Lenses in Arrangements of Pseudo-circles and their Applications*

Eran Nevo† János Pach‡ Rom Pinchasi§ Micha Sharir¶ Shakhar Smorodinsky‖

ABSTRACT

A collection of simple closed Jordan curves in the plane is called a family of *pseudo-circles* if any two of its members intersect at most twice. A closed curve composed of two subarcs of distinct pseudo-circles is said to be an *empty lens* if it does not intersect any other member of the family. We establish a linear upper bound on the number of empty lenses in an arrangement of n pseudo-circles with the property that any two curves intersect precisely twice. Enhancing this bound in several ways, and combining it with the technique of Tamaki and Tokuyama [16], we show that any collection of n pseudo-circles can be cut into $O(n^{3/2}(\log n)^{O(\alpha^s(n))})$ arcs so that any two intersect at most once, provided that the given pseudo-circles are x-monotone and admit an algebraic representation by three real parameters; here $\alpha(n)$ is the inverse Ackermann function, and s is a constant that depends on the algebraic degree of the representation of the pseudo-circles ($s = 2$ for circles and parabolas). For arbitrary collections of pseudo-circles, any two of which intersect twice, the number of necessary cuts reduces to $O(n^{4/3})$. As applications, we obtain improved bounds for the number of point-curve incidences, the complexity of a single level, and the complexity of many faces in arrangements of circles, pairwise intersecting pseudo-circles, parabolas, and families of homothetic copies of a fixed convex curve. We also obtain a variant of the Gallai-Sylvester theorem for arrangements of pairwise intersecting pseudo-circles, and a new lower bound for the number of distinct distances among n points in the plane under any simply-defined norm or convex distance function.

*Work on this paper by János Pach and Micha Sharir has been supported by NSF Grants CCR-97-32101 and CCR-00-98246, and by a joint grant from the U.S.–Israel Binational Science Foundation. Work by Micha Sharir and Shakhar Smorodinsky has also been supported by a grant from the Israeli Academy of Sciences for a Center of Excellence in Geometric Computing at Tel Aviv University, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.

†Institute of Mathematics, Hebrew University, Jerusalem, Israel; eranevo@math.huji.ac.il

‡Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA; pach@cims.nyu.edu

§Institute of Mathematics, Hebrew University, Jerusalem, Israel. Current address: Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA; room@math.mit.edu

¶School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel, and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA; sharir@cs.tau.ac.il

‖School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; shakhar@cs.tau.ac.il

1. INTRODUCTION

The *arrangement* of a finite collection C of geometric curves in $\mathbb{R}^2$, denoted as $\mathcal{A}(C)$, is the planar subdivision induced by C, whose vertices are the intersection points of the curves of C, whose edges are the maximal connected portions of curves in C not containing a vertex, and whose faces are maximal connected portions of $\mathbb{R}^2 \setminus \bigcup C$. Because of numerous applications and the rich geometric structure that they possess, arrangements of curves, especially of lines and segments, have been widely studied [2].

A family of unbounded Jordan curves (resp., arcs) is called a family of *pseudo-lines* (resp., *pseudo-segments*) if every pair of curves intersect in at most one point and they cross at that point. A collection C of closed Jordan curves is called a family of *pseudo-circles* if every pair of them cross at most twice. If the curves of C are graphs of continuous functions everywhere defined on the set of real numbers, such that every two cross at most twice, we call them *pseudo-parabolas.*[1] Although many combinatorial results on arrangements of lines and segments extend to pseudo-lines and pseudo-segments, as they rely on the fact that any two curves intersect in at most one point, they rarely extend to arrangements of curves in which a pair intersect in more than one point. In the last few years, progress has been made on analyzing arrangements of circles, pseudo-circles, or pseudo-parabolas by "cutting" the curves into subarcs so that the resulting set is a family of pseudo-segments and by applying results on pseudo-segments to the new arrangement; see [1, 5, 6, 7, 14, 16]. This paper continues this line of study—it improves a number of previous results on arrangements of

[1]For simplicity, we assume that every tangency counts as two crossings, i.e., if two pseudo-circles or pseudo-parabolas are tangent at some point, but they do not properly cross there, they do not have any other point in common.

pseudo-circles, and extends a few of the recent results on arrangements of circles (e.g., those presented in [5, 6, 14]) to arrangements of pseudo-circles.

Let C be a finite set of pseudo-circles in the plane. Let c and c' be two pseudo-circles in C, intersecting at two points u, v. A *lens* λ formed by c and c' is the union of two arcs, one of c and one of c', both delimited by u and v. If λ is a face of $\mathcal{A}(C)$, we call λ an *empty* lens; λ is called a *lens-face* if it is contained in the interiors of both c and c', and a *lune-face* if it is contained in the interior of one of them and in the exterior of the other. See Figure 1. (We ignore the case where λ lies in the exteriors of both pseudo-circles, because there can be only one such face in $\mathcal{A}(C)$.) Let $\mu(C)$ denote the number of empty lenses in C. A family of lenses formed by the curves in C is called *pairwise nonoverlapping* if the arcs forming any two of them do not overlap. Let $\nu(C)$ denote the maximum size of a family of nonoverlapping lenses in C. We define the *cutting number* of C, denoted by $\chi(C)$, as the minimum number of arcs into which the curves of C have to be cut so that any pair of resulting arcs intersect at most once (i.e., these arcs form a collection of pseudo-segments). In this paper, we obtain improved bounds on $\mu(C), \nu(C)$, and $\chi(C)$ for several special classes of pseudo-circles, and apply them to obtain bounds on various substructures of $\mathcal{A}(C)$.

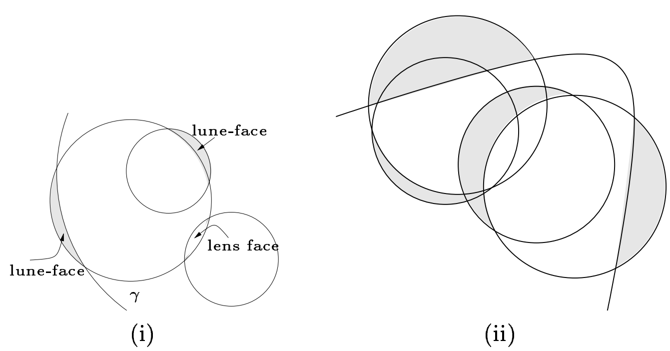

Figure 1: (i) A pseudo-circle γ supporting one lens-face and two lune-faces. (ii) A family of (shaded) nonoverlapping lenses.

Previous results. Tamaki and Tokuyama [16] proved that $\nu(C) = O(n^{5/3})$ for a family C of n pseudo-parabolas or pseudo-circles, and exhibited a lower bound of $\Omega(n^{4/3})$. In fact, their construction gives a lower bound on the number of empty lenses in an arrangement of circles or parabolas. Subsequently, improved bounds on $\mu(C)$ and $\nu(C)$ have been obtained for arrangements of circles. Alon *et al.* [5] and Pinchasi [14] proved that $\mu(C) = \Theta(n)$ for a set of n pairwise intersecting circles. If C is an arbitrary collection of circles, then $\nu(C) = O(n^{3/2+\varepsilon})$, for any $\varepsilon > 0$, as shown by Aronov and Sharir [6]. No better bound is known for the number of empty lenses in a family of circles. However, we have $\mu(C) = O(n^{4/3})$ for a set of n unit circles, though no superlinear lower bound is known for this special case.

The analysis in [16] shows that the cutting number $\chi(C)$ is proportional to $\nu(C)$ for collections of pseudo-parabolas or of pseudo-circles. Therefore one has $\chi(C) = O(n^{5/3})$ for pseudo-parabolas and pseudo-circles [16], and $\chi(C) = O(n^{3/2+\varepsilon})$ for circles. Using this bound on $\chi(C)$, Aronov and Sharir [6] proved that the maximum number of incidences between a set C of n circles and a set P of m points is $O(m^{2/3}n^{2/3} + m^{6/11+3\varepsilon}n^{9/11-\varepsilon} + m + n)$, for any $\varepsilon > 0$. Recently, following a similar but more involved argument, Agarwal *et al.* [1] proved a similar bound on the complexity of m distinct faces in an arrangement of n circles in the plane. An interesting consequence of the results in [5, 14] is the following generalization of the Sylvester-Gallai theorem: In an arrangement of pairwise intersecting circles, there always exists a vertex incident to at most three circles, provided that the number of circles is sufficiently large. For pairwise intersecting unit circles, the property holds when the number of circles is at least 5 [5, 14].

New results. In this paper we first obtain improved bounds on $\mu(C)$ and $\nu(C)$ for various special classes of pseudo-circles, and then apply these bounds to several problems involving arrangements of such pseudo-circles. Let C be a collection of n pseudo-parabolas such that any two have at least one point in common. We show that the number of tangencies in C is at most $2n - 4$ (for $n \geq 3$). In fact, we prove the stronger result that the tangency graph for such a collection C is bipartite and planar. Using this result, we prove that $\mu(C) = \Theta(n)$ for a set C of n pairwise intersecting pseudo-circles. Next, we show that $\nu(C) = O(n^{4/3})$ for collections C of pairwise intersecting pseudo-parabolas. Then, in Section 3, we study a somewhat artificial extension of the analysis of lenses to certain kinds of bichromatic lenses (where C is the disjoint union of two subsets A, B, and we only consider lenses formed by a curve in A and a curve in B). This extension is needed as a crucial component for the analysis of arbitrary arrangements of pseudo-circles. In the general case, we can no longer just assume the pseudo-circle property, and we need to make some additional assumptions on the geometric shape of the given curves. Specifically, we assume that they are all x-monotone, and that they admit a 3-parameter algebraic representation (a notion defined more precisely in Section 4). Then $\nu(C) = O(n^{3/2}(\log n)^{O(\alpha^s(n))})$, where $\alpha(n)$ is the inverse Ackermann function and s is a constant depending on the algebraic parametrization ($s = 2$ for circles and for parabolas of the form $y = ax^2 + bx + c$). This bound gives a slightly improved bound on $\nu(C)$, compared to the bound proved in [6], for a family of circles.

In Section 5, we present several applications of the above results. The results imply that $\chi(C) = O(n^{4/3})$ for a family of n pairwise intersecting pseudo-circles. If C is a family of n x-monotone pseudo-circles with 3-parameter algebraic representation, then $\chi(C) = O(n^{3/2}(\log n)^{O(\alpha^s(n))})$, where s is as above. The better bounds on the cutting number lead to improved bounds on the complexity of levels, on the number of incidences between points and pseudo-circles, and on the complexity of many faces, in arrangements of several classes of pseudo-circles, including the cases of circles, parabolas, pairwise-intersecting pseudo-circles, and homothetic copies of a fixed convex curve. Our results also yield a new lower bound for the number of distinct distances in the plane under norms or convex distance functions other than the Euclidean norm. Finally, we obtain a generalization of the Sylvester-Gallai-type results of [5] to the case of pairwise-intersecting pseudo-circles. The exact bounds and detailed results are stated in Section 5.

We close the introduction by mentioning recent work by Chan [7], in which nontrivial and improved bounds for $\chi(C)$ are obtained for families C of graphs of polynomials of any constant maximum degree. Interestingly, although his analysis is based on cutting pseudo-parabolic arcs into pseudo-segments, it does not (and probably cannot) exploit the new bounds obtained in this paper. In the full version of the paper, we also present applications of his bounds for obtaining improved bounds on the number of incidences between points and graphs of polynomials of a fixed degree.

2. THE CASE OF PAIRWISE INTERSECTING PSEUDO-CIRCLES

2.1 Tangencies in arrangements of pairwise intersecting pseudo-parabolas

Before going into the analysis of the general case of pairwise intersecting pseudo-circles, which requires several topological transformations and reductions of a more technical nature, we begin with a more specialized result, which is interesting in its own right, and which constitutes the main tool in the derivation of all the other results of this paper.

Let Γ be a set of n pairwise intersecting *pseudo-parabolas*, i.e., graphs of totally defined continuous functions, each pair of which intersect, either in exactly two crossing points or in exactly one point of tangency. Assume also that no three of these curves have a point in common. Let T denote the set of all pairs of tangent curves in Γ. We regard T as the edge set of a 'tangency graph' $G = (\Gamma, T)$.

THEOREM 2.1. *The tangency graph $G = (\Gamma, T)$, as defined above, is a bipartite planar graph. Consequently, $|T| \leq 2n - 4$, for $n \geq 3$.*

Proof: We first show that G is bipartite. A pseudo-parabola in Γ is called *lower* (resp., *upper*) if it forms a tangency with another curve that lies above (resp., below) it. We observe that a curve $\gamma \in \Gamma$ cannot be both upper and lower, or else the two other curves forming the respective tangencies with γ would have to be disjoint. Hence, G is bipartite.

The drawing rule. Let ℓ be some fixed vertical line that lies to the left of all the vertices of $\mathcal{A}(\Gamma)$. We draw G in the plane as follows. Each $\gamma \in \Gamma$ is represented by the point $\gamma^* = \gamma \cap \ell$. Each edge $(\gamma_1, \gamma_2) \in G$ is drawn as a y-monotone curve that connects the points γ_1^*, γ_2^*. This path has to navigate to the left or to the right of each of the intermediate vertices δ^* between γ_1^* and γ_2^* along ℓ.

The rule for drawing an edge (γ_1^*, γ_2^*) is: Assume that γ_1^* lies below γ_2^* along ℓ. Let $W(\gamma_1, \gamma_2)$ denote the *left wedge* formed by γ_1 and γ_2, consisting of all points that lie above γ_1 and below γ_2 and to the left of the tangency between them. Let $\delta \in \Gamma$ be a curve so that δ^* lies on ℓ between γ_1^* and γ_2^*. The curve δ has to exit $W(\gamma_1, \gamma_2)$. If its first exit point (i.e., its leftmost intersection with $\partial W(\gamma_1, \gamma_2)$) lies on γ_1 then we draw (γ_1^*, γ_2^*) to pass to the right of δ^*. Otherwise we draw it to pass to the left of δ^*. Except for these requirements, the edge (γ_1^*, γ_2^*) can be drawn in an arbitrary y-monotone manner. See Figure 2 for an illustration.

LEMMA 2.2. *Suppose that the following conditions hold for each quadruple $\gamma_1, \gamma_2, \gamma_3, \gamma_4$ of pseudo-parabolas in Γ, listed in the order of their intersections with the line ℓ:*
(a) If (γ_1, γ_4) and (γ_2, γ_3) are edges of G then the drawing of (γ_1^*, γ_4^*) does not pass to the left of γ_2^* and to the right of γ_3^*, nor does it pass to the right of γ_2^* and to the left of γ_3^*.
(b) If (γ_1, γ_3) and (γ_2, γ_4) are edges of G then, if the drawing of (γ_1^*, γ_3^*) passes to the left (resp., to the right) of γ_2^* then the drawing of (γ_2^*, γ_4^*) passes to the right (resp., to the left) of γ_3^*.
Then G is planar.

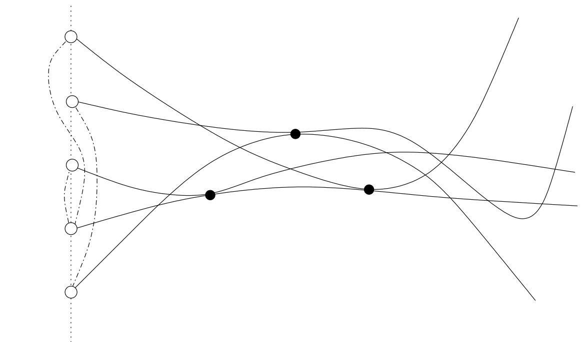

Figure 2: Illustrating the drawing rule.

Proof: See Figure 3 for the configurations allowed and disallowed by conditions (a) and (b). We show that the drawings of each pair of edges of G cross an even number of times. This, combined with Hanani-Tutte's theorem [17] (see also [9, 12]), implies that G is planar. Clearly, it suffices to check this for pairs of edges for which the y-projections of their drawings have a nonempty intersection. In this case, the projections are either *nested*, as in case (a) of the condition in the lemma, or *partially overlapping*, as in case (b).

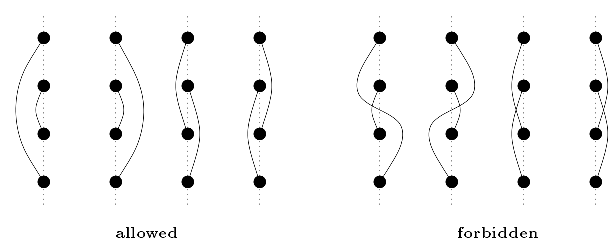

Figure 3: The allowed and forbidden configurations in conditions (a) and (b).

Consider first a pair of edges $e = (\gamma_1, \gamma_4)$ and $e' = (\gamma_2, \gamma_3)$, with nested projections, as in case (a). Regard the drawing of e as the graph of a continuous partial function $x = e(y)$, defined over the interval $[\gamma_1^*, \gamma_4^*]$, and similarly for e'. Part (a) of the condition implies that either e is to the left of e' at both γ_2^* and γ_3^*, or e is to the right of e' at both these points. Since e and e' correspond to graphs of functions that are defined and continuous over $[\gamma_2^*, \gamma_3^*]$, it follows that e and e' intersect in an even number of points.

Consider next a pair of edges $e = (\gamma_1, \gamma_3)$ and $e' = (\gamma_2, \gamma_4)$, with partially overlapping projections, as in case (b). Here, too, part (b) of the condition implies that either e is to the left of e' at both γ_2^* and γ_3^*, or e is to the right of e' at both these points. This implies, as above, that e and e' intersect in an even number of points. □

We next show that the conditions in Lemma 2.2 do indeed hold for our drawing of G. The argument is that, in any forbidden pattern of lemma 2.2, two of the curves $\gamma_1, \ldots, \gamma_4$ must be disjoint, which contradicts our assumption.

Lemma 2.3. *Let $\gamma_1, \gamma_2, \gamma_3, \gamma_4$ be four pseudo-parabolas in Γ, whose intercepts with ℓ appear in this increasing order, and suppose that (γ_1, γ_4) and (γ_2, γ_3) are tangent pairs. Then it is impossible that the first exit points of γ_2 and γ_3 from the wedge $W(\gamma_1, \gamma_4)$ are at opposite sides of the wedge.*

Proof: Suppose to the contrary that such a configuration exists. Then, except for the respective points of tangency, γ_3 always lies above γ_2, and γ_4 always lies above γ_1. This implies that if the first exit point of γ_2 from $W(\gamma_1, \gamma_4)$ lies on γ_4, then the first exit point of γ_3 also has to lie on γ_4, contrary to assumption. Hence, the first exit point of γ_2 lies on γ_1 and the first exit point of γ_3 lies on γ_4. See Figure 4. Let v_{14} denote the point of tangency of γ_1 and γ_4. We distinguish between two cases:
(a) *γ_2 passes below v_{14} and γ_3 passes above v_{14}:* See Figure 4(i). In this case, the second intersection point of γ_1 and γ_2 must lie to the right of v_{14}, for otherwise γ_2 could not have passed below v_{14}. Similarly, the second intersection point of γ_3 and γ_4 also lies to the right of v_{14}. This also implies that γ_2 and γ_4 do not intersect to the left of v_{14}, and that γ_1 and γ_3 also do not intersect to the left of v_{14}. Let u_{13} (resp., u_{24}) denote the leftmost intersection point of γ_1 and γ_3 (resp., of γ_2 and γ_4), both lying to the right of v_{14}. Suppose, without loss of generality, that u_{13} lies to the left of u_{24}. In this case, the second intersection of γ_1 and γ_2 must lie to the right of u_{13}. Indeed, otherwise γ_2 would become "trapped" inside the wedge $W(\gamma_1, \gamma_3)$ because γ_2 cannot cross γ_3 and it has already crossed γ_1 at two points. The second intersection of γ_3 and γ_4 occurs to the left of u_{13}. Now, γ_2 and γ_4 cannot intersect to the left of u_{13}: γ_2 does not intersect γ_4 to the left of its first exit w_{12} from $W(\gamma_1, \gamma_4)$. To the right of w_{12} and to the left of u_{13}, γ_2 remains below γ_1, which lies below γ_4. Finally, to the right of u_{13}, γ_2 lies below γ_3, which lies below γ_4. This implies that γ_2 cannot intersect γ_4 at all, a contradiction, which shows that case (a) is impossible.

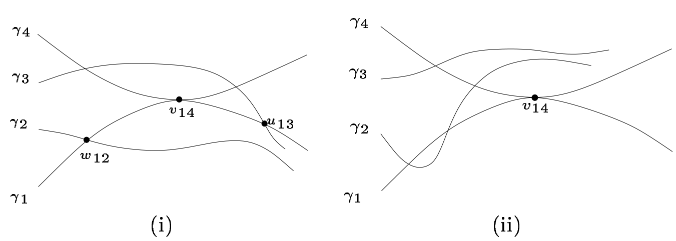

Figure 4: Edges of G with nested projections: (i) γ_2 passes below v_{14} and γ_3 passes above v_{14}; (ii) both γ_2 and γ_3 pass on the same side of v_{14}.

(b) *Both γ_2 and γ_3 pass on the same side of v_{14}:* Without loss of generality, assume that they pass above v_{14}. See Figure 4(ii). Then γ_2 must cross γ_1 again and then cross γ_4, both within $\partial W(\gamma_1, \gamma_4)$. In this case, γ_3 cannot cross γ_1 to the left of v_{14}, because to do so it must first cross γ_4 again, and then it would get 'trapped' inside the wedge $W(\gamma_2, \gamma_4)$. But then γ_1 and γ_3 cannot intersect at all: We have argued that they cannot intersect to the left of v_{14}. To the right of this point, γ_3 lies above γ_2, which lies above γ_1. This contradiction rules out case (b), and thus completes the proof of the lemma. □

Lemma 2.4. *Let $\gamma_1, \gamma_2, \gamma_3, \gamma_4$ be four curves in Γ, whose intercepts with ℓ appear in this increasing order, and suppose that (γ_1, γ_3) and (γ_2, γ_4) are tangent pairs. Then it is impossible that the first exit point of γ_2 from the wedge $W(\gamma_1, \gamma_3)$ and the first exit point of γ_3 from the wedge $W(\gamma_2, \gamma_4)$ both lie on the bottom sides of the respective wedges, or both lie on the top sides.*

Proof. Suppose to the contrary that such a configuration exists. By symmetry, we may assume, without loss of generality, that both exit points lie on the bottom sides. That is, the exit point u_{12} of γ_2 from $W(\gamma_1, \gamma_3)$ lies on γ_1 and the exit point u_{23} of γ_3 from $W(\gamma_2, \gamma_4)$ lies on γ_2. See Figure 5. By definition, γ_2 and γ_3 do not intersect to the left of u_{12}. So, u_{23} occurs to the right of u_{12} and, in fact, also to the right of the second intersection point of γ_1 and γ_2. Again, by assumption, γ_3 and γ_4 do not intersect to the left of u_{23}. Hence γ_1 and γ_4 also do not intersect to the left of u_{23}, because γ_1 lies below γ_3. But then γ_1 and γ_4 cannot intersect at all, because to the right of u_{23}, γ_4 lies above γ_2, which lies above γ_1. This contradiction completes the proof of the lemma. □

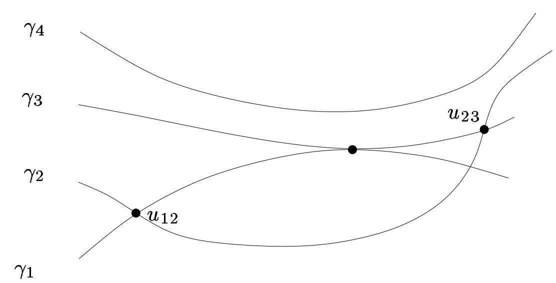

Figure 5: Edges of G with partially overlapping projections.

Lemmas 2.3 and 2.4 show that the conditions in Lemma 2.2 hold, so G is planar and bipartite and thus has at most $2n-4$ edges, for $n \geq 3$. This completes the proof of Theorem 2.1. □

2.2 Empty lenses in arrangements of pairwise intersecting pseudo-circles

We next extend Theorem 2.1 to families C of n pseudo-circles, any two of which intersect each other in two points. The extension is rather technical, and aims to reduce this case to the case of pseudo-parabolas. Here is a sketch of the process; some of the details are left out due to lack of space.

We refer to the interiors of the pseudo-circles in C as *pseudo-disks*. Using planarity, we first show that, among any five pseudo-disks bounded by the elements of C, there are at least three that have a point in common: A configuration that violates this property leads to an impossible plane drawing of K_5.

We then apply a topological variant of Helly's theorem [10], due to Molnár [13], which asserts that any finite family of at least three simply connected regions in the plane has a nonempty simply connected intersection, provided that any two of its members have a connected intersection and any three have a nonempty intersection. Consequently, the intersection of any subfamily of pseudo-disks bounded by elements of C is either empty or simply connected and hence contractible.

For any $p \geq q \geq d+1$, a finite collection F of open regions in d-space is said to have the (p,q)-*property* if among any p members of F there are q that have a point in common. Alon et al. [3] have recently extended a celebrated result of Alon and Kleitman [4], by showing that there exists a constant $k = k(p,q,d)$ such that, if F satisfies the (p,q)-property and the intersection of every subfamily of F is either empty or contractible, then there are k points so that every member of F contains at least one of them. Such a set is often called a k-element *transversal* or *piercing set*.

All this implies that there is an absolute constant k such that any family of pseudo-disks bounded by pairwise intersecting pseudo-circles can be pierced by at most k points.

Fix a set $\mathcal{O} = \{o_1, o_2, \ldots, o_k\}$ piercing the family of all pseudo-disks bounded by the elements of C. Let C_i consist of all elements of C that contain o_i, for $i = 1, 2, \ldots, k$.

It suffices to derive an upper bound on the number of empty lenses formed by pairs of pseudo-circles belonging to the same class C_i, and on the number of empty lenses formed by pairs of pseudo-circles belonging to two fixed classes C_i, C_j. Using an inversion with respect to the piercing point of one of the two classes, one can show that the second task can be reduced to the first one (details omitted here), so we focus our attention on lenses formed within a fixed class C_i.

Consider now a fixed class C_i. We next show that, by deforming the plane without changing the combinatorial structure of the arrangement of C_i, we can transform the elements of C_i into sets that are star-shaped with respect to o_i. This is accomplished using a topological sweeping argument, akin to that due to Hershberger and Snoeyink [11]. Specifically, we show that the union of any subset of pseudo-disks bounded by the pseudo-circles in C_i is simply connected. This allows us to draw a curve $\bar{r}$ that starts at o_i and extends to infinity, crossing each $c \in C_i$ exactly once; see Figure 6(i). We then sweep $\bar{r}$ around o_i. This is based on the following crucial property, whose proof is omitted here: There always exist two consecutive edges of $\mathcal{A}(C_i)$ crossed by $\bar{r}$ which have a common endpoint w counterclockwise to $\bar{r}$ (such as the edges e_1, e_2 in Figure 6(i)). This allows us to advance $\bar{r}$ past w, and to continue the sweep in this manner until we perform a complete revolution about o_i. We then simulate the sweep by replacing $\bar{r}$ by a straight ray r emanating from o_i. Each step of sweeping $\bar{r}$ past a vertex of $\mathcal{A}(C_i)$ is simulated by a swap of the modified pseudo-circles at an appropriate position of r. We omit further details due to lack of space; an illustration of the process is depicted in Figure 6.

We then regard each pseudo-circle in C_i as a graph of a function in polar coordinates; the collection of these graphs (with an appropriate 'stretching' of the θ-axis to the full line) is the desired collection of pairwise-intersecting pseudo-parabolas. We still need to go through a few technical steps: First, we dispose of cases where all the pseudo-parabolas pass through the same pair of points (they form a so-called *pencil*), or where two empty lenses share an arc. Then we deform the arrangement slightly, so as to ensure that no three pseudo-parabolas have a common point, without destroying any empty lens. Finally, we shrink each empty lens to a point of tangency between the respective curves, thereby reducing the setup to that assumed in Theorem 2.1. Applying all these steps, and handling in this manner all pairs of subfamilies pierced by two respective points of $\mathcal{O}$, we obtain the main technical result of the paper:

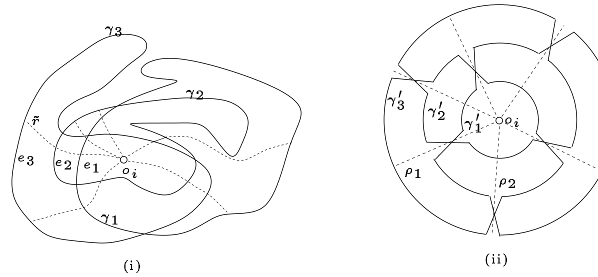

Figure 6: Converting C into a star-shaped family by a counterclockwise topological sweep: (i) original curves, (ii) transformed curves.

THEOREM 2.5. *The number of empty lenses in an arrangement of n pairwise intersecting pseudo-circles is $O(n)$.*

2.3 Pairwise nonoverlapping lenses

Let C be a family of n pairwise-intersecting pseudo-parabolas or pseudo-circles, and let L be a family of pairwise nonoverlapping lenses in $\mathcal{A}(C)$. In this subsection, we obtain the following bound for the size of L.

THEOREM 2.6. *Let C be a family of n pairwise-intersecting pseudo-parabolas or pseudo-circles. Then the maximum size of a family of pairwise nonoverlapping lenses in $\mathcal{A}(C)$ is $O(n^{4/3})$.*

Proof: We only consider the case of pseudo-parabolas; the other case can be reduced to this case, using the analysis given in the preceding subsections. The proof proceeds through the following sequence of lemmas.

LEMMA 2.7. *Let C and L be as above, and assume further that the lenses in L have pairwise disjoint interiors. Then $|L| = O(n)$.*

Proof: For each lens $\lambda \in L$, let σ_λ denote the number of edges of $\mathcal{A}(C)$ that lie in the interior of λ (i.e., the region bounded by λ), and set $\sigma_L = \sum_{\lambda \in L} \sigma_\lambda$. We prove the lemma by induction on the value of σ_L. If $\sigma_L = 0$, i.e., all lenses in L are empty, then the lemma follows from Theorem 2.5. Suppose $\sigma_L \geq 1$.

Let λ_0 be a lens in L with $\sigma_{\lambda_0} \geq 1$, and let K_0 be the interior of λ_0. Let $\gamma, \gamma' \in C$ be the pseudo-parabolas forming λ_0, and let $\delta \subset \gamma$ and $\delta' \subset \gamma'$ be the two arcs forming λ_0. Let $\zeta \in C$ be a curve that intersects K_0; clearly, $\zeta \in C$ cannot be fully contained in the interior of K_0. Therefore, up to symmetry, there are two possible kinds of intersection between ζ and λ_0:

(i) $|\zeta \cap \delta'| = 2$, and $\zeta \cap \delta = \emptyset$.

(ii) ζ intersects both δ and δ'. In this case, either ζ intersects each of δ, δ' at a single point, or intersects each of them at two points.

Suppose K_0 contains a curve $\zeta \in C$ of type (i). Let λ_1 be the lens formed by ζ and γ'. Let L' be the family obtained from L by replacing λ_0 with λ_1. See Figure 7(a). The interior of λ_1 is strictly contained in K_0 and contains fewer edges of $\mathcal{A}(C)$ than K_0, so $\sigma_{L'} < \sigma_L$. The lemma now holds by the induction hypothesis. We may thus assume that no curve of type (i) crosses K_0, so all these curves are of type (ii). In this case, we can shrink K_0 to an empty lens between γ and γ'. For example, we can replace δ' by an arc that proceeds parallel to δ and outside K_0, and connects two points on γ' close to the endpoints of δ', except for a small region where the new δ' crosses δ twice, forming a small empty lens; see Figure 7(b). Since only curves of type (ii) cross K_0, it is easy to check that C is still a collection of pairwise-intersecting pseudo-parabolas. The lens λ_0 is replaced by the new lens λ_1 formed between δ and the modified δ'. Since $\sigma_{\lambda_1} = 0$, we have reduced the size of σ_L, and the claim follows by the induction hypothesis. This completes the proof of the lemma. □

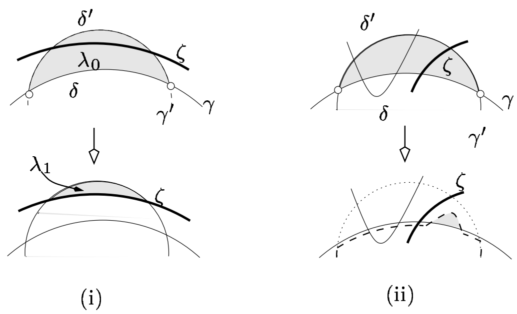

Figure 7: Replacing λ_0 by s 'smaller' lens: (a) The case of a type (i) curve ζ. (b) Shrinking λ_0 to an empty lens when it is crossed only by type (ii) curves.

A pair (λ, λ') of lenses in L is said to be *crossing* if an arc of λ intersects an arc of λ'. (Note that a pair of lenses may be nonoverlapping and yet crossing.) A pair (λ, λ') of lenses in L is said to be *nested* if both arcs of λ' are fully contained in the interior of λ. Let X (resp., Y) be the number of crossing (resp., nested) pairs of lenses in L.

LEMMA 2.8. *Let C, L, X and Y be as above. Then*

$$|L| = O(n + X + Y). \qquad (1)$$

Proof: If L contains a pair of crossing or nested lenses, remove one of them from L. This decreases $|L|$ by 1 and $X + Y$ by at least 1, so if (1) holds for the new L, it also holds for the original set. Repeat this step until L has no pair of crossing or nested lenses. Every pair of lenses in (the new) L must have disjoint interiors. The lemma is then an immediate consequence of Lemma 2.7. □

We next derive upper bounds for X and Y. The first bound is easy:

LEMMA 2.9. $X = O(n^2)$.

Proof: We charge each crossing pair of lenses (λ, λ') in L to an intersection point of some arc bounding λ and some arc bounding λ'. Since the lenses of L are pairwise nonoverlapping, it easily follows that such an intersection point can be charged at most $O(1)$ times (it is charged at most once if the crossing occurs at a point in the relative interior of arcs of both lenses), and this implies the lemma. □

We next derive an upper bound for Y, with the following twist:

LEMMA 2.10. *Let $k < n$ be some threshold integer parameter, and suppose that each lens of L is crossed by at most k curves of C. Then $Y = O(k|L|)$.*

Proof: Fix a lens $\lambda' \in L$. Let $\lambda \in L$ be a lens that contains λ' in its interior, i.e., (λ, λ') is a nested pair. Pick any point q on λ' (e.g., its left vertex), and draw an upward vertical ray ρ from q; ρ must cross the upper boundary of λ. It cannot cross more than k other curves before hitting λ because any such curve has to cross λ. Because of the nonoverlap of the lenses of L, the crossing point $\rho \cap \lambda$ uniquely identifies λ (unless it is a vertex of λ, in which case there is a constant number of possible lenses λ). This implies that at most $O(k)$ lenses in L can contain λ', thereby implying that the number of nested pairs of lenses in L is $O(k|L|)$. □

We are now ready to complete the proof of Theorem 2.6. Let L be a family of pairwise nonoverlapping lenses in $\mathcal{A}(C)$. Let k be any fixed threshold parameter, which will be determined later. First, remove from L all lenses which are intersected by at least k curves of C. Any such lens contains points of intersection of at least k pairs of curves of C. Since these lenses are pairwise nonoverlapping, and there are $n(n-1)$ intersection points, the number of such 'heavily intersected' lenses is at most $O(n^2/k)$. So, we may assume that each remaining lens in L is crossed by at most k curves of C.

Draw a random sample R of curves from C, where each curve is chosen independently with probability p, to be determined shortly. The expected number of curves in R is np, and the expected size $|L'|$ of the subset L' of lenses of L that materialize in R is $|L|p^2$ (where L refers to the set after removal, within $\mathcal{A}(C)$, of the 'heavily intersected' lenses). The expected number Y' of nested pairs (λ, λ') in L' is Yp^4 (any such pair must be counted in Y for the whole arrangement, and its probability of materializing in R is p^4). Similarly, the expected number X' of crossing pairs (λ, λ') in L' is Xp^4. By Lemmas 2.8 (applied to $\mathcal{A}(R)$), 2.9, and 2.10, we have

$$|L|p^2 \le c(np + n^2p^4 + k|L|p^4),$$

for an appropriate constant c. That is, we have

$$|L|(1 - ckp^2) \le c\left(\frac{n}{p} + n^2p^2\right).$$

Choose $p = 1/(2ck)^{1/2}$, to obtain $|L| = O(nk^{1/2} + n^2/k)$. Adding the bound on the number of heavy lenses, we conclude that the size of the whole L is

$$|L| = O\left(nk^{1/2} + \frac{n^2}{k}\right).$$

By choosing $k = n^{2/3}$, we obtain $|L| = O(n^{4/3})$, thereby completing the proof of the theorem. □

3. BICHROMATIC LENSES

In this section we consider the following *bichromatic* extensions of the problems involving empty lenses and pairwise nonoverlapping lenses. These extensions are somewhat artificial, but they are required as a main technical tool in the analysis of the general case, treated in the next section, where not all pairs of the given pseudo-circles necessarily intersect.

We begin the study by assuming that we have a collection C of n pseudo-parabolas, which is the disjoint union of two subsets A, B, so that each pseudo-parabola of A intersects every pseudo-parabola of B twice; a pair of pseudo-parabolas within A (or B) may be disjoint. A lens formed by a pseudo-parabola belonging to A and by one belonging to B is called *bichromatic*.

THEOREM 3.1. *Let C be a collection of n pseudo-parabolas, so that C is the disjoint union of two subsets A, B, such that each element of A intersects every element of B. Then the number of bichromatic empty lenses in $\mathcal{A}(C)$ is $O(n)$.*

Proof: It suffices to estimate the number of empty bichromatic lenses formed by some $a \in A$ and by some $b \in B$, so that a lies above b within the lens. The complementary set of empty bichromatic lenses is analyzed in a fully symmetric manner.

We apply the following pruning process to the curves of C. Let a, a' be two disjoint curves in A, so that a' lies fully below a. Then no empty bichromatic lens of the kind under consideration can be formed between a and any pseudo-parabola $b \in B$, because then a' and b would have to be disjoint; see Figure 8(i). Hence, we may remove a from A without affecting the number of empty bichromatic lenses under consideration. A fully symmetric process (depicted in Figure 8(ii)) prunes away curves from B.

We keep applying this pruning process untill all pairs of remaining curves in $A \cup B$ intersect each other. By Theorem 2.1, the number of empty lenses in $\mathcal{A}(A \cup B)$ is $O(n)$. As discussed above, this completes the proof of the theorem. $\square$

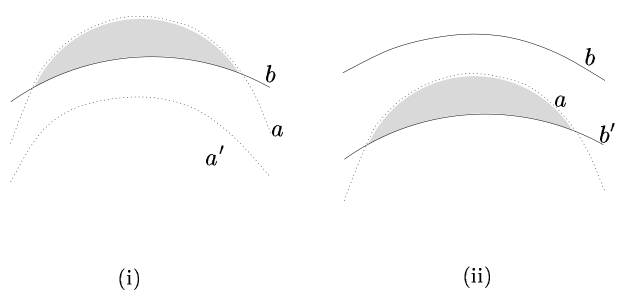

Figure 8: Discarding one of the nested pseudo-parabolas: (i) a is discarded, (ii) b' is discarded.

THEOREM 3.2. *Let C be a collection of n pseudo-parabolas, so that C is the disjoint union of two subsets A, B, such that each element of A intersects every element of B. Let L be a family of pairwise nonoverlapping bichromatic lenses in $\mathcal{A}(C)$. Then the size of L is $O(n^{4/3})$.*

Proof: The proof proceeds by adapting the analysis given in Section 2.3. In fact, the only modification required in the proof is that of lemma 2.7. The modified variant is:

LEMMA 3.3. *Let C and L be as in the theorem, and assume further that the lenses in L have pairwise disjoint interiors. Then $|L| = O(n)$.*

Proof: As in the proof of Theorem 3.1, it suffices to estimate the number of lenses in L that are formed by some $a \in A$ and by some $b \in B$, so that a lies above b within the lens. Let a, a' be two disjoint curves in A, so that a' lies fully below a. We argue that a can be pruned away, as follows. Let $\lambda \in L$ be a lens formed by a and by some $b \in B$. Let $\delta \subset b$ be the arc of b forming λ. Clearly, b must also intersect a', and the two points of intersection must lie on δ, since $b \setminus \delta$ lies fully above a and thus above a'. Replace λ by the lens λ', formed between a' and b. Since the lenses in L have disjoint interiors, λ' is not a member of L, and, after the replacement, L is still a family of bichromatic lenses with pairwise disjoint interiors, of the same size. Hence, by applying this replacement rule to each lens in L formed along a, we may prune away a without affecting the size of L. We keep applying this pruning rule, as well as the symmetric rule for pruning away curves of B, until every pair of remaining curves intersect each other twice. The lemma now follows from Theorem 3.1. $\square$

By plugging the modified lemma into the analysis in the preceding section, Theorem 3.2 follows. $\square$

3.1 The case of pseudo-circles

Let C be a collection of bounded closed x-monotone pseudo-circles. For each $c \in C$, denote by λ_c (resp., ρ_c) the leftmost (resp., rightmost) point of c. To simplify the analysis, we assume that the 2n points λ_c, ρ_c, for $c \in C$ are all well defined, and that their x-coordinates are all distinct. The analysis can be easily extended to handle degenerate cases as well.

Fix a curve $c \in C$. The points λ_c, ρ_c partition c into two complementary x-monotone arcs c^+, c^-, so that c^+ lies above c^-. Extend c^+ (resp., c^-) to an unbounded curve by two downward-directed (resp., upward-directed) rays emanating from λ_c and from ρ_c, so that the absolute values of the slopes of all these rays is some fixed, sufficiently large value, and so that the rays extend the curves c^+, c^- into graphs of continuous totally-defined functions. See Figure 9 for an illustration. We have thus obtained a new collection C^* of $2n$ unbounded x-monotone curves. The following lemma is easy to establish, by examining the few cases that can arise:

LEMMA 3.4. *C^* is a collection of pseudo-parabolas.*

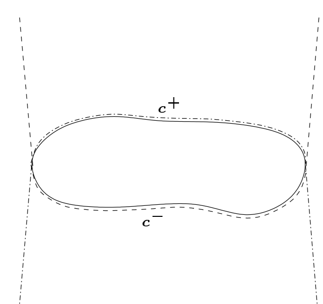

Figure 9: Transforming an x-monotone pseudo-circle into two graphs of totally defined continuous functions.

Theorems 3.1 and 3.2 thus yield the following extensions:

THEOREM 3.5. *Let A and B be two disjoint subsets of C^* with the property that each curve in A intersects every curve in B. Then:*

(i) The number of bichromatic empty lenses in $\mathcal{A}(A \cup B)$ is $O(n)$.
(ii) The maximum size of a family of pairwise nonoverlapping bichromatic lenses in $\mathcal{A}(A \cup B)$ is $O(n^{4/3})$.

4. THE GENERAL CASE

Here we consider the case where C is a collection of n pseudo-circles, not every pair of which intersect. It follows from the construction given in [16] that the number of empty lenses in C can be $\Omega(n^{4/3})$. We conjecture that this is also an upper bound for the number of empty lenses. We obtain a weaker upper bound for the potential larger quantity $\nu(C)$, defined, as above, to be the maximal size of a family of pairwise nonoverlapping lenses in C.

Tamaki and Tokuyama [16] have shown that $\nu(C) = O(n^{5/3})$. Aronov and Sharir [6] have recently improved this bound to $O(n^{3/2+\varepsilon})$, for any $\varepsilon > 0$, for the case of circles. Theorem 2.6 asserts that $\nu(C) = O(n^{4/3})$ for the case of pairwise intersecting pseudo-circles or pseudo-parabolas. Unfortunately, at the moment we can only obtain the improved bound of [6] (in fact, as a consequence of Theorem 2.6, we will even be able to slightly improve it further) in the following special case.

We say that C has a *3-parameter algebraic representation*, if C is a finite subset of some infinite family $\mathbf{C}$ of x-monotone curves, so that every curve in $\mathbf{C}$ can be represented by a triple of real parameters (ξ, η, ζ), such that, for each curve $\gamma_0 \in \mathbf{C}$, the locus of all curves in $\mathbf{C}$ whose top boundary is tangent to the top boundary γ_0^+ of γ_0 is a 2-dimensional surface patch which is a semialgebraic set of constant description complexity. Moreover, this surface partitions 3-space into two subsets, one consisting of all points that represent pseudo-circles of $\mathbf{C}$ whose top boundary intersects γ_0^+, and the other consisting of points representing curves whose top boundaries are disjoint from γ_0^+. Fully analogous conditions are assumed to hold for each of the three other combinations of top vs. bottom boundaries. Moreover, for any point $q \in \mathbb{R}^2$, the locus of all curves of $\mathbf{C}$ that pass through q is a 2-dimensional semi-algebraic surface of constant description complexity.

Three important classes of pseudo-circles that admit 3-parameter algebraic representations are the class of circles, the class of parabolas, given by equations of the form $y = ax^2 + bx + c$, and the class of homothetic copies of any fixed convex curve of constant description complexity.

Suppose then that C is a collection of n x-monotone curves that admit a 3-parameter algebraic representation, as above. Let C^* be the collection of the $2n$ extended top and bottom boundaries of the curves of C. Let L be a family of pairwise nonoverlapping lenses in C. Only $O(n)$ of them can contain the leftmost or rightmost point of any curve in C, so we may assume that each arc of every lens in L is fully contained in the top or bottom portion of the corresponding pseudo-circle. Our plan of attack, similar to those employed in [5, 6], is to decompose the *intersection graph* G of C^* (whose edges represent all intersecting pairs of curves in C^*) into a union of complete bipartite graphs $\{A_i \times B_i\}_i$. We then estimate the number of lenses in L that are formed between a curve in A_i and a curve in B_i, using Theorem 3.5(ii), and add up these bounds to obtain an upper bound for $\nu(C)$.

In more details, we proceed as follows. Without loss of generality, it suffices to consider the task of decomposing the portion of G corresponding to intersections between top boundaries of pairs of curves in C. Put $C^+ = \{c^+ \mid c \in C\}$. For each $\gamma \in C^+$, let p_γ denote the point in 3-space that represents γ (or, rather, the curve of C that contains γ), and let τ_γ denote the tangency surface associated with γ, representing tangencies of top boundaries of curves in $\mathbf{C}$ with γ. Put $\hat{C} = \{p_\gamma \mid \gamma \in C^+\}$, and let Σ denote the set of tangency surfaces $\{\tau_\gamma \mid \gamma \in C^+\}$.

We fix a parameter $r = n^\gamma$, for some sufficiently small constant fraction γ, and construct a $(1/r)$-*cutting* of the arrangement $\mathcal{A}(\Sigma)$. This is a decomposition of 3-space into cells, each having constant description complexity, so that each cell is crossed by at most n/r surfaces of Σ. The cutting is constructed as in [8], using the *vertical decomposition* (as defined, e.g., in [15]) of an arrangement of some r randomly sampled surfaces of Σ. (More precisely, the technique of [8] constructs first a 'master arrangement' $\mathcal{A}$ of r such surfaces, and then constructs additional arrangements of sampled surfaces within each cell of the vertical decomposition of $\mathcal{A}$ that is still crossed by more than n/r surfaces.) The decomposition consists of $O(r^3\beta(r))$ cells of constant description complexity, where $\beta(r) = \lambda_q(r)/r$ for some constant q depending on the algebraic degree (and other properties) of the representation of the curves in C, and where $\lambda_q(r)$ is the maximum length of Davenport-Schinzel sequences of order q composed of r symbols [15]. Thus $\beta(r)$ is a very slowly growing function of r. By cutting cells further as necessary, we may also assume that each cell contains at most n/r^3 points of $\hat{C}$.

Each cell ξ of the cutting induces two subproblems. One involves the points in ξ and the surfaces τ_γ that avoid the cell, but are such that all points in the cell represent pseudo-circles whose top boundaries intersect γ; this subproblem yields right away a complete bipartite graph for the output, consisting of the points in the cell and these avoiding surfaces. The second subproblem involves the surfaces that intersect ξ, and is handled recursively.

To simplify the recurrence, we apply one more recursive round (for each cell ξ), in which the roles of points and surfaces are interchanged. We sum over all resulting subproblems, and handle in an analogous fashion the otherthree types of interaction between top and bottom boundaries. Omitting further details, we obtain the recurrence

$$\nu(n) = O(r^6\beta^2(r)) \cdot \nu\left(\frac{n}{r^4}\right) + O(n^{4/3}r^6\beta^2(r)).$$

With an appropriate choice of $r = n^\gamma$, this solves to

$$\nu(n) = O(n^{3/2}(\log n)^{O(\log\beta(n))}) = O(n^{3/2}(\log n)^{O(\alpha^s(n))}),$$

using the explicit bound $\beta(n) = 2^{O(\alpha^s(n))}$, for an appropriate constant s [15]. We put

$$\kappa_s(n) = (\log n)^{O(\alpha^s(n))}.$$

Following the analysis of [16] (see also [6]), we then obtain:

THEOREM 4.1. *Let C be a collection of n x-monotone pseudo-circles that admit a 3-parameter algebraic representation. Then $\nu(C), \chi(C) = O(n^{3/2}\kappa_s(n))$, where s is a constant that depends (in the manner outlined above) on the algebraic representation of the curves in C.*

The case of circles. We next apply Theorem 4.1 to the case of circles, to obtain a slight improvement in the previous bound of Aronov and Sharir [6]. We first show that the constant s, for the case of circles, is 2 (details are routine, and omitted here). Hence, putting $\kappa(n) = (\log n)^{O(\alpha^2(n))}$, we then have the following improved bound.

COROLLARY 4.2. *n arbitrary circles in the plane can be cut into $O(n^{3/2}\kappa(n))$ subarcs, so that each pair of arcs intersect at most once.*

The case of parabolas. Theorem 4.1 can also be applied to the case of vertical parabolas, given by equations of the form $y = ax^2 + bx + c$. Omitting the routine details, due to lack of space, we show that $s = 2$ here too, and obtain:

THEOREM 4.3. *n parabolas can be cut into $O(n^{3/2}\kappa(n))$ arcs, so that each pair of these arcs intersect at most once. In particular, the number of empty lenses in such a collection of parabolas is $O(n^{3/2}\kappa(n))$.*

The case of homothetic copies. Here one can show that n homothetic copies of a fixed convex curve can be cut into $O(n^{3/2}\kappa_s(n))$ pseudo-segment arcs, where s depends on the shape of the fixed curve. The easy details are given in the full version.

5. APPLICATIONS

The preceding results have numerous applications to problems involving incidences, many faces, and levels, which extend (and also slightly improve) similar applications obtained for the case of circles in [1, 5, 6].

5.1 Levels

Given a collection C of curves, the *level* of a point $p \in \mathbb{R}^2$ is defined to be the number of intersection points between the relatively-open downward vertical ray emanating from p and the curves of C. The *k-th level* of $\mathcal{A}(C)$ is the (closure of the) locus of all points on the curves of C, whose level is exactly k. The k-th level consists of portions of edges of $\mathcal{A}(C)$, delimited either at vertices of $\mathcal{A}(C)$ or at points that lie above an x-extremal point of some curve. The complexity of a level is the number of edge portions that constitute the level.

The main tool for establishing bounds on the complexity of levels in arrangements of curves is an upper bound, given by Chan [7, Theorem 2.1], on the complexity of a level in an arrangement of *extendible* pseudo-segments, which is a collection of x-monotone bounded curves, each of which is contained in some unbounded x-monotone curve, so that the collection of these extensions is a family of pseudo-lines (in particular, each pair of the original curves intersect at most once).

Chan showed that the complexity of a level in an arrangement of m extendible pseudo-segments with ξ intersecting pairs is $O(m + m^{2/3}\xi^{1/3})$. Chan also showed that a collection of m x-monotone pseudo-segments can be turned, by further cutting the given pseudo-segments, into a collection of $O(m \log m)$ extendible pseudo-segments.

Thus, Theorem 4.1 leads to the following result (where the extra logarithmic factor incurred in turning our pseudo-segments into extendible pseudo-segments, as well as the power 2/3 to which we raise the number of pseudo-segments, are absorbed in the factor $\kappa(n)$).

THEOREM 5.1. *Let C be a set of n x-monotone pseudo-circles that admit a 3-parameter algebraic representation. Then the maximum complexity of a level in $\mathcal{A}(C)$ is $O(n^{5/3}\kappa_s(n))$ pseudo-segments, where s is a constant that depends on the algebraic representation of the curves in C; $s = 2$ for circles and vertical parabolas. If all pseudo-circles in C are pairwise intersecting, then, with no further assumption on these curves, the bound improves to $O(n^{14/9}\log n)$.*

The above theorem implies the following result in the area of kinetic geometry (improving a previous bound of [16]).

COROLLARY 5.2. *Let P be a set of n points in the plane, each moving along some line with a fixed velocity. For each time t, let $p(t)$ and $q(t)$ be the pair of points of P whose distance is the median distance at time t. The number of times in which this median pair changes is $O(n^{10/3}\kappa(n))$. The same bound applies to any fixed quantile.*

5.2 Incidences and marked faces

Let C be a set of n curves in the plane, and let P be a set of m points in the plane. Let $I(C, P)$ denote the number of incidences between P and C, i.e., the number of pairs $(c, p) \in C \times P$ such that $p \in c$. Let $K(C, P)$ denote the sum of the complexities of the faces of $\mathcal{A}(C)$ that contain at least one point of P; the complexity of a face f is the number of edges on its boundary. The results in [1, 6] imply the following

LEMMA 5.3. *Let C be a set of n curves in the plane, and let P be a set of m points in the plane. Then*

$$I(C, P) = O(m^{2/3}n^{2/3} + m + \chi(C)),$$

$$K(C, P) = O(m^{2/3}n^{2/3} + \chi(C)\log^2 n).$$

Hence, Theorem 4.1 implies the following.

THEOREM 5.4. *Let C be a set of n pairwise-intersecting pseudo-circles, and P a set of m points in the plane. Then*

$$I(C, P) = O(m^{2/3}n^{2/3} + m + n^{4/3}),$$

$$K(C, P) = O(m^{2/3}n^{2/3} + n^{4/3}\log^2 n).$$

If the pseudo-circles in C are not pairwise intersecting but are x-monotone and admit a 3-parameter algebraic representation, then we can obtain the following weaker bound by plugging Theorem 4.1 into Lemma 5.3.

$$I(C, P) = O(m^{2/3}n^{2/3} + m + n^{3/2}\kappa_s(n)),$$
$$K(C, P) = O(m^{2/3}n^{2/3} + n^{3/2}\kappa_s(n)). \quad (2)$$

However, we can obtain an improved bound on $I(C, P)$ and $K(C, P)$ following the approach in [1, 6]. In that approach, the circles are mapped to points in $\mathbb{R}^3$, and the points in P to planes in $\mathbb{R}^3$, so that one halfspace bounded by the plane corresponds to circles that contain the point. The arrangement of these dual planes is partitioned, via a cutting, into subcells, and one applies a weaker bound within each subcell separately. We follow their approach verbatim. That is, we map pseudo-circles in C into points in $\mathbb{R}^3$, and points in P into surfaces in $\mathbb{R}^3$. We then compute a cutting of the resulting arrangement of surfaces, and apply the bounds in (2) within each cell. Of course, if the pseudo-circles in C are pairwise intersecting, then we use instead the bounds from Theorem 5.4. The only difference is that the size of a $(1/r)$-cutting is $O(r^3\beta_1(r))$, instead of $O(r^3)$, where $\beta_1(r) = 2^{O(\alpha^{s_1}(r))}$, for an appropriate constant s_1. Following this analysis, we obtain the following.

THEOREM 5.5. *Let C be a set of n x-monotone pseudo-circles that admit a 3-parameter algebraic representation, and let P be a set of m points in the plane. Then*

(i) $I(C,P) = O(m^{2/3}n^{2/3} + m^{6/11}n^{9/11}\kappa^*(m,n) + m + n)$;

(ii) $K(C,P) = O(m^{2/3}n^{2/3} + m^{6/11}n^{9/11}\kappa^*(m,n) + n\log n)$,

where $\kappa*(m,n) = 2^{O(\alpha^{s_1}(n^5\kappa_s(n)/m^4))} \cdot \kappa_s(m^3/n)$. *In addition, if the pseudo-circles in C are pairwise intersecting, then*

(i) $I(C,P) = O(m^{2/3}n^{2/3} + m^{1/2}n^{5/6} \cdot 2^{O(\alpha^{s_1}(n/m))} + m + n)$.

(ii) $K(C,P) = O(m^{2/3}n^{2/3} + m^{1/2}n^{5/6} \cdot 2^{O(\alpha^{s_1}(n/m))}\log^{1/2} n + n\log n)$.

Since circles and vertical parabolas can be linearized in $\mathbb{R}^3$ (i.e., mapped to planes in $\mathbb{R}^3$), we have $s_1 = 0$, and thus:

THEOREM 5.6. *Let C be a set of n circles or vertical parabolas and let P be a set of m points in the plane. Then*

(i) $I(C,P) = O(m^{2/3}n^{2/3} + m^{6/11}n^{9/11}\kappa(m^3/n) + m + n)$.

(ii) $K(C,P) = O(m^{2/3}n^{2/3} + m^{6/11}n^{9/11}\kappa(m^3/n) + n\log n)$.

In addition, if the curves in C are pairwise intersecting, then

(i) $I(C,P) = O(m^{2/3}n^{2/3} + m^{1/2}n^{5/6} + m + n)$.

(ii) $K(C,P) = O(m^{2/3}n^{2/3} + m^{1/2}n^{5/6}\log^{1/2} n + n\log n)$.

5.3 Distinct distances under arbitrary norms

THEOREM 5.7. *Let Q be a compact convex centrally symmetric semi-algebraic region in the plane, of constant description complexity, which we regard as the unit ball of a norm $\|\cdot\|_Q$. Then any set P of n distinct points in the plane determines at least $\Omega(n^{7/9}/\kappa_s(n))$ distinct $\|\cdot\|_Q$-distances, where s is a constant that depends on Q. If Q is not centrally symmetric, it defines a convex distance function, and the same lower bound applies in this case too. In both cases, this is also a lower bound on the number of distinct $\|\cdot\|_Q$-distances that can be attained from a single point of P.*

Proof: The proof proceeds by considering nt homothetic copies of Q, shifted to each point of P and scaled by the t possible distinct $\|\cdot\|_Q$-distances. There are n^2 incidences between these curves and the points of P. Using (5.5), the bound follows easily. □

5.4 Gallai-Sylvester theorem

In our final application, similar to Theorem 4.1 in [5], the following theorem is a consequence of Theorem 2.5.

THEOREM 5.8. *Let C be a family of n pairwise intersecting pseudo-circles in the plane. If n is sufficiently large and C is not a pencil, then there exists an intersection point incident to at most three pseudo-circles.*

6. OPEN PROBLEMS

The paper leaves many open problems unanswered. We mention a few of the more significant ones:
(i) Obtain tight (or improved) bounds for the number of pairwise nonoverlapping lenses in an arrangement of n pairwise intersecting pseudo-circles. We conjecture that the upper bound of $O(n^{4/3})$, given in Theorem 2.6, is not tight, and that the correct bound is $O(n)$ or near-linear.
(ii) Obtain tight (or improved) bounds for the number of empty lenses in an arrangement of n arbitrary circles. There is a gap between the lower bound $\Omega(n^{4/3})$ and the upper bound of $O(n^{3/2}\kappa(n))$, given in Theorem 4.1 and Corollary 4.2. Even improving the upper bound to $O(n^{3/2})$ seems a challenging open problem. A related problem is to obtain an improved bound for the number of pairwise nonoverlapping lenses in an arrangement of n arbitrary circles.
(iii) One annoying aspect of our analysis is the difference between the analysis of pairwise intersecting pseudo-circles, which is purely topological and requires no further assumptions concerning the shape of the pseudo-circles, and the analysis of the general case, in which we require x-monotonicity and 3-parameter algebraic representation. It would be interesting and instructive to find a purely topological way of tackling the general problem. For example, can one obtain a bound close to $O(n^{3/2})$, or even any bound smaller than the general bound $O(n^{5/3})$ of [16] (which *is* purely topological), for the number of empty lenses in an arbitrary arrangement of pseudo-circles, without having to make any assumption concerning their shape? Work in progress by Pinchasi provides an initial affirmative answer to this question.

7. REFERENCES

[1] P. K. Agarwal, B. Aronov and M. Sharir, On the complexity of many faces in arrangements of pseudo-segments and of circles, *Proc. 42nd IEEE Symposium on Foundations of Computer Science*, 2001, pp. 74–83.

[2] P. K. Agarwal and M. Sharir, Arrangements and their applications, in: *Handbook of Computational Geometry* (J.-R. Sack and J. Urrutia, eds.), Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000, pp. 49–119.

[3] N. Alon, G. Kalai, J. Matoušek, and R. Meshulam, Transversal numbers for hypergraphs arising in geometry, *Adv. Appl. Math.*, to appear.

[4] N. Alon and D. J. Kleitman, Piercing convex sets and the Hadwiger-Debrunner (p,q)-problem, *Adv. Math.* 96 (1992), 103–112.

[5] N. Alon, H. Last, R. Pinchasi, and M. Sharir, On the complexity of arrangements of circles in the plane, *Discrete Comput. Geom.*, 26 (2001), 465–492.

[6] B. Aronov and M. Sharir, Cutting circles into pseudo-segments and improved bounds for incidences, *Discrete Comput. Geom.*, in press.

[7] T.M. Chan, On levels in arrangements of curves, to appear in *Discrete Comput. Geom.* (A preliminary version appeared in *Proc. 41st IEEE Symp. Found. Comput. Sci.* (2000), 219–227.)

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

[9] Ch. Chojnacki (A. Hanani), Über wesentlich unplättbare Kurven im dreidimensionalen Raume, *Fund. Math.* 23 (1934), 135–142.

[10] E. Helly, Über Systeme abgeschlossener Mengen mit gemeinschaftlichen Punkten, *Monatshefte d. Mathematik* 37 (1930), 281–302.

[11] J. Hershberger and J. Snoeyink, Sweeping arrangements of curves, *DIMACS Series in Discrete Mathematics, Discrete and Computational Geometry, the DIMACS Special Year* 6 (1991), 309–349.

[12] L. Lovász, J. Pach and M. Szegedy, On Conway's thrackle conjecture, *Discrete Comput. Geom.* 18 (1997), 369–376.

[13] J. Molnár, Über eine Verallgemeinerung auf die Kugelfläche eines topologischen Satzes von Helly, *Acta Math. Acad. Sci.* 7 (1956), 107–108.

[14] R. Pinchasi, *Problems in Combinatorial Geometry in the Plane*, Ph.D. Dissertation, Dept. of Mathematics, Hebrew University, 2001.

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

[16] H. Tamaki and T. Tokuyama, How to cut pseudo-parabolas into segments, *Discrete Comput. Geom.* 19 (1998), 265–290.

[17] W.T. Tutte, Toward a theory of crossing numbers, *J. Combinat. Theory* 8 (1970), 45–53.