Lenses in Arrangements of Pseudo-circles and Their Applications*

Pankaj K. Agarwal† Eran Nevo‡ János Pach§ Rom Pinchasi¶ Micha Sharir‖
Shakhar Smorodinsky**

November 26, 2002

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. This bound implies that any collection of n x-monotone pseudo-circles can be cut into $O(n^{8/5})$ arcs so that any two intersect at most once; this improves a previous bound of $O(n^{5/3})$ due to Tamaki and Tokuyama. If, in addition, the given collection admits an algebraic representation by three real parameters that satisfies some simple conditions, then the number of cuts can be further reduced to $O(n^{3/2}(\log n)^{O(\alpha^s(n))})$, where $\alpha(n)$ is the inverse Ackermann function, and s is a constant that depends on the the representation of the pseudo-circles. For arbitrary collections of pseudo-circles, any two of which intersect exactly twice, the number of necessary cuts reduces still further to $O(n^{4/3})$. As applications, we obtain improved bounds for the number of incidences, the complexity of a single level, and the complexity of many faces in arrangements of circles, of pairwise intersecting pseudo-circles, of arbitrary x-monotone pseudo-circles, of parabolas, and of homothetic copies of any fixed simply-shaped convex curve. We also obtain a variant of the Gallai-Sylvester theorem for arrangements of pairwise intersecting pseudo-circles, and a new lower bound on the number of distinct distances under any well-behaved norm.

*Work on this paper by Pankaj Agarwal, János Pach and Micha Sharir has been supported by a joint grant from the U.S.–Israel Binational Science Foundation. Work by Pankaj Agarwal has also been supported by NSF grants EIA-98-70724, EIA-99-72879, ITR-333-1050, CCR-97-32787, and CCR-00-86013. Work by Micha Sharir has also been supported by NSF Grants CCR-97-32101 and CCR-00-98246, 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. A preliminary version of this paper has appeared in *Proc. 18th ACM Annu. Sypos. Comput. Geom.*, 2002, pp. 123–132.

†Department of Computer Science, Duke University, Durham, NC 27708-0129, USA; pankaj@cs.duke.edu

‡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; michas@post.tau.ac.il

**School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; smoro@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 [4].

A family of 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 intersect 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 intersect 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, 7, 8, 11, 24, 27]. This paper continues this line of study—it improves a number of previous results on arrangements of pseudo-circles, and extends a few of the recent results on arrangements of circles (e.g., those presented in [7, 8, 24]) 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 (relative interiors of 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); thus $\chi(C) = |C|$ when no cuts need to be made. 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 $\mathcal{A}(C)$.

Previous results. Tamaki and Tokuyama [27] 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.* [7] and Pinchasi [24] 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 [8]. No better bound is known for the number of empty lenses in an arbitrary

[1] For simplicity, we assume that every tangency counts as two intersections, 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.

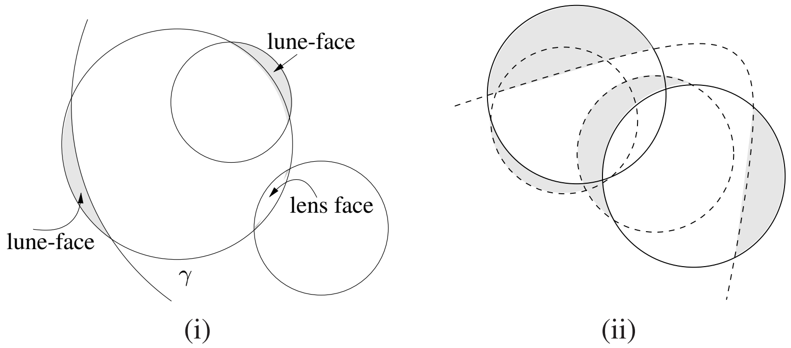

Figure 1. (i) A pseudo-circle γ supporting one lens-face and two lune-faces. (ii) A family of (shaded) nonoverlapping lenses.

family of circles. However, $\mu(C) = O(n^{4/3})$ for a set of n unit circles, though no superlinear lower bound is known for this special case.

The analysis in [27] 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 [27], and $\chi(C) = O(n^{3/2+\varepsilon})$ for circles. Using this bound on $\chi(C)$, Aronov and Sharir [8] 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.[2] An interesting consequence of the results in [7, 24] is the following generalization of the Sylvester-Gallai theorem: In an arrangement of pairwise intersecting circles, there always exists a vertex incident upon at most three circles, provided that the number of circles is sufficiently large and that they do not form a pencil. For pairwise intersecting unit circles, the property holds when the number of circles is at least five [7, 24].

New results. In this paper we first obtain improved bounds on $\mu(C)$, $\nu(C)$, and $\chi(C)$ 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 $\chi(C) = O(n^{4/3})$ for collections C of n pairwise intersecting pseudo-parabolas. We then go on to study the general case, in which not every pair of curves intersect. We first show, in Section 4, that $\chi(C) = O(n^{8/5})$ for *arbitrary* collections of n pseudo-parabolas and for collections of n x-monotone pseudo-circles. This improves the general bound of Tamaki and Tokuyama [27], and is based on a recent result of Pinchasi and Radoičić [25] on the size of graphs drawn in the plane so that any pair of edges in a cycle of length 4 intersect an even number of times. In order to improve this bound further, we need to make a few additional assumptions on the geometric shape of the given curves. Specifically, we assume, in Section 5, that, in addition to x-monotonicity, the n given curves admit a 3-parameter algebraic representation that satisfies some simple conditions (a notion defined more precisely in Section 5). Three important

[2] Actually, the paper [1], having been written alongside with the present paper, already exploits the slightly improved bound derived here.

classes of curves that satisfy these assumptions are the classes of circles, vertical parabolas (of the form $y = ax^2 + bx + c$), and of homothetic copies of any fixed simply-shaped convex curve. We show that, in the case of such a representation, $\chi(C) = O(n^{3/2}(\log n)^{O(\alpha^s(n))})$, where $\alpha(n)$ is the inverse Ackermann function and s is a constant depending on the algebraic parametrization; $s = 2$ for circles and vertical parabolas. This bound gives a slightly improved bound on $\chi(C)$, compared to the bound proved in [8], for a family of circles.

In Section 6, we apply the above results to several problems. The better bounds on the cutting number $\chi(C)$ lead to improved bounds on the complexity of levels, on the number of incidences between points and curves, 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, homothetic copies of a fixed convex curve, and general pseudo-parabolas and x-monotone pseudo-circles. The exact bounds are stated in Section 6. We also obtain a generalized Gallai-Sylvester result for arrangements of pairwise-intersecting pseudo-circles, and a new lower bound for the number of distinct distances determined by n points in the plane and induced by an arbitrary well-behaved norm.

2 Pairwise Intersecting Pseudo-Circles

Let C be a set of n pseudo-circles, any two of which intersect in two points. We prove that $\mu(C)$, the number of empty lenses in $\mathcal{A}(C)$, is $O(n)$. The proof proceeds in three stages: First, we reduce the problem to $O(1)$ instances of counting the number of empty lenses in an arrangement of at most n pairwise intersecting pseudo-circles, all of whose interiors are star shaped with respect to a fixed point o. Next, we reduce the latter problem to counting the number of tangencies in a family of pairwise intersecting pseudo-parabolas. Finally, we prove that the number of such tangencies is $O(n)$. For simplicity, we provide the proof in the reverse order: Section 2.1 proves a bound on the number of tangencies in a family of pairwise intersecting pseudo-parabolas; this provides the main geometric insight of this paper, on which all other results are built. Section 2.2 proves a bounds on $\mu(C)$ for a family C of pairwise-intersecting star-shaped pseudo-circles, by using the result in the previous subsection; Section 2.3 supplies the final reduction, and shows that the number of empty lenses in a family of arbitrary pairwise-intersecting pseudo-circles can be counted using the result obtained in Section 2.2.

2.1 Tangencies of pseudo-parabolas

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, where no crossing occurs.[3] We also assume that no three of these curves have a point in common. This general position assumption is made in order to simplify our analysis. Later on, we will show how to extend our analysis to sets of curves that are not in general position. Note also that considering tangencies, rather than empty lenses, is just another simplifying step: Since no three curves are concurrent, any tangency can be deformed into a small empty lens and vice versa. Let T denote the set of all tangencies between pairs of curves in Γ. Our goal is to bound the size of

[3]The requirement that the number of intersections of every pair be exactly two can be relaxed to that of requiring that every pair intersect at least once: A family satisfying the latter condition can easily be extended to a family that satisfies the former condition.

T.

We associate a graph G with T, whose vertices are the curves of Γ and whose edges connect pairs of tangent curves. 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 because the two other curves forming the respective tangencies with γ would have to be disjoint, contrary to assumption. Hence, G is bipartite. In the remainder of this subsection we show that G is planar, and this will establish a linear upper bound on the size of T.

The drawing rule. Let ℓ be a 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^*. We use (γ_1^*, γ_2^*) to denote the arc drawn for (γ_1, γ_2). The arc 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.

We use the following rule for drawing an edge (γ_1, γ_2): 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^*; see Figure 2(i). Note that a tangency also counts as an exit point (with immediate re-entry back into the wedge). Except for these requirements, the edge (γ_1, γ_2) can be drawn in an arbitrary y-monotone manner.

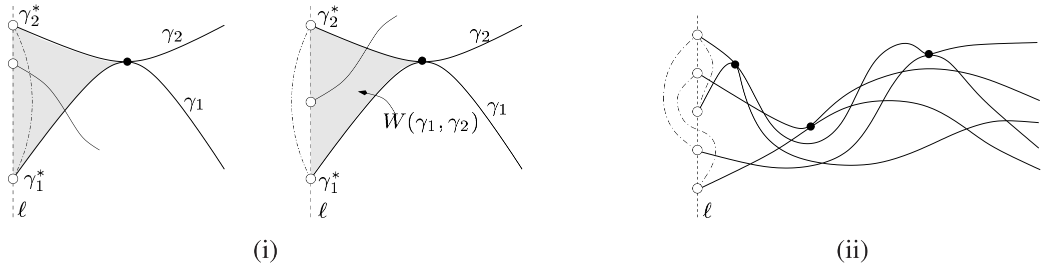

Figure 2. (i) Illustrating the drawing rule. (ii) Drawing the graph G for an arrangement of five pairwise intersecting pseudo-parabolas with three tangencies.

Lemma 2.1 *Suppose that the following conditions hold for each quadruple $\gamma_1, \gamma_2, \gamma_3, \gamma_4$ of distinct curves in Γ, whose intercepts with ℓ appear in this y-increasing order:*

(a) If (γ_1, γ_4) and (γ_2, γ_3) are edges of G, then both γ_2^* and γ_3^* lie on the same side of the arc (γ_1^*, γ_4^*).

(b) If (γ_1, γ_3) and (γ_2, γ_4) are edges of G and the arc (γ_1^*, γ_3^*) passes to the left (resp., right) of γ_2^*, then the arc (γ_2^*, γ_4^*) passes to the right (resp., left) of γ_3^*.

Then G is planar.

Proof: Figure 3 shows the configurations allowed and forbidden by conditions (a) and (b). We show that the drawings of each pair of edges of G without a common endpoint cross an even number of times. (With additional care, this property can also be enforced for pairs of edges with a common endpoint, as will be shown later. This extension is not needed for the main result, Theorem 2.4, but is needed for the analysis in Section 4 involving general pseudo-parabolas and x-monotone pseudo-circles.) This, combined with Hanani-Tutte's theorem [29] (see also [16] and [22]), implies that G is planar. Clearly, it suffices to check this for pairs of edges (with distinct endpoints) 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).

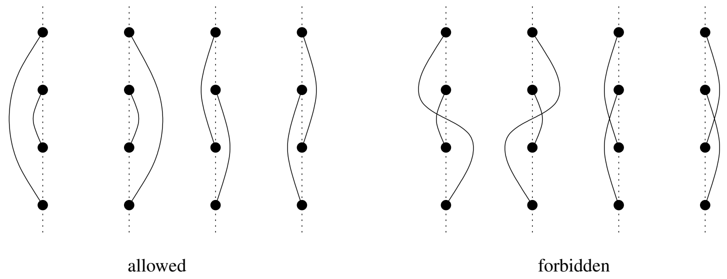

Figure 3. The allowed and forbidden configurations in conditions (a) and (b).

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.

This completes the proof of the lemma. $\square$

We next show that the conditions in Lemma 2.1 do indeed hold for our drawing of G.

Lemma 2.2 *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_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, by symmetric reasoning, 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 (since it has already intersected γ_4 twice). This implies that γ_2 cannot intersect γ_4 at all, a contradiction, which shows that case (a) is impossible.

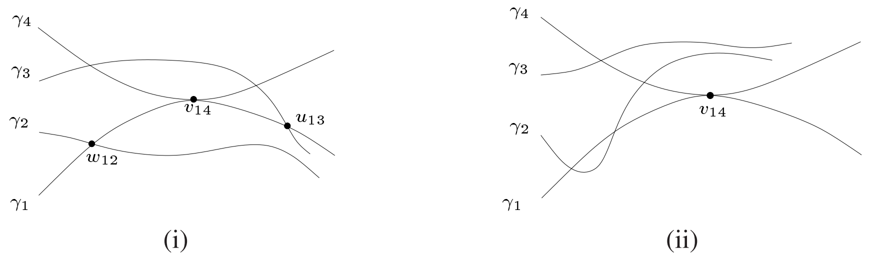

Figure 4. Edges of G with nested projections: (i) γ_2 passes below v_{14} and γ_3 passes above v_{14}; (ii) both γ_2 and γ_3 pass on the same side of v_{14}.

(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 2.3 *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. □

Lemmas 2.2 and 2.3 show that the conditions in Lemma 2.1 hold, so G is planar and bipartite and thus has at most $2n - 4$ edges, for $n \geq 3$. Hence, we obtain the following.

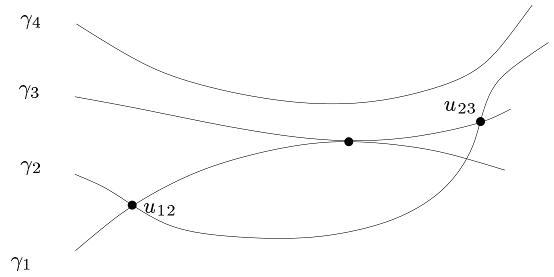

Figure 5. Edges of G with partially overlapping projections.

Theorem 2.4 *Let Γ be a family of n pairwise intersecting pseudo-parabolas in the plane, i.e., each pair intersect either in exactly two crossing points or in exactly one point of noncrossing tangency. Assume also that no three curves of Γ meet at a common point. Then there are at most $2n-4$ tangencies between pairs of curves in Γ, for $n \geq 3$.*

2.2 Empty lenses in star-shaped pseudo-circles

The main result of this subsection is:

Theorem 2.5 *The number of empty lenses in an arrangement of $n \geq 3$ pairwise intersecting pseudo-circles, no pair of whicch are tangent and no three concurrent, so that all their interiors are star shaped with respect to a point o, is at most $2n-3$. This number is 3 for $n = 2$. Both bounds are tight in the worst case.*

The lower bound, for $n = 5$, is illustrated in Figure 6. It is easy to generalize this construction for any $n \geq 3$. The case $n = 2$ is trivial: A pair of intersecting circles form three empty lenses (ignoring the unbounded face), of which two are lune-faces and one is a lens-face, containing o.

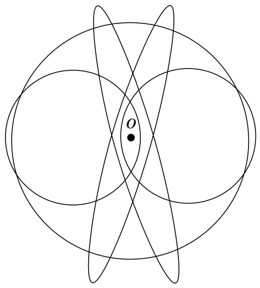

Figure 6. Lower-bound construction: Five circles with a common interior point forming seven empty lenses.

Assume then that $n \geq 3$. At most one empty lens contains o. We will show that the number of empty lenses not containing o is at most $2n-4$. By definition, each of these lenses is a lune-face (whereas the empty lens containing o, if any, is a lens-face).

We deform the pseudo-circles of C, so as to turn each lune-face into a *tangency* between the two corresponding pseudo-circles. This is easy to do, by deforming the two pseudo-circles bounding such an empty lens, using the facts that no two empty lenses share an arc or a vertex; see Figure 7 for

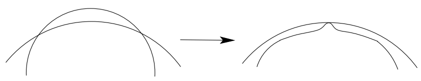

Figure 7. Transforming an empty lens into a tangency.

an illustration. We can deform the pseudo-circles in this manner without losing the star-shapedness property.

Draw a generic ray ρ that emanates from o and does not pass through any vertex of $\mathcal{A}(C)$; in particular, it does not pass through any empty lens, each now reduced to a point of tangency between the respective pseudo-circles. Without loss of generality, assume that ρ has orientation 0, i.e., it points to the direction of the positive x-axis. Regard each curve of C as the graph of a function in polar coordinates, and map the open interval $(0, 2\pi)$ of orientations onto the real line (e.g., by $x = -\cot\theta/2$). This transforms C into a collection Γ of pairwise intersecting *pseudo-parabolas*, that is, graphs of totally defined continuous functions, each pair of which intersect exactly twice. The ray ρ is mapped to the vertical lines at $x = \pm\infty$.

The problem has thus been reduced to that of bounding the number of tangencies among n pairwise intersecting pseudo-parabolas, no three of which are concurrent. By Theorem 2.4, the number of tangencies is at most $2n - 4$, for $n \geq 3$, so the number of lune-faces is at most $2n - 4$. This completes the overall inductive proof of the theorem.

2.3 Reduction to pairwise intersecting star-shaped pseudo-circles

Let C be a family of n pseudo-circles, any two of which intersect each other in two points. We refer to the interiors of these pseudo-circles as *pseudo-disks*. We bound $\mu(C)$ by reducing the problem to a constant number of subproblems, each of which is ultimately reduced to counting the number of empty lenses in a family of pairwise intersecting star-shaped pseudo-circles. We continue to assume that the curves in C are in general position, as in the preceding subsection.

We need the following easy observation.

Lemma 2.6 *Among any five pseudo-disks bounded by the elements of C, there are at least three that have a point in common.*

Proof: Indeed, if this were false, then there would exist five pseudo-disks such that any two of them intersect in an empty lens (in the arrangement of the five corresponding boundary curves), which would give rise to a forbidden planar drawing of K_5, the complete graph with five vertices.
$\square$

The following topological variant of Helly's theorem [18] was found by Molnár [23]. It can be proved by a fairly straightforward induction.

Lemma 2.7 *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.*

Let $p \geq q \geq 2$ be integers. We say that a family F of sets has the (p, q) *property* if among every p members of F there are q that have a point in common. We say that a family of sets F is *pierced* by a set T if every member of F contains at least one element of T. The set T is often called a *transversal* of F. Fix $p \geq q \geq d+1$. Alon and Kleitman [6] proved that there exists a transversal of size at most $k = k(p, q, d)$ for any finite family of convex sets in $\mathbb{R}^d$ with the (p, q)-property. Recently, Alon *et al.* [5] extended this result to any finite family F of open regions in d-space with the property that the intersection of every subfamily of F is either empty or contractible. Their result, combined with Lemmas 2.6 and 2.7, implies the following.

Corollary 2.8 *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 $\mathcal{O} = \{o_1, o_2, \ldots, o_k\}$ of k points that pierces all pseudo-disks bounded by the elements of C. Let C_i consist of all elements of C that contain o_i in their interior, 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. We begin by considering the first case and then reduce the second case to the first one.

Let C be a family of pseudo-circles, so that any two of them intersect and each of them contains the origin o in its interior. We wish to bound $\mu(C)$. Obviously, there exists at most one empty lens-face formed by elements of C, namely, the face containing o. Therefore, it is sufficient to bound the number of lune-faces determined by C. The *combinatorial structure* of an arrangement is its face lattice. We call two arrangements *c*ombinatorially equivalent if the face lattices of their arrangements are isomorphic. For a face f, we say that an edge e bounding f is pointing *inside* (resp., *outside*) if f is in the interior (resp., the exterior) of the pseudo-disk whose boundary includes e.

We need the following technical lemma to prove the main result.

Lemma 2.9 *Let C be a family of pseudo-circles such that all of them have an interior point o in common. Then the union of any set of pseudo-disks bounded by the elements of C is simply connected.*

Proof: For any $\gamma_i \in C$, let D_i denote the pseudo-disk bounded by γ_i. Using stereographic projection, we can map each D_i into a simply connected region D'_i of a sphere $\mathbb{S}^2$ touching the plane at o, where the center of projection is the point $o' \in \mathbb{S}^2$ antipodal to o. Clearly, we have

$$\mathbb{S}^2 \setminus \bigcup_{1 \leq i \leq k} D'_i = \bigcap_{1 \leq i \leq k} (\mathbb{S}^2 \setminus D'_i).$$

The sets $\overline{D'_i} = \mathbb{S}^2 \setminus D'_i$ form a collection of pseudo-disks in the "punctured" sphere $\mathbb{S}^2 \setminus \{o\}$, isomorphic to the plane, and they all contain o'. Thus, applying Lemma 2.7 (clearly, the intersection of two pseudo-disks is always connected), we obtain that the right-hand side of the above equation is simply connected. Therefore, $\mathbb{S}^2 \setminus \bigcup_{1 \leq i \leq k} D'_i$ is also simply connected, which implies that the union of C is simply connected. □

By Lemma 2.9, $\mathbb{R}^2 \setminus \bigcup_i D_i$ consists of only one (unbounded) cell in $\mathcal{A}(C)$. An immediate corollary of the above lemma is the following.

Corollary 2.10 *Every bounded face of $\mathcal{A}(C)$ has an edge that points inside.*

Proof: Let f be a bounded face of $\mathcal{A}(C)$. Denoting by s_i and D_i, for $i = 1, 2, \ldots, k$, the edges of f and the respective pseudo-disks whose boundaries contain these edges, and assuming that every s_i is pointing outside, we obtain that f lies in the exterior of all pseudo-disks D_i, for $i = 1, 2, \ldots, k$. However, this would imply that f is a bounded cell of the complement of $\bigcup_{1 \le i \le k} D_i$, contradicting Lemma 2.9, which states that $\bigcup_{1 \le i \le k} D_i$ is a simply connected bounded set. $\square$

We now prove the main technical result of this subsection.

Lemma 2.11 *Let C be a finite family of pseudo-circles in general position, such that all of them have an interior point o in common. Then there exists a combinatorially equivalent family C' of pseudo-circles, all of which are star-shaped with respect to o.*

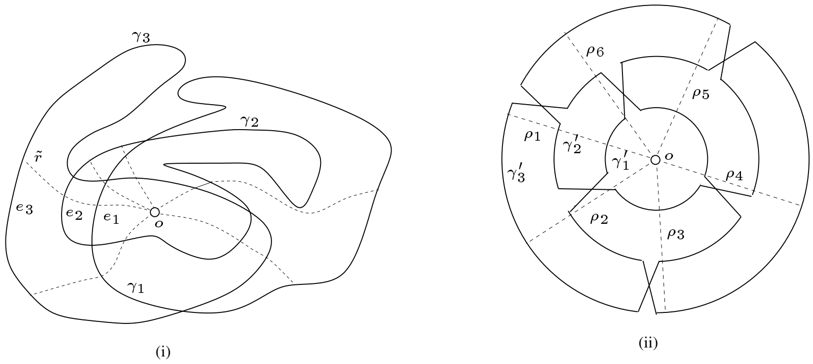

Figure 8. Converting C into a star-shaped family by a counterclockwise topological sweep: (i) The original curves; (ii) The transformed curves. $\Pi = (123, 213, 231, 321, 312, 132, 123)$.

Proof: We perform an "angular" topological sweep of $\mathcal{A}(C)$ with respect to o by a semi-infinite arc $\tilde{r}$ that has o as an endpoint, and intersects, at any time, each pseudo-circle of C exactly once. The ordering of the intersections of $\tilde{r}$ with the members of C gives a permutation of C, and the sweep produces a circular sequence Π of permutations, each differing from the preceding one by a swap of two adjacent elements. We then construct a family C' of pseudo-circles, all of which are star-shaped with respect to o, so that the angular sweep of $\mathcal{A}(C')$ by a *ray* emanating from o produces the same sequence Π; this will imply that C' is combinatorially equivalent to C.

First we show how to construct an initial instance of the curve $\tilde{r}$. Let f_1 be the cell of $\mathcal{A}(C)$ containing o. Clearly, all edges of f_1 point inside. Start drawing a curve $\tilde{r}$ from o so that it first crosses an edge e_1 of f_1, pointing inside f_1. Let f_2 denote the cell on the other side of e_1, and let e_2 be an edge of this cell pointing inside; clearly, $e_2 \neq e_1$. Extend $\tilde{r}$ through f_2 until it crosses e_2. Proceeding in this way, we reach, after n steps, the unique unbounded cell f_{n+1}; see Figure 8(i). This follows by noting that at each step we exit a different pseudo-disk, and never enter into any pseudo-disk. Let $\gamma_i \in C$ denote the pseudo-circle whose boundary contains e_i. Clearly, the sequence $\pi_1 = (\gamma_1, \ldots, \gamma_n)$, where γ_i is the curve containing the edge e_i, is a permutation of C.

The following claim shows that there always exists a "local" move that advances the sweep of the curve $\tilde{r}$ around o. It is reminiscent of a similar result given in [20].

Claim A *There exist two consecutive edges e_i, e_{i+1} that are crossed by $\tilde{r}$ and have a common endpoint counterclockwise to $\tilde{r}$, i.e., the triangular region enclosed by e_i, e_{i+1}, and $\tilde{r}$ is contained in a face of $\mathcal{A}(C)$ and lies (locally) on the counterclockwise side of $\tilde{r}$.*

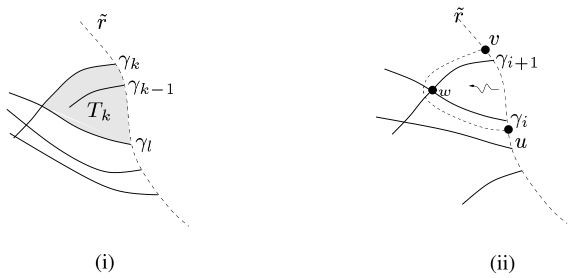

Figure 9. (i) e_l and e_k have a common endpoint counterclockwise to $\tilde{r}$; (ii) advancing the sweep curve.

Proof: Let $j(i)$, for each $1 \leq i \leq n$, denote the index of the first element of C that intersects γ_i counterclockwise to $\tilde{r}$. Let T_i denote the triangular region bounded by $\gamma_i, \gamma_{j(i)}$, and $\tilde{r}$. We say that T_i is *positive* (resp., *negative*), if $j(i) < i$ (resp., $j(i) > i$). Let k be the smallest integer for which T_k is positive, and put $l = j(k)$; see Figure 9(i). Observe that T_n is positive, so k is well defined. No curve whose index is greater than k can intersect T_k because such a curve would have to intersect γ_l at more than two points (it has to "enter" and "leave" T_k through γ_l, but to reach the entry point it has to cross γ_l once more, counterclockwise to T_k). Since $j(l) > l$, it follows that if $l = k - 1$ then e_l and e_k satisfy the property in the claim. The proof is completed by noting that this is the only possible case: If $l < k - 1$ then γ_{k-1} cannot exit T_k at all, which is impossible. Indeed, γ_{k-1} cannot intersect any curve of C in the interior of T_k, because then T_{k-1} would be positive, as the index of any curve intersecting the interior of T_k is smaller than k. If γ_{k-1} exits T_k by intersecting γ_l, then again T_{k-1} would be positive. Finally, γ_{k-1} cannot exit T_k by crossing γ_k because $k - 1 \neq l = j(k)$. This contradiction implies that $l = k - 1$, and the claim holds with e_l, e_k. □

Assume that e_i and e_{i+1} share an endpoint w counterclockwise to $\tilde{r}$. Now fix a pair of points $u, v \in \tilde{r}$, close to the points where $\tilde{r}$ crosses ∂T_i and lying outside T_i, and continuously sweep the portion of the curve $\tilde{r}$ between u and v, keeping the other parts fixed, pushing the crossing points with ∂T_i towards w, and finally pull it through w, so that $\tilde{r}$ no longer intersects T_i; see Figure 9(ii). In this new position, $\tilde{r}$ meets γ_{i+1} before it meets γ_i. We obtain a new permutation π_2, which is the same as π_1 except that the positions of γ_i and γ_j are swapped.

We repeat the above procedure for the new curve $\tilde{r}$. Continuing in this manner, we obtain a sequence $\Pi = (\pi_1, \pi_2, \dots)$ of permutations of the elements of C, corresponding to the different orders in which $\tilde{r}$ crosses the curves.

We now construct a family of pseudo-circles that realize the same sequence Π if we sweep their arrangement by a ray around o. This is done similar to the procedure described by Goodman and Pollack [17] for realizing an *allowable* sequence by an arrangement of pseudo-lines. Roughly speaking, we draw n concentric circles $\sigma_1, \sigma_2, \dots, \sigma_n$ around o, and draw a ray ρ_i from o for each permutation π_i in Π. If π_{i+1} is obtained from π_i by swapping γ_j and γ_{j+1}, we erase small arcs of σ_j and σ_{j+1} near their intersection points with ρ_{i+1} and connect the endpoints of the two erased arcs by two crossing segments; see Figure 8(ii). Let C' denote the set of n curves, obtained by modifying the circles $\sigma_1, \dots, \sigma_n$ in this manner. By construction, each curve in C' is star-shaped with respect to o and C' produces the sequence Π if we sweep it around o with a ray. By induction on the length of Π, one can show that C and C' are combinatorially equivalent, which implies that C' is a family of pseudo-circles, any pair of which intersect in exactly two points. $\square$

Lemma 2.11 implies that the number of empty lenses in C is the same as that in C'. Hence, by Theorem 2.5, we obtain the following.

Corollary 2.12 *Let C be a family of $n \geq 3$ pairwise-intersecting pseudo-circles in general position whose common interior is not empty. Then $\mu(C) \leq 2n - 3$. For $n = 2$, $\mu(C) = 3$.*

We are now ready to prove the main result of this section.

Theorem 2.13 *Let C be a family of n pairwise-intersecting pseudo-circles in general position. Then $\mu(C) = O(n)$.*

Proof: By Corollary 2.8, there exists a partition $\{C_1, \dots, C_k\}$ of C into $O(1)$ subsets, so that all the pseudo-circles in C_i contain a point o_i in their common interior, for $i = 1, \dots, k$. Corollary 2.12 implies that the number of empty lenses induced by two pseudo-circles within the same family C_i is at most $2|C_i| - 1$, for a total of at most $2n - k$. It thus remains to consider the case in which the given family of pairwise intersecting pseudo-circles is the union of two subfamilies C, C', such that the interiors of all pseudo-circles in C (resp., in C') contain a common point o (resp., o'). We wish to bound the number of "bichromatic" empty lenses, i.e., empty lenses in $\mathcal{A}(C \cup C')$ formed by a pseudo-circle in C and a pseudo-circle in C'. We may assume that none of the pseudo-circles of C' contains o in its interior. Indeed, each pseudo-circle of C' whose interior contains o can be added to C, and every bichromatic empty lens it determines is counted among the empty lenses in $\mathcal{A}(C)$, using Theorem 2.5. Similarly, we may assume that none of the pseudo-circles of C contains o' in its interior. Any bichromatic lune-face in $\mathcal{A}(C \cup C')$ must contain either o or o', so there can be at most two such faces. Thus, it suffices to bound the number of bichromatic *lens-faces*.

Apply an inversion of the plane with respect to o. Then each bichromatic lens-face is mapped into a lune-face, which lies outside the incident pseudo-circle of C and inside the incident pseudo-circle of C'. Moreover, all the pseudo-circles of both families now contain o' in their interior. Hence, by Theorem 2.5, the number of these lune-faces (that is, the original lens-faces) is at most $2n - 4$, for $n \geq 3$; it is 2 for $n = 2$. Summing this bound over all pairs of sets in the partition, the theorem follows. $\square$

2.4 Pairwise nonoverlapping lenses

Let C be a family of n pairwise-intersecting pseudo-parabolas or pseudo-circles in general position, 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 2.14 *Let C be a family of n pairwise-intersecting pseudo-parabolas or pseudo-circles in general position. Then the maximum size of a family of pairwise nonoverlapping lenses in $\mathcal{A}(C)$ is $O(n^{4/3})$.*

We begin by considering the case of pseudo-parabolas; we then show that the other case can be reduced to this case, using the analysis given in the preceding subsections. We first prove several lemmas.

Lemma 2.15 *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 2.13. 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, so it must cross λ_0. 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 it intersects each of them at two points.

Suppose K_0 is crossed by a curve $\zeta \in C$ of type (i). Let λ_1 be the lens formed by ζ and γ'. We replace λ_0 with λ_1 in L. See Figure 10(i). The new set L' still consists of lenses with pairwise disjoint interiors, so in particular the lenses in L' are still pairwise nonoverlapping. Moreover, 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 deform γ or γ', thereby shrinking 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 10(ii). 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. Moreover, since the lenses in L are pairwise nonoverlapping and no pair of them share an endpoint, the deformation of δ' can be done in such a way that no other lens in L is affected. 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. $\square$

A pair (λ, λ') of lenses in L is called *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 be the number of crossing pairs of lenses in L, and let Y be the number of nested pairs of lenses in L.

Lemma 2.16 *Let C, L, X and Y be as above. Then*

$$|L| = O(n + X + Y). \tag{1}$$

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.

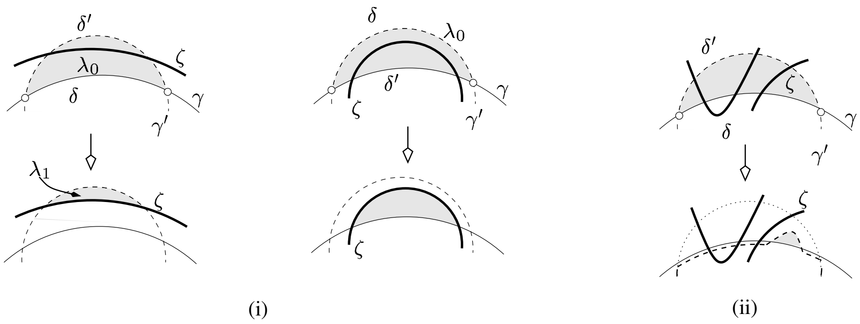

Figure 10. (i) Replacing λ_0 by a "smaller" lens if it intersects a type (i) curve. (ii) Shrinking λ_0 to an empty lens when it is crossed only by type (ii) curves.

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 2.15. $\square$

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

Lemma 2.17 $X = O(n^2)$.

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. $\square$

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

Lemma 2.18 *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 (as mentioned in the proof of Lemma 2.15, no curve can be fully contained in the interior of a lens of L). Because of the nonoverlap of the lenses of L and the general position assumption, the crossing point $\rho \cap \lambda$ uniquely identifies 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 $O(k|L|)$. $\square$

Proof of Theorem 2.14: Continue to assume that L is a collection of pseudo-parabolas, and 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 survive in R (i.e., both curves bounding the lens are chosen in R) is $|L|p^2$. Here 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 surviving in R is p^4). Similarly, the expected number X' of crossing pairs (λ, λ') in L' is Xp^4. By Lemmas 2.16 (applied to $\mathcal{A}(R)$), 2.17, and 2.18, we have

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

for an appropriate constant c. That is, we have

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

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 for the case of pseudo-parabolas.

Suppose next that C is a collection of pairwise intersecting pseudo-circles. We apply the sequence of reductions used in Section 2, and keep track of the "fate" of each lens in L, ensuring that they remain pairwise nonoverlapping. The transformations effected by Lemma 2.11 and Theorem 2.13 clearly do not violate this property. Moreover, when we pass to the subcollections C_i or $C_i \cup C_j$, the remaining lenses continue to be pairwise nonoverlapping. Finally, "opening-up" the pseudo-circles into pseudo-parabolas by cutting them with a ray may destroy some lenses of L, but the number of lenses of L that are cut by the ray is clearly only $O(n)$, so we can remove them from L and consider only the surviving lenses, to which the analysis just presented can be applied. $\square$

2.5 Cutting pairwise intersecting pseudo-circles into pseudo-segments

Let C be a family of n pairwise intersecting pseudo-parabolas or pseudo-circles that are not necessarily in general position. (This is the first time that we treat degenerate situations as well.) Recall that $\chi(C)$ denotes the minimum number of subarcs into which the curves in C need to be cut so that any two arcs intersect at most once. As noted, the analysis of Tamaki and Tokuyama [27] implies that $\chi(C) = O(\nu(C))$. Hence, if the curves in C are in general position, Theorem 2.14 implies that $\chi(C) = O(n^{4/3})$.

Remark. For the analysis of [27] to apply, one has to assume that the properties of C that are needed for the derivation of a bound on $\nu(C)$ also hold for any (random) sample of C. For example, here we assume that every pair of curves in C intersect, and this clearly holds for any subset of C. In later applications similar hereditary behavior also has to be verified, but we will not do it explicitly, as it will trivially hold in all cases.

Handling degeneracies. Suppose that the curves in C are in degenerate position. For technical reasons, we assume that, for the case of pseudo-circles, the curves are x-monotone. We will first deform them into a collection of curves in general position, then apply Theorem 2.14 to obtain the bound $O(n^{4/3})$ on $\nu(C')$, for the deformed collection C', then apply the analysis of Tamaki and Tokuyama to cut the curves of C' into $O(n^{4/3})$ pseudo-segments, and finally deform the cut curves of C', together with the cutting points, back to their original position.

In more detail, we proceed as follows. Let p be a point at which at least three curves of C are incident or at least two curves of C are tangent; any number of pairs of curves incident to p may be tangent to each other at p.[4] Draw a small axis-parallel rectangle $\gamma = \gamma_p$ centered at p, so that (i) the interior of γ does not contain any vertex of $\mathcal{A}(C)$ except for p; (ii) each curve incident to p intersects γ in exactly two points, which lie on the left and right edges of γ; and (iii) no curve that is not incident to p intersects γ. The x-monotonicity and continuity of the curves of C are easily seen to imply that such a γ exists. For each curve c that is incident to p, we replace the (connected) portion of c inside γ by the pair of straight segments connecting p to the two points of $c \cap \gamma$. See Figure 11(i).

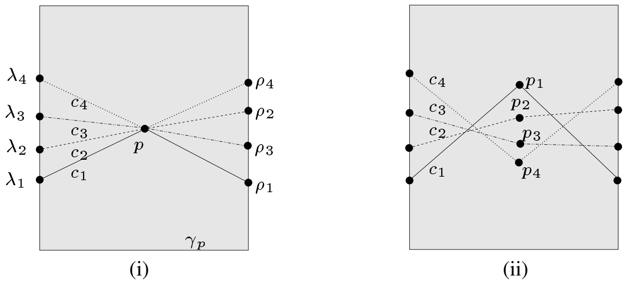

Figure 11. Perturbing arrangements in degenerate position: (i) Straightening the curves in the vicinity of a degenerate point p. (ii) Deforming the curves near p. (Note that c_2 and c_3 cross at p, while every other pair is tangent at p.)

For each curve $c_i \in C$ passing through p, let λ_i (resp., ρ_i) denote the intersection of c_i with the left (resp., right) edge of γ. Order the curves incident to p as $c_1, \ldots, c_j$, so that $\lambda_1, \ldots, \lambda_j$ appear in this increasing y-order along the left edge of γ. Replace p by a sequence of j distinct points $p_1, \ldots, p_j$ lying on the vertical line passing through p, and arranged along it in this decreasing y-order. For each $i = 1, \ldots, j$, replace the portion of c_i within γ by the two straight segments connecting λ_i and ρ_i to p_i; see Figure 11(ii).

It is easily verified that (i) each pair of original curves that were tangent at p are replaced by a pair of curves that cross twice within γ and (ii) each pair of original curves that crossed at p are replaced by a pair of curves that cross once within γ. This implies that the resulting curves are still a family of pairwise-intersecting pseudo-parabolas or x-monotone pseudo-circles, and, with an appropriate choice of the points $p_1, \ldots, p_j$, the portions of these curves within γ are in general position.

We repeat this perturbation in the neighborhood of each point that is incident to at least three curves or to at least one tangent pair. The final perturbed collection C' is still a family of pairwise intersecting pseudo-parabolas or x-monotone pseudo-circles, and they are now in general position. Applying, as above, the analysis of Tamaki and Tokuyama and Theorem 2.14, we can cut the curves in C' into $O(n^{4/3})$ pseudo-segments. Moreover, the cuts can be made in such a way that, for any curve c incident to a degenerate point p, its perturbed version c' is cut within the corresponding

[4]Note that it may be the case that (c_1, c_2) and (c_1, c_3) are two pairs of tangent curves at p, but c_2 and c_3 are *not* tangent; see Figure 11(i).

surrounding rectangle γ_p only if c' participates in a lens that is fully contained in γ_p, which is equivalent to the original curve c being tangent to some other curve(s) at p.

Finally, after having cut the perturbed curves, we deform them back to their original positions. If a perturbed curve c' was cut within some rectangle γ_p, we cut the original curve c at the center p itself. It is easily verified that the resulting collection of arcs is indeed a family of pseudo-segments. No two arcs are tangent to each other (in their relative interiors), but an endpoint of an arc may lie on (the relative interior of) another arc. We summarize this analysis in the following theorem.

Theorem 2.19 *Let C be a collection of n pairwise intersecting pseudo-parabolas or x-monotone pseudo-circles, not necessarily in general position. Then $\chi(C) = O(n^{4/3})$. (x-monotonicity need not be assumed for pseudo-circles in general position.)*

3 Bichromatic Lenses in Pseudo-Parabolas and Their Elimination

In this section we consider the following *bichromatic* extension of the problems involving empty and pairwise-nonoverlapping lenses, which is required as a main technical tool in the analysis of the general case, treated in Section 5, where not all pairs of the given pseudo-circles necessarily intersect.

We consider in this section only the case of pseudo-parabolas, which is simpler to handle. The case of pseudo-circles will be treated indirectly in Section 5. Moreover, we return to our initial assumption that the given curves are in general position. Degenerate cases will be treated later on. Let $\Gamma = A \cup B$ be a family of n pseudo-parabolas in general position, where $A \cap B = \emptyset$ and 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 another belonging to B is called *bichromatic*.

We first extend Theorem 2.4 to the bichromatic case, and show that the number of empty bichromatic lenses, in the setup assumed above, is $O(n)$. Then we obtain a bound of $O(n^{4/3})$ on the maximum size of a family of bichromatic pairwise nonoverlapping lenses. These results are obtained by pruning away some curves from Γ, so that the remaining curves are pairwise intersecting, and no lens in the family under consideration is lost. More specifically, we proceed as follows.

Theorem 3.1 *Let $\Gamma = A \cup B$ be a family of n pseudo-parabolas in general position, where $A \cap B = \emptyset$ and each pseudo-parabola of A intersects every pseudo-parabola of B twice. Then the number of empty bichromatic lenses in $\mathcal{A}(\Gamma)$ 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 Γ. 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 12(i). Hence, we may remove a from A without affecting the number of empty bichromatic lenses under consideration. Similarly, if b and b' are two disjoint curves in B, with b lying fully below b', then, for similar reasons, no empty bichromatic lens of the kind under consideration can be formed between b and any pseudo-parabola $a \in A$; see Figure 12 (ii). Hence, b may be removed from B without affecting the number of lenses that we are after.

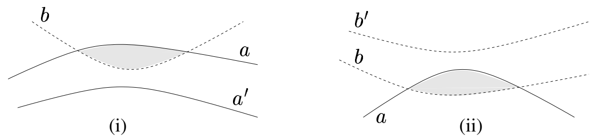

Figure 12. Discarding one of the nested pseudo-parabolas: (i) a is discarded, (ii) b is discarded.

We keep applying this pruning process until all pairs of remaining curves in $A \cup B$ intersect each other. By Theorem 2.4, the number of empty lenses in $\mathcal{A}(A \cup B)$ is $O(n)$. As discussed above, this completes the proof of the theorem. $\square$

In order to bound the maximum number of bichromatic pairwise-nonoverlapping lenses in Γ, we need the following lemma.

Lemma 3.2 *Let $\Gamma = A \cup B$ be a family of n pseudo-parabolas in general position, where $A \cap B = \emptyset$ and each pseudo-parabola of A intersects every pseudo-parabola of B twice. Let L be a family of pairwise-nonoverlapping lenses in $\mathcal{A}(\Gamma)$ that have pairwise disjoint interiors. Then $|L| = O(n)$.*

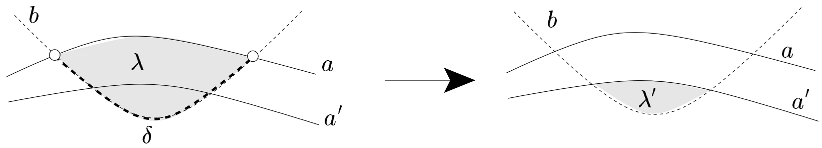

Figure 13. Transforming a lens into an empty lens.

Proof: As earlier, 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. As in the proof of Theorem 3.1, we argue that if there are two disjoint curves $a, a' \in A$ so that a' lies fully below a, then a can be pruned away. Let $\lambda \in L$ be a lens formed by a and by some curve $b \in B$. Let $\delta \subset b$ be the arc of b forming λ. Since $b \setminus \delta$ lies fully above a and thus above a', the curve a' must intersect δ at two points. 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 (and thus pairwise nonoverlapping), of the same size. Hence, by applying this replacement rule to each lens in L formed along a, we construct a family of pairwise-nonoverlapping lenses in which no lens is bounded by a, so we delete a from A. Hence, we can assume that all pairs of curves in A intersect. By applying a symmetric rule for pruning the curves of B, we can assume that every pair in B also intersect. Since every two curves in Γ intersect, the lemma follows from Theorem 2.4. $\square$

By proceeding as in Section 2.4 but using the above lemma instead of Lemma 2.15, we obtain the following result.

Lemma 3.3 *Let $\Gamma = A \cup B$ be a family of n pseudo-parabolas in general position, where $A \cap B = \emptyset$ and each pseudo-parabola of A intersects every pseudo-parabola of B twice. Let L be a family of pairwise-nonoverlapping bichromatic lenses in $\mathcal{A}(\Gamma)$. Then the size of L is $O(n^{4/3})$.*

As a result, we obtain the main result of this section.

Theorem 3.4 *Let $\Gamma = A \cup B$ be a family of n pseudo-parabolas, not necessarily in general position, where $A \cap B = \emptyset$ and each pseudo-parabola of A intersects every pseudo-parabola of B twice. Then one can cut the curves in Γ into $O(n^{4/3})$ arcs, so that each arc lying on a curve of A intersects every arc lying on a curve of B at most once.*

Proof: If the curves are in general position, this is an immediate corollary of the analysis of [27], in a similar manner to the application in Section 2.5. (As remarked there, we need to verify that the conditions assumed in the theorem also hold for subsets of A, B, which is clearly the case.) If A and B are in degenerate position, we apply the perturbation scheme used in Section 2.5. It is easily checked that this scheme maintains the property that each curve in A intersects every curve in B, so the bound on the number of cuts remains $O(n^{4/3})$ in this case too. $\Box$

4 Improving the Tamaki-Tokuyama Bound

In this section we improve the bound of Tamaki and Tokuyama [27] for arbitrary collections C of pseudo-parabolas or x-monotone pseudo-circles, and show that $\nu(C) = O(n^{8/5})$ in these cases.

4.1 The case of pseudo-parabolas

Theorem 4.1 *Let Γ be a family of n pseudo-parabolas (not necessarily in general position). Then $\chi(\Gamma) = O(n^{8/5})$.*

Proof: Let us first assume that the given collection is in general position, and handle the degenerate case towards the end of the proof, as in the preceding sections. Let Γ be a collection of n pseudo-parabolas in general position, and let L be a family of pairwise nonoverlapping lenses in Γ. Consider the graph $G = (\Gamma, L)$ as in Section 2.1. We draw G in the plane using the same drawing rule described in Section 2.1,[5] We partition Γ into two subsets Γ_1, Γ_2 of size at most $\lceil n/2 \rceil$ each so that for all $(\gamma_1, \gamma_2) \in \Gamma_1 \times \Gamma_2$, γ_1^* lies above γ_2^*. Let G' be the bipartite subgraph of G in which $E(G') = E(G) \cap (\Gamma_1 \times \Gamma_2)$. Then $|L| \le \nu(\Gamma_1) + \nu(\Gamma_2) + |E(G')|$.

By refining the rule described in Section 2.1 we draw G' so that the drawings of every pair of edges in G' that belong to a cycle of length 4 intersect an even number of times. By a result of Pinchasi and Radoičić [25], a graph on n vertices with this property has at most $O(n^{8/5})$ edges. Put $\nu(n) = \max_\Gamma \nu(\Gamma)$, where the maximum is taken over all sets Γ of n pseudo-parabolas in general position. Since $|\Gamma_1|, |\Gamma_2| \le \lceil n/2 \rceil$, we obtain the recurrence

$$\nu(n) \le 2\nu\left(\left\lceil \frac{n}{2} \right\rceil\right) + O(n^{8/5}),$$

whose solution is $\nu(n) = O(n^{8/5})$. This implies that $|L| = O(n^{8/5})$. This, plus the analysis in [27] implies that $\chi(\Gamma) = O(n^{8/5})$.

[5] We make a small technical modification in the statement of the rule: the wedge $W(\gamma_1, \gamma_2)$ is now defined to terminate on the right at the left intersection point of γ_1 and γ_2 (rather than at their tangency, as in Section 2.1).

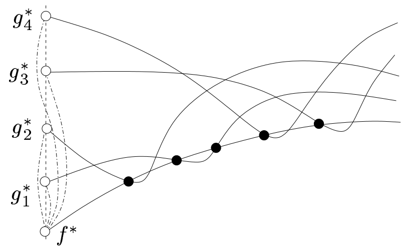

Figure 14. Illustrating the refined drawing rule for the plane embedding of G'. The lenses of L all appear along the bottommost curve, and each empty circle designates the left endpoint of a lens, and the apex of the corresponding wedge.

We first describe how to refine the drawing of G'. The drawing rule of Section 2.1 only specifies how the edges of G' have to "navigate" around intermediate vertices along the vertical line ℓ, but the rule does not specify the order in which edges emanate from a vertex. Let f^* be a vertex of the drawn graph G'. Let $g_1^*, \ldots, g_k^*$ be all the vertices above f^* that are connected to it by an edge. For each $1 \le i \le k$, let x_i be the x-coordinate of the leftmost intersection point between f and g_i. Order the g_i's so that $x_i < x_j$ whenever $i < j$. We then draw the edges $(f, g_1), \ldots, (f, g_k)$ so that they emanate from f^* upward in this clockwise order. See Figure 14.[6]

Symmetrically, for any given vertex f^* let $h_1^*, \ldots, h_m^*$ denote all the vertices below f^* that are connected to it by an edge. Order them, as above, in the left-to-right order of the leftmost intersection points between $h_1, \ldots, h_m$ and f. We draw the edges $(f, h_1), \ldots, (f, h_m)$ so that they emanate from f^* downward in this counterclockwise order. We call two edges of G' *adjacent* if they share an endpoint.

Claim A *The drawings of every pair of adjacent edges in G' cross an even number of times.*

Proof: We prove this only for two adjacent edges whose drawings go upward from a common vertex f^*; the argument for edges that go downward is fully symmetric. Let the other endpoints of these edges be g^* and h^*, and assume, without loss of generality, that h^* lies above g^*.

If the arc (f^*, h^*) passes to the left of g^*, then the leftmost intersection v_{gh} between h and g is to the left of the leftmost intersection v_{fh} between h and f (clearly, both intersections exist); see Figure 15(i). We claim that in this case v_{fh} lies to the left of the leftmost intersection v_{fg} between f and g. Indeed, assume to the contrary that v_{fh} lies to the right of v_{fg}. Then g must intersect h twice to the left of v_{fg} and then intersect f at least once to the left of v_{fg}. Moreover, since the lenses (f, g) and (f, h) are nonoverlapping, the rightmost intersection v'_{fg} of f and g must also lie to the left of v_{fh}; see Figure 15(i). But then, immediately to the right of v'_{fg}, the curve g is "trapped" in the wedge $W(f, h)$, since it has already intersected each of these curves twice. This contradiction implies that v_{fh} lies to the left of v_{fg}, and our modified drawing rule thus implies that (f^*, g^*) lies clockwise to (f^*, h^*) near f^*. Regarding the two edges as graphs of functions of y, and using the mean-value theorem, as in Section 2.1, we conclude that (f^*, g^*) and (f^*, h^*) intersect an even number of times.

If the arc (f^*, h^*) passes to the right of g^* then the leftmost intersection v_{fg} of f and g lies to the left of the leftmost intersection v_{fh} of f and h. See Figure 15(ii). Then our modified drawing rule implies that (f^*, g^*) lies counterclockwise to (f^*, h^*) near f^*. Arguing as above, this implies that these two edges intersect an even number of times, thus completing the proof of our claim. $\square$

[6]Note that in this figure, unlike Figure 2(ii), we do not draw the lenses as tangencies, since they need not be empty.

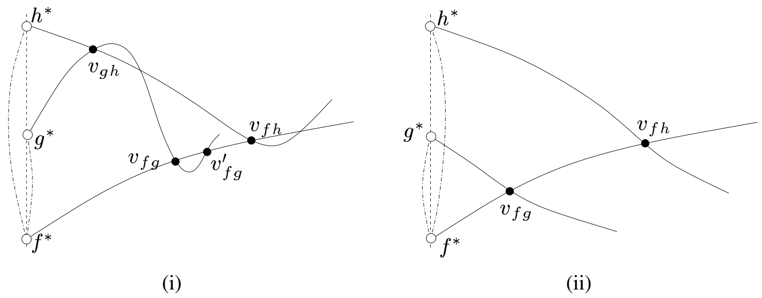

Figure 15. Illustrating the proof that adjacent edges of G' intersect an even number of times. (i) The case where (f^*, h^*) passes to the left of g^*. (ii) The case where (f^*, h^*) passes to the right of g^*.

Claim B *If (f, p, g, q) is a cycle of length four in G', then the curves f, p, g, and q are pairwise intersecting.*

Proof: This clearly holds for each pair of curves whose corresponding vertices are adjacent in the cycle, so the only pairs that need to be analyzed are the pair f, g and the pair p, q. We show that f, g must intersect each other, and the argument for p, q is similar. Assume to the contrary that f and g are disjoint and, without loss of generality, that f lies always above g. Trace the curve p from left to right. It starts above f, g and it creates a lens with each of f and g. Clearly, p must first intersect f, but then it cannot intersect g before it intersects f again, for otherwise the lenses (p, f) and (p, g) would be overlapping. However, after p intersects f for the second time, it cannot intersect g anymore, since f now separates these two curves. See Figure 16 (i). This contradiction implies that f, p, g, q are pairwise intersecting. $\square$

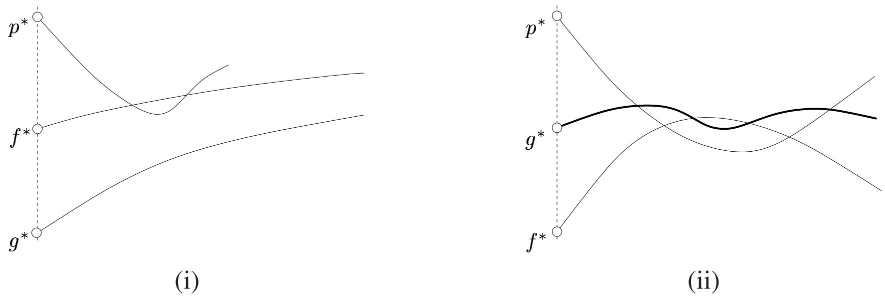

Figure 16. (i) All the pairs of curves that correspond to the given 4-cycle must intersect. (ii) The lenses that correspond to the 4-cycle are all empty relative to the four curves f, p, g, q.

Claim C *If (f, p, g, q) is a cycle of length four in G', then the four lenses corresponding to the cycle are empty with respect to the arrangement of these four curves.*

Proof: Consider any of these four lenses, say (f, p), and assume that either g or q intersects it. Since the two cases are similar, we only consider the case where g intersects (f, p). g cannot intersect the arc of (f, p) that belongs to p, for then (f, p) and (g, p) would be overlapping. It follows that g must intersect twice the arc of (f, p) that belongs to f; see Figure 16 (ii). In this case, since g starts below p, g must intersect p once to the left of the lens (f, p) and once to its right, in which case the two lenses (f, p) and (g, p) are overlapping, a contradiction that implies the claim. $\square$

Finally, let (f, p, g, q) be a cycle of length four in G'. By Claim A, the drawings of each of the four pairs of adjacent edges intersect an even number of times. By Claims B and C, the lenses (f, p) and (g, q) are empty in the family of the four pairwise intersecting pseudo-parabolas f, p, g, q. It now follows from the analysis of Section 2.1 that the drawings of (f, p) and (g, q) intersect an even number of times. Similarly, we can argue that the drawings of (f, q) and (g, p) intersect an even number of times, thereby implying that the drawings of every pair of edges in the above cycle intersect in an even number of times. Hence, $|E(G')| = O(n^{8/5})$, by the result in [25].

This completes the proof of the theorem for curves in general position. In the degenerate case we proceed exactly as in Section 2.5, concluding that $\chi(\Gamma) = O(n^{8/5})$ in these cases too. $\square$

4.2 The case of pseudo-circles

We next extend Theorem 4.1 to the case of *x-monotone* pseudo-circles. The corresponding extension to the case of arbitrary pseudo-circles remains an open problem, although we expect it to hold just as well. Let C be a family of n x-monotone pseudo-circles. For any closed and bounded x-monotone Jordan curve c in the plane, denote by λ_c (resp., ρ_c) the leftmost (resp., rightmost) point of c, assuming these points to be well defined. The points λ_c, ρ_c partition c into two x-monotone arcs, called *upper* and *lower* arcs and denoted as c^+, c^-, respectively; see Figure 17 (i).

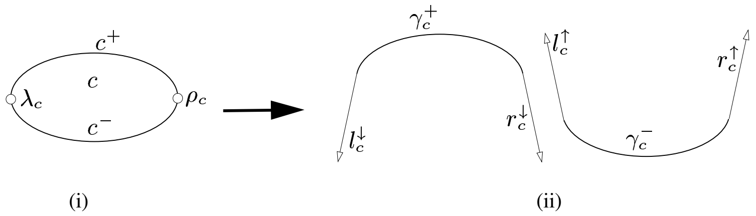

Figure 17. Converting a pseudo-circle into two pseudo-parabolas.

We convert C into a family of pseudo-parabolas. For each $c \in C$, we extend its upper arc c^+ to an x-monotone curve γ_c^+ by adding a downward (almost vertical) ray $l_c^\downarrow$ (resp., $r_c^\downarrow$) of sufficiently large positive (resp., negative) slope from λ_c (resp., ρ_c); all rays emanating from the left (resp., right) endpoints of the pseudo-circles are parallel. Similarly we extend every c^- to an x-monotone curve γ_c^- by attaching upward (almost vertical) rays $l_c^\uparrow$ and $r_c^\uparrow$ to λ_c and ρ_c, respectively. We assume that the rays are chosen sufficiently steep so that a downward (resp., upward) ray intersects a pseudo-disk of C only if it lies vertically below (resp., above) the apex of the ray. If x-coordinates of the left (or right) endpoints are are not all distinct, then we draw the rays as earlier, but they have slightly different slopes. For example, we draw the rays $\lambda_c^\downarrow$ as follows. We sort the left endpoints of all the curves in C in nondecreasing order of their x-coordinates. If two endpoints have the same x-coordinates, then we sort them in nonincreasing order of their y-coordinates. If two curves have the same left endpoint, i.e., they are tangent at their left endpoints and one of the curves lies inside the other, then the left endpoint of the outer curve appears first. Let Λ^- be the resulting sequence of left endpoints. We choose a sufficiently large slope σ, as above, and a sufficiently small parameter δ. For the ith left endpoint λ_c in Λ, we draw a downward ray $l_c^\downarrow$ of slope $\sigma + i\varepsilon$. The interiors of these rays are pairwise disjoint, and they are parallel for all practical purposes. We do the same for the other three types of rays to handle degeneracies. We now prove that the resulting curves form a family of pseudo-parabolas.

Lemma 4.2 *Let C be a finite family of x-monotone pseudo-circles. Then $\Gamma = \{\gamma_c^+, \gamma_c^- \mid c \in C\}$ is a family of pseudo-parabolas.*

Proof: For simplcity we prove the lemma for the case in which the x-coordinates of the extremal points on the curves of C are all distinct. With a little care, the proof can be extended to the general case. Let a and b be two pseudo-circles in C. We first prove that γ_a^+ and γ_b^+ intersect in at most two points. For simplicity, for a curve $c \in C$, we will use l_c, r_c to denote the rays $l_c^\downarrow$ and $r_c^\downarrow$, respectively. Without loss of generality assume that λ_a lies to the left of λ_b; then the ray l_a does not intersect γ_b^+. There are three cases to consider:

Case (A): λ_b lies to the right of ρ_a: In this case the only intersection between γ_a^+ and γ_b^+ is between the rays l_b and r_a (see Figure 18 (A)).

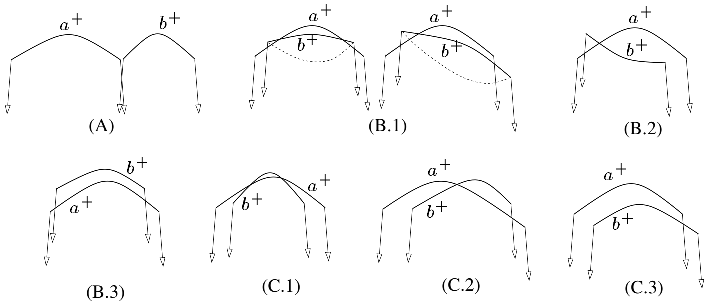

Figure 18. Two extended upper arcs intersect at most twice: (A) ρ_a lies to the left of λ_b; (B) λ_b lies above a^+: (B.1) a^+, b^+ intersect at two points or they intersect at one point but ρ_b lies to the right of ρ_a; (B.2) a^+ and b^+ intersect at one point and ρ_b lies to the left of ρ_a; (B.3) a^+ and b^+ do not intersect. (C) λ_b lies below a^+: (C.1) a^+ and b^+ intersect at two points and ρ_b lies to the left of ρ_a; (C.2) a^+ and b^+ intersect at one point; (C.3) a^+ and b^+ do not intersect.

Case (B): λ_b lies above a^+. In this case l_b intersects a^+, so we show that there is at most one additional intersection point between γ_a^+ and γ_b^+. If a^+ and b^+ intersect at two points or if a^+ and b^+ intersect at one point but ρ_b lies to the right of ρ_a, then a and b intersect in at least four points (see Figure 18 (B.1)), contradicting the assumption that C is a family of pseudo-circles. If a^+ and b^+ intersect at one point and ρ_b lies to the left of ρ_a (and, necessarily, below a^+), then neither r_a intersects γ_b^+ (r_a lies to the right of b^+) nor r_b intersects γ_a^+ (r_b lies below a^+); see Figure 18 (B.2). Hence, there are only two intersection points between γ_a^+ and γ_b^+.

If a^+ and b^+ do not intersect, then r_a cannot intersect γ_b^+, as it lies below b^+. Hence, only r_b may intersect γ_a^+ (if ρ_a lies to the right of ρ_b), thereby showing that there are at most two intersection points between γ_a^+ and γ_b^+; see Figure 18 (B.3).

Case (C): λ_b lies below a^+. In this case l_b does not intersect a^+. If a^+ intersects b^+ at two points and ρ_b lies to the right of ρ_a, then a and b intersect in at least four points, a contradiction (the situation is similar to that shown in Figure 18 (B.1)). If they intersect at two points but ρ_b lies to the left of ρ_a, then neither r_a intersects b^+ nor r_b intersects a^+, so there are at most two intersection points between γ_a^+, γ_b^+; see Figure 18 (C.1).

If a^+ and b^+ intersect at one point, then r_a cannot intersect γ_b^+ (see Figure 18 (C.2)), so the number of intersection points between γ_a^+ and γ_b^+ is easily seen to be at most two. Finally, if a^+ and b^+ do not intersect, then there is at most one intersection between γ_a^+ and γ_b^+, namely between r_a and b^+ (if ρ_b lies to the right of ρ_a); see Figure 18 (C.3).

Hence, in all cases, there are at most two intersection points between γ_a^+ and γ_b^+. A symmetric argument shows that γ_a^- and γ_b^- also intersect at most twice. Finally, a similar case analysis, depicted in Figure 19, shows that γ_a^- and γ_b^+ also intersect at most twice. We leave it to the reader to fill in the fairly straightforward details, similar to those given above. $\square$

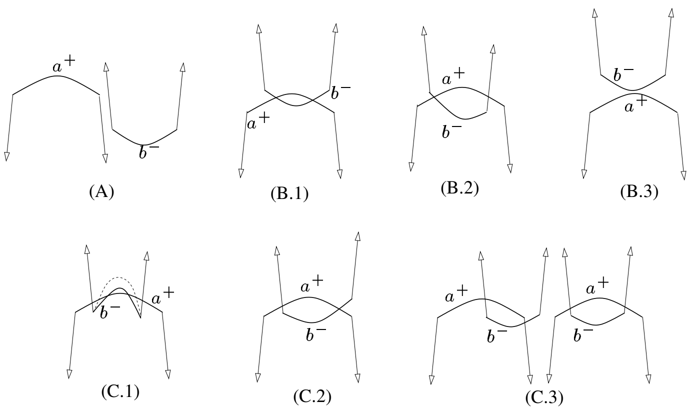

Figure 19. An extended upper arc and an extended lower arc intersect at most twice: (A) ρ_a lies to the left of λ_b; (B) λ_b lies above a^+: (B.1) a^+, b^- intersect at two points; (B.2) a^+ and b^- intersect at one point; (B.3) a^+ and b^- do not intersect. (C) λ_b lies below a^+: (C.1) a^+ and b^- intersect at two points (an impossible configuration); (C.2) a^+ and b^- intersect at one point; (C.3) a^+ and b^- do not intersect.

Theorem 4.3 *Let C be an arbitrary family of n x-monotone pseudo-circles in the plane. Then $\chi(C) = O(n^{8/5})$.*

Proof: Assume first that the curves in C are in general position. Let L be a family of pairwise-nonoverlapping lenses in C. We convert C into a family $\Gamma = \{\gamma_c^+, \gamma_c^- \mid c \in C\}$ of $2n$ pseudo-parabolas, as described above. There are at most $2n$ lenses in L that contain λ_c, ρ_c of a curve $c \in C$ on its boundary, as the lenses in L are nonoverlapping. Any remaining lens lies on the upper or the lower arc of a pseudo-circle in C, and therefore it lies in the transformed collection Γ of pseudo-parabolas. By Theorem 4.1, the number of such lenses is $O(n^{8/5})$. Hence, $|L| = O(n^{8/5})$, which implies the claim for curves in general position. The case of degenerate position is handled exactly as in Section 2.5. $\square$

5 Curves with 3-Parameter Algebraic Representation

In this section we further improve the bound obtained in the previous section, and derive a bound close to $n^{3/2}$ for a few important special cases, in which the curves possess what we term as a 3-*parameter algebraic representation*. As in Sections 2 and 4, we first prove the bound for pseudo-parabolas and then reduce the case of pseudo-circles to that of pseudo-parabolas.

5.1 The case of pseudo-parabolas

Let Γ be a family of n pseudo-parabolas. We say that Γ has a 3-*parameter algebraic representation* if Γ is a finite subset of some infinite family $\mathbf{P}$ of curves so that each curve $\gamma \in \mathbf{P}$ can be represented by a triple of real parameters (ξ, η, ζ), which we regard as a point $\gamma^* \in \mathbb{R}^3$, so that the following three conditions are satisfied.

(AP1) For each point q in the plane, the locus of all curves in $\mathbf{P}$ that pass through q is, under the assumed parametrization, a 2-dimensional surface patch in $\mathbb{R}^3$, which is a semialgebraic set of *constant description complexity*, i.e., it is defined as a Boolean combination of a constant number of polynomial equations and inequalities of constant maximum degree. For any two distinct points p and q in the plane, the locus of all curves in $\mathbf{P}$ that pass through both p and q is, under the assumed parametrization, a 1-dimensional semialgebraic curve of constant description complexity.

(AP2) For each curve $\gamma \in \mathbf{P}$, the set of all curves $g \in \mathbf{P}$ that intersect γ maps to a 3-dimensional semialgebraic set K_γ of constant description complexity. The boundary of K_γ, denoted by τ_γ, is the locus of all curves in $\mathbf{P}$ that are tangent to γ (and, being pseudo-parabolas, do not meet γ at any other point); τ_γ partitions $\mathbb{R}^3$ into two regions, one of which is K_γ and the other consists of points representing curves that are disjoint from γ.

(AP3) Each curve in $\mathbf{P}$ is a semialgebraic set of constant description complexity in the plane, and the family $\mathbf{P}$ is closed under translations.

We remark that condition (AP1) is not needed for obtaining bounds on $\nu(\Gamma)$ and $\chi(\Gamma)$. It is used for obtaining improved bounds for the number of incidences between points and the curves in Γ, and for the complexity of many faces in $\mathcal{A}(\Gamma)$; see Section 6 for details. The class of *vertical* parabolas, given by equations of the form $y = ax^2 + bx + c$, is an example of pseudo-parabolas having a 3-parameter algebraic representation, where each parabola is represented by the triple of its coefficients.

Suppose then that $\mathbf{P}$ is a fixed collection of pseudo-parabolas that have a 3-parameter algebraic representation, and let $\Gamma \subset \mathbf{P}$ be a family of n pseudo-parabolas.

Our plan of attack, similar to those employed in [7, 8], is to decompose the *intersection graph* H of Γ (whose edges represent all intersecting pairs of curves in Γ) into a union of complete bipartite graphs $\{A_i \times B_i\}_i$, so that, for each $a \in A_i$, $b \in B_i$, a intersects b. We then use Theorem 3.4 to derive an upper bound on the number of cuts needed to eliminate all bichromatic lenses in $A_i \times B_i$. We repeat this process for each complete bipartite graph $A_i \times B_i$, and add up the numbers of cuts to derive the overall bound on $\chi(\Gamma)$.

In more detail, we proceed as follows. Let $\Gamma^* = \{\gamma^* \mid \gamma \in \Gamma\}$, and $\hat{\Gamma} = \{\tau_\gamma \mid \gamma \in \Gamma\}$. We describe a recursive scheme to generate the desired bipartite decomposition of the intersection graph of Γ. At each step, we have two families $A, B \subseteq \Gamma$, of size m and n, respectively. Let $\chi(A, B)$ denote the minimum number of cuts needed to eliminate all bichromatic lenses in $\mathcal{A}(A \cup B)$. Set $\chi(m, n) = \max\ \chi(A, B)$ where the maximum is taken over all families of m and n pseudo-parabolas of $\mathbf{P}$, respectively. Set $\chi(m) = \chi(m, m)$. We need to introduce a few concepts before beginning with the analysis of $\chi(m)$.

For any constant integer q, let $\lambda_q(r)$ denote the maximum length of Davenport-Schinzel sequences of order q composed of r symbols [26]. Put $\beta_q(r) = \lambda_q(r)/r$. In what follows, we sometimes drop the parameter q, and write $\beta_q(r)$ simply as $\beta(r)$. Assuming q to be even, we have $\beta_q(r) = 2^{O(\alpha(r)^{(q-2)/2})}$, where $\alpha(r)$ is the extremely slowly growing inverse Ackermann function. See [26] for more details. Let $\tau \subseteq \mathbb{R}^3$ be a simply connected region of constant description complexity. For a set G of surfaces in $\mathbb{R}^3$, we define the conflict list $G_\tau \subseteq G$ with respect to G to be the set of surfaces that intersect τ but do not contain τ. Each surface in G_τ either crosses τ, or it is tangent to τ.

Lemma 5.1 *For any m, n and for any given parameter $1 \le r \le \min\{m^{1/3}, n\}$,*

$$\chi(m,n) \le cr^3\beta_q(r)\left[\chi\left(\frac{m}{r^3},\frac{n}{r}\right) + O((m+n)^{4/3})\right], \tag{2}$$

where q is a constant that depends on the family $\mathbf{P}$, and c is an absolute constant.

Proof: Let $A, B \subset \mathbf{P}$ be two families of m and n pseudo-parabolas, respectively. Let $\hat{B} = \{\tau_b \mid b \in B\}$. For a parameter $1 \le r \le n$, a $(1/r)$-*cutting* Ξ of the arrangement $\mathcal{A}(\hat{B})$ is a decomposition of $\mathbb{R}^3$ into relatively open and simply connected cells of dimensions $0, 1, 2, 3$, each having constant description complexity, so that the size of the conflict list of each cell with respect to $\hat{B}$ is at most n/r. Since each τ_b is a two-dimensional algebraic set of constant description complexity, it follows from the results in [2, 3] that there exists a $(1/r)$-cutting Ξ of size $O(r^3\beta_q(r))$, where q is 2 plus the maximum number $s' = s'(\gamma_1,\gamma_2,\gamma_3,\gamma_4)$, over all quadruples of curves $\gamma_1,\gamma_2,\gamma_3,\gamma_4$ in $\mathbf{P}$, of vertical lines ℓ that pass through both intersection curves $\tau_{\gamma_1} \cap \tau_{\gamma_2}$ and $\tau_{\gamma_3} \cap \tau_{\gamma_4}$ in $\mathbb{R}^3$. More precisely, $s'(\gamma_1,\gamma_2,\gamma_3,\gamma_4)$ is the number of connected components of the union of all these vertical lines; equivalently, it is the number of connected components of the intersection of the vertical projections of $\tau_{\gamma_1} \cap \tau_{\gamma_2}$ and $\tau_{\gamma_3} \cap \tau_{\gamma_4}$.

We construct such a $(1/r)$-cutting Ξ of $\hat{B}$. For each cell $\Delta \in \Xi$, let $A_\Delta = \{\gamma \in A \mid \gamma^* \in \Delta\}$. If $|A_\Delta| > m/r^3$, we cut Δ further into subcells (e.g., by planes parallel to some generic direction), each containing at most m/r^3 points. The number of cells remain asymptotically $O(r^3\beta_q(r))$. For each (new) cell Δ, let $\tilde{B}_\Delta = \{b \in B \mid \Delta \subseteq K_b\}$, i.e., any curve in $\tilde{B}_\Delta$ intersects all curves of A_Δ (if $\Delta \subseteq \partial K_b$, then b is tangent to all curves in A_Δ), and let B_Δ be the set of curves corresponding to the conflict list of Δ with respect to $\hat{B}$.

It follows by construction that

$$\chi(A,B) \le \sum_{\Delta\in\Xi}[\chi(A_\Delta,B_\Delta) + \chi(A_\Delta,\tilde{B}_\Delta)].$$

Since every pair of pseudo-parabolas in $A_\Delta \times \tilde{B}_\Delta$ intersect, by Theorem 3.4, $\chi(A_\Delta,\tilde{B}_\Delta) = O((|A_\Delta| + |\tilde{B}_\Delta|)^{4/3}) = O((m+n)^{4/3})$. Since $|A_\Delta| \le m/r^3$ and $|B_\Delta| \le n/r$ (the latter inequality holds for the original cells of Ξ, before any cell with two many points of A^* has been split, and it thus also holds for each split cell), we have $\chi(A_\Delta,B_\Delta) \le \chi(m/r^3, n/r)$. This completes the proof of the lemma. □

Flipping the roles of A and B, i.e., mapping B to a set of points and A to a set of surfaces in $\mathbb{R}^3$, and applying the same decomposition scheme, we obtain

$$\chi(m,n) \le cr^3\beta_q(r)\left[\chi\left(\frac{m}{r},\frac{n}{r^3}\right) + O((m+n)^{4/3})\right]. \tag{3}$$

Substituting (3) into the right-hand side of (2), we obtain

$$\chi(m) \le c^2r^6\beta_q^2(r)\chi\left(\frac{m}{r^4}\right) + O(m^{4/3}r^6\beta_q^2(r)).$$

Choosing $r = m^{1/36}$, we obtain

$$\chi(m) \le c_1 m^{1/6}\beta_q^2(m)) \cdot \chi(m^{8/9}) + c_1 m^{3/2}\beta_q^2(m) \tag{4}$$

for an appropriate constant $c_1 \geq 1$. We claim that the solution of this recurrence is

$$\chi(m) \leq m^{3/2}(\log m)^{c' \log \beta_q(m)} \tag{5}$$

where $c' \geq 1$ is a sufficiently large constant. This can be proved by induction on m, as follows. We may assume that (5) holds for all $m \leq m_0$, where m_0 is a sufficiently large constant that satisfies $(\log m)^{c' \log \beta_q(m)} \geq 2c_1\beta_q^2(m)$ for all $m > m_0$. Plugging (5) into (4), we obtain, for $m > m_0$,

$$\begin{aligned}
\chi(m) &\leq c_1 m^{1/6}\beta_q^2(m)m^{4/3}\left(\log(m^{8/9})\right)^{c' \log \beta_q(m)} + c_1 m^{3/2}\beta_q^2(m) \\
&\leq c_1 m^{3/2}(\log m)^{c' \log \beta_q(m)}\beta_q^2(m)\left(\frac{8}{9}\right)^{c' \log \beta_q(m)} + c_1 m^{3/2}\beta_q^2(m) \\
&\leq m^{3/2}(\log m)^{c' \log \beta_q(m)}\left(c_1\beta_q^{2+c' \log(8/9)}(m) + \frac{1}{2}\right) \\
&\leq c_1 m^{3/2}(\log m)^{c' \log \beta_q(m)},
\end{aligned}$$

provided that the constant c' is chosen sufficiently large. This establishes the induction step and thus proves (5). Recall that $\beta_q(n) = 2^{O(\alpha^s(n))}$, where $\alpha(n)$ is the inverse Ackermann function and $s = \lceil (q-2)/2 \rceil$ is a constant. Putting

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

and using the fact that, initially, $|A|, |B| \leq n$, we obtain the following main result of this section:

Theorem 5.2 *Let* $\mathbf{P}$ *be a collection of pseudo-parabolas that admits a 3-parameter algebraic representation. Then* $\chi(\Gamma) = O(n^{3/2}\kappa_s(n))$*, for any subset* Γ *of* n *elements of* $\mathbf{P}$*, and for some constant parameter* s *that depends on the algebraic representation of the curves in* $\mathbf{P}$.

Remark. In what follows, we will sometimes raise $\kappa_s(n)$ to some fixed power, or multiply it by a polylogarithmic factor, or replace n by some fixed power of n. These operations do not change the asymptotic form of the expression—they merely affect the constant of proportionality in the exponent. For the sake of simplicity, we use the notation $\kappa_s(n)$ to denote these modified expressions as well. We allow ourselves this freedom because we strongly believe that the factor $\kappa_s(n)$ is just an esoteric artifact of our analysis, and has nothing to do with the real bound, which we conjecture to be $o(n^{3/2})$.

5.2 The case of vertical parabolas

As a first application of Theorem 5.2, consider the family $\mathbf{V}$ of *vertical* parabolas, each of which is given by an equation of the form $y = ax^2 + bx + c$. Every vertical parabola has a natural 3-parameter representation, by the triple (a, b, c) of its coefficients, and $\mathbf{V}$ trivially satisfies (AP3).

For a fixed point $p = (\alpha, \beta) \in \mathbb{R}^2$, the set of vertical parabolas $y = \xi x^2 + \eta x + \zeta$ passing through p is the plane

$$\alpha^2\xi + \alpha\eta + \zeta = \beta,$$

which is obviously a two-dimensional semialgebraic set of constant description complexity. Similarly, the locus of parabolas that pass through two distinct points p, q is either empty or a 1-dimensional curve of constant description complexity. Thus (AP1) is satisfied.

Finally, for a fixed parabola $\gamma : y = ax^2 + bx + c$, another vertical parabola $y = \xi x^2 + \eta x + \zeta$ is tangent to γ if and only if

$$(\eta - b)^2 - 4(\xi - a)(\zeta - c) = 0.$$

Hence, the surface τ_γ is given by the equation

$$(\eta^2 - 4\xi\zeta) - 2b\eta + 4c\xi + 4a\zeta + (b^2 - 4ac) = 0, \tag{6}$$

which is a quadric in $\mathbb{R}^3$, and thus (AP2) is also satisfied. In order to estimate the value of $s = \lceil s'/2 \rceil$, recall that s' satisfies the following condition: Given any four curves $\gamma_1, \ldots, \gamma_4 \in \mathbf{P}$, there are at most s' intersection points between the $\xi\eta$-projections of the intersection curves $\sigma_{12} = \tau_{\gamma_1} \cap \tau_{\gamma_2}$ and $\sigma_{34} = \tau_{\gamma_3} \cap \tau_{\gamma_4}$.

It follows from (6) that the intersection curve σ_{12} of two surfaces τ_{γ_1} and τ_{γ_2} is a planar curve, whose projection on the $\xi\eta$-plane ($\zeta = 0$) is a quadric. Hence, the projections of σ_{12} and σ_{34} on the $\xi\eta$-plane intersect in at most four points, implying that $s' \le 4$ and $s \le 2$. (These bounds also apply in case of partial overlap between the projections.) Letting

$$\kappa(n) = \kappa_2(n) = (\log n)^{O(\alpha^2(n))},$$

we obtain the following.

Theorem 5.3 *Let Γ be a set of n vertical parabolas in the plane; then $\chi(\Gamma) = O(n^{3/2}\kappa(n))$.*

5.3 The case of pseudo-circles

We now prove a near $n^{3/2}$-bound on the maximum number of pairwise-nonoverlapping lenses for a few special classes of pseudo-circles. In addition to the condition of 3-parameter algebraic representation, which we define in a slightly different manner, we also require, as in Section 4, that the pseudo-circles be *x-monotone*. We say that an infinite family $\mathbf{C}$ of x-monotone pseudo-circles has a 3-*parameter algebraic representation* if every curve c can be represented by a triple of real parameters (ξ, η, ζ), which we regard as a point $c^* \in \mathbb{R}^3$, so that the following three conditions are satisfied.

(AC1) For each point q in the plane, the locus of all curves in $\mathbf{C}$ that pass through q is, under the assumed parametrization, a 2-dimensional semialgebraic set σ_q of constant description complexity. For any two distinct points p and q in the plane, the locus of all curves in $\mathbf{C}$ that pass through both p and q is, under the assumed parametrization, a 1-dimensional semialgebraic curve of constant description complexity.

(AC2) For each curve $c \in \mathbf{C}$, the locus of all curves $g \in \mathbf{C}$ whose upper (resp., lower) arc intersects the upper arc c^+ of c at two points is a 3-dimensional semialgebraic set $K^+_{c^+}$ (resp., $K^-_{c^+}$) of constant description complexity. The same also holds for the lower arc c^- of c.

(AC3) Each curve in $\mathbf{C}$ is a semialgebraic set of constant description complexity in the plane, and the family $\mathbf{C}$ is closed under translations.

Let $\mathbf{C}$ be a family of x-monotone pseudo-circles having a 3-parameter algebraic representation, and let $C \subseteq \mathbf{C}$ be a subset of n pseudo-circles. We replace C by the collection $\Gamma = \{c^+, c^- \mid$

$c \in C\}$, which by Lemma 4.2, is a collection of pseudo-parabolas. By Theorem 5.2, $\chi(\Gamma) = O(n^{3/2}\kappa_s(n))$, for an appropriate constant parameter s. We now cut the curves in C at the same points where their top or bottom boundaries have been cut in Γ, and, in addition, cut each curve $c \in C$ at the two extreme points λ_c, ρ_c. It follows trivially that the resulting subarcs form a collection of pseudo-segments. We thus have:

Theorem 5.4 *Let* $\mathbf{C}$ *be a collection of pseudo-circles that satisfies (AC1)–(AC3). Then* $\chi(C) = O(n^{3/2}\kappa_s(n))$*, for any subset* C *of* n *elements of* $\mathbf{C}$*, and for some constant parameter* s *that depends on* $\mathbf{C}$*.*

5.4 The case of circles

The most obvious application of Theorem 5.4 is to the family $\mathbf{C}$ of all circles in the plane. $\mathbf{C}$ trivially satisfies condition (AC3). We map each circle $c : (x - \xi)^2 + (y - \eta)^2 = \zeta^2$ to the point $c^* = (\xi, \eta, \zeta) \in \mathbb{R}^3$. The set of points $c^* = (\xi, \eta, \zeta) \in \mathbb{R}^3$ corresponding to circles c that pass through a fixed point $p = (\alpha, \beta)$ is the region

$$\sigma_p = \{(\xi, \eta, \zeta) \mid (\xi - \alpha)^2 + (\eta - \beta)^2 = \zeta^2\},$$

which is a 2-dimensional cone in 3-space. Moreover, using a standard transformation [14], we can map these surfaces into planes, without changing the incidence pattern between points and surfaces. Similarly, the locus of circles that pass through two distinct points p, q is, in the new representation, the line of intersection of the two corresponding planes. Hence, (AC1) is satisfied.

Concerning condition (AC2), it is straightforward to verify that the set of (points in $\mathbb{R}^3$ representing) circles that satisfy the condition that their upper arc, say, intersect the upper arc of a fixed given circle at two points, is a semialgebraic set of constant description complexity (an explicit expression for this set is given in Appendix A). However, naïve calculations that exploit this condition to derive a recurrence similar to that in Lemma 5.1, yield (bounds on the) constants s' and s that are somewhat high. Using a more sophisticated, but somewhat tedious, analysis, one can lower the constants to $s' = 4$ and $s = 2$. The details of this analysis are given in Appendix A.

Writing, as above, $\kappa(n)$ for $\kappa_2(n)$, we thus obtain:

Theorem 5.5 *Let* C *be a set of* n *circles in the plane; then* $\chi(C) = O(n^{3/2}\kappa(n))$.

5.5 The case of homothetic copies of a strictly convex curve

Theorem 5.4 can also be applied to the family $\mathbf{C}$ of homothetic copies of a fixed strictly convex curve γ_0 having constant description complexity. First, as already noted in [21], $\mathbf{C}$ is indeed a family of pseudo-circles (this does not necessarily hold if γ_0 is not strictly convex). Clearly, condition (AC3) is satisfied. Each homothetic copy of γ_0 has the form

$$(\xi, \eta) + \lambda\gamma_0 \equiv \{(\xi, \eta) + \lambda(x, y) \mid (x, y) \in \gamma_0\},$$

for some triple of real parameters $\xi, \eta \in \mathbb{R}$, $\lambda \in \mathbb{R}^+$. We represent each copy by the corresponding triple $(\xi, \eta, \lambda) \in \mathbb{R}^3$. Condition (AC1) is easy to establish: For a fixed point p, the condition $p \in (\xi, \eta) + \lambda\gamma_0$ is equivalent to $\frac{1}{\lambda}(p - (\xi, \eta)) \in \gamma_0$, which clearly defines a semialgebraic surface patch of constant description complexity.

For a pair p, q of distinct points, each homothetic copy of γ_0 that passes through p and q satisfies $\frac{1}{\lambda}(p-(\xi,\eta)) \in \gamma_0$, $\frac{1}{\lambda}(q-(\xi,\eta)) \in \gamma_0$. Hence $(p-q)/\lambda$ is a chord of γ_0. Since λ_0 is strictly convex, for each fixed λ there is a *unique* chord equal to $(p-q)/\lambda$, so ξ, η are also uniquely determined. Hence the locus of copies of γ_0 that pass through p and q is a 1-dimensional curves, which clearly has constant description complexity.

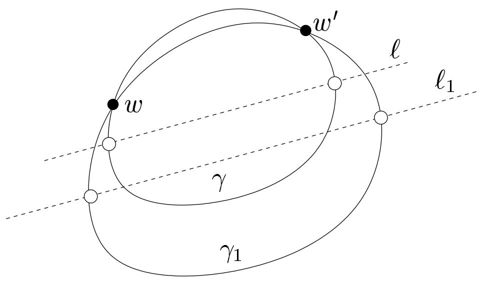

Figure 20. Upper arcs of two homothetic copies of γ_0 intersecting at two points.

Establishing condition (AC2) is a bit more technical. For a fixed homothetic copy $\gamma_1 = (\alpha, \beta, \mu)$ of γ_0, the condition that another homothetic copy $\gamma = (\xi, \eta, \lambda)$ be such that, say, its upper arc meets the upper arc of γ_1 at two points, can be expressed by the following predicate:

> There exist $w, w' \in \mathbb{R}^2$ such that $\{w, w'\} = \gamma_1 \cap \gamma$ and each of w, w' lies above both lines ℓ_1 and ℓ, where ℓ_1 (resp., ℓ) is the line connecting the leftmost and rightmost points of γ_1 (resp., γ).

See Figure 20. Using the fact that γ_0 is a semialgebraic set of constant description complexity, it follows that the above predicate also defines a semialgebraic set of constant description complexity; see [9, 10] for properties of real semialgebraic sets that imply this claim. Theorem 5.4 thus implies the following.

Theorem 5.6 *Let γ_0 be a convex curve of constant description complexity, and let C be a set of n homothetic copies of γ_0. Then $\chi(C) = O(n^{3/2}\kappa_s(n))$, for some constant s that depends on γ_0.*

6 Applications

The preceding results have numerous applications to problems involving incidences, many faces, levels, distinct distances, and results of the Gallai-Sylvester type, which extend (and also slightly improve) similar applications obtained for the case of circles in [1, 7, 8].

6.1 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 *kth level* of $\mathcal{A}(C)$, for a fixed parameter k, is the (closure of the) locus of all points on the curves of C, whose level is exactly k. The k-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 the k-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 [11, 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 subsegments, into a collection of $O(m \log m)$ extendible pseudo-segments.

Thus, the bounds on $\chi(n)$ lead to the following result (where, in part (b), 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 $\kappa_s(n)$).

Theorem 6.1 *(a) Let C be a set of n pseudo-parabolas or n x-monotone pseudo-circles. Then the maximum complexity of a level in $\mathcal{A}(C)$ is $O(n^{26/15}\log^{2/3} n)$.*

(b) If, in addition, C admits a 3-parameter algebraic representation that satisfies (AP1)–(AP3) for the case of pseudo-parabolas, or (AC1)–(AC3) for the case of pseudo-circles, then the maximum complexity of a single level is $O(n^{5/3}\kappa_s(n))$, where s is a constant that depends on the algebraic representation of the curves in C; $s = 2$ for circles and vertical parabolas.

(c) If all curves in C are pairwise intersecting, then the bound improves to $O(n^{14/9}\log^{2/3} n)$ (with no further assumption on these curves).

Remark. Recently, Chan [11] has studied the complexity of levels in arrangements of graphs of polynomials of constant maximum degree $s \geq 3$. His bound relies on cutting the given graphs into subarcs that constitute a collection of pseudo-segments, which is achieved by repeated differentiation of the given polynomials, eventually reducing to the problem of cutting an arrangement of pseudo-parabolas (actually, of pseudo-parabolic arcs) into pseudo-segments. In the earlier conference version of his paper, the bound on the number of the desired cuts was obtained by applying the Tamaki-Tokuyama result as a "black box." In the new version Chan uses a more sophisticated variant of the Tamaki-Tokuyama technique, which leads to improved bounds on the number of cuts. It is not clear whether our new bounds can be used to further improve his new bounds.

The above theorem implies the following result in the area of kinetic geometry, which improves upon an earlier bound given in [27]. This problem was one of the motivations for the initial study of Tamaki and Tokuyama [27].

Corollary 6.2 *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.*

6.2 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. Two closely related and widely studied problems concern two kinds of interaction between C and P: (i) Assuming that the points of P lie on curves of C, 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 $p \in c$. (ii) Assuming that no point of P lies on any curve of 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 is the number of edges of $\mathcal{A}(C)$ on its boundary. The results in [1, 8] imply the following bounds.

Lemma 6.3 *Let C be a set of n curves in the plane, and let P be a set of m points in the plane. Then*

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

Hence, Theorems 3.4, 4.3, 5.2, and 5.4 imply the following.

Theorem 6.4 *(a) Let C be a set of n pairwise-intersecting pseudo-circles, and P a set of m points in the plane. Then*

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

(b) Let C be a set of n pseudo-parabolas or n x-monotone pseudo-circles, and P a set of m points in the plane. Then

$$I(C,P) = O(m^{2/3}n^{2/3} + m + n^{8/5}), \qquad K(C,P) = O(m^{2/3}n^{2/3} + n^{8/5}\log^2 n).$$

We note that these bounds are worst-case tight when the first term dominates the last term, which is the case when m is larger than n or $n\log^3 n$ in part (a), and larger than $n^{7/5}$ or $n^{7/5}\log^3 n$ in part (b).

Similarly, if C is a set of n pseudo-parabolas or n x-monotone pseudo-circles that are not pairwise intersecting but admit a 3-parameter algebraic representation with corresponding parameter s, as above, then we can obtain the following bounds by plugging Theorems 3.4 and 4.3 into Lemma 6.3.

$$I(C,P) = O(m^{2/3}n^{2/3} + m + n^{3/2}\kappa_s(n)), \qquad K(C,P) = O(m^{2/3}n^{2/3} + n^{3/2}\kappa_s(n)). \quad (7)$$

As above, these bounds are worst-case tight when m is sufficiently large (larger than roughly $n^{5/4}$) [1, 8]. We can improve these bounds for smaller values of m, by exploiting properties (AP1) or (AC1) of the definition of 3-parameter algebraic representation, following the approaches in [1, 8]. We describe the argument for the case of incidences and briefly discuss how to handle the case of marked faces.

We map the pseudo-circles $\gamma \in C$ to points γ^* in $\mathbb{R}^3$, and the points in P to surfaces σ_p in $\mathbb{R}^3$, so that incidences between points and curves correspond to incidences between the dual surfaces and points, and so that one halfspace bounded by the surface σ_p corresponds to pseudo-circles that contain the point p in their interior. Let P^* be the resulting set of surfaces in $\mathbb{R}^3$, and let C^* be the resulting set of points in $\mathbb{R}^3$.

We fix a parameter $r > 1$. Roughly speaking, as in [1, 8], we wish to compute a $(1/r)$-cutting of P^*. However, since we are dealing with an arrangement of surfaces instead an arrangement of planes, a $(1/r)$-cutting for P^* is not a cell complex and the incidence structure between C^* and P^* is more involved. Consequently we rely on a random-sampling argument similar to the one in [13].

Sampling lemma. For a subset $R \subseteq P^*$, we define a partition $\Xi = \Xi(R)$ of $\mathbb{R}^3$ into relatively open and simply connected 0-, 1-, 2-, and 3-dimensional cells, which is very similar to the vertical decomposition of $\mathcal{A}(R)$ [13, 15]. Specifically, we add all vertices and edges of $\mathcal{A}(R)$ into Ξ. For each (open) 2-face f of $\mathcal{A}(R)$, we compute the vertical decomposition f^* of f, as described in [13], and add the relatively open edges and pseudo-trapezoids to Ξ. (The newly created vertices, which lie on the edges of f, are not added to Ξ.) Finally, for each (open) 3-face ϕ of $\mathcal{A}(R)$, we compute its vertical decomposition as described in [13], and we add the vertical edges, 2-faces, and 3-dimensional pseudo-prisms to Ξ; none of these cells lie in any surface of R. Let $\Xi_A \subseteq \Xi$ be the set of vertices and edges of $\mathcal{A}(R)$, which were added to Ξ, let $\Xi_E \subseteq \Xi$ be the set of 1-dimensional cells that lie in exactly one surface of R, and let $\Xi_\Box \subseteq \Xi$ be the set of vertical edges that were added to Ξ in the last step. For each cell $\Delta \in \Xi$, let $C_\Delta = \{c \in C \mid c^* \in \Delta\}$, $P_\Delta = \{p \in P \mid p^* \in P^*_\Delta\}$, where P^*_Δ is the conflict list of Δ (with respect to P^*), and $\tilde{P}_\Delta = \{p \in P \mid \Delta \subseteq p^*\}$. Set $n_\Delta = |C_\Delta|$, $m_\Delta = |P_\Delta|$, and $\tilde{m}_\Delta = |\tilde{P}_\Delta|$. The result in [15] implies that $|\Xi| = O(r^3\beta_q(r))$, where $\beta_q(r)$ is the function defined in Section 5.1.

Lemma 6.5 *For a given parameter $r > 1$, there exists a set $R \subseteq P^*$ of $O(r)$ surfaces with the following properties:*

(i) $\displaystyle\sum_{\Delta\in\Xi} n_\Delta^{2/3} = n$ *and* $m_\Delta \le \dfrac{m}{r}\log r$*, for any* $\Delta \in \Xi$.

(ii) $\sum_{\Delta\in\Xi_A} \tilde{m}_\Delta = O(mr^2)$.

(iii) $\tilde{m}_\Delta \le \dfrac{m}{r}\log r$*, for any* $\Delta \in \Xi_E \cup \Xi_\Box$.

Proof: We choose a random subset $R \subseteq P^*$ of size cr, for a sufficiently large constant parameter c, where each subset is chosen with equal probability. Since Ξ is a partition of $\mathbb{R}^3$, $\sum_\Delta n_\Delta = n$. By the theory of ε-nets, an appropriate choice of c guarantees that, with high probability, $m_\Delta \le (m/r)\log r$, for any $\Delta \in \Xi$ [19]. This proves part (i). As for (ii), observe that if $p \in \tilde{P}_\Delta$, for a vertex or edge Δ in $\mathcal{A}(R)$, then Δ is also a vertex or an edge, respectively, in the arrangement of the intersection curves $\{p^* \cap r^* \mid r^* \in R\}$. Since this arrangement has $O(r^2)$ vertices and edges, the bound in part (ii) follows. A vertical edge $\Delta \in \Xi_\Box$ does not lie in any surface of R, therefore by the theory of ε-nets and with an appropriate choice of c, $\tilde{m}_\Delta \le (m/r)\log r$ with high probability, for all such Δ's. Similarly, one can argue that $\tilde{m}_\Delta \le (m/r)\log r$ for each cell $\Delta \in \Xi_E$, as such a cell lies in exactly one surface of R. See [13, 19] for details. This completes the proof of the lemma. □

Bounding incidences. Let R be a subset of P^* satisfying the conditions of Lemma 6.5. We compute Ξ as defined above. Then

$$I(C, P) = \sum_{\Delta\in\Xi} I(C_\Delta, P_\Delta) + I(C_\Delta, \tilde{P}_\Delta).$$

Since each point in $\tilde{P}_\Delta$ lies on every curve in C_Δ and two curves in C intersect in at most two points, $\tilde{m}_\Delta > 2$ implies that $n_\Delta \le 1$. Hence,

$$I(C_\Delta, \tilde{P}_\Delta) = O(n_\Delta + \tilde{m}_\Delta),$$

Note that $\sum_\Delta n_\Delta = n$, $\tilde{m}_\Delta = 0$ for any 3-dimensional cell $\Delta \in \Xi$, and $\tilde{m}_\Delta \le 1$ for any 2-dimensional cell $\Delta \in \Xi$ because, by conditions (AC1) and (AP1), two surfaces intersect along a 1-dimensional curve. Hence,

$$\sum_{\Delta \in \Xi} I(C_\Delta, \tilde{P}_\Delta) = O(n + mr^2 \beta_q(r) \log r).$$

In order to bound $\sum_\Delta I(C_\Delta, P_\Delta)$, we refine the cells of Ξ as follows. If $n_\Delta > n/(r^3\beta_q(r))$ for a cell $\Delta \in \Xi$, we split it further so that each new cell contains at most $n/(r^3\beta_q(r))$ points. The number of refined cells in the resulting partition Ξ' is still $O(r^3\beta_q(r))$. Therefore, using the bound (7) for $I(C_\Delta, P_\Delta)$, we obtain

$$\begin{aligned}
\sum_{\Delta \in \Xi'} I(C_\Delta, P_\Delta) &= \sum_{\Delta \in \Xi'} O(m_\Delta^{2/3} n_\Delta^{2/3} + m_\Delta + n_\Delta^{3/2} \kappa_s(n_\Delta)) \\
&= O(r^3\beta(r)) \left(\left(\frac{m \log r}{r} \right)^{2/3} \left(\frac{n}{r^3 \beta(r)} \right)^{2/3} + \frac{m \log r}{r} + \left(\frac{n}{r^3\beta(r)} \right)^{3/2} \kappa_s \left(\frac{n}{r^3} \right) \right) \\
&= O(m^{2/3} n^{2/3} r^{1/3} \beta^{1/3}(r) \log^{2/3} r + mr^2 \beta(r) \log r + (n/r)^{3/2} \kappa_s(n/r^3)).
\end{aligned}$$

Hence,

$$I(C,P) = O(m^{2/3} n^{2/3} r^{1/3} \beta^{1/3}(r) \log^{2/3} r + mr^2\beta(r) \log r + (n/r)^{3/2} \kappa_s(n/r^3) + n).$$

We choose $r = \lceil n^{5/11}/m^{4/11} \rceil$, which is in the range $1 \le r \le m$ when $n^{1/3} \le m \le n^{5/4}$. If $m > n^{5/4}$ we take $r = 1$, and if $m < n^{1/3}$ we take $r = m$. It follows easily, as in [8], that

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

where s is a constant depending on the representation of C.

Bounding the complexity of marked faces. We use the approach in [1] to prove an improved bound on the complexity of marked faces. There is one significant difference in the proof for this case compared with the case of incidences. Here we need a *hierarchical cutting*[7] of $\mathcal{A}(R)$. The best known algorithm for computing such a hierarchical $(1/r)$-cutting returns a cutting of size $O(r^{3+\varepsilon})$, for any $\varepsilon > 0$. Plugging this weaker bound on the size of hierarchical cuttings in the analysis of [1], the bound on the marked faces increases by a factor $O(m^\varepsilon)$. We refer the reader to the papers just cited for further details, and omit the description of the modifications of the analyses given there that need to be performed.

Putting everything together, we obtain the following results on the number of incidences and the complexity of marked faces.

Theorem 6.6 *Let C be a set of n pseudo-parabolas or n x-monotone pseudo-circles that admit a 3-parameter algebraic representation, and let P be a set of m points in the plane.*

(i) $I(C,P) = O(m^{2/3} n^{2/3} + m^{6/11} n^{9/11} \kappa_s(m^3/n) + m + n)$, where s is a constant depending on the representation, and

[7] For a set Γ of surfaces, a $(1/r)$-cutting Ξ of Γ is called *hierarchical* if there exist a constant r_0 and a sequence of cuttings $\Xi_0, \Xi_1, \ldots, \Xi_u = \Xi$, for $u = \lceil \log_{r_0} r \rceil$, where Ξ_i is a $(1/r_0^i)$-cutting of Γ and each cell of Ξ_i lies inside a cell of Ξ_{i-1}.

(ii) $K(C,P) = O(m^{2/3}n^{2/3} + m^{6/11+\varepsilon}n^{9/11} + n\log n)$*, for any* $\varepsilon > 0$.

If the pseudo-parabolas or pseudo-circles in C *are also pairwise intersecting, then (we do not need to require that the pseudo-circles be* x*-monotone in this case)*

(iii) $I(C,P) = O(m^{2/3}n^{2/3} + m^{1/2}n^{5/6}\beta(n/m) + m + n)$*, and*

(iv) $K(C,P) = O(m^{2/3}n^{2/3} + m^{1/2+\varepsilon}n^{5/6}\log^{1/2} n + n\log n)$*, for any* $\varepsilon > 0$.

For the cases of circles and of vertical parabolas, the relevant surfaces are (or can be transformed into) planes, so there is no extra $\beta(r)$ factor, and efficient hierarchical cuttings can be constructed (for the analysis of many faces). Hence, the analysis in [1, 8] yields the following improved bounds. (The bound in Theorem 6.7(ii) has actually been proven in [1] for the case of circles; we state it here for the sake of completeness.)

Theorem 6.7 *Let* C *be a set of* n *circles or* n *vertical parabolas and* P *a set of* m *points in the plane. Then*

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

(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*

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

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

Remark. Using a standard sampling technique, such as the one used in [1, 8, 11], we can also obtain versions of these bounds that are sensitive to the number of intersecting pairs of the given curves (for parts (i) and (ii) of both theorems).

6.3 Distinct distances under arbitrary norms

An interesting application of Theorem 6.6(i) is the following result.

Theorem 6.8 *Let* Q *be a compact strictly 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.) 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 that the points in P determine. There are n^2 incidences between these curves and the points of P. Using Theorem 6.6(i), the bound follows easily (here too the constant in the exponent of the expression for $\kappa_s(n)$ is changed). □

Remarks. (1) The proof technique is identical to an older proof for distinct distances under the Euclidean metric, given in [13, Section 5.4]. Meanwhile, the bound for the Euclidean case has been substantially improved (see [28] for the current "record"), but, as far as we know, the problem has not been considered at all for more general metrics.

(2) Theorem 6.8 is false if Q is not strictly convex. For example, let Q be the unit ball of the L_1-norm, and let P be the set of vertices of the $\sqrt{n} \times \sqrt{n}$ integer lattice. There are only $2\sqrt{n}$ distinct L_1-distances among the points of P.

6.4 A generalized Gallai-Sylvester theorem

Similar to Theorem 4.1 in [7], the following theorem is a consequence of Theorem 2.13.

Theorem 6.9 *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 of C.*

7 Conclusion and Open Problems

In this paper we obtained a variety of results involving lenses in arrangements of pseudo-circles, with numerous applications to incidences, levels, and complexity of many faces in arrangements of circles, vertical parabolas, homothetic copies of a fixed convex curve, pairwise intersecting pseudo-circles, and arbitrary pseudo-parabolas and x-monotone pseudo-circles. We also obtained a Gallai-Sylvester result for arrangements of pairwise-intersecting pseudo-circles, and a new lower bound on the number of distinct distances in the plane under fairly arbitrary norms. The main tool that facilitated the derivation of all these results is the somewhat surprising property that the tangency graph in a family of pairwise intersecting pseudo-parabolas is planar (Theorem 2.4).

The paper leaves many 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 2.14, 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 or more general classes of pseudo-circles. There is a gap between the lower bound $\Omega(n^{4/3})$, which follows from the construction of $\Omega(n^{4/3})$ incidences between n points and n lines, and which can be realized by circles, and the upper bound of $O(n^{3/2}\kappa(n))$, given in Theorem 5.2 and Corollary 5.5. Even improving the upper bound to $O(n^{3/2})$, for the case of circles, seems a challenging open problem. A related and harder problem is to obtain an improved bound for the number of pairwise nonoverlapping lenses (and for the cutting number) 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. (At least for pseudo-parabolas, the weaker bound of $O(n^{8/5})$ holds in general.) It would be interesting and instructive to find a purely topological way of tackling the general problem involving pseudo-circles. 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 [27] (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? Assuming x-monotonicity, can the bound $O(n^{8/5})$ in Theorem 4.1 be further improved?

References

[1] P. K. Agarwal, B. Aronov and M. Sharir, On the complexity of many faces in arrangements of pseudo-segments and of circles, in *Discrete and Computational Geometry: The Goodman-Pollack Festschrift* (B. Aronov, S. Basu, J. Pach, and M. Sharir, eds.), Springer Verlag, Berlin, to appear.

[2] P. K. Agarwal, A. Efrat and M. Sharir, Vertical decomposition of shallow levels in 3-dimensional arrangements and its applications, *SIAM J. Comput.* 29 (2000), 912–953.

[3] P. K. Agarwal and J. Matoušek, On range searching with semialgebraic sets, *Discrete Comput. Geom.* 11 (1994), 393–418.

[4] 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.

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

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

[7] 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.

[8] B. Aronov and M. Sharir, Cutting circles into pseudo-segments and improved bounds for incidences, *Discrete Comput. Geom.* 29 (2003), in press.

[9] S. Basu, R. Pollack and M.-F. Roy, *Algorithms in Real Algebraic Geometry*, in preparation.

[10] J. Bochnak, M. Coste and M.-F. Roy, *Real Algebraic Geometry*, Springer Verlag, Berlin-Heidelberg, 1998.

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

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

[15] B. Chazelle, H. Edelsbrunner, L. J. Guibas, and M. Sharir, A singly-exponential stratification scheme for real semi-algebraic varieties and its applications, *Proc. 16th Internat. Colloq. Automata Lang. Program.*, *Lecture Notes Comput. Sci.*, Vol. 372, Springer-Verlag, 1989, pp. 179–192.

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

[17] J. E. Goodman and R. Pollack, Allowable sequences and order types in discrete and computational geometry, in: *New Trends in Discrete and Computational Geometry* (J. Pach, ed.), Springer-Verlag, 1993, pp. 103–134.

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

[19] D. Haussler and E. Welzl, Epsilon-nets and simplex range queries, *Discrete Comput. Geom.*, 2 (1987), 127–151.

[20] 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.

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

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

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

[25] R. Pinchasi and R. Radoičić, On the number of edges in geometric graphs with no self-intersecting cycle of length 4, manuscript, 2002.

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

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

[28] G. Tardos, On distinct sums and distinct distances, *Adv. Math.*, to appear.

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

Appendix A: Analysis of the Case of Circles

In this appendix, we show how to refine the upper bound on $\chi(C)$, in the case of circles, so that the associated constant s' is 4, and thus $s = 2$ and $q = 4$.

Lemma A.1 *Let c_1 and c_2 be two circles in the plane, with $c_1^* = (a_1, b_1, r_1)$ and $c_2^* = (a_2, b_2, r_2)$ and $r_1 \geq r_2$. The upper arcs c_1^+ and c_2^+ intersect at two points if and only if the following condition holds (see Figure 21 (i)):*

(UU) $b_2 \geq b_1$, λ_{c_2} and ρ_{c_2} lie inside c_1, and c_1 intersects c_2.

Proof: If c_1^+ and c_2^+ intersect at two points u, v then both centers lie below the line ℓ passing through u and v. Moreover, the portion of the smaller disk (the disk bounding the smaller circle) below ℓ is contained in the corresponding portion of the bigger disk, and the center of the smaller disk is closer to ℓ. This is easily seen to imply (UU). Conversely, if (UU) holds then both intersection points lie on c_2^+ or both lie on c_2^- (because the endpoints of both arcs lie inside c_1). Translate c_2 vertically downward until its center has the same y-coordinate as that of c_1. In this position λ_{c_2} and ρ_{c_2} continue to lie inside c_1, and the two circles must be disjoint (any intersection point on c_2^- must have a matching symmetric point on c_2^+, which would produce at least 4 intersection points). This is easily seen to imply that the original c_2^- is also disjoint from c_1, so the two intersection points must lie on c_2^+, and, since $b_2 \geq b_1$, they must also lie on c_1^+. $\square$

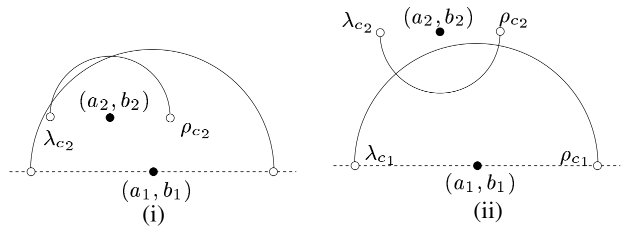

Figure 21. (i) Illustration of condition (UU). (ii) Illustration of condition (UL).

Lemma A.2 *Let c_1 and c_2 be two circles in the plane, with $c_1^* = (a_1, b_1, r_1)$ and $c_2^* = (a_2, b_2, r_2)$. The arcs c_1^+ and c_2^- intersect at two points if and only if the following condition holds (see Figure 21 (ii)):*

(UL) $b_2 \geq b_1$, λ_{c_2} and ρ_{c_2} lie outside c_1, λ_{c_1} and ρ_{c_1} lie outside c_2, and c_1 intersects c_2.

Proof: Suppose that c_1^+ and c_2^- intersect at two points u, v. Then the portion of c_1^+ between u and v lies inside c_2, and the portion of c_2^- between u and v lies inside c_1. This is easily seen to imply that each of the x-extreme points λ_{c_1}, ρ_{c_1}, λ_{c_2} and ρ_{c_2} lies outside the other circle. Moreover, the center of c_1 (resp., c_2) lies below (resp., above) the line passing through u and v, implying that $b_2 \geq b_1$. Hence (UL) holds. Conversely, if (UL) holds then both intersection points must lie on the same arc (upper or lower) of c_1, and on the same arc (upper or lower) of c_2. However, in view of Lemma A.1, it cannot be the case that both arcs are upper or that both arcs are lower. Hence one arc is upper and one is lower, and the condition $b_2 \geq b_1$ is easily seen to imply that the upper arc is of c_1 and the lower arc is of c_2. $\square$

Fix a circle $c : (x-a)^2 + (y-b)^2 = r^2$. Then by Lemma A.1, the locus $K_{c^+}^+$ of circles whose upper arc intersects c^+ at two points is given by $K^+_{c^+} = K^+_{c^+,>} \cup K^+_{c^+,<}$, where[8]

$$\begin{aligned}
K^+_{c^+,>} &= \{(\xi,\eta,\zeta) \mid (\zeta \le r) \wedge (\eta \ge b) \wedge [(\xi \pm \zeta - a)^2 + (\eta - b)^2 \le r^2] \wedge \\
&\qquad\qquad\qquad\qquad\qquad\qquad [(\xi - a)^2 + (\eta - b)^2 \ge (r-\zeta)^2]\} \\
K^+_{c^+,<} &= \{(\xi,\eta,\zeta) \mid (\zeta \ge r) \wedge (\eta \le b) \wedge [(\xi - a \pm r)^2 + (\eta - b)^2 \le \zeta^2] \wedge \\
&\qquad\qquad\qquad\qquad\qquad\qquad [(\xi - a)^2 + (\eta - b)^2 \ge (r-\zeta)^2]\}.
\end{aligned}$$

This implies that $K^+_{c^+}$ is a semialgebraic set of constant description complexity. Symmetrically, it follows that $K^-_{c^+}$ and the corresponding regions for c^- are also semialgebraic sets of constant description complexity. We thus conclude that $\chi(C) = O(n^{3/2}\kappa_s(n))$ for some integer s. However, the surfaces bounding these regions are quadrics, so their intersection curves are in general of degree four, and a naive bound on the number of intersection points between the $\xi\eta$-projections of a pair of such curves is $s_1 \le 4^2 = 16$, yielding $s = 8$. For mostly aesthetic reasons, we set out to improve this bound to $\chi(C) = O(n^{3/2}\kappa(n))$, where $\kappa(n) = \kappa_2(n)$.

Let $\psi(A,B)$ denote the minimum number of cuts needed to eliminate all bichromatic upper-upper lenses in $A \cup B$ (lenses formed by the upper arcs of one circle in A and one in B). Put $\psi(A) = \psi(A,A)$. For $k = 0,1,2$, set $\psi^{(k)}(u,v) = \max\ \psi(A,B)$, where the maximum is taken over all pairs of families of circles A and B of sizes at most u and v, respectively, so that

- for $k = 0$, no constraint is imposed on A and B;
- for $k = 1$, we require that the radius of each circle in A be greater than or equal to the radius of each circle in B; and
- for $k = 2$, we require the same condition on the radii as for $k = 1$, and also that the y-coordinate of the center of each circle in A be smaller than or equal to the y-coordinate of the center of each circle in B.

We set $\psi^{(k)}(m) = \psi^{(k)}(m,m)$, and our task is to bound $\psi^{(0)}(n)$.

Sort the circles in C in increasing order of their radii, and let C_1, C_2 be the subsets of the circles with the $n/2$ smallest and $n/2$ largest radii, respectively. We clearly have

$$\psi(C) \le \psi(C_1) + \psi(C_2) + \psi(C_2, C_1),$$

from which we deduce the recurrence

$$\psi^{(0)}(n) \le 2\psi^{(0)}\left(\frac{n}{2}\right) + \psi^{(1)}\left(\frac{n}{2},\frac{n}{2}\right). \tag{8}$$

Next we estimate $\psi^{(1)}$. Let A and B be two sets of m and n circles, respectively, so that the radius of each circle in A is greater than or equal to the radius of every circle in B. Sort the circles in $C = A \cup B$ in increasing order of the y-coordinate of their centers, and split C into two subsets C^-, C^+, consisting respectively of the circles with the $(m+n)/2$ lowest and the $(m+n)/2$ highest

[8] The condition for the intersection of two circles is that the distance between their centers be larger than the difference between the radii and smaller than their sum. In what follows, we only use the first inequality, because the second is implied by the additional condition that one circle contains points of the other in its interior. This simplification does not hold, though, when we consider intersections between lower and upper arcs.

y-coordinates. Put $A^- = A \cap C^-$, $A^+ = A \cap C^+$, $B^- = B \cap C^-$, and $B^+ = B \cap C^+$. We clearly have

$$\psi(A,B) \le \psi(A^-,B^-) + \psi(A^+,B^+) + \psi(A^-,B^+);$$

the fourth term, $\psi(A^+,B^-)$, is 0, because all pairs of circles in $A^+ \times B^-$ violate condition (UU). Put $k = |A^-|$, $\ell = |B^+|$. Hence, we obtain the recurrence

$$\psi^{(1)}(m,n) \le \max_{\substack{k,\ell \le \frac{m+n}{2} \\ k-l=\frac{m-n}{2}}} \left\{ \psi^{(1)}\left(k, \frac{m+n}{2} - k\right) + \psi^{(1)}\left(\frac{m+n}{2} - \ell, \ell\right) + \psi^{(2)}(k,\ell) \right\}, \quad (9)$$

where the conditions on k and ℓ follow from the construction.

We next bound $\psi^{(2)}$, where a more complex recurrence is needed. Let A and B be two sets of m and n circles, respectively, so that for any $(c_1,c_2) \in A \times B$, with $c_1 = (a_1,b_1,r_1)$ and $c_2 = (a_2,b_2,r_2)$, the following condition holds:

(C0) $r_1 \ge r_2$ and $b_2 \ge b_1$.

If the upper arc of a circle $c_1 = (a_1,b_1,r_1) \in A$ intersects the upper arc of $c_2 = (a_2,b_2,r_2) \in B$ at two points, then by Lemma A.1, the following two conditions also hold:

(C1) $\lambda_{c_2} = (a_2 - r_2, b_2)$ and $\rho_{c_2} = (a_2 + r_2, b_2)$ lie inside c_1;

(C2) c_1 and c_2 intersect.

Fix a circle $c = (a,b,r)$ in A. The locus $K_1(c)$ of all circles $(\xi,\eta,\zeta) \in B$ that satisfy (C1) with c is the region

$$\{(\xi,\eta,\zeta) \mid (\xi-\zeta-a)^2 + (\eta-b)^2 \le r^2 \text{ and } (\xi+\zeta-a)^2 + (\eta-b)^2 \le r^2\},$$

which is bounded by the pair of surfaces

$$\pi_1(c) : (\xi-\zeta)^2 + \eta^2 - 2a(\xi-\zeta) - 2b\eta + a^2 + b^2 - r^2 = 0, \quad (10)$$

$$\pi_2(c) : (\xi+\zeta)^2 + \eta^2 - 2a(\xi+\zeta) - 2b\eta + a^2 + b^2 - r^2 = 0. \quad (11)$$

On the other hand, if we fix a circle $c' = (a,b,r)$ in B, then the locus $\bar{K}_1(c')$ of all circles $(\xi,\eta,\zeta) \in A$ that satisfy (C1) with c' is the region

$$\{(\xi,\eta,\zeta) \mid (\xi-(a-r))^2 + (\eta-b)^2 \le \zeta^2 \text{ and } (\xi-(a+r))^2 + (\eta-b)^2 \le \zeta^2\},$$

which is bounded by the pair of surfaces

$$\bar{\pi}_1(c') : \xi^2 + \eta^2 - \zeta^2 - 2(a-r)\xi - 2b\eta + (a-r)^2 + b^2 = 0, \quad (12)$$

$$\bar{\pi}_2(c') : \xi^2 + \eta^2 - \zeta^2 - 2(a+r)\xi - 2b\eta + (a+r)^2 + b^2 = 0. \quad (13)$$

Finally, for a fixed circle $c = (a,b,r)$ in A or B, the locus $K_2(c)$ of all circles (ξ,η,ζ) that satisfy (C2) with c, given that they already satisfy (C1), is bounded by the surface (as already remarked, only one of the two inequalities that represent intersection between circles need to be considered)

$$(\xi - a)^2 + (\eta - b)^2 = (\zeta - r)^2, \quad \text{or}$$

$$\pi_3(c) : \xi^2 + \eta^2 - \zeta^2 - 2a\xi - 2b\eta + 2r\zeta + a^2 + b^2 - r^2 = 0. \tag{14}$$

An important observation is that the bound on the parameter s is large because we consider intersection curves of "mixed" pairs of surfaces from among the possible types (10)–(14). However, if we only consider pairs of surfaces of the same type, say of type (14), the corresponding intersection curves are *plane quadrics*, so the number of intersection points between the projections of two such curves is at most 4, as in the case of vertical parabolas (Section 5.2). Our approach is thus to enforce the conditions (C1)–(C2) in two stages, where the first stage enforces (C1) and the second enforces (C2). This will suffice to reduce s to 2.

In more detail, we proceed as follows. For $k = 3, 4$, set $\psi^{(k)}(u, v) = \max\ \psi(A, B)$, where the maximum is taken over all pairs of families of circles A and B of sizes at most u and v, respectively, that satisfy (C0)–(C($k-2$)). We set $\psi^{(k)}(m) = \psi^{(k)}(m, m)$. Recall that our task is to bound $\psi^{(2)}(m)$.

Bounding $\psi^{(4)}(m)$. We first observe that $\psi^{(4)}(m) = O(m^{4/3})$. Indeed, if every pair of circles in $A \times B$ satisfy (C0)–(C2), i.e., the upper arcs of every pair intersect at two points, then the bound follows by considering the collection of extended upper arcs of the circles in $A \cup B$, and applying Lemma 4.2 and Theorem 3.4, as argued in Section 5.3.

Bounding $\psi^{(3)}(m)$. Next, we apply the analysis in the proof of Lemma 5.1 to the arrangement of the surfaces $\pi_3(c)$, for $c \in A$ or $c \in B$. Choosing a parameter $1 \le r \le m^{1/4}$, we obtain the recurrence

$$\psi^{(3)}(m) \le cr^6\beta_q^2(r)\left[\psi^{(3)}\left(\frac{m}{r^4}\right) + \psi^{(4)}(m)\right] \le cr^6\beta_q^2(r)\left[\psi^{(3)}\left(\frac{m}{r^4}\right) + O(m^{4/3})\right],$$

with $q = 4$. Indeed, the overhead term bounds the minimum number of cuts needed to eliminate all bichromatic upper-upper lenses between pairs of subfamilies of circles that satisfy (C2) (where one subfamily corresponds to all circles in, say, A, whose representing points lie in some cell Δ of the relevant cutting, and the other subfamily corresponds to all circles $c \in B$ whose associated surface $\pi_3(c)$ fully encloses Δ), in addition to (C0)–(C1) which are satisfied, by assumption, by all pairs of circles in $A \times B$. Here $q = 4$, because we are dealing here only with surfaces of the form $\pi_3(c)$, and, as already remarked, the intersection curve of two such surfaces is a plane quadric, so, as argued in Section 5.2, the projections of two such intersection curves on the $\xi\eta$-plane intersect in at most four points, thereby implying that $q = 4$ and $\beta_q(r) = 2^{O(\alpha^2(r))}$. The same analysis as in Section 5.1 now shows that

$$\psi^{(3)}(m) = O(m^{3/2}\kappa(m)). \tag{15}$$

Bounding $\psi^{(2)}(m)$. This is achieved by a similar process of interleaved recursion, in which we keep flipping the roles of A and B. However, this can be done so that one of the two recursive steps is performed in the plane (and only one in three dimensions). Specifically, we have:

Lemma A.3 *For any m, n and for any parameter $1 \le r_1 \le \min\{m, n^{1/2}\}$,*

$$\psi^{(2)}(m,n) \le c_2 r_1^2 \psi^{(2)}\left(\frac{m}{r_1}, \frac{n}{r_1^2}\right) + c_2 r_1^4 \psi^{(3)}\left(m, \frac{n}{r_1^2}\right), \tag{16}$$

for some positive constant c_2.

Proof: Let A and B be two families of circles of size m and n, respectively, so that every pair in $A \times B$ satisfy condition (C0). We need to "enforce" condition (C1), namely, that the leftmost and rightmost points of a circle in B lie inside a circle in A. This can be done via the following cutting-based partitioning in the plane, where each circle $g = (\xi, \eta, \zeta) \in B$ is mapped to the two respective points $\lambda_g = (\xi - \zeta, \eta)$, $\rho_g = (\xi + \zeta, \eta)$, and the circles of A remain as they are.

We compute a $(1/r_1)$-cutting Ξ of A of size $O(r_1^2)$. For each $\Delta \in \Xi$, let $B_\Delta = \{g \in B \mid \lambda_g \in \Delta \text{ or } \rho_g \in \Delta\}$. If $|B_\Delta| > n/r_1^2$, we partition Δ into subcells, each of which contains at most n/r_1^2 points. The number of new cells remains $O(r_1^2)$. For each new cell Δ, let $A_\Delta = \{c \in A \mid c \cap \Delta \ne \emptyset\}$ and $\tilde{A}_\Delta = \{c \in A \mid \Delta \subseteq \text{int}(c)\}$. Since Ξ is a cutting, we have $|A_\Delta| \le m/r_1$ for each Δ.

To bound $\psi(A, B)$, we first sum up the recursive terms $\sum_\Delta \psi(A_\Delta, B_\Delta)$. Let (c, g) be a pair that needs to be counted in $\psi(A, B)$ but has not been counted in this recursive manner. Let Δ, Δ' be the cells of the cutting that contain λ_g, ρ_g, respectively. Then both cells Δ, Δ' are fully contained in the interior of c. This suggests the following approach to completing the count: Take each pair (Δ, Δ') of cells of the cutting, and put $B_{(\Delta,\Delta')} = \{g \in B \mid \lambda_g \in \Delta \text{ and } \rho_g \in \Delta'\}$, $A_{(\Delta,\Delta')} = \{c \in A \mid \Delta, \Delta' \subseteq \text{int}(c)\}$. The number of remaining pairs that need to be counted is thus bounded by

$$\sum_{(\Delta,\Delta')} \psi\left(A_{(\Delta,\Delta')}, B_{(\Delta,\Delta')}\right).$$

However, every pair of sets in this sum also satisfy (C1), so the sum is at most $O(r^4 \psi^{(3)}(m, n/r_1^2))$. This completes the proof of the lemma. $\square$

We also need a dual partitioning scheme for the "flipped" version of the recursion, in which the circles of A are mapped into points and those of B into surfaces. Here, unlike the preceding partition, we need to use the 3-dimensional representation of the circles:

Lemma A.4 *For any m, n and for any parameter $1 \le r_2 \le \min\{m^{1/3}, n\}$,*

$$\psi^{(2)}(m,n) \le c_3 r_2^3 \beta_q(r_2) \left[\psi^{(2)}\left(\frac{m}{r_2^3}, \frac{2n}{r_2}\right) + \psi^{(3)}\left(\frac{m}{r_2^3}, n\right)\right], \tag{17}$$

for some integer constant q and some positive constant c_3.

Proof: Let A and B be two families of circles of size m and n, respectively, which satisfy condition (C0). We now map each circle $g \in A$ to the point $g^* = (\xi, \eta, \zeta) \in \mathbb{R}^3$, using the 3-parameter representation of C. Let $\Sigma = \{\bar{\pi}_1(c), \bar{\pi}_2(c) \mid c \in B\}$. We compute a $(1/r_2)$-cutting Ξ of Σ of size $O(r_2^3 \beta_q(r_2))$, for some appropriate constant q.[9] For each cell $\tau \in \Xi$, set $A_\tau = \{c \in A \mid c^* \in \tau\}$ and partition τ further, as needed, to ensure that, for any resulting subcell τ', $|A_{\tau'}| \le m/r_2^3$; this does not

[9]Curiously, $q = 4$ for the collection of surfaces $\bar{\pi}_1(c), \bar{\pi}_2(c)$, which follows by the same reasoning used for the surfaces $\pi_3(c)$. However, this extra property is not needed in this step of our analysis.

change the asymptotic bound on the number of cells. Set $B_\tau = \{c \in B \mid (\bar{\pi}_1(c) \cup \bar{\pi}_2(c)) \cap \tau \neq \emptyset\}$ and $\tilde{B}_\tau = \{c \in B \mid \tau \subseteq \bar{K}_1(c)\}$. Hence, we obtain the following recurrence

$$\psi(A, B) = \sum_{\tau \in \Xi} [\psi(A_\tau, B_\tau) + \psi(A_\tau, \tilde{B}_\tau)].$$

By construction, every pair $(c_1, c_2) \in A_\tau \times \tilde{B}_\tau$ satisfies (C0)–(C1), which implies that $\psi(A_\tau, \tilde{B}_\tau) \leq \psi^{(3)}(|A_\tau|, |\tilde{B}_\tau|)$. Since $|A_\tau| \leq m/r_2^3$ and $|B_\tau| \leq 2n/r_2$ for each τ, we thus obtain, summing over all cells of the cutting,

$$\psi^{(2)}(m, n) \leq c_3 r_2^3 \beta_q(r_2) \left[\psi^{(2)}\left(\frac{m}{r_2^3}, \frac{2n}{r_2}\right) + \psi^{(3)}\left(\frac{m}{r_2^3}, n\right)\right],$$

as asserted. □

Combining (16) and (17), choosing $r_2 = r$ and $r_1 = 2r^2$ for an appropriate parameter $r > 1$, and substituting the bound (15) on $\psi^{(3)}(\cdot)$, we obtain the recurrence for appropriate values of constants c, c':

$$\psi^{(2)}(m) \leq c r^7 \beta_q(r) \psi^{(2)}\left(\frac{m}{2r^5}\right) + c' r^8 m^{3/2} \kappa(m).$$

Since the overhead term in the recurrence dominates its homogeneous solution, it can be shown (by induction on m) that if we choose r to be a sufficiently large constant, then the solution to the recurrence is

$$\psi^{(2)}(m) = O(m^{3/2}\kappa(m)).$$

Bounding $\psi^{(1)}(m)$ and $\psi^{(0)}(m)$. We now return to the first two stages of divide and conquer. Substituting the bound for $\psi^{(2)}(\cdot)$ in (9), we obtain a recurrence in which each instance involving a total of $m+n$ circles is replaced by two instances, each involving a total of $(m+n)/2$ circles. This readily implies that the recurrence solves to

$$\psi^{(1)}(m) = O(m^{3/2}\kappa(m)).$$

Substituting this bound into (8), we again obtain a simple recurrence for $\psi^{(0)}(\cdot)$ which also solves to

$$\psi^{(0)}(m) = O(m^{3/2}\kappa(m)).$$

We have thus shown that the minimum number of cuts needed to eliminate all upper-upper lenses in a set of n circles is $O(n^{3/2}\kappa(n))$. A fully symmetric argument yields the same bound for the number of cuts needed to eliminate all lower-lower lenses, and it remains to bound the number of cuts needed to eliminate upper-lower lenses. For this we need to carry out a similar analysis, based on the condition (UL) in Lemma A.2. The analysis is indeed rather similar, and we do not spell it out in detail. We only comment on several technical differences that arise:

(1) At the bottommost recursive stage, we enforce the condition that a pair of circles $c = (a, b, r)$ and $c' = (\xi, \eta, \zeta)$ intersect. Here we need to enforce both inequalities, that the distance between the centers be at least the difference between the radii and at most their sum. The corresponding surfaces, with c fixed and c' varying, are

$$\begin{aligned} \pi_3(c):&\quad \xi^2 + \eta^2 - \zeta^2 - 2a\xi - 2b\eta + 2r\zeta + a^2 + b^2 - r^2 = 0 \\ \bar{\pi}_3(c):&\quad \xi^2 + \eta^2 - \zeta^2 - 2a\xi - 2b\eta - 2r\zeta + a^2 + b^2 - r^2 = 0. \end{aligned}$$

Fortunately, the intersection curve of any pair of these surfaces is still a plane quadric, and the preceding analysis can be easily adapted to keep the parameter q equal to 4 (and s to 2) in this case too.

(2) We now need only one stage of a simple divide-and-conquer, to enforce the condition $b_2 \geq b_1$, but we need two stages to enforce the conditions concerning the points λ_{c_1}, ρ_{c_1}, λ_{c_2} and ρ_{c_2}, one stage enforcing that λ_{c_1}, ρ_{c_1} lie outside c_2, and the other stage enforcing that λ_{c_2}, ρ_{c_2} lie outside c_1. Both stages are carried out exactly as above.

The modified analysis thus yields a bound of $O(n^{3/2}\kappa(n))$ for the minimum number of cuts needed to eliminate all upper-lower lenses in a set C of n circles, showing, at long last, that $\chi(C) = O(n^{3/2}\kappa(n))$.