

On the Complexity of Many Faces in Arrangements of Pseudo-Segments and of Circles

Pankaj K. Agarwal
Boris Aronov
Micha Sharir

Abstract

We obtain improved bounds on the complexity of m distinct faces in an arrangement of n pseudo-segments, n circles, or n unit circles. The bounds are worst-case optimal for unit circles; they are also worst-case optimal for the case of pseudo-segments, except when the number of faces is very small, in which case our upper bound is a polylogarithmic factor from the best-known lower bound. For general circles, the bounds nearly coincide with the best-known bounds for the number of incidences between m points and n circles, recently obtained in [9].

1 Introduction

Problem statement and motivation. The *arrangement* $\mathcal{A}(\Gamma)$ of a finite collection Γ of curves or surfaces in \mathbb{R}^d is the decomposition of the space into relatively open connected cells of dimensions $0, \dots, d$ induced by Γ , where each cell is a maximal connected set of points lying in the intersection of a fixed subset of Γ and avoiding all other elements of Γ . The *combinatorial complexity* (or *complexity* for short) of a cell ϕ in $\mathcal{A}(\Gamma)$, denoted as $|\phi|$, is the number of faces of $\mathcal{A}(\Gamma)$ of all dimensions that lie on the boundary of ϕ . Besides being interesting in their own right, due to the rich geometric, combinatorial, algebraic, and topological structure that they possess, arrangements also lie at the heart of numerous geometric problems arising in a wide range of applications, including robotics, computer graphics, and molecular modeling. The study of arrangements of lines and hyperplanes has a long, rich history, but most of the work until the 1980s dealt with the combinatorial structure of the entire arrangement or of a single cell in the arrangement (which, in this case, is a convex polyhedron); see [16] for a summary of early work. More recently, motivated by problems in computational and combinatorial geometry, various substructures of, and algorithmic issues involving arrangements of hyperplanes, and, more generally, of hypersurfaces, have received considerable attention, mostly during the last two decades; see [5] for a recent survey.

This paper studies the so-called *many-faces* problem for arrangements of pseudo-segments or of circles in the plane. (A set of arcs is called a family of *pseudo-segments* if every pair of arcs intersect in at most one point.) More precisely, given a set Γ of n arcs and a set P of m points in the plane, none lying on any arc in Γ , let $K(P, \Gamma)$ be the combined combinatorial complexity of the cells of $\mathcal{A}(\Gamma)$ that contain at least one point of P . We wish to obtain an

upper bound for the maximum value of $K(P, \Gamma)$, as a function of n and m , for the cases where Γ is a set of pseudo-segments or a set of circles. The study of the complexity of many faces, and the accompanying algorithmic problem of computing many faces, in planar arrangements (as studied, e.g., in [3, 15]) has several motivations: (i) It arises in a variety of problems involving 3-dimensional arrangements [8, 18]. (ii) It is closely related to the classical problem in combinatorial geometry of bounding the number of *incidences* between points and curves, as studied in numerous papers, including [13, 23, 26, 27, 28]. Informally, in both cases we have points and curves; in the case of incidences, the points lie on the curves and an incidence is a pair (p, γ) , where point p lies on curve γ . In the case of many faces, the points lie “in between” the curves, and we are essentially interested in “extended incidences,” involving pairs (p, γ) , where point p can reach curve γ without crossing any other curve (i.e., γ appears on the boundary of the face containing p). The incidence problem for points and curves has attracted considerable attention in combinatorial and computational geometry; see the papers cited above. The problem of many faces is typically much harder than the (already quite hard) corresponding incidence problem. (iii) The many-faces problem is the “loosest” (i.e., least restricted) of all problems that study substructures in arrangements. It poses the biggest challenge because there is less structure to exploit. Tackling this problem has led to the derivation of various tools, such as the Combination Lemma [19, 25], which are interesting in their own right, and have many algorithmic applications; see, e.g., [2] for a recent such application.

Previous results. An early paper by Canham [10] initiated the study of the many-faces problem for line arrangements. After a number of intermediate results, tight bounds on the complexity of many faces in line and pseudo-line arrangements were obtained by Clarkson *et al.* [13], using an approach based on random sampling. This work, and a series of subsequent papers, proved near-optimal or nontrivial bounds on the complexity of many faces in arrangements of line segments, of circles, and of other classes of curves in the plane, and in arrangements of hyperplanes in higher dimensions; see [5] and the references therein. Aronov *et al.* [7] showed, using a fairly involved analysis, that the complexity of m distinct faces in an arrangement of n segments in the plane is $O(m^{2/3}n^{2/3} + n\alpha(n) + n \log m)$, which is optimal in the worst case except for a small range of m near the value $n^{1/2}$. Unlike the case of lines, their proof does not immediately extend to the case of pseudo-segments. In fact, some key properties of segments that are used in the proof do not hold for arbitrary collections of pseudo-segments, but do hold if we assume that the pseudo-segments are *extendible*; see below for details and further discussion. The best-known bound on the complexity of m distinct faces in an arrangement of n circles in the plane is $O(m^{3/5}n^{4/5}4^{\alpha(n)/5} + n)$. If all circles are congruent, then the bound is $O(m^{2/3}n^{2/3}\alpha^{1/3}(n) + n)$; here $\alpha(n)$ is the extremely slowly growing inverse of the Ackermann’s function [25]. These bounds were obtained in [13].

As mentioned above, the many-faces problem is closely related to the incidence problem, which, given a set Γ of curves and a set P of points in the plane, asks for bounding the number of pairs $(p, \gamma) \in P \times \Gamma$ such that $p \in \gamma$. For example, the tight bounds on the maximum number of incidences between points and lines (or segments, or pseudo-lines) are asymptotically the same as the maximum complexity of m distinct faces in an arrangement of n lines, viz., $\Theta(m^{2/3}n^{2/3} + m + n)$ [13]. (Note, though, that the best-known bound for the *complexity* of many faces in an arrangement of line segments, mentioned above, is

slightly weaker [7].) The same was true for arrangements of circles (except for the tiny $4^{\alpha(n)/5}$ factor in the leading term) until recently, when Aronov and Sharir [9] obtained an improved bound of $O(m^{2/3}n^{2/3} + m^{6/11+3\varepsilon}n^{9/11-\varepsilon} + m + n)$, for any $\varepsilon > 0$, on the number of incidences between points and circles. They raised the question whether a similar bound can be obtained for the complexity of many faces in circle arrangements, which, after the cases of lines, segments, and pseudo-lines, is one of the natural next problem instances to be tackled.

Our results: The case of extendible and general pseudo-segments. We first study the complexity of many faces in an arrangement of *extendible pseudo-segments*. A set S of x -monotone Jordan arcs is called a family of extendible pseudo-segments if there exists a family Γ of pseudo-lines, such that each $s \in S$ is contained in some $\gamma \in \Gamma$. See a recent work of Chan [11], where extendible pseudo-segments are discussed. In particular, not every family of pseudo-segments is a family of extendible pseudo-segments; the simplest demonstration of this fact is depicted in Figure 1. Chan has shown that a family of n x -monotone pseudo-segments

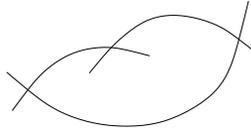


Figure 1: Three pseudo-segments that do not form an extendible family.

can be transformed into a family of $O(n \log n)$ extendible pseudo-segments by cutting each of the given pseudo-segments into at most $O(\log n)$ pieces, in a “segment-tree” fashion.

We prove that the complexity of m distinct faces in an arrangement of n extendible pseudo-segments with X intersecting pairs is $O(m^{2/3}X^{1/3} + n \log n)$. The best lower bound, which is constructed using straight segments, is $\Omega(m^{2/3}X^{1/3} + n\alpha(n))$ [7]. Hence, our bound is worst-case tight when the first term dominates, and is otherwise within a logarithmic factor of the lower bound. Thus, since $X = O(n^2)$, the bound is $O(m^{2/3}n^{2/3} + n \log n)$, which is worst-case optimal for $m = \Omega(n^{1/2} \log^{3/2} n)$.

A closer inspection of the argument in [7] shows that, with some obvious modifications, it also applies to the case of extendible pseudo-segments thus yielding the bound $O(m^{2/3}X^{1/3} + n\alpha(n) + n \log m)$, which beats our bound only when $\log m = o(\log n)$ and $\log m = \omega(\alpha(n))$. Nevertheless, our proof is simpler than that of [7].

Using Chan’s observation, this bound implies an upper bound of $O(m^{2/3}X^{1/3} + n \log^2 n)$ for the complexity of m faces in an arrangement of arbitrary x -monotone pseudo-segments; this bound also holds when the pseudo-segments are not x -monotone, but each of them has only $O(1)$ locally x -extremal points. Again, this is worst-case optimal, unless m is small. For example, substituting $X = O(n^2)$, the bound becomes $O(m^{2/3}n^{2/3} + n \log^2 n)$, which is worst-case optimal for $m = \Omega(n^{1/2} \log^3 n)$.

The analysis of the cases of extendible and general pseudo-segments is important for two independent reasons. First, we obtain nontrivial bounds (which are worst-case optimal or near-optimal) for these cases. In doing so, we obtain a proof that is much simpler than the one given in [7] and, of course, applies also to the case of segments. As mentioned, our bounds are known to be worst-case optimal, unless the value of m is small (about $n^{1/2}$ or smaller). The bound in [7] (modified for the case of extendible pseudo-segments) is slightly better, but these two bounds differ from each other by at most a logarithmic factor. Second, the result for extendible pseudo-segments is used as a major tool in our derivation of the bounds for the case of circles.

The case of circles. Next, we (almost) answer the question raised by Aronov and Sharir affirmatively: Let C be a set of n circles in the plane and P a set of m points, not lying on any circle. As defined earlier, we will use $K(P, C)$ to denote the combined combinatorial complexity of the faces of $\mathcal{A}(C)$ that contain at least one point of P . Set $K(m, n) = \max K(P, C)$, with the maximum taken over all families C of n circles and all families P of m points. We prove that

$$K(m, n) = O\left(m^{2/3}n^{2/3} + m^{6/11+3\varepsilon}n^{9/11-\varepsilon} + n \log n\right),$$

where $\varepsilon > 0$ is an arbitrarily small constant.

Let $K'(m, n)$ denote the maximum value of $K(P, C)$ with the added assumption that all pairs of circles intersect. In this case, following the analysis by Agarwal *et al.* [4], we obtain the following improved bound:

$$K'(m, n) = O\left(m^{2/3}n^{2/3} + m^{1/2}n^{5/6} \log^{1/2} n + n \log n\right).$$

If not all pairs of circles intersect, we obtain a bound that depends on X , the number of intersecting pairs of circles. Let $K(m, n, X) = \max K(P, C)$, with the maximum taken over all families P of m points and C of n circles with X intersecting pairs. We show that

$$K(m, n, X) = O\left(m^{2/3}X^{1/3} + m^{6/11+3\varepsilon}X^{4/11+2\varepsilon}n^{1/11-5\varepsilon} + n \log n\right),$$

where $\varepsilon > 0$ is an arbitrarily small constant. These three bounds are nearly the same as the new corresponding bounds for incidences, given in [9], apart from polylogarithmic factors. Note that the bound of $O(m^{3/5}n^{4/5}4^{\alpha(n)/5} + n)$, obtained by Clarkson *et al.* [13], is slightly better than the ones stated here for $m \leq (n^{1/3}/4^{\alpha(n)/3}) \log^{5/3} n$. For example, $K(m, n) = O(n)$ if $m \leq n^{1/3}$.

Face-curve incidences. Our general technique is similar to the one used in [9], i.e., we first prove a weaker bound, which is almost optimal for large values of m , by cutting the circles of C into extendible pseudo-segments and using the bound for extendible pseudo-segments that we derive separately. Next, to handle small values of m , we use a partitioning scheme in the “dual space,” decompose the problem into many subproblems, bound the complexity for each subproblem using the weaker bound, and estimate the overall complexity as we merge the subproblems.

However, several new ideas are needed to carry out each of these steps. First, we introduce the notion of *face-curve incidences* between the given collection of arcs and a set of marked faces in its arrangement, where a face-curve incidence is a pair (f, γ) , where f is a marked face and γ appears along ∂f . Thus, even if a curve appears many times along a face boundary, we count it only once in this new measure. We show that it suffices to bound the number of those face-curve incidences in order to bound the complexity of the given marked faces. The advantage of incidences is that, with a careful extension of their definition to arrangements of subsets of the given set of curves, this measure is *additive* with respect to both the number of face-marking points and the number of curves. This makes it considerably easier to partition the set of arcs into various subsets, bound the complexity of marked faces in each subarrangement, and then merge (i.e., add the face-curve incidence counts for) the subarrangements. Once a bound on the number of face-curve incidences is obtained, it can be converted to a bound on the actual complexity of these faces (see Lemmas 2.1, 2.2, 2.3).

Previous techniques (e.g., that of [7]) have faced the same problem of merging subarrangements into the whole arrangement, and solved it using *combination lemmas*, which provide relations between the complexity of the marked faces in the subarrangements and the complexity of the marked faces in the whole arrangement. These combination lemmas (presented, e.g., in [19, 25]) are more involved, and generally yield weaker bounds when the partition into subarrangements consists of many recursive levels, as is the case in the analysis presented in this paper.

The case of unit circles. Finally, for the case where all circles in C are congruent (the case of “unit circles”), we show that the complexity of m distinct faces in an arrangement of n congruent circles with X intersecting pairs, is $O(m^{2/3}X^{1/3} + n)$. This bound is asymptotically tight in the worst case, in contrast with the same asymptotic upper bound for the case of incidences [13, 26, 27], which is far away from the best-known, near-linear lower bound. Note that the improvement here is rather marginal—we only remove the factor $\alpha(n)^{1/3}$ from the leading term, appearing in the previous bound of [13].

The paper is organized as follows. In Section 2, we introduce the notion of face-curve incidences, and establish several general properties of this measure, including a relationship between the number of face-curve incidences and the actual complexity of the corresponding faces. Next, we establish in Section 3 complexity bounds for the case of extendible (and general) pseudo-segments. Section 4 derives the bounds for general circles, and Section 5 establishes an optimal bound for congruent circles.

2 Incidences between Curves and Faces

Let Γ be a set of n Jordan arcs in the plane, each pair of which intersect in at most s points. For a point p not lying on any arc in Γ , let f_p denote the face of $\mathcal{A}(\Gamma)$ that contains p . Let P be a finite set of points so that no point lies on any arc in Γ . For a subset $G \subseteq \Gamma$, we define $I_\Gamma(P, G)$ to be the number of pairs $(p, \gamma) \in P \times G$ such that an arc of γ appears on ∂f_p . Note that f_p is defined as a face of the entire arrangement $\mathcal{A}(\Gamma)$ rather than a face of $\mathcal{A}(G)$; it is in fact a subset of the face of $\mathcal{A}(G)$ that contains p . Note also that a pair (p, γ)

is counted only once, even if γ contains more than one edge of ∂f_p .

Lemma 2.1. *Let Γ be a set of Jordan arcs in the plane such that every pair of arcs intersect in at most s points. Let P be a set of points so that none of them lies on any arc of Γ and so that no face of $\mathcal{A}(\Gamma)$ contains more than one point of P . Then*

$$I_\Gamma(P, \Gamma) \leq K(P, \Gamma) = O(\lambda_t(I_\Gamma(P, \Gamma))),$$

where $t = s$ if every arc in Γ is either an unbounded curve that separates the plane or a closed curve, and $t = s + 2$ otherwise; $\lambda_t(n)$ is the maximum length of an (n, t) -Davenport–Schinzel sequence.

Proof. Let n_p be the number of arcs of Γ that appear on the boundary of f_p , for a point $p \in P$. Then $I_\Gamma(P, \Gamma) = \sum_{p \in P} n_p$, and this is clearly a lower bound for $K(P, \Gamma)$. By a result of Guibas *et al.* [17], the complexity of f_p is $O(\lambda_t(n_p))$, where $t = s$ if every arc in Γ is either an unbounded curve separating the plane or a closed curve, and $t = s + 2$ otherwise. Since each face of $\mathcal{A}(\Gamma)$ contains at most one point of P , $K(P, \Gamma) = \sum_{p \in P} O(\lambda_t(n_p)) = O(\lambda_t(I_\Gamma(P, \Gamma)))$. \square

The quantity $K(P, \Gamma) - I(P, \Gamma)$ is closely related to the notion of *excess* introduced by Aronov and Sharir [8]. Specifically, the excess of a face ϕ is the number of edges bounding ϕ minus the number of distinct arcs of Γ that appear on the boundary of ϕ . A result of Sharir [24] implies the following:

Lemma 2.2. *Let Γ be a set of n line segments in the plane, and let P be a set of points, none lying on any segment, so that no face of $\mathcal{A}(\Gamma)$ contains more than one point of P . Then*

$$K(P, \Gamma) = I(P, \Gamma) + O(n \log \log n).$$

A close inspection of the proof given in [24] shows that it also holds for extendible pseudo-segments:

Lemma 2.3. *Let Γ be a set of n extendible pseudo-segments in the plane, and let P be a set of points, none lying on any pseudo-segment, so that no face of $\mathcal{A}(\Gamma)$ contains more than one point of P . Then*

$$K(P, \Gamma) = I(P, \Gamma) + O(n \log \log n).$$

If the set Γ is obvious from the context, we will simply use $I(P, G)$ to denote $I_\Gamma(P, G)$. The following lemma will be crucial in proving the bounds on the complexity of many faces.

Lemma 2.4. *Let $G \subseteq \Gamma$ be a subset of g arcs, and let P be a set of m points, none lying on any arc, so that no face of $\mathcal{A}(\Gamma)$ contains more than one point of P . Then*

$$I_\Gamma(P, G) \leq 2m + 2g + K(P, G).$$

Note the difference between this lemma, which deals with the case where G is a proper subset of Γ , and Lemma 2.1, which deals only with the case $G = \Gamma$. The difference lies in the fact that a face in $\mathcal{A}(G)$ may contain many points of P , and each of its edges may appear on the boundary of many marked faces of $\mathcal{A}(\Gamma)$. Lemma 2.4 shows that the number of these additional multiple occurrences of edges is bounded by $2m + 2g$. (The lemma also holds for $G = \Gamma$, but then the bound in Lemma 2.1 is better.)

Proof. Let F be the set of faces of $\mathcal{A}(G)$ that contain points of P . Let f be a face in F that contains $m_f > 0$ points of P , say, p_1, \dots, p_{m_f} . The corresponding faces f_{p_j} of $\mathcal{A}(\Gamma)$, for $j = 1, \dots, m_f$, are pairwise-disjoint connected regions within f (because each face of $\mathcal{A}(\Gamma)$ is assumed to contain at most one point of P). Suppose ∂f has ξ_f connected components. For each connected component, we choose a point q_j , for $1 \leq j \leq \xi_f$, that lies in the complement of f bounded by that component. We decompose each connected component of ∂f into maximal connected portions, so that each portion overlaps with the boundary of a single face f_{p_i} of $\mathcal{A}(\Gamma)$; such a portion might appear on ∂f_{p_i} in many disconnected pieces; see Figure 2. Let $\gamma_1, \dots, \gamma_{h_f}$ denote the resulting partition of ∂f . Then the points of P lying in f contribute at most $h_f + |f|$ to $I(P, G)$, where $|f|$ is the number of edges in ∂f . Hence,

$$I(P, G) \leq \sum_{f \in F} (h_f + |f|) = K(P, G) + \sum_{f \in F} h_f.$$

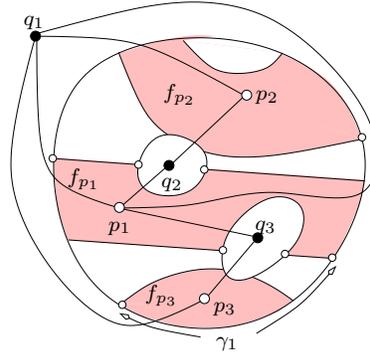


Figure 2: Construction of the bipartite graph to bound $I(P, G)$ within a single face of $\mathcal{A}(G)$; small white circles denote the partition of ∂f into $\gamma_1, \gamma_2, \dots$.

In order to bound h_f , we construct a planar bipartite graph whose vertices are the points p_j , for $j = 1, \dots, m_f$, on one side, and the points q_j , for $j = 1, \dots, \xi_f$, on the other side. For each γ_l , if γ_l is a portion of the j th connected component of ∂f and overlaps with f_{p_i} , we connect p_i to q_j by an edge; we draw the edge as an arc passing through γ_l ; see Figure 2. This can easily be done so that these edge drawings are pairwise disjoint (except at their endpoints). The resulting graph is planar and has no faces of degree two (although there may be multiple edges between a pair of vertices). Hence, the number h_f of edges in the graph is at most $2(m_f + \xi_f) - 4$.

The points of P are partitioned among the faces of F , so $\sum_{f \in F} m_f = m$. Moreover, $\sum_{f \in F} (\xi_f - 1) \leq |G| = g$. Indeed, $\xi_f - 1$ is the total number of “islands” (inner boundary components) inside the face f , and an arc of G cannot belong to more than one island. This completes the proof of the lemma. \square

A useful property of $I_\Gamma(P, G)$, which justifies its introduction, is given in the following lemma; its proof is immediate from the definition.

Lemma 2.5. $I_\Gamma(\cdot, \cdot)$ is additive in both variables: If $P = P_1 \cup P_2$, where P_1 and P_2 are disjoint subsets of marking points, so that no face of $\mathcal{A}(\Gamma)$ contains more than one point of P , and if $G = G_1 \cup G_2$, where G_1 and G_2 are disjoint subsets of G , then

$$\begin{aligned} I_\Gamma(P_1 \cup P_2, G) &= I_\Gamma(P_1, G) + I_\Gamma(P_2, G) \\ I_\Gamma(P, G_1 \cup G_2) &= I_\Gamma(P, G_1) + I_\Gamma(P, G_2). \end{aligned}$$

3 The Case of Pseudo-Segments

A collection Γ of bounded Jordan arcs (resp., unbounded Jordan curves, each separating the plane) in the plane is called a family of *pseudo-segments* (resp., *pseudo-lines*) if every pair of them intersect in at most one point (resp., in exactly one point), where they cross each other. A collection Γ of n *x-monotone* pseudo-segments is called a family of *extendible pseudo-segments* if there exists a family Γ_0 of *x-monotone* pseudo-lines, so that each arc in Γ is contained in some pseudo-line of Γ_0 . The basic properties of extendible pseudo-segments are mentioned in the introduction, and presented in more detail in [11].

In the main result of this section, we bound the complexity of m faces in an arrangement of n extendible pseudo-segments. By combining this bound with the machinery of [11], we also obtain an upper bound for the complexity of many faces in an arrangement of arbitrary *x-monotone* pseudo-segments. This bound also holds for non-*x-monotone* pseudo-segments, provided that each of them has only $O(1)$ locally *x-extremal* points. As far as we know, this is the first study of these cases. Besides being interesting in its own right, the case of extendible pseudo-segments will be used, as already indicated in the introduction, as a main tool in our derivation of the bounds for the case of general circles, presented in the next section.

A weaker bound for extendible pseudo-segments. Let Γ be a set of n *x-monotone* extendible pseudo-segments, and let P be a set of m points in the plane so that no point lies on any pseudo-segment or on any vertical line passing through an endpoint of a pseudo-segment. Let X denote the number of intersecting pairs in Γ .

Lemma 3.1. *The maximum complexity of m distinct faces in an arrangement of n extendible pseudo-segments in the plane is $O(m^{2/3}n^{2/3} + n^{4/3})$.*

Proof. For each $p \in P$ for which f_p is not *x-monotone*, partition f_p into *x-monotone* subfaces by erecting vertical segments up and down from each pseudo-segment endpoint that lies on ∂f_p , until they meet another pseudo-segment (or extend all the way to $\pm\infty$). The number of resulting subfaces is at most $m + 4n$. Let $P_0 \supseteq P$ be a new set of marking points, one in each of the new subfaces. We apply Székely's technique [27] to bound the complexity of these $O(m + n)$ *x-monotone* subfaces, using the same approach as in [14]. Namely, we define a graph G with the set P_0 of marking points as vertices. Two points $p, p' \in P_0$ are connected by an edge if there exists a pseudo-segment $s \in \Gamma$ that appears on the top boundaries of these two faces (resp., on their bottom boundaries), so that s does not appear along the top (resp., bottom) boundary of any other marked subface between these two appearances. An illustration of a portion of such a graph is given in Figure 3. As shown in [14], one can

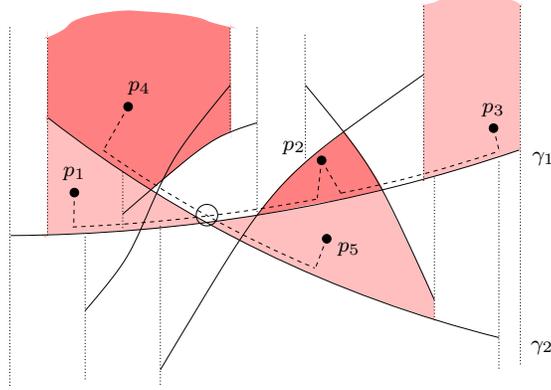


Figure 3: Three edges of the graph G . The edge (p_1, p_2) connecting p_1 and p_2 along the upper side of the arc γ_1 and the edge (p_4, p_5) connecting p_4 to p_5 along the upper side of the arc γ_2 cross at an intersection point of γ_1 and γ_2 .

draw the edges of G as arcs in the plane, so that they intersect only at points of intersection between the curves of Γ . The graph G may have multiple edges connecting the same pair of points. However, the facts that the pseudo-segments are extendible and that the subfaces are x -monotone imply that the edge multiplicity in the resulting graph is at most four. This is shown as follows.

Define, as in [11], a relation on Γ , so that, for $s, s' \in \Gamma$, $s \prec_{\Gamma} s'$ if s and s' intersect, and, slightly to the left of their intersection point, s lies below s' . As noted in [11], Γ is a collection of extendible pseudo-segments if and only if \prec_{Γ} is a partial order. Let L be a family of pseudo-lines, so that each $s \in \Gamma$ is contained in some $\tilde{s} \in L$. Define a total order on L so that, for $\gamma, \gamma' \in L$, $\gamma \prec_L \gamma'$ if γ lies below γ' to the left of their intersection point. By construction, \prec_L is a linear extension of \prec_{Γ} .

Claim. *Let f be an x -monotone subface, as constructed above, and let $s \in \Gamma$ be a pseudo-segment appearing on the top (resp., bottom) boundary of f . Then f lies fully below (resp., above) the pseudo-line $\tilde{s} \in L$ containing s .*

Proof. Suppose to the contrary that the top portion of ∂f , which is a connected x -monotone curve, crosses \tilde{s} , say to the right of s (clearly, the boundary cannot cross s itself). Consider the leftmost such crossing. Let $t \in \Gamma$ be the pseudo-segment along which the crossing takes place, and let \tilde{t} be the pseudo-line in L containing t . By definition, we have $\tilde{t} \prec_L \tilde{s}$. On the other hand, follow the top boundary of f from s to the right, and let $s = s_1, s_2, \dots, s_j = t$ be the sequence of pseudo-segments that we encounter between s and t . See Figure 4. By definition, we have $s_i \prec_{\Gamma} s_{i+1}$ and thus $\tilde{s}_i \prec_L \tilde{s}_{i+1}$, for each $i = 1, \dots, j - 1$. Therefore $\tilde{s} \prec_L \tilde{t}$, a contradiction that establishes the asserted claim. \square

Now let f and f' be two (x -monotone sub-)faces that are connected by at least five edges in G . Then there exist three distinct pseudo-segments, s_1, s_2, s_3 that appear, say, along the top boundaries of both f and f' . Let E denote the lower envelope of the three

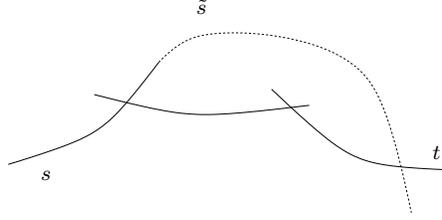
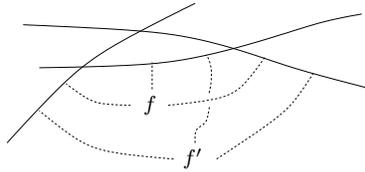


Figure 4: Illustration to the claim in the proof of Lemma 3.1

corresponding pseudo-lines $\tilde{s}_1, \tilde{s}_2, \tilde{s}_3$. The above claim implies that f and f' lie fully below E , and each of them touches E at three distinct points. Since E consists of three connected arcs, each contained in a different pseudo-line, it follows easily that this configuration yields an impossible planar drawing of $K_{3,3}$. See Figure 5 for an illustration.

Figure 5: Impossible drawing of $K_{3,3}$ when 3 distinct pseudo-segments bound f and f' on their top sides.

Hence the edge multiplicity of G is at most 4, and the lemma now follows exactly as in [14, 27]. \square

Next, we obtain an improved bound on $K(P, \Gamma)$ using a decomposition in dual space. It suffices to obtain a bound on $I(P, \Gamma)$ since, by Lemma 2.2, $K(P, \Gamma) = I(P, \Gamma) + O(n \log n)$.

Cuttings. Although the following discussion applies to (and is presented for) any dimension d , we only need it for $d = 2$ (in this section), and for $d = 3$ (when treating the case of circles).

Let H be a set of m hyperplanes in \mathbb{R}^d , and let S be a set of n points in \mathbb{R}^d . For a simplex Δ , we use $H_\Delta \subseteq H$ to denote the set of hyperplanes that cross (i.e., meet the interior of) Δ , and S_Δ to denote $S \cap \Delta$. Set $m_\Delta = |H_\Delta|$ and $n_\Delta = |S_\Delta|$. Let k_Δ be the number of vertices of $\mathcal{A}(H)$ that lie inside Δ .

Let $1 \leq r \leq m$ be a parameter and Δ a simplex. A simplicial subdivision Ξ of Δ is called a $(1/r)$ -cutting of H (with respect to Δ) if at most m/r hyperplanes of H cross any simplex of Ξ . We will use Chazelle's *hierarchical cuttings* [12] to compute a $(1/r)$ -cutting Ξ of H . In this approach, one chooses a sufficiently large constant r_0 , and sets $\nu = \lceil \log_{r_0} r \rceil$. One then computes a sequence of cuttings $\Xi_0, \Xi_1, \dots, \Xi_\nu = \Xi$, where Ξ_i is a $(1/r_0^i)$ -cutting of H . The initial cutting Ξ_0 is simply Δ itself. The cutting Ξ_i is obtained from Ξ_{i-1} by computing, for each $\tau \in \Xi_{i-1}$, a $(1/r_0)$ -cutting Ξ_i^τ of H_τ within τ . It is shown in [12] that $|\Xi_i| \leq cr_0^{di}$, for some constant $c > 0$ that only depends on d . Hence, $|\Xi| = O(r^d)$.

Chazelle's technique is presented only for the case of hyperplanes. However, in the planar case it also applies to other families of curves. In particular, it holds for families of pseudo-lines. The only technical difference is that, instead of simplices (i.e., triangles), one needs to use vertical pseudo-trapezoids (see, e.g., [1] for details).

A stronger bound for extendible pseudo-segments. We now return to our analysis of pseudo-segments. Let L be the set of pseudo-lines containing the pseudo-segments of Γ . As above, we assume that no two pseudo-segments lie on the same pseudo-line. We apply the recent duality transform of Agarwal and Sharir [6], which maps L into a set L^* of n dual points, and maps P to a set P^* of m dual x -monotone pseudo-lines, so that the above/below relationships between points and pseudo-lines are preserved.

We fix a parameter $r \geq 1$, to be determined later, and construct a hierarchical $(1/r)$ -cutting Ξ of $\mathcal{A}(P^*)$, as just described. Ξ consists of $O(r^2)$ pseudo-trapezoids, and it is constructed in $\lceil \log_{r_0} r \rceil$ phases, for some constant $r_0 > 1$. Let Ξ_i be the i -th layer of the cutting; we have $|\Xi_i| \leq cr_0^{2i}$. If a pseudo-trapezoid Δ of the final cutting Ξ contains more than n/r^2 points of L^* , then we split it further into subtrapezoids, each of which contains at most n/r^2 points of L^* . Let $\tilde{\Xi}$ denote the resulting $(1/r)$ -cutting. Using the notation introduced above, $|\tilde{\Xi}| = O(r^2)$, $n_\Delta \leq n/r^2$, and $m_\Delta \leq m/r$, for every pseudo-trapezoid $\Delta \in \tilde{\Xi}$. By choosing the marking points of P generically, and by exploiting the flexibility available in drawing the dual family P^* of pseudo-lines (see [6]), we may assume that all the points of L^* lie in the interiors of the cells of each of the cuttings in the hierarchy. Thus $\sum_\Delta n_\Delta = n$.

Lemma 3.2. *Let τ be a pseudo-trapezoid in one of the cuttings Ξ_i . Let $P' \subset P$ be a subset of the marking points. Then $I(P' \setminus P_\tau, \Gamma_\tau) = O(|P'| \log^* |P'| + n_\tau)$.*

Proof. By definition, for any point $p \in P' \setminus P_\tau$, the dual pseudo-line p^* does not cross the pseudo-trapezoid τ , and therefore passes below all the points of L_τ^* , or above all these points. In primal space, p lies below all the pseudo-lines in L_τ or above all these pseudo-lines. In particular, every such p lies in the unbounded face φ of $\mathcal{A}(\Gamma_\tau)$.

Let γ be a connected component of $\partial\varphi$, and let n_γ be the number of pseudo-segments of Γ_τ that appear on γ . Partition γ into maximal connected portions (referred to as *blocks*), each overlapping the boundary of a single face f_p (in the entire $\mathcal{A}(\Gamma)$), for some $p \in P' \setminus P_\tau$; let m_γ be the number of blocks into which γ is partitioned. We proved in Lemma 2.4 that

$$\sum_\gamma n_\gamma = n_\tau \quad \text{and} \quad \sum_\gamma m_\gamma \leq 2(|P'| + \xi_\tau) - 4, \quad (1)$$

where ξ_τ is the number of connected components of $\partial\varphi$. As a matter of fact, the proof of Lemma 2.4 shows that this serves as an upper bound on the number of blocks δ along any subset of components, with ξ_τ replaced by the size of that subset.

Fix a connected component γ of φ , and let $m_\gamma \geq 3$ denote the number of blocks δ into which γ has been partitioned. (Components with $m_\gamma \leq 2$ will be handled separately.) Enumerate these blocks as $\delta_1, \dots, \delta_{m_\gamma}$ in their circular counterclockwise order along γ . For each block δ_i , encode it as a sequence of the pseudo-segments that appear along δ_i in order, but (i) use different symbols for the two different sides of each pseudo-segment, and also use,

if necessary, two different symbols for a side of a pseudo-segment, to account for the possible ‘wrap-around’ of that side when the circular sequence is being linearized; see [25, Chap. 5.2] for details), and (ii) record only one appearance of each of these symbols in a block, even if it appears there several times. Let σ_γ denote the concatenation of these ‘block-sequences.’ If the last symbol of a block is the same as the first symbol of the next block, we delete one of them; at most m_γ symbols are deleted. The resulting sequence is a Davenport-Schinzel sequence of order three, composed of at most $4n_\gamma$ symbols, and consisting of m_γ blocks, each composed of distinct symbols. The analysis of Davenport-Schinzel sequences of order three, as presented in [25, Chap. 2.2], implies that the length of σ_γ is $O(km_\gamma\alpha_k(m_\gamma) + kn_\gamma)$, for any integer k , where α_k is the inverse of the k -th Ackermann’s function (see [25, Chap. 2.2] for details). Choosing $k = 3$, we obtain $|\sigma_\gamma| = O(m_\gamma \log^* m_\gamma + n_\gamma)$.

We sum this bound over all connected components γ for which $m_\gamma \geq 3$. Let t denote the number of such components. Then they contribute at least $3t$ to the left-hand side of the second inequality of (1), implying that $3t \leq 2|P'| + 2t - 4$, or $t \leq 2|P'| - 4$. Hence, $\sum_{m_\gamma \geq 3} m_\gamma \leq 2(|P'| + t) - 4 = O(|P'|)$. For components γ with $m_\gamma \leq 2$, the total length of their associated sequences σ_γ is at most $8 \sum_\gamma n_\gamma$. Hence, using (1), the total length of all the sequences σ_γ is

$$O\left(\sum_{m_\gamma \geq 3} (m_\gamma \log^* m_\gamma + n_\gamma) + \sum_{m_\gamma \leq 2} n_\gamma\right) = O(|P'| \log^* |P'| + n_\tau).$$

□

By Lemma 2.5,

$$I(P, C) = \sum_{\Delta \in \Xi} I(P, C_\Delta) = \sum_{\Delta \in \Xi} (I(P_\Delta, C_\Delta) + I(P \setminus P_\Delta, C_\Delta)). \quad (2)$$

Instead of bounding the right-hand side directly, we use a recursive approach, based on the hierarchy of cuttings $\Xi_0, \Xi_1, \dots, \Xi_\nu = \Xi$ that underlies the construction of Ξ .

$$\begin{aligned} I(P, \Gamma) &\leq \sum_{\Delta \in \Xi_1} \left(I(P_\Delta, \Gamma_\Delta) + I(P \setminus P_\Delta, \Gamma_\Delta) \right) \\ &\leq \sum_{\Delta \in \Xi_1} \sum_{\tau \in \Xi_2^\Delta} (I(P_\tau, \Gamma_\tau) + I(P_\Delta \setminus P_\tau, \Gamma_\tau)) + \sum_{\Delta \in \Xi_1} a(m \log^* m + n_\Delta) \\ &\hspace{15em} \text{(by Lemma 3.2)} \\ &\leq \sum_{\tau \in \Xi_2} (I(P_\tau, \Gamma_\tau) + a(n_\tau + (m/r_0) \log^* m)) + a(n + c'r_0^2 m \log^* m) \\ &\dots \\ &\leq \sum_{\tau \in \Xi_i} I(P_\tau, \Gamma_\tau) + ian + a'm \log^* m \sum_{j=0}^{i-1} r_0^j, \end{aligned}$$

for all $i = 1, \dots, \lceil \log_{r_0} r \rceil$, where $a' = acr_0^2$. We thus obtain

$$\begin{aligned} I(P, \Gamma) &\leq \sum_{\tau \in \Xi} I(P_\tau, \Gamma_\tau) + O(n \log r + mr \log^* m) \\ &\leq \sum_{\tau \in \tilde{\Xi}} I(P_\tau, \Gamma_\tau) + O(n \log r + mr \log^* m) \\ &\leq \sum_{\tau \in \tilde{\Xi}} (K(P_\tau, \Gamma_\tau) + 2m_\tau + 2n_\tau) + O(n \log r + mr \log^* m), \end{aligned} \quad (3)$$

where the last inequality follows from Lemma 2.4. Substituting the value of $K(P_\tau, \Gamma_\tau)$ from Lemma 3.1, and using the fact that $m_\tau \leq m/r$ and $n_\tau \leq n/r^2$, for each τ , and that $|\tilde{\Xi}| = O(r^2)$, we obtain

$$\begin{aligned} I(P, \Gamma) &= O(n \log r + mr \log^* m) + \sum_{\tau \in \tilde{\Xi}} O\left(m_\tau^{2/3} n_\tau^{2/3} + n_\tau^{4/3}\right) \\ &= O\left(n \log r + mr \log^* m + m^{2/3} n^{2/3} + \frac{n^{4/3}}{r^{2/3}}\right). \end{aligned}$$

Choosing $r = \lceil n/m \rceil$ and using Lemma 2.2, we have

$$K(P, \Gamma) = O\left(m^{2/3} n^{2/3} + n(\log(n/m) + \log^* n + \log \log n)\right).$$

We note that the near-linear terms dominate only when m is smaller than, or is very close to $n^{1/2}$. For such values of m , the first near-linear term is $O(n \log n)$ and thus dominates all the others. Hence we obtain the following bound, which coincides with the one in [7] for all but very small values of m .

Theorem 3.3. *The maximum complexity of m distinct faces in an arrangement of n extendible pseudo-segments in the plane is $O(m^{2/3} n^{2/3} + n \log n)$.*

We next refine Theorem 3.3, to obtain a bound that depends on the number X of intersections between the pseudo-segments of Γ . This is done using the following fairly-standard approach. Put $s = \lceil n^2/X \rceil$, and construct a $(1/s)$ -cutting of $\mathcal{A}(\Gamma)$ that consists of $O(s + s^2 X/n^2) = O(s)$ vertical pseudo-trapezoids, each crossed by at most n/s pseudo-segments (see, e.g., [20]). We apply Theorem 3.3 to bound the complexity of the marked faces within each cell, add up the resulting complexity bounds, and also add the complexity of the zones of the cell boundaries to account for faces not confined to a single cell (as in [13]). The overall complexity of the zones is $O(s) \cdot O(\frac{n}{s} \alpha(\frac{n}{s})) = O(n \alpha(n))$. This leads to the following result.

Theorem 3.4. *The maximum complexity of m distinct faces in an arrangement of n extendible pseudo-segments in the plane is $O(m^{2/3} X^{1/3} + n \log n)$.*

The case of arbitrary pseudo-segments. We next extend the analysis to the case of arbitrary x -monotone pseudo-segments. This is an easy consequence of Chan's analysis [11]. Namely, we cut the n given pseudo-segments into $O(n \log n)$ subarcs, which constitute a family of extendible pseudo-segments, and then apply Theorem 3.4 to the new collection, observing that the cuts do not change X . We thus obtain:

Theorem 3.5. *The maximum complexity of m distinct faces in an arrangement of n x -monotone pseudo-segments in the plane is $O(m^{2/3} X^{1/3} + n \log^2 n)$.*

As already noted, the same bound also holds for collections of pseudo-segments that are not x -monotone, provided that each of them has only $O(1)$ locally x -extremal points. By cutting each pseudo-segment at its x -extremal points, we obtain a family of $O(n)$ x -monotone pseudo-segments, and can then apply Theorem 3.5 to the new collection.

4 The Case of Circles

In this section, we derive an improved bound on the complexity of many faces in an arrangement of circles in the plane. Let C be a set of n circles in the plane with X intersecting pairs, and let P be a set of m points, none of which lies on any input circle. By Lemma 2.1,

$$K(P, C) = O(I(P, C)). \quad (4)$$

We first prove a weak bound on $K(P, C)$ by cutting the circles into pseudo-segments and using the results of the preceding section, and then derive an improved bound by decomposing the problem into subproblems using cuttings in dual space, similar to the approach used for pseudo-segments.

A weaker bound. Aronov and Sharir [9] showed that any family of n circles in the plane, with X intersecting pairs, can be cut into $O(n^{1/2-\varepsilon} X^{1/2+\varepsilon} + n)$ x -monotone pseudo-segments, for any arbitrarily small constant $\varepsilon > 0$. Using the result of Chan [11], mentioned above, we can decompose each of the resulting pseudo-segments into $O(\log n)$ subarcs, that collectively constitute a family Γ of extendible pseudo-segments. Then $I_C(P, C) \leq I_\Gamma(P, \Gamma)$. Using Theorem 3.4, the inequality (4), and the fact that $X = O(n^2)$ in the worst case, we obtain the following lemma (where the two logarithmic factors, one incurred by cutting the arcs further into extendible pseudo-segments, and one appearing in the bound of Theorem 3.4, are both subsumed by the factor n^ε).

Lemma 4.1. *The maximum complexity of m distinct faces in an arrangement of n circles in the plane is $O(m^{2/3} n^{2/3} + n^{3/2+\varepsilon})$, for any $\varepsilon > 0$.*

If every pair of circles in C intersect, then a recent result by Agarwal *et al.* [4] shows that C can be cut into $O(n^{4/3})$ pseudo-segments, and thus into $O(n^{4/3} \log n)$ extendible pseudo-segments, which implies the following bound.

Lemma 4.2. *The maximum complexity of m distinct faces in an arrangement of n pairwise-intersecting circles in the plane is $O(m^{2/3} n^{2/3} + n^{4/3} \log^2 n)$.*

In Lemma 4.1, the term $n^{3/2+\varepsilon}$ becomes dominant when m is smaller than roughly $n^{5/4}$. In order to obtain an improved bound for smaller values of m , we (i) choose a parameter r , depending on n and m , (ii) partition C into $O(r^3)$ subsets, each of size at most n/r^3 , so that the points of P lie in at most m/r distinct faces of the arrangement of each subset, excluding faces in the common exterior or in the common interior of the circles in the subset, (iii) use Lemma 4.1 to bound the complexity of the faces in question in each subarrangement, and (iv) analyze the cost of overlaying all the subarrangements. Although this technique is similar in spirit to an analogous approach used in [9] for the case of incidences, it is considerably more involved when analyzing the complexity of many faces.

Decomposing into subproblems. Using hierarchical cuttings in dual space, we decompose the problem of estimating $I(P, C)$ into subproblems, each involving appropriate subsets of P and C . We use the standard lifting transformation, as in [9], to map circles to points, and points to planes, in \mathbb{R}^3 : A circle γ of radius ρ and center (a, b) in the plane is mapped to the point $\gamma^* = (a, b, a^2 + b^2 - \rho^2) \in \mathbb{R}^3$, and a point $p = (\xi, \eta)$ in the plane is mapped to the plane $p^* : z = 2\xi x + 2\eta y - (\xi^2 + \eta^2)$ in \mathbb{R}^3 . As is easily verified, a point p lies on (resp., inside, outside) a circle γ if and only if the dual plane p^* contains (resp., passes above, below) the dual point γ^* . Let P^* denote the set of planes dual to the points of P and let C^* denote the set of points dual to the circles of C . No three planes of P^* pass through a common line, as all planes of P^* are tangent to the paraboloid $\Pi : z = x^2 + y^2$.

We apply the hierarchical cutting procedure, reviewed in the preceding section, to P^* and C^* in the dual 3-dimensional space, with respect to a sufficiently large simplex that contains C^* and all vertices of $\mathcal{A}(P^*)$, with a value of r that will be fixed later. Let Ξ denote the resulting hierarchical $(1/r)$ -cutting. If $|C_\Delta^*| > n/r^3$ for any simplex $\Delta \in \Xi$, then we split it further into a set $\tilde{\Xi}_\Delta$ of simplices so that each simplex of $\tilde{\Xi}_\Delta$ contains at most n/r^3 points of P^* . This step creates at most r^3 new simplices. Let $\tilde{\Xi}$ denote the resulting cutting. The size of $\tilde{\Xi}$ is also $O(r^3)$. For each simplex $\Delta \in \tilde{\Xi}$, we have $|P_\Delta^*| \leq m/r$ and $|C_\Delta^*| \leq n/r^3$; put $m_\Delta = |P_\Delta^*|$ and $n_\Delta = |C_\Delta^*|$. Finally, for a simplex $\Delta \in \tilde{\Xi}$, let C_Δ be the subset of circles in C that are dual to the points of C_Δ^* , and let P_Δ denote the set of points of P dual to the planes of P_Δ^* . Since no point of P lies on any circle, and we can choose them generically, we may assume that all points of C^* lie in the interiors of the simplices of the cutting. We thus have $\sum_\Delta n_\Delta = n$. We define similar quantities for the simplices of intermediate cuttings.

Obtaining the improved bound. We will follow the notation introduced above for computing a $(1/r)$ -cutting. We first prove the following lemma.

Lemma 4.3. *Let Δ be a simplex in one of the cuttings Ξ_i . Let $P' \subseteq P$ be a subset of the marking points. Then $I(P' \setminus P_\Delta, C_\Delta) \leq a(|P'| + n_\Delta)$, for an absolute constant $a \geq 1$.*

Proof. For any point $p \in P' \setminus P_\Delta$, the dual plane p^* does not cross the simplex Δ . If p^* lies below (resp., above) Δ , and therefore below (resp., above) all points of C_Δ^* , then p lies in the common exterior (resp., common interior) of the circles in C_Δ . Since the complexity of the common exterior or common interior of n_Δ circles in the plane is $O(n_\Delta)$ [21], we obtain that $K(P' \setminus P_\Delta, C_\Delta) = O(n_\Delta)$. The claim now follows from Lemma 2.4. \square

We proceed now in a manner similar to the case of pseudo-segments. Applying (2) to Ξ_1 and noticing that $\sum_{\Delta \in \Xi_1} n_\Delta = n$, we have

$$\begin{aligned} I(P, C) &\leq \sum_{\Delta \in \Xi_1} (I(P_\Delta, C_\Delta) + I(P \setminus P_\Delta, C_\Delta)) \\ &\leq \sum_{\Delta \in \Xi_1} I(P_\Delta, C_\Delta) + \sum_{\Delta \in \Xi_1} a(m + n_\Delta) \quad (\text{by Lemma 4.3}) \\ &\leq \sum_{\Delta \in \Xi_1} \left(\sum_{\tau \in \Xi_\Delta^2} I(P_\tau, C_\tau) + I(P_\Delta \setminus P_\tau, C_\tau) \right) + a(n + cr_0^3 m), \end{aligned}$$

where c is the constant of proportionality in the bound for the size of the cutting Ξ_i . Setting $a' = acr_0^3$ and using Lemma 4.3 again to bound $I(P_\Delta \setminus P_\tau, C_\tau)$, we obtain

$$\begin{aligned} I(P, C) &\leq \sum_{\tau \in \Xi_2} \left(I(P_\tau, C_\tau) + a \left(n_\tau + \frac{m}{r_0} \right) \right) + an + a'm \\ &\leq \sum_{\tau \in \Xi_2} I(P_\tau, C_\tau) + 2an + a'm(1 + r_0^2), \end{aligned}$$

because $\sum_\tau n_\tau = n$ and $|\Xi_2| \leq c'r_0^6$. Continuing in this manner and recalling that for any simplex $\tau \in \Xi_{j-1}$, $m_\tau \leq m/r_0^{j-1}$, $|\Xi_j| \leq cr_0^{3j}$, and that $\sum_{\tau \in \Xi_j} n_\tau = n$, we obtain

$$\begin{aligned} I(P, C) &\leq \sum_{\tau \in \Xi_i} I(P_\tau, C_\tau) + ian + a'm \sum_{j=0}^{i-1} r_0^{2j} \\ &= \sum_{\tau \in \Xi} I(P_\tau, C_\tau) + O(n \log r + mr^2). \end{aligned}$$

Since $\sum_{\tau \in \Xi} n_\tau = n$, and $m_\Delta \leq m/r$ for all $\Delta \in \Xi$, we have

$$\begin{aligned} I(P, C) &\leq \sum_{\Delta \in \Xi} I(P_\Delta, C_\Delta) + O(n \log r + mr^2) \\ &\leq \sum_{\Delta \in \Xi} \sum_{\tau \in \Xi_\Delta} (I(P_\tau, C_\tau) + I(P_\Delta \setminus P_\tau, C_\tau)) + O(n \log r + mr^2) \\ &= \sum_{\tau \in \Xi} I(P_\tau, C_\tau) + O(n \log r + mr^2) \\ &= \sum_{\tau \in \Xi} O(K(P_\tau, C_\tau) + m_\tau + n_\tau) + O(n \log r + mr^2). \end{aligned} \tag{5}$$

Substituting the bound from Lemma 4.1 in (5) and using the inequalities $m_\tau \leq m/r$, $n_\tau \leq n/r^3$, and $|\Xi| = O(r^3)$, we obtain

$$\begin{aligned} I(P, C) &= \sum_{\tau \in \Xi} O \left(m_\tau^{2/3} n_\tau^{2/3} + n_\tau^{3/2+\varepsilon} \right) + O(mr^2 + n \log r) \\ &= O \left(m^{2/3} n^{2/3} r^{1/3} + n^{3/2+\varepsilon} / r^{3/2+3\varepsilon} + mr^2 + n \log r \right). \end{aligned}$$

Choose $r = \lceil n^{(5+6\varepsilon)/(11+18\varepsilon)} / m^{4/(11+18\varepsilon)} \rceil$. Note that $1 \leq r \leq m$ when $n^{1/3} \leq m \leq n^{5/4+3\varepsilon/2}$. If $m < n^{1/3}$ then $K(m, n) = O(n)$, as follows, e.g., from [13]. For $m > n^{1/3}$, the term mr^2 is dominated by $(n/r)^{3/2}$. Using this, substituting the value of r , and including the bounds obtained when r does not fall into the required range, we have, as in [9],

$$I(P, C) = O\left(m^{2/3}n^{2/3} + m^{6/11+3\varepsilon'}n^{9/11-\varepsilon'} + n \log n\right),$$

for $\varepsilon' = 8\varepsilon/(11(11+18\varepsilon)) > 0$, which can also be made arbitrarily small. Using (4), we obtain the following main result of the paper.

Theorem 4.4. *The maximum complexity of m distinct faces in an arrangement of n arbitrary circles in the plane is $O(m^{2/3}n^{2/3} + m^{6/11+3\varepsilon}n^{9/11-\varepsilon} + n \log n)$, for any arbitrarily small constant $\varepsilon > 0$.*

We can extend Theorem 4.4 to obtain an upper bound for $K(m, n, X)$, which takes into account the number X of intersecting pairs of circles in C . This is done exactly as in the case of pseudo-segments. That is, put $s = \lceil n^2/X \rceil$, and construct a $(1/s)$ -cutting of $\mathcal{A}(C)$ that consists of $O(s + s^2X/n^2) = O(s)$ cells, each crossed by at most n/s circles [20]. Apply Theorem 4.4 to bound the complexity of the marked faces within each cell, add up the resulting complexity bounds, and also add the complexity of the zones of the cell boundaries to account for faces not confined to a single cell (as in [13]). The complexity of the zones is

$$O(s) \cdot O\left(\lambda_4\left(\frac{n}{s}\right)\right) = O(s) \cdot O\left(\frac{n}{s} \cdot 2^{\alpha(n/s)}\right) = O(n \cdot 2^{\alpha(n)}).$$

This leads to the following result.

Theorem 4.5. *The maximum complexity of m distinct faces in an arrangement of n arbitrary circles in \mathbb{R}^2 with X intersecting pairs is $O(m^{2/3}X^{1/3} + m^{6/11+3\varepsilon}X^{4/11+2\varepsilon}n^{1/11-5\varepsilon} + n \log n)$, for an arbitrarily small constant $\varepsilon > 0$.*

The case of pairwise intersecting circles can be handled in a similar manner, using Lemma 4.2 to substitute the value of $K(P_\tau, C_\tau)$ in (5). Omitting the straightforward details, we obtain:

Theorem 4.6. *The maximum complexity of m distinct faces in an arrangement of n arbitrary pairwise-intersecting circles in the plane is $O\left(m^{2/3}n^{2/3} + m^{1/2}n^{5/6} \log^{1/2} n + n \log n\right)$.*

Remark. Note that for $m \leq (n^{1/3}/4^{\alpha(n)/3}) \log^{5/3} n$, the bound of $O(m^{3/5}n^{4/5}4^{\alpha(n)/5} + n)$, obtained by Clarkson et al. [13] on the complexity of m distinct faces in an arrangement of n circles in \mathbb{R}^2 , is slightly better than the ones stated in the above three theorems. For example, as mentioned above, $K(m, n) = O(n)$ if $m \leq n^{1/3}$.

5 The Case of Unit Circles

In this section we prove the following worst-case optimal bound on the maximum complexity of many faces in an arrangement of unit circles in the plane.

Theorem 5.1. *The combinatorial complexity of m distinct faces in an arrangement of n unit circles in the plane with X intersecting pairs is $O(m^{2/3}X^{1/3} + n)$. This bound is tight in the worst case.*

Proof. Let C be a collection of n unit circles in the plane and P a collection of m points marking (lying in the interior of) distinct faces in $\mathcal{A}(C)$. We aim to bound the total complexity of the marked faces. By Lemma 2.1, it suffices to bound the number $I = I(P, C)$ of incidences between the marked faces and the circles.

Note that $m = O(X + n)$, as the total number of faces in the arrangement is at most $2X + n + 1$, since the rightmost vertex of every bounded face is either one of the at most $2X$ arrangement vertices or one of the n rightmost points of the circles, and each point can be used only once in this manner. In the remainder of the proof we assume, without loss of generality, that the union of the circles of C is connected, so $X = \Omega(n)$ and $m = O(X)$. The analysis can easily be extended to the case in which the union is disconnected.

The analysis begins in a manner similar to that for the case of a line arrangement, as presented in [14], and its variant used in Section 3. For each circle $\gamma \in C$, we distinguish between faces touching γ “from the inside” and those that touch γ “from the outside.” We construct two separate (multi)graphs G_- and G_+ to encode the two types of face-circle incidences. The graphs are drawn as prescribed in [14], and briefly reviewed in Section 3.

More precisely, the graph G_- has P as its set of vertices. For each face-circle incidence along the “inside” of a circle $c \in C$ of $\mathcal{A}(C)$, fix a point of c on the face boundary to represent the incidence. Two consecutive representative points are connected along c and each of them is connected to the point marking the face it is incident to. The graph G_+ is constructed similarly, and encodes the “outer” incidences. The total number $|G_-| + |G_+|$ of edges in the two graphs is exactly I , by definition.

The analysis of Clarkson *et al.* [13] implies that the multiplicity of any edge of G_- is at most two. Actually, a stronger property holds: It is impossible for two distinct faces to touch three distinct unit circles on their interior sides (the argument is essentially the same as the one illustrated in Figure 5). Hence, arguing as in [14, 27], $|G_-|$ is $O(m^{2/3}X_-^{1/3} + m)$, where X_- is the number of edge crossings in G_- . Since, by construction, an edge crossing in G_- is also an intersection point of a pair of circles in C , and no two edge crossings can use the same intersection point of the same pair of circles, it follows that $X_- \leq 2X$ and $|G_-| = O(m^{2/3}X^{1/3} + m) = O(m^{2/3}X^{1/3})$ (the latter estimate follows from the fact that $m = O(X)$).

Handling the graph G_+ is somewhat more involved. It is shown in [13] that G_+ can be manipulated as follows. We first disregard the faces of the arrangement that lie outside all circles of C , if any of them are marked, because they can contribute at most $6n - 12$ (for $n \geq 3$) to $K(P, C)$ [21]. Each remaining marked face is enclosed by at least one circle of C and thus has diameter at most 2. We overlay the arrangement of the circles of C with the unit grid. Each circle meets the grid lines at most 8 times, so the total number of circle arcs that are part of the drawing of G_+ and are met by the grid lines is at most $8n$ —we remove the edges corresponding to these arcs from G_+ . It can now be shown (adapting the analysis given in [13]) that in what remains of G_+ , the edge multiplicities are all bounded by a constant. Hence we can apply an analysis similar to that above to conclude that $|G_+|$, and thus also the overall face-circle incidence count, are $O(m^{2/3}X^{1/3} + m + n) = O(m^{2/3}X^{1/3} + n)$

(the latter estimate follows, as above, from the fact that $m = O(X)$). This completes the proof of the upper bound.

To see that the bound is tight in the worst case, consider an arrangement of n lines which has m faces whose combined complexity is $\Theta(m^{2/3}n^{2/3} + n)$ (see [22] for details). We can then “bend” the lines slightly into large but congruent circles without changing the combinatorial structure of any face. This shows that the bound is worst-case tight when $X = \Theta(n^2)$. For smaller values of X , put $k = \lceil n^2/X \rceil$, and take k copies of the preceding construction, placed far away from each other, each involving $\lfloor n/k \rfloor$ circles and $\lfloor m/k \rfloor$ faces, of combined complexity (within a single copy)

$$\Theta\left(\left(\frac{m}{k}\right)^{2/3}\left(\frac{n}{k}\right)^{2/3} + \frac{n}{k}\right).$$

Together, we have at most n congruent circles and at most m faces in their arrangement. The number of intersecting pairs is at most $k \cdot (n/k)^2 \leq X$, and the overall complexity of all the marked faces is

$$k \cdot \Theta\left(\left(\frac{m}{k}\right)^{2/3}\left(\frac{n}{k}\right)^{2/3} + \frac{n}{k}\right) = \Theta\left(\frac{m^{2/3}n^{2/3}}{k^{1/3}} + n\right) = \Theta(m^{2/3}X^{1/3} + n).$$

□

References

- [1] P.K. Agarwal, T. Hagerup, R. Ray, M. Sharir, M. Smid and E. Welzl, Translating a planar object to maximize point containment: Exact and approximation algorithms, manuscript, 2001.
- [2] P.K. Agarwal, R. Klein, C. Knauer and M. Sharir, Computing the detour of polygonal curves, manuscript, 2001.
- [3] P.K. Agarwal, J. Matoušek and O. Schwarzkopf, Computing many faces in arrangements of lines and segments, *SIAM J. Comput.* 27 (1998), 491–505.
- [4] P.K. Agarwal, E. Nevo, J. Pach, R. Pinchasi, M. Sharir and S. Smorodinsky, Lenses in arrangements of pseudocircles and their applications, in preparation.
- [5] P.K. Agarwal and M. Sharir, Arrangements and their applications, in: *Handbook of Computational Geometry* (J.-R. Sack and J. Urrutia, eds.), North-Holland, Amsterdam, 2000, pp. 49–119.
- [6] P.K. Agarwal and M. Sharir, Pseudo-line arrangements: Duality, algorithms and applications, to appear in SODA 2002.
- [7] B. Aronov, H. Edelsbrunner, L. Guibas and M. Sharir, Improved bounds on the number of edges of many faces in arrangements of line segments, *Combinatorica* 12 (1992), 261–274.

- [8] B. Aronov and M. Sharir, Triangles in space, or: Building (and analyzing) castles in the air, *Combinatorica* 10 (1990), 137–173.
- [9] B. Aronov and M. Sharir, Cutting circles into pseudo-segments and improved bounds on incidences, *Discrete Comput. Geom.*, to appear.
- [10] R.J. Canham, A theorem on arrangements of lines in the plane, *Israel J. Math.* 7 (1969), 393–397.
- [11] T.M. Chan, On levels in arrangements of curves, *Proc. 41st IEEE Symp. Found. Comput. Sci.* (2000), 219–227.
- [12] B. Chazelle, Cutting hyperplanes for divide-and-conquer, *Discrete Comput. Geom.* 9 (1993), 145–158.
- [13] K. Clarkson, H. Edelsbrunner, L. Guibas, M. Sharir, and E. Welzl, Combinatorial complexity bounds for arrangements of curves and spheres, *Discrete Comput. Geom.* 5 (1990), 99–160.
- [14] T. Dey and J. Pach, Extremal problems for geometric hypergraphs, *Discrete Comput. Geom.* 19 (1998), 473–484.
- [15] H. Edelsbrunner, L. Guibas and M. Sharir, The complexity and construction of many faces in arrangements of lines and of segments, *Discrete Comput. Geom.* 5 (1990), 161–196.
- [16] B. Grünbaum, Arrangements of hyperplanes, *Congr. Numer.* 3 (1971), 41–106.
- [17] L. J. Guibas, M. Sharir, and S. Sifrony, On the general motion planning problem with two degrees of freedom, *Discrete Comput. Geom.*, 4 (1989), 491–521.
- [18] D. Halperin and M. Sharir, Improved combinatorial bounds and efficient techniques for certain motion planning problems with three degrees of freedom, *Comput. Geom. Theory Appl.* 1(5) (1992), 269–303.
- [19] S. Har-Peled, *The Complexity of Many Cells in the Overlay of Many Arrangements*, M.Sc. Thesis, Dept. Computer Science, Tel Aviv University, 1995. See also: Multicolor combination lemmas, *Comput. Geom. Theory Appl.* 12 (1999), 155–176.
- [20] S. Har-Peled, Constructing cuttings in theory and practice, *SIAM J. Comput.* 29 (2000), 2016–2039.
- [21] 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.
- [22] J. Pach and P.K. Agarwal, *Combinatorial Geometry*, Wiley-Interscience, New York, 1995.
- [23] J. Pach and M. Sharir, On the number of incidences between points and curves. *Combinatorics, Probability and Computing* 7 (1998), 121–127.

- [24] M. Sharir, Excess in arrangements of segments, *Inform. Process. Lett.* 58 (1996), 245–247.
- [25] M. Sharir and P.K. Agarwal, *Davenport Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.
- [26] J. Spencer, E. Szemerédi and W. T. Trotter, Unit distances in the Euclidean plane, in *Graph Theory and Combinatorics* (B. Bollobás, Ed.), Academic Press, New York, NY, 1984, pp. 293–303.
- [27] L. Székely, Crossing numbers and hard Erdős problems in discrete geometry, *Combinatorics, Probability and Computing* 6 (1997), 353–358.
- [28] E. Szemerédi and W. Trotter, Jr., Extremal problems in discrete geometry, *Combinatorica* 3 (1983), 381–392.

About Authors

Pankaj K. Agarwal is at the Department of Computer Science, Duke University, Durham, NC 27708-0129, USA; pankaj@cs.duke.edu. Boris Aronov is at the Department of Computer and Information Science, Polytechnic University, Brooklyn, NY 11201-3840, USA; aronov@ziggy.poly.edu. Micha Sharir is at the 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. sharir@cs.tau.ac.il.

Acknowledgments

Work on this paper has been supported by a joint grant from the U.S.-Israeli Binational Science Foundation. Work by Pankaj Agarwal was also supported by Army Research Office MURI grant DAAH04-96-1-0013, by a Sloan fellowship, by NSF grants ITR-333-1050, EIA-98-70724, EIA-99-72879, CCR-97-32787, and CCR-00-86013. Work by Boris Aronov has also been supported by NSF Grants CCR-99-72568 and ITR-00-81964. Work by Micha Sharir was also supported by NSF Grants CCR-97-32101 and CCR-00-98246, by a grant from the Israel Science Fund (for a Center of Excellence in Geometric Computing), and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.