermann Minkowski – Minerva Center for Geometry at Tel Aviv University. Part of the work on this paper was carried out when Boris Aronov was visiting Tel Aviv University in March 1997. A preliminary version of the paper has appeared in *Proc. 6th Scandinavian Workshop on Algorithm Theory (SWAT '98)*, 1998, pp. 322–334.

the sets of \mathcal{C} . (i) If each set of \mathcal{C} is convex, then $|R(\mathcal{C})| = O(n^{1.5+\varepsilon})$ for any $\varepsilon > 0$.¹ (ii) If no further assumptions are made on the sets of \mathcal{C} , then we show that there is a positive integer r that depends only on s such that $|R(\mathcal{C})| = O(n^{2-1/r})$. (iii) If \mathcal{C} consists of two collections \mathcal{C}_1 and \mathcal{C}_2 where \mathcal{C}_1 is a collection of m convex *pseudo-disks* in the plane (closed Jordan regions with the property that the boundaries of any two of them intersect at most twice), and \mathcal{C}_2 is a collection of polygons with a total of n sides, then $|R(\mathcal{C})| = O(m^{2/3}n^{2/3} + m + n)$, and this bound is tight in the worst case.

1. Introduction

Let \mathcal{C} be a collection of n Jordan regions (the interiors of closed Jordan curves) in the plane, with the property that the boundaries of any pair of regions intersect in at most some constant number s of points. We assume that the sets of \mathcal{C} are in *general position*, so that no point is incident to more than two boundaries, and that the boundaries cross transversally at each intersection point. Let U denote the *union* of \mathcal{C} . We consider the *arrangement* $\mathcal{A}(\mathcal{C})$, formed by the boundaries of the sets in \mathcal{C} , and define a *vertex* of $\mathcal{A}(\mathcal{C})$ formed by an intersection point of two boundaries $\partial C_1, \partial C_2$ to be *regular* if ∂C_1 and ∂C_2 cross exactly twice; all other vertices are called *irregular*. The goal is to obtain sharp bounds on the maximal number of regular vertices that appear on ∂U . We denote the set of regular vertices on ∂U by $R(\mathcal{C})$.

The interest in this problem goes back to the work of Kedem et al. [12], where it was shown that if all vertices of $\mathcal{A}(\mathcal{C})$ are regular (such a collection \mathcal{C} is called a family of *pseudo-disks*), then the number of (regular) vertices of ∂U is at most $6n - 12$, for $n \geq 3$, and this bound is tight in the worst case. Recently, Pach and Sharir [16] have shown that if \mathcal{C} is an arbitrary collection of n convex sets (so that any two of their boundaries can intersect in an arbitrary number of points), then $|R(\mathcal{C})| \leq 2|I(\mathcal{C})| + 6n - 12$, where $I(\mathcal{C})$ is the set of irregular vertices on ∂U . This result was instrumental in a recent paper by Efrat and Sharir [8], showing that the complexity of the union of n planar “fat” convex sets, each pair of whose boundaries intersect in at most some constant number of points, is nearly linear in n . However, since $I(\mathcal{C})$ can be $\Omega(n^2)$ for a general collection \mathcal{C} , even when no pair of boundaries cross at more than four points, the bound of [16] only yields the trivial $O(n^2)$ upper bound on $|R(\mathcal{C})|$. As an example of such a construction, consider a collection of n narrow rectangles arranged in an $n/2 \times n/2$ grid.

Pach and Sharir [16] also construct a set \mathcal{C} of n rectangles and m unit disks for which $|R(\mathcal{C})| = \Omega(m^{2/3}n^{2/3} + m + n)$. This is the best known lower bound for the general problem stated above, with a constant number of intersections between any pair of boundaries. If two regions are allowed to intersect in an arbitrary number of points, then it is easy to obtain examples with $\Omega(n^2)$ regular vertices on the boundary of the union; see, e.g., [13], [16], and Fig. 1.

For deriving the main results of the paper, we first present in Section 2 a technique for transforming the family of regions so that every regular vertex in $R(\mathcal{C})$ becomes a point of tangency between the two regions, and so that the number of intersections between any pair of boundaries does not increase (see Lemma 1).

Next, in Section 3, we consider the case of general convex regions, and show:

¹ Throughout this paper, ε stands for an arbitrarily small positive constant; the constants of proportionality in bounds that involve ε also depend on ε and, generally, tend to infinity as ε approaches zero.




Fig. 1. Union of polygons with a quadratic number of regular vertices on its boundary: the middle rectangles represent $n/2$ polygons, each with $n/2$ vertices; adding $n/2$ triangles as shown, we get the asserted quadratic lower bound.

Theorem 1. *The number of regular vertices on the boundary of the union of a family of n convex Jordan regions in the plane, where any two boundaries intersect in at most a constant number of points, is $O(n^{1.5+\epsilon})$ for any constant $\epsilon > 0$.*

Next, in Section 4, we study the case of general Jordan regions, and show:

Theorem 2. *The number of regular vertices on the boundary of the union of a family of n Jordan regions in the plane, where any two boundaries intersect in at most a constant number s of points, is $O(n^{2-1/r})$, where r is a positive integer that depends only on s .*

In other words, we show that in fairly general settings, the number of regular vertices on ∂U is subquadratic.

Finally, in Section 5 we show that the $\Omega(m^{2/3}n^{2/3} + m + n)$ lower bound is tight for the special class (iii) of collections \mathcal{C} as in the abstract. Specifically, we show:

Theorem 3. *The number of regular vertices on the boundary of the union of a family of m convex pseudo-disks and a family of polygons with a total of n edges is $\Theta(m^{2/3}n^{2/3} + m + n)$.*

Besides being an interesting collection of results in itself, the study in this paper is likely to have implications on the analysis of the complexity of the union of geometric objects in two and three dimensions. Moreover, our experience has been that improved combinatorial bounds on the complexity of the union of geometric objects often entails efficient algorithms for computing such unions, a task that often arises in several application areas, such as robot motion planning [11], solid modeling, and others.

Remark. Note that using a straightforward perturbation scheme (such as the one in [17]), one can show that the maximum number of regular vertices on the union is achieved when the given regions are in general position (i.e., when no three boundaries have a common point and each crossing between two boundaries is transversal), so this assumption involves no loss of generality.

2. Transforming the Regions

The next lemma is used as a first step in the proofs of Theorems 1 and 2.

Lemma 1. *Let \mathcal{C} be a collection of n Jordan regions in the plane, so that each pair of boundaries intersect in at most a finite number s of points. Then we can transform this collection so that if (c, c') is a pair of regions in \mathcal{C} whose boundaries originally crossed regularly, with at least one of these two crossing points lying on ∂U , then after the transformation c and c' are openly disjoint and touch at a single point that lies on the new union boundary. Moreover, the number of intersections between any two region boundaries does not increase after the transformation. (If two region boundaries cross each other irregularly, then after the transformation they either continue to cross irregularly, or cross regularly, or become disjoint.) Finally, if all original regions in \mathcal{C} are convex, then they remain so after the transformation.*

Proof. Before starting the transformation process, we first remove from \mathcal{C} any region that is fully contained in the interior of the union U and any region that is disjoint from all the other sets (and so does not contribute any vertex to the union). This can be done without loss of generality. We thus assume from now on that \mathcal{C} does not contain any such region.

The transformation process is iterative: Order the regions in \mathcal{C} arbitrarily as (c_1, \dots, c_n) . Let \mathcal{C}_i denote the collection after the first i steps of the transformation, with $\mathcal{C}_0 = \mathcal{C}$. We now describe how to transform \mathcal{C}_{i-1} into \mathcal{C}_i .

Put $c = c_i$, and apply a homeomorphism τ_c of the plane that maps c onto the closed unit disk D ; the existence of such a homeomorphism is a consequence of Schönfliess' theorem [14]. If all the sets are convex, then we take τ_c to be the identity and for uniformity put $D = c$. We now apply the following steps:

- (i) For each $c' \neq c \in \mathcal{C}$, we replace each maximal connected arc of $D \cap \partial\tau_c(c')$ by the line segment connecting its endpoints. See Fig. 2(ii). (It is possible that such an arc is a single point, in which case this step does not modify it.)
- (ii) We shortcut each maximal connected arc γ of $\partial D \cap \text{int}(\tau_c(U))$ as follows. Let u and v denote the endpoints of γ , and let $a, b \in \mathcal{C}$, $a, b \neq c$, be the sets whose transformed boundaries contain u, v , respectively. If $a \neq b$, we replace γ by the line segment uv ; see the left-hand side of Fig. 2(iii). If $a = b$, we choose a point v' on ∂D that lies outside $\tau_c(a)$ very close to v (and does not lie inside any other transformed set), and again replace the portion of ∂D between u and v' by the line segment uv' ; see the right-hand side of Fig. 2(iii) (where the relevant set is d).
- (iii) We now transform the plane back using τ_c^{-1} .

We iterate these three stages, applying them to each $c_i \in \mathcal{C}$ in the above order.

We claim that this transformation does not increase the number of intersections between any pair of boundaries. Indeed, consider the step where a set c_i is processed, and let a, b be two distinct sets in \mathcal{C}_{i-1} . If $c_i \neq a, b$, then the portions of ∂a and ∂b outside c_i do not change, while their portions inside c_i , after applying τ_c , and our "straightening" step (step (ii)), consist of straight segments. If two such segments uv and wz , lying on




Fig. 2. Demonstration of the transformation rules: (i) The disk D with the images of four other sets, referred to as a , b , d , and f . (ii) The modified sets a , b , d , and f ; the arcs γ of ∂D shortcut in the next stage are highlighted. (iii) D after applying stage (ii) of the transformation; the shaded region is a new connected component of the complement of the union (a new “hole”).

the modified respective boundaries ∂a , ∂b inside D , cross at some point x , then, since u , w , v , z all lie on ∂D and must appear there in this cyclic order, and since the original respective portions of $\partial\tau_{c_i}(a)$ and $\partial\tau_{c_i}(b)$ connecting u to v and w to z lie fully inside D , these original portions must cross each other at some point x' . We charge x to x' , and note that this charging is unique, implying that the number of intersections has not increased.

Suppose then that $c_i = a$, say, and let x be a new intersection between ∂D and $\partial\tau_{c_i}(b)$, after the straightening step. Then x must lie (in the τ_{c_i} image) on one of the new straight segments uv or uv' on ∂a in the transformed plane, and on some new straight portion wz of the transformed ∂b . If x is a common endpoint of these two straight segments, then it must have been an intersection point between ∂a and ∂b in C_{i-1} , so no new intersection arises in this case. Otherwise, as above, it follows that one of w , z , say w , must lie on the portion of ∂D that has been replaced by uv or uv' , so we can charge x to w , which was an old intersection point of ∂a and ∂b (in the transformed plane), and this charging is unique, again implying that the number of intersections has not increased.

In particular, every pair of new boundaries intersect (transversally) at most s times, and every pair of boundaries that originally intersected in two (regular) vertices either continue to do so, or just touch each other in a single point, or do not intersect at all after the transformation. Specifically, in the case of regular intersection between two boundaries ∂a , ∂b , if at least one of the two intersections lies on ∂U , then the two transformed boundaries touch at a single point that lies on ∂U ; see Fig. 2(iii). Indeed, suppose that ∂a and ∂b intersect at exactly two points u , v , at least one of which lies on ∂U , and suppose that a is processed before b . When a is processed, ∂a and ∂b become touching by construction. It is easy to verify that this situation does not change after any subsequent transformation step (one needs to verify this only for the step that processes b because no other step affects the neighborhood of this touching point that lies outside all other regions).

All these considerations complete the proof of the lemma. □

Observe that U might be changed by these transformations: The original ∂U is still a portion of the new boundary (not taking into account the slight perturbation introduced in the second case of step (ii)), but U could have gained additional “holes,” as shown in

Fig. 2(iii). Note also that the result of the transformation depends on the order in which the sets are processed. Finally, note that the transformation used here is different from that used in [12], although they do have some common features.

3. Convex Regions: Proof of Theorem 1

We first apply the transformation described in the previous section to the given family \mathcal{C} of convex sets. As already noted, if the sets are convex, then there is no need to apply a homeomorphism to the plane at each step of this process, and the sets remain convex after the transformation. The transformed sets have the following properties. They are convex. No set forms a single component of the union. Any two boundaries intersect at most s times. Any two sets that intersected regularly become disjoint or touch at a single point. If two sets intersected regularly with at least one point of intersection of their boundaries on ∂U , the transformed sets are openly disjoint and touch on ∂U . If they intersected regularly without creating vertices on ∂U , they are now disjoint. From now on, we assume that \mathcal{C} has the properties just noted.

For each $c \in \mathcal{C}$, let e_c denote its “equator,” namely, the segment connecting the leftmost point and the rightmost point in c . (By an appropriate general position assumption, or by appropriate tilting of the plane, e_c is uniquely defined for each $c \in \mathcal{C}$; these segments were called “sentinels” in [2].) In what follows we ignore boundary touchings that occur at endpoints of equators—there are at most $2n$ such touchings.

We first construct a *hereditary segment tree* Q on the x -projections of the equators (which are the same as the x -projections of the sets in \mathcal{C}), as in [5] (consult [5] for more details, and for the terminology that we use below). Each node v of Q stores the standard segment-tree list L_v of sets with so-called “long” equators, and also a list S_v of sets with “short” equators, those that are stored in some list L_w , for a proper descendant w of v , and thus have at least one equator endpoint in the interior of the vertical strip σ_v associated with v . It is easily verified that for any *boundary touching* between two sets $a, b \in \mathcal{C}$, not occurring at an endpoint of any equator, there is a unique node v of Q (on the path to the leaf w whose strip σ_w contains the touching), such that the touching point lies in σ_v and either both a and b belong to L_v or one of them belongs to L_v and the other to S_v . Also, any such touching occurs between the upper boundary of one set and the lower boundary of the other. We have $\sum_v (|L_v| + |S_v|) = O(n \log n)$.

We now fix a node v and bound the number of boundary touchings within σ_v formed between two sets in $L_v \cup S_v$, at least one of which lies in L_v . We only describe the case where the other set lies in S_v , because the case where both of them lie in L_v is simpler and can be handled by a similar approach. Moreover, with no loss of generality, it suffices to consider only boundary touchings where the set in L_v lies above the set in S_v .

We apply fairly standard range-searching techniques, described in some detail below, to obtain a finite collection $\{A_i \times B_i\}_i$ of complete bipartite graphs, such that:

- (a) For each i , $A_i \subseteq S_v$ and $B_i \subseteq L_v$.
- (b) For each i there is a substrip $\sigma^{(i)} \subseteq \sigma_v$, such that for each $a \in A_i$, $b \in B_i$, the x -projection of a contains that of $\sigma^{(i)}$, the equator e_b lies fully above a within σ_v , and the equator e_a lies fully below b within $\sigma^{(i)}$. See Fig. 3.




Fig. 3. Illustrating property (b).

- (c) For each pair of sets $a \in S_v, b \in L_v$, such that ∂a touches ∂b at a point ζ within σ_v and a lies below b , there is an index i such that $a \in A_i, b \in B_i$, and $\zeta \in \sigma^{(i)}$.
- (d) $\sum_i (|A_i| + |B_i|) = O(|S_v|^{1+\epsilon}|L_v|^{1/2+\epsilon} + |L_v|^{1+\epsilon}|S_v|^{1/2+\epsilon})$, for any $\epsilon > 0$, where the constant of proportionality depends on ϵ and on s .

Suppose this has been done. Then fix an index i , and note that any boundary touching $\zeta \in \partial U$ that occurs within $\sigma^{(i)}$ between a set a in A_i and a set b in B_i must occur along the upper envelope of the upper boundaries of the sets of A_i , and along the lower envelope of the lower boundaries of the sets of B_i . Indeed, suppose to the contrary that, say, ζ does not lie on this upper envelope. Since ζ lies on the boundary of the union of A_i , the set a' appearing on the envelope at the x -coordinate of ζ must be such that $e_{a'}$ lies above ζ , which contradicts the property that $e_{a'}$ has to lie fully below b within $\sigma^{(i)}$. The argument for the lower envelope is fully symmetric. The number of such touchings is therefore proportional to the combined complexity of these envelopes within $\sigma^{(i)}$, which is at most $\lambda_s(|A_i|) + \lambda_s(|B_i|)$, where $\lambda_s(n)$ denotes the maximum length of (n, s) -Davenport–Schinzel sequences [17]. Summing over all i 's, we obtain that the total number of boundary touchings “associated” with v is proportional to $O(|S_v||L_v|^{1/2+\epsilon} + |L_v||S_v|^{1/2+\epsilon})$, for a slightly larger, but still arbitrarily small $\epsilon > 0$. Summing these bounds over all nodes v of \mathcal{Q} , we obtain the overall bound $O(n^{3/2+\epsilon})$, again for slightly larger, but still arbitrarily small $\epsilon > 0$, with the constant of proportionality depending on ϵ and on s . This therefore completes the proof of the theorem.

To obtain the decomposition $\{A_i \times B_i\}$, we use a multilevel range-searching structure, where each node w at each level of the structure will store a complete bipartite graph $A'_w \times B'_w$, such that $A'_w \subseteq S_v$ and $B'_w \subseteq L_v$. Each subsequent level of the structure enforces some more of the desired constraints, and the bipartite graphs within each subsequent level form a refinement of the graphs obtained at the previous level. See [1] and [3] for more details concerning multilevel range-searching structures.

In the first level we enforce the property that the segments e_b , for $b \in L_v$, lie above the sets a in S_v (for each of the bipartite graphs to be generated at this step). In what follows,

we clip all the relevant a 's and b 's to within σ_v . Actually, we want to have the property that the equator e_b of any such b lies above the upper boundary of any such a . We may replace any e_b by the line containing it (this has no effect on what happens within σ_v), and replace any a by the portion of σ_v lying below the upper boundary of the original a (so we first make a smaller by ignoring its portion outside σ_v , and then make a larger by allowing it to expand downward within σ_v). We apply a standard duality transform to the plane (as in [7]), which preserves incidences and the above/below relationship. This duality maps the upper boundary of any $a \in S_v$ to a convex x -monotone curve γ_a , and the (extended) equators e_b of sets in L_v are mapped to points e_b^* . An equator e_b lies above the upper boundary of a if and only if the dual point e_b^* lies above γ_a . Note also that any pair of curves $\gamma_a, \gamma_{a'}$ intersect each other at most s times, because any such intersection point is the dual of a common tangent to the upper boundaries of a and a' , and there can be at most s such common tangents, because ∂a and $\partial a'$ intersect in at most s points.

Thus, in this dual setting, the desired first-level decomposition $\{A'_w \times B'_w\}_w$ of $S_v \times L_v$ has to satisfy the following properties:

- (e) For each w , the point e_b^* dual to the equator e_b of any set $b \in B'_w$ lies above γ_a , for every $a \in A'_w$.
- (f) For any $a \in S_v, b \in L_v$, such that e_b^* lies above γ_a , there is a w such that $a \in A'_w$ and $b \in B'_w$.

Put $m_v = |S_v|, n_v = |L_v|$. To obtain this decomposition, we fix some sufficiently large constant parameter ξ , draw a random sample R of ξ sets $a \in S_v$, consider the arrangement \mathcal{A}_R of the corresponding curves γ_a , and apply a vertical decomposition to \mathcal{A}_R that produces $O(\xi^2)$ pseudo-trapezoidal cells. Since each pseudo-trapezoid is determined by at most four curves γ_a , it follows from [6] that with high probability, no pseudo-trapezoid is crossed by more than $(cm_v/\xi) \log \xi$ curves, for some appropriate constant c . We may assume that our sample R does indeed have this property.

For each pseudo-trapezoid τ , let B_τ denote the subset of L_v consisting of those sets b whose dual points e_b^* lie inside τ , and let A_τ (resp. C_τ) denote the subset of S_v consisting of those sets a whose dual curves γ_a pass below τ (resp. cross τ). Put $k_\tau = |A_\tau|, n_\tau = |B_\tau|$, and $m_\tau = |C_\tau|$. We have $\sum_\tau n_\tau = n_v, k_\tau \leq m_v$, and $m_\tau \leq (cm_v/\xi) \log \xi$. By partitioning τ vertically into subcells, if necessary, we may also assume that $n_\tau \leq n_v/\xi^2$ for each τ , while the total number of subcells remains $O(\xi^2)$.

We add to the first-level output collection of complete bipartite graphs all the products $A_\tau \times B_\tau$, and repeat the whole process recursively within each cell τ , with the sets C_τ and B_τ . We stop the recursion when the size of C_τ or of B_τ falls below some specified constant, and then output all appropriate singleton products $\{a\} \times \{b\}$, for $a \in C_\tau, b \in B_\tau$.

It is clear from the construction that the resulting decomposition satisfies the required properties (e) and (f). We next estimate its total size $\sum_\tau (|A_\tau| + |B_\tau|)$. At the top level of the recursion we have $\sum_\tau |B_\tau| = n_v$. Since the number of levels of recursion is $O(\log n_v)$ and the points e_b^* are partitioned among the recursive subproblems, it follows that in total we have $\sum_\tau |B_\tau| = O(n_v \log n_v)$. Similarly, at the top level of the recursion we have $\sum_\tau |A_\tau| = O(m_v \xi^2)$, and $\sum_\tau |C_\tau| = O(m_v \xi \log \xi)$ (where both constants of proportionality depend only on s). The maximum depth j of the recursion satisfies $\xi^{2j} \leq n_v$, or $\xi^j \leq n_v^{1/2}$. It follows that, for an appropriate constant c' that depends on s ,

the overall sum $\sum_{\tau} |A_{\tau}|$ is at most proportional to

$$m_v \xi^2 (1 + c' \xi \log \xi + \dots + (c' \xi \log \xi)^{j-1}) = O(m_v (c' \xi \log \xi)^j) = O(m_v n_v^{1/2+\varepsilon}),$$

where we can make $\varepsilon > 0$ arbitrarily small by choosing ξ sufficiently large (as a function of ε and s). Hence we have

$$\sum_{\tau} |A_{\tau}| = O(m_v n_v^{1/2+\varepsilon}) \quad \text{and} \quad \sum_{\tau} |B_{\tau}| = O(n_v \log n_v). \tag{1}$$

Next, fix a pair $A = A_{\tau}, B = B_{\tau}$ in this decomposition. For each $a \in A, b \in B$ the equator e_b lies above the upper boundary of a (within σ_v). This already implies (arguing as above) that any boundary touching between ∂a and ∂b , for any pair $a \in A, b \in B$ that lies on the union boundary, must lie on the lower envelope E of the lower boundaries of the sets in B . The complexity of E is at most $\lambda_x(|B|)$, which is nearly linear in $|B|$.

For each $a \in A$, let a_x denote its x -projection, clipped to within σ_v . We construct a secondary segment tree T on the intervals a_x . Each node u of T is associated with a vertical strip $\sigma^{(u)} \subseteq \sigma_v$. In addition to the standard segment-tree list $A^{(u)}$ of x -projections of sets in A that is stored at u , we also store there a list $B^{(u)}$ of the sets in B whose lower boundaries appear in $E \cap \sigma^{(u)}$. We clearly have that $\sum_u |A^{(u)}| = O(|A| \log |A|)$ and that $\sum_{|B^{(u)}|>1} |B^{(u)}| = O(\lambda_x(|B|) \log |A|)$.

Indeed, the first bound is a standard property of segment trees. The second bound is a consequence of the following observations: (a) The number of breakpoints of E is at most $\lambda_x(|B|)$. (b) Each breakpoint belongs to at most $\log |A|$ strips $\sigma^{(u)}$. (c) For any u , the size of $B^{(u)}$ is upper bounded by 1 plus the number of breakpoints of E in $\sigma^{(u)}$.

Now suppose that there is a pair $a \in A, b \in B$, with a touching ζ between the upper part of ∂a and the lower part of ∂b that lies on the union boundary. There is a unique node u of T such that $a \in A^{(u)}$ and $\zeta \in \sigma^{(u)}$. Hence ∂b appears on E within $\sigma^{(u)}$, so $b \in B^{(u)}$. Note also that the line containing e_a passes fully below $b \cap \sigma^{(u)}$ (because the x -projection of a fully contains the projection of $\sigma^{(u)}$).

We now fix a node u for which $|B^{(u)}| > 1$, and apply a symmetric version of the first-level decomposition to $A^{(u)} \times B^{(u)}$, to obtain a collection $\{A'_j \times B'_j\}_j$ of complete bipartite graphs, such that:

- (g) For each j , we have $A'_j \subseteq A^{(u)}$ and $B'_j \subseteq B^{(u)}$; we also associate the strip $\sigma^{(u)}$ with j .
- (h) For each j , each $a \in A'_j$ and each $b \in B'_j$, the equator e_a lies fully below b within $\sigma^{(u)}$.
- (i) For any boundary touching ζ as above, there is an index j such that $a \in A'_j$ and $b \in B'_j$ (and ζ lies in the strip of j , that is, in $\sigma^{(u)}$).
- (j) $\sum_j (|A'_j| + |B'_j|) = O(|A^{(u)}| \log |A^{(u)}| + |B^{(u)}| |A^{(u)}|^{1/2+\varepsilon})$, for any $\varepsilon > 0$.

If $|B^{(u)}| = 1$ we output only one bipartite graph $\tilde{A}^{(u)} \times B^{(u)}$ where $\tilde{A}^{(u)}$ consists of all $a \in A^{(u)}$ such that e_a lies fully below the unique $b \in B^{(u)}$ within $\sigma^{(u)}$.

The grand collection of complete bipartite graphs $A'_j \times B'_j$, gathered over all nodes u of T , and over all first-level pairs $A \times B$ (recall that we originally denoted such a pair by $A_{\tau} \times B_{\tau}$ for some trapezoid τ), is the desired output collection. It clearly satisfies

properties (a)–(c). Concerning (d), we first sum the bounds (j) over all $u \in T$, to obtain the bound

$$O(|A| \log^2 |A| + \lambda_s(|B|)|A|^{1/2+\varepsilon}),$$

for a slightly larger but still arbitrarily small $\varepsilon > 0$, where the constant of proportionality depends on ε and on s . Clearly, this bound also subsumes the case $|B^{(u)}| = 1$.

Finally, we sum these bounds over all pairs $A = A_\tau$, $B = B_\tau$, in the first-level decomposition, and make use of (1), to conclude that this sum is at most proportional to

$$\sum_\tau O(|A_\tau| \log^2 |A_\tau| + \lambda_s(|B_\tau|)|A_\tau|^{1/2+\varepsilon}) = O(m_v^{1+\varepsilon} n_v^{1/2+\varepsilon} + n_v^{1+\varepsilon} m_v^{1/2+\varepsilon}),$$

again, for a slightly larger but still arbitrarily small $\varepsilon > 0$, where the constant of proportionality depends on ε and on s . Hence (d) is also satisfied. As already noted, this completes the proof of the theorem. \square

Remark. An obvious open problem is to close the gap between this upper bound and the lower bound $\Omega(n^{4/3})$ noted earlier.

4. Arbitrary Regions: Proof of Theorem 2

We prove the theorem using the following “forbidden subgraph” argument. We first transform \mathcal{C} as described in Lemma 1, but continue to denote the transformed collection by \mathcal{C} . Now we define a graph H whose vertices are the regions in \mathcal{C} and whose edges connect pairs of regions that touch at a point on the boundary of the union. We claim that H does not contain a complete bipartite graph $K_{r,g}$ for some sufficiently large constants $r = r(s)$, $g = g(s)$, where s is the maximum number of intersections between any two region boundaries.

Suppose to the contrary that there exist subsets \mathcal{R}, \mathcal{G} of \mathcal{C} of sizes r, g , respectively, such that for each $(c, c') \in \mathcal{R} \times \mathcal{G}$, c and c' touch at only one point and that point lies on the boundary of the union U . With no loss of generality, we may assume that $\mathcal{C} = \mathcal{R} \cup \mathcal{G}$. We refer to sets of \mathcal{R} (resp. \mathcal{G}) as “red” (resp. “green”).

Consider the arrangement $\mathcal{A}(\mathcal{G})$ and let $U_{\mathcal{G}}$ denote the union of \mathcal{G} . We claim that any $c \in \mathcal{R}$ is fully contained in the closure of a single *hole* (connected component of the complement) of $U_{\mathcal{G}}$. Indeed, if this were false, ∂c would have to cross some boundary of a green region c' , which is impossible by construction.

Consider the collection of holes of $U_{\mathcal{G}}$ that contain red regions. We call them *interesting green holes*. Since each red region must touch the boundary of every green set at a point that lies on $\partial U_{\mathcal{G}}$, it follows that all interesting green holes are part of the *zone* in $\mathcal{A}(\mathcal{G})$ of any green boundary, which is the collection of cells in the arrangement whose closure is met by the boundary. Hence the overall complexity, i.e., the number of vertices and edges, of all these holes is at most $a\lambda_{s+2}(g)$, for some absolute constant a (see [15]). Symmetrically, the overall complexity of interesting red holes, defined analogously for \mathcal{R} , is at most $a\lambda_{s+2}(r)$.

We now construct a planar bipartite graph G , whose nodes are the arcs forming the boundaries of interesting green and red holes; each edge of G connects a green arc ζ to




Fig. 4. Both pairs u, v are pseudo-regular.

a red arc η if these two arcs touch at a (regular) vertex of the overall union. It is easy to verify that G is indeed planar, and that it has no multiple edges. Since the graph is bipartite, the number of its edges is at most twice the number of its nodes, that is, at most $2a(\lambda_{s+2}(g) + \lambda_{s+2}(r))$. On the other hand, the number of edges of G must be rg , because every green set and every red set touch at some (regular) vertex on ∂U . We thus obtain: $rg \leq 2a(\lambda_{s+2}(g) + \lambda_{s+2}(r))$, which is impossible if r and g are sufficiently large constants (that depend on s).

Using standard results in extremal graph theory (see [15]), it follows that the number of regular vertices on the boundary of the union of \mathcal{C} is $O(n^{2-1/r})$. This completes the proof of the theorem. \square

Remarks. (1) The following is a natural extension of the concept of regular intersection: We say that the boundaries of two Jordan regions a and b meet pseudo-regularly at u and v if there is a connected component of $a \cap b$ that contains u and v and no other intersection point of $\partial a \cap \partial b$ on its boundary. In other words, we allow ∂a and ∂b to cross more than twice, but require that their intersections u and v behave “locally” as regular intersections; see Fig. 4. Unfortunately, it is possible to construct families of n Jordan regions, each pair of whose boundaries cross at most six times, such that the boundary of their union has $\Theta(n^2)$ pseudo-regular vertices. Such a construction is depicted in Fig. 5. We start with a grid-like arrangement of thin rectangles and deform the horizontal rectangle near each hole to create two pseudo-regular vertices. Thus Theorem 2 fails for pseudo-regularity.

In the lower bound construction just given, even though any two boundaries cross at most six times, the shape of some of the boundaries is quite complicated. We do not know whether a similar construction can be obtained for semialgebraic Jordan regions of constant description complexity,² so the question whether Theorem 2 holds in this case remains open.

² That is, regions that are each described by a Boolean combination of a constant number of polynomial equalities and inequalities of constant maximum degree.




Fig. 5. Regions with at most six intersections between any pair of boundaries, whose union has $\Theta(n^2)$ pseudo-regular vertices.

(2) Nevertheless, if any two boundaries cross at most four times, then if two boundaries cross four times but meet pseudo-regularly, then the four intersections come in two pairs, each of which is a pseudo-regular pair; see Fig. 4(right). It is easy to verify that Theorem 2 continues to hold in this case, using the same proof with slight and obvious modifications. In fact, this suggests an alternative definition of pseudo-regularity: we say that a and b intersect pseudo-regularly if the intersection $a \cap b$ consists of a finite number of connected components, each with at most two points of $\partial a \cap \partial b$ on it. With this modification, the proof of Theorem 2 carries through, for any fixed s .

5. Pseudo-Disks and Polygons: Proof of Theorem 3

We finally turn to the special case considered in Theorem 3. Let $\mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2$ where \mathcal{C}_1 is a collection of m convex pseudo-disks in the plane (closed Jordan regions with the property that the boundaries of any two of them intersect at most twice), and \mathcal{C}_2 is a collection of simple polygons with a total of n sides. Let U_1, U_2 be the union of $\mathcal{C}_1, \mathcal{C}_2$, respectively, and put $U = U_1 \cup U_2$. By [12], the number of regular vertices on ∂U that are incident to two boundaries of sets in \mathcal{C}_1 is at most $6m$.

Next, we show that the number of regular vertices on U incident to two boundaries of polygons in \mathcal{C}_2 is $O(n)$. Let v be such a vertex, and let c, c' be the two polygons in \mathcal{C}_2 whose boundaries contain v . Let K be the connected component of $c \cap c'$ that contains v . The connected component of ∂K that contains v is a polygonal cycle that contains only two vertices (one of which is v) where ∂c and $\partial c'$ meet. Hence it must also contain at least one vertex w of c or of c' , such that vw is an edge of K . We charge v to such a vertex w , and note that any vertex w can be charged at most twice in this manner. Hence the total number of vertices of ∂U of this type is $O(n)$.

It thus remains to bound the number of *mixed* regular vertices of ∂U , namely, those that are incident to the boundaries of a pseudo-disk in \mathcal{C}_1 and of a polygon in \mathcal{C}_2 . Without loss of generality, we count only those mixed vertices that lie on the top portion of the boundary of pseudo-disks in \mathcal{C}_1 . Let Γ denote the set of arcs of ∂U_1 that lie on the top

portion of the boundaries of pseudo-disks; arcs that contain the leftmost or rightmost points are split at those points. Let $M = O(m)$ denote the number of these arcs.

Let $\gamma \in \Gamma$, let $d \in \mathcal{C}_1$ be the pseudo-disk whose boundary contains γ , and consider its interaction with a polygon c of \mathcal{C}_2 that forms a regular vertex on γ . Since γ lies on the top boundary of d , it follows that either c contains the extreme left or right point of d , or $c \cap \partial d$ lies completely in the top portion of ∂d . Polygons $c \in \mathcal{C}_2$ that contain the leftmost (rightmost) point of d contribute at most one mixed vertex of ∂U on γ (two if c contains both extreme points of d), for a total of at most $2n$ vertices. Similarly we can eliminate polygons c which contain one of the endpoints of γ , as those produce at most $2M = O(m)$ mixed vertices on ∂U . From this point on we restrict our attention to polygons c that form regular vertices on γ and satisfy $c \cap \partial d \subset \gamma$. Let c be such a polygon, and let v, w be the two points of intersection of γ with ∂c . If v, w lie on different edges of c , then the portion of c inside d must contain a vertex of c , and we can then charge v and w to such a vertex, in the same manner as in a preceding paragraph. It follows that the number of such vertices v, w is $O(n)$, and we can thus ignore such cases in subsequent analysis. We can thus assume that γ crosses the boundary of each such polygon c at two points that lie on the same edge e of c . We refer to these remaining regular vertices as *edge-touching* vertices. Note that since γ lies on the top boundary of d , the edge e must lie on the bottom boundary of c (meaning that c lies locally above e).

We adapt the analysis technique of [9] and [10]. First, we derive a weaker upper bound on the number of mixed edge-touching regular vertices on ∂U . To this end, we construct a bipartite graph H whose nodes are the arcs in Γ and the polygon edges, and each of its edges connects an arc γ to a polygon edge e if γ crosses e twice, at two regular vertices, at least one of which lies on ∂U . We claim that H does not contain a $K_{2,3}$ as a subgraph (composed of two arcs in Γ and of three polygon edges). Indeed, suppose that H did contain such a subgraph, consisting of two arcs $\gamma_1, \gamma_2 \in \Gamma$, and of three polygon edges e_1, e_2, e_3 of polygons c_1, c_2, c_3 , respectively. Suppose $\gamma_1 \subset \partial d_1, d_1 \in \mathcal{C}_1$. Since γ_1 intersects each of the three edges twice and each edge e_i lies on the lower boundary of its polygon, the portion of γ_1 outside $c_1 \cup c_2 \cup c_3$ must lie below the lower envelope L of the lines containing e_1, e_2, e_3 . Moreover, the regular vertices of ∂U formed by γ_1 with the three edges must lie on L . Assume, with no loss of generality, that e_1, e_2, e_3 appear in this left-to-right order along L ; see Fig. 6.

L is partitioned by γ_1 into at most seven sections: the three segments $d_1 \cap e_i \cap L, i = 1, 2, 3$, and at most four maximal connected complementary sections of L where




Fig. 6. The graph H does not contain any $K_{2,3}$ consisting of two arcs γ_1, γ_2 , and three polygon edges e_1, e_2, e_3 .

the edge-touching vertices of γ_2 and the edges e_i that appear on ∂U may show up. Since γ_1 and γ_2 are openly disjoint and $\gamma_2 \setminus (c_1 \cup c_2 \cup c_3)$ has to lie below L , all edge-touching vertices of ∂U induced by γ_2 must appear along only one of the latter four complementary sections of the envelope (refer to Fig. 6). However, none of these sections contain portions of all three edges. Hence γ_2 cannot connect to all of $e_1, e_2,$ and e_3 in the graph.

Since the graph H is bipartite, with at most M and n nodes, respectively, in each class, and since it contains no $K_{2,3}$ as a subgraph, it follows from standard extremal graph-theoretic arguments (see Corollary 9.7 of [15]) that the number of its edges, and hence the number of mixed edge-touching regular vertices on ∂U , is $O(Mn^{1/2} + n)$.

We next choose an integer parameter r , to be fixed below, and construct a $(1/r)$ -cutting of the arrangement of the edges of the polygons in \mathcal{C}_2 (see [4] for details). This yields a tiling of the plane by $O(r^2)$ pairwise openly disjoint vertical trapezoids, each crossed by at most n/r edges. For each trapezoid τ , consider the set Γ_τ of all the arcs in Γ that cross τ , clipped to within τ ; some of these arcs may intersect τ in two connected portions, and we regard each such portion as a separate arc.

We classify the arcs in Γ_τ according to the pairs of sides of τ that they cross. One class consists of *short* arcs that have at least one endpoint inside τ ; any other *long* arc meets $\partial\tau$ exactly twice. There are at most $2M$ short arcs in total. One class of long arcs is referred to as the class of *bottom-edge* arcs; these are the arcs that have both endpoints on the bottom side of τ .

Let $CH(X)$ denote the convex hull of the set X . We claim that, for any class $\Gamma' \subseteq \Gamma_\tau$ of arcs, other than those of the short arcs or of the bottom-edge arcs, the number of mixed edge-touching regular vertices on ∂U that are formed within τ by arcs of Γ' is $O(n/r)$. The proof is similar to that in [9] and [10]. It is based on the observation that none of the at most n/r polygon edges that cross τ can form mixed edge-touching regular vertices on ∂U with more than one arc of Γ' . Indeed, if this could happen for one such edge e and two such arcs γ, γ' , then $CH(e \cap \gamma)$ and $CH(e \cap \gamma')$ must be disjoint, with, say, $CH(e \cap \gamma)$ lying to the left of $CH(e \cap \gamma')$; see Fig. 7. (Indeed, the convexity of the pseudo-disks rules out the case that these two intervals are nested within each other; on the other hand, if these intervals overlapped without being nested, then γ and γ' would have to cross, which is impossible.) If we trace γ from the right endpoint of $CH(e \cap \gamma)$ to the right, and trace γ' from the left endpoint of $CH(e \cap \gamma')$ to the left, then these




Fig. 7. e cannot form mixed edge-touching regular vertices with both γ and γ' , unless the family they belong to is the family of short arcs or of bottom-edge arcs.

curves must cross within τ —a contradiction. Indeed, for this not to happen, either one of these curves has to end inside τ (so Γ' is the class of short arcs), or both extensions must cross the bottom side of τ (so Γ' is the class of bottom-edge arcs). Since Γ' is neither of these classes, the claim follows; see [9] and [10] for a similar argument. We have thus shown that the number of mixed edge-touching regular vertices on ∂U that are formed within τ by arcs that are neither short nor bottom-edge is $O(n/r)$, for an overall bound of $O(r^2 \cdot (n/r)) = O(nr)$, over all trapezoids τ .

We next claim that the total number of bottom-edge arcs, over all trapezoids τ , is $O(r^2 + m)$. More precisely, we first discard any bottom-edge arc that does not form a mixed edge-touching regular vertex of ∂U within its containing trapezoid, and then claim the above bound for the number of remaining bottom-edge arcs. If τ is a trapezoid with bottom edge e , then any two remaining bottom-edge arcs γ, γ' within τ are such that $CH(e \cap \gamma), CH(e \cap \gamma')$ are disjoint. This follows from an argument similar to that in the preceding paragraph.

We establish our claim using a graph-planarity argument, similar to that used in [10]. We construct a plane embedding of a planar graph G as follows. The nodes of G are the bottom edges of the trapezoids of our cutting. The edges of G are defined and drawn as follows.

Let γ' be a bottom-edge subarc of some arc $\gamma \in \Gamma$, formed within some trapezoid τ . Let e denote the bottom edge of τ , and let u and v be the two intersection points of γ with e , where u lies to the left of v —these are the endpoints of γ' . Now follow γ from v to the right until another bottom edge e' of some trapezoid τ' is encountered; denote the traced subarc of γ by γ^R . Note that γ can hit e' either from above or from below. (If we do not meet any bottom edge, we can charge γ' uniquely to the right endpoint of γ , so the overall number of these bottom-edge arcs is at most $M = O(m)$. We disregard them in the following argument.) We distinguish between two cases: If the portion of γ to the left of u does not intersect e' , then we connect e and e' in the graph G along γ^R . If, on the other hand, the portion of γ to the left of u does intersect e' , then we claim that both u and v must lie on the top edge of τ' . Indeed, by construction, γ^R does not meet any nonvertical edge of the cutting. Moreover, γ^R must hit e' from above, for otherwise the convexity of γ implies that the entire portion of γ to the left of the hitting point lies below the line supporting e' , so it cannot meet e' again. Now if we follow γ^R from e' to the left, the first vertical edge of the cutting that we meet must be the left edge of τ' , but then γ could not have met e' again further to the left. Hence γ^R is fully contained in τ' , so v lies on the top side of τ' . It follows that u also lies on the top side of τ' , for otherwise e' would have ended to the right of u , so γ could not have intersected it to the left of u . A similar, slightly modified argument implies that the portion of γ between u and e' is also fully contained in τ' . Let u' and v' be the intersection points of γ with e' , where u' lies to the left of v' . We now apply the same analysis to the portions of γ extending to the left of u' and to the right of v' , respectively. We iterate this process until those portions do not end up on the same bottom side. When this is the case, we add the last right portion of γ as an edge of the graph G . (Again, if during this process we reach the right endpoint of γ , we charge γ' uniquely to that endpoint; the overall number of these bottom-edge subarcs is at most $M = O(m)$.) See Fig. 8. Note that the entire portion of γ between (the final) points u' and v' contains only one bottom-edge subarc, namely γ' . Hence the number of edges of G is equal to the number of bottom-edge subarcs




Fig. 8. The portion of γ between v' and v'' is added as an edge to the graph G .

excluding those that we have already charged to the right endpoint of their containing arcs in Γ .

Since the relative interiors of the arcs of Γ are pairwise disjoint, G is indeed a plane embedding of a planar graph. Let f be a face of G of degree 2. By construction, f is bounded by the subarcs δ, δ' of two respective distinct arcs γ, γ' of Γ , and by portions of the bottom edges e, e' of two respective trapezoids τ, τ' . If the interior of f , as a planar region, contains an endpoint of e or of e' , then we can charge f to this endpoint (see Fig. 9 for an illustration of this case). Since there are $O(r^2)$ such endpoints and each of them is charged at most once, the number of such faces f is $O(r^2)$. Otherwise, both δ, δ' must leave e from the same (top or bottom) side of it, and similarly for e' . (We say that δ leaves the top (bottom) side of e , if in a small neighborhood of the intersection point of δ and e , δ is on the top (bottom) side of e .) Since, by construction, each of δ, δ' must leave at least one of e, e' from its bottom side, there are only two possible cases:

(i) δ and δ' leave both e and e' from their bottom sides (see Fig. 10(i)). It is easily seen that in this case the left endpoints of δ, δ' lie on the same edge, say e . Suppose, without loss of generality, that $v = \delta \cap e$ lies to the right of $w = \delta' \cap e$. Let u be the other intersection of γ with e ; by construction, such a point must exist. By convexity of the pseudo-disks in C_1 , u must lie between w and v , and the extension of γ to the left of u lies in f locally near u . Moreover, this extension lies fully below the lines containing e and e' , implying that it cannot intersect any of the four sides of f , and thus the left endpoint of γ must lie inside f . The number of such faces is therefore at most $M = O(m)$.

(ii) δ and δ' leave, say, e from its bottom side and e' from its top side (see Fig. 10(ii)). In this case, by construction, the left endpoints of δ and δ' lie on e . Again, we assume that $v = \delta \cap e$ lies to the right of $w = \delta' \cap e$. Let u be the other intersection of γ with e ;




Fig. 9. A face of G of degree 2 containing an endpoint of a bottom side.




Fig. 10. The portion of γ to the left of $CH(\gamma \cap e)$ must lie fully within f .

as above, by construction, such a point must exist, it must lie between w and v , and the extension of γ to the left of u must lie in f locally near u . As above, we claim that this extension lies fully within f . As in the previous case, this extension cannot intersect e , δ , or δ' . Moreover, it also cannot meet e' (from its top side), since this would contradict our assumption that δ is an edge of G . Hence, in this case too the left endpoint of γ must lie inside f , so the overall number of faces of G of degree 2 is at most $M = O(m)$.

Hence, a straightforward application of Euler's formula for planar graphs implies that the number of edges of G is $O(r^2 + m)$.

Let m_τ denote the number of short and (undiscarded) bottom-edge arcs in τ . The preceding analysis implies that $\sum_\tau m_\tau = O(r^2 + m)$. The weaker bound obtained above implies that the number of mixed edge-touching regular vertices on ∂U that are formed within a trapezoid τ by its m_τ short and bottom-edge arcs is $O(m_\tau(n/r)^{1/2} + n/r)$, for a total of

$$O\left(\sum_\tau \left[m_\tau \left(\frac{n}{r}\right)^{1/2} + \frac{n}{r}\right]\right) = O\left((m + r^2) \left(\frac{n}{r}\right)^{1/2} + nr\right).$$

As argued above, the overall number of all other mixed edge-touching regular vertices on ∂U is $O(m + nr)$. We now choose

$$r = \begin{cases} \lceil m^{2/3}/n^{1/3} \rceil & \text{if } n^{1/2} \leq m \leq n^2, \\ 1 & \text{if } m < n^{1/2}, \\ n & \text{if } m > n^2. \end{cases}$$

The overall bound then becomes $O(m^{2/3}n^{2/3} + m + n)$. This, combined with the lower bound construction in [16], completes the proof of Theorem 3. □

References

1. 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.), AMS Press, Providence, RI, 1998, pp. 1–56.
2. B. Aronov and M. Sharir, The common exterior of convex polygons in the plane, *Comput. Geom. Theory Appl.* **8** (1997), 139–149.
3. M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf, *Computational Geometry: Algorithms and Applications*, Springer-Verlag, Berlin, 1997.

4. B. Chazelle, Cutting hyperplanes for divide-and-conquer, *Discrete Comput. Geom.* **9** (1993), 145–158.
5. B. Chazelle, H. Edelsbrunner, L. Guibas, and M. Sharir, Algorithms for bichromatic line segment problems and polyhedral terrains, *Algorithmica* **11** (1994), 116–132.
6. K. Clarkson, New applications of random sampling in computational geometry, *Discrete Comput. Geom.* **2** (1987), 195–222.
7. H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg, 1987.
8. A. Efrat and M. Sharir, The complexity of the union of fat objects in the plane, *Discrete Comput. Geom.* **23** (2000), 171–189.
9. D. Halperin and M. Sharir, On disjoint concave chains in arrangements of (pseudo)lines, *Inform. Process. Lett.* **40** (1991), 189–192.
10. D. Halperin and M. Sharir, Corrigendum: On disjoint concave chains in arrangements of (pseudo) lines, *Inform. Process. Lett.* **51** (1994), 53–56.
11. D. Halperin and M. Sharir, Arrangements and their applications in robotics: recent developments, in *The Algorithmic Foundations of Robotics*, K. Goldberg, D. Halperin, J. C. Latombe and R. Wilson, eds., A. K. Peters, Boston, MA, 1995, pp. 495–511.
12. 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.
13. J. Matoušek, J. Pach, M. Sharir, S. Sifrony, and E. Welzl, Fat triangles determine linearly many holes, *SIAM J. Comput.* **23** (1994) 154–169.
14. E. E. Moise, *Geometric Topology in Dimension 2 and 3*, Springer-Verlag, New York, 1977.
15. J. Pach and P. K. Agarwal, *Combinatorial Geometry*, Wiley, New York, 1995.
16. J. Pach and M. Sharir, On the boundary of the union of planar convex sets, *Discrete Comput. Geom.* **21** (1999), 321–328.
17. M. Sharir and P. K. Agarwal, *Davenport Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.

Received December 4, 1998, and in revised form June 3, 2000. Online publication February 1, 2001.