

Arrangements of pseudocircles and circles

Ross J. Kang* Tobias Müller†

December 27, 2013

Abstract

An arrangement of pseudocircles is a finite collection of Jordan curves in the plane with the additional properties that **(i)** every two curves meet in at most two points; and **(ii)** if two curves meet in a point p , then they cross at p .

We say that two arrangements $\mathcal{C} = (c_1, \dots, c_n)$ and $\mathcal{D} = (d_1, \dots, d_n)$ are equivalent if there is a homeomorphism φ of the plane onto itself such that $\varphi[c_i] = d_i$ for all $i \in \{1, \dots, n\}$. Linhart and Ortner (2005) gave an example of an arrangement of five pseudocircles that is not equivalent to an arrangement of circles, and they conjectured that every arrangement of at most four pseudocircles is equivalent to an arrangement of circles. Here we prove their conjecture.

We consider two related recognition problems. The first is the problem of deciding, given a (combinatorial description of a) pseudocircle arrangement, whether it is equivalent to an arrangement of circles. The second is deciding whether it is equivalent to an arrangement of convex pseudocircles. We prove that both problems are NP-hard, answering questions of Bultena, Grünbaum and Ruskey (1998) and of Linhart and Ortner (2008).

We also give an example of an arrangement of convex pseudocircles with the property that its intersection graph (i.e. the graph with one vertex for each pseudocircle and an edge between two vertices if and only if the corresponding pseudocircles intersect) cannot be realised as the intersection graph of a family of circles. This disproves a folklore conjecture communicated to us by Pyatkin.

1 Introduction and statement of results

An arrangement of pseudocircles is a finite list $\mathcal{C} = (c_1, \dots, c_n)$ of Jordan curves in the plane satisfying the following two conditions:

- (i)** every two curves intersect in at most two points; and
- (ii)** if two curves meet in a point p , then they cross at p .

We speak of a *simple* arrangement of pseudocircles if in addition the following holds:

- (iii)** no three curves intersect in a point.

*Utrecht University, Utrecht, the Netherlands. E-mail: ross.kang@gmail.com. Part of the work in this paper was done while this author was at Durham University, supported by EPSRC grant EP/G066604/1. He is now supported by a VENI grant from Netherlands Organisation for Scientific Research (NWO).

†Utrecht University, Utrecht, the Netherlands. E-mail: t.muller@uu.nl. Part of the work in this paper was done while this author was supported by a VENI grant from Netherlands Organisation for Scientific Research (NWO).

Naturally, an arrangement of pseudocircles $\mathcal{C} = (c_1, \dots, c_n)$ is called an *arrangement of circles* if each c_i is a circle. Arrangements of circles and pseudocircles have been studied previously by several groups of authors including Agarwal, Aronov and Sharir [1], Alon et al. [2], Linhart and Ortner [15, 16, 17] and Linhart and Yang [18].

We will say that two arrangements $\mathcal{C} = (c_1, \dots, c_n)$ and $\mathcal{D} = (d_1, \dots, d_n)$ are *equivalent* if there exists a homeomorphism φ from the plane onto itself with the property that $\varphi[c_i] = d_i$ for all $i \in \{1, \dots, n\}$. We will say that an arrangement of pseudocircles is *circleable* if it is equivalent to an arrangement of circles. A natural question is whether every arrangement of pseudocircles is circleable. As it turns out the answer to this question is negative: Edelsbrunner and Ramos [8] gave an example of an arrangement of six pseudocircles that is not circleable. Later, Linhart and Ortner [16] showed that the arrangement of five pseudocircles shown in Figure 1 is not circleable.




Figure 1: An arrangement of five pseudocircles that is not circleable.

They also conjectured this to be the smallest non-circleable pseudocircle arrangement. Some evidence in favour of this conjecture was reported in the PhD thesis of Ortner [22]. There, a computer enumeration was implemented which verified the circleability of all simple arrangements on at most four pseudocircles satisfying the additional condition that every two pseudocircles intersect. Here we confirm the full conjecture.

Theorem 1.1 *Every arrangement of at most four pseudocircles is circleable.*

It should be mentioned that Linhart and Ortner used a more restrictive notion of arrangement of pseudocircles (corresponding to arrangements that we have called simple above) and a more general notion of equivalence than ours. So Theorem 1.1 not only proves their conjecture, but also strengthens it slightly.

Theorem 1.1 provides a natural analogue of a celebrated result of Goodman and Pollack [11], stating that every arrangement of up to eight pseudolines is equivalent to an arrangement of lines. Prior to this an example was already known of an arrangement of nine pseudolines inequivalent to any arrangement of lines.

Given that not all pseudocircle arrangements are circleable, one may wonder how easy it is to tell whether a pseudocircle arrangement is circleable. One way to frame this as a

mathematically precise question is by asking the complexity of an appropriate computational recognition problem. Let us define the *combinatorial description* of a pseudocircle arrangement as the associated labelled cell-complex — we describe this precisely in the next section. The next theorem shows that the computational problem of deciding, given a combinatorial description of a pseudocircle arrangement \mathcal{C} , whether \mathcal{C} is equivalent to a circle arrangement is NP-hard, even if we restrict the input to simple arrangements of convex pseudocircles.

Theorem 1.2 *It is NP-hard to decide, given a combinatorial description of a simple arrangement of convex pseudocircles, whether the arrangement is circleable.*

The proof of this theorem is given in Section 5.

It is also natural to consider the analogous problem for convex pseudocircle arrangements. We say that an arrangement of pseudocircles is *convexible* if it is equivalent to an arrangement of convex pseudocircles. Bultena, Grünbaum and Ruskey [4] have asked about the complexity of the computational problem of deciding, given a combinatorial description of a pseudocircle arrangement, whether the arrangement is equivalent to an arrangement of convex pseudocircles. Later Linhart and Ortner [17] asked the indeed weaker question of whether there exists a pseudocircle arrangement that is not convexible. The next result answers both questions.

Theorem 1.3 *It is NP-hard to decide, given a combinatorial description of an arrangement of pseudocircles, whether the arrangement is convexible.*

The proof of this theorem is given in Section 6.

If $\mathcal{A} = (A_1, \dots, A_n)$ is a list of sets, then the *intersection graph* of \mathcal{A} is the graph $G = (V, E)$ with vertex set $V = \{1, \dots, n\}$ and an edge $ij \in E$ if and only if $A_i \cap A_j \neq \emptyset$. A folklore conjecture that was communicated to us by Artem Pyatkin [23] states that every intersection graph of an arrangement of convex pseudocircles is also the intersection graph of a list of circles. (We do not use the word “arrangement” here because we do not necessarily want to impose the (ii) above.) This conjecture was apparently inspired by the work of Dobrynin and Mel’nikov [5, 6, 7] on the chromatic number of “arrangement graphs” of Jordan curves in the plane (i.e. graphs whose vertices are the intersection points of the curves and whose edges are the curve segments between these intersection points). To get a feel for the conjecture observe for instance that, while all the pseudocircles of the arrangement in Figure 1 are convex curves and the arrangement is not equivalent to any arrangement of circles, one can easily construct a family of five circles in the plane with the same intersection graph. We are however able to produce a counterexample to the folklore conjecture by adding a “scaffolding” of additional pseudocircles as in Figure 2.

Theorem 1.4 *The intersection graph of the convex pseudocircles in Figure 2 cannot be realised as the intersection graph of a list of circles.*




Figure 2: An arrangement of convex pseudocircles such that no collection of circles defines the same intersection graph.

2 Preliminaries

We will write $[n] := \{1, \dots, n\}$. If $p \in \mathbb{R}^2$ and $r > 0$ then $B(p, r)$ denotes the open disk of radius r around p .

Throughout the paper we identify the Euclidean plane \mathbb{R}^2 with \mathbb{C} . Often it will be convenient to work in the *extended complex plane* $\overline{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$, the one-point compactification of \mathbb{C} . Recall that the extended complex plane $\overline{\mathbb{C}}$ is homeomorphic to the sphere S^2 . (Stereographic projection provides the canonical example of a homeomorphism between S^2 and $\overline{\mathbb{C}}$ if we adopt the convention that the north pole $N \in S^2$ is mapped to $\infty \in \overline{\mathbb{C}}$.) Working in the extended complex plane has a number of advantages. There is for instance no “outer face” that needs separate treatment and we can view lines as circles passing through the point ∞ .

Two basic topological facts we shall use frequently in this paper are the Jordan curve theorem and the Jordan-Schoenflies theorem. The *Jordan curve theorem* states that if c is a Jordan curve in the plane (i.e. c is a homeomorphic image of the unit circle S^1), then $\mathbb{C} \setminus c$ has precisely two arcwise connected components, a bounded one (the “inside”) and an unbounded one (the “outside”). The *Jordan-Schoenflies theorem* states that if $\varphi : S^1 \rightarrow c$ is a homeomorphism, then there also exists a homeomorphism $\bar{\varphi} : \mathbb{C} \rightarrow \mathbb{C}$ that agrees with φ on S^1 . We shall make frequent use of these basic topological facts without mentioning them explicitly each time.

If C is a circle in the plane with centre p and radius r , then the *circle inversion in C* is the map $\varphi : \mathbb{C} \setminus \{p\} \rightarrow \mathbb{C} \setminus \{p\}$ that assigns to a point $z \neq p$ the unique point $\varphi(z)$ that satisfies $|p - \varphi(z)| = r^2/|p - z|$ and lies on the “ray” starting from p and going through z . See Figure 3 for a depiction. The point $\varphi(z)$ can be written in an explicit expression as

$$\varphi(z) = \frac{p \cdot \bar{z} + r^2 - |p|^2}{\bar{z} - \bar{p}}.$$




Figure 3: Inversion in the circle C .

To obtain a map from $\overline{\mathbb{C}}$ to itself, we also define $\varphi(p) := \infty$ and $\varphi(\infty) := p$. A circle inversion has the convenient properties that it maps circles not going through p to circles; it maps circles through p to lines; it maps lines through p to lines; it maps lines not through p to circles; and it exchanges the inside and outside of C (in particular, it is the identity on C and swaps p and ∞). Circle inversions are *conformal*, meaning that they preserve angles (i.e. if two curves meet at a point p and make angle α there, then their images also make angle α at $\varphi(p)$). Note also that φ is a self-homeomorphism of the extended complex plane $\overline{\mathbb{C}}$. More background on circle inversions can for instance be found in [25].

If $\mathcal{C} = (c_1, \dots, c_n)$ is an arrangement of pseudocircles, then we also call the intersection points of the curves the *vertices* of \mathcal{C} and we denote the set of vertices by $I(\mathcal{C})$. A *segment* of \mathcal{C} is the piece of curve connecting two consecutive intersection points. The *faces* of \mathcal{C} are the connected components of $\mathbb{C} \setminus \bigcup \mathcal{C}$, each of which is an open region. There is exactly one unbounded face, also called the *outer face*. The unbounded face is homeomorphic to an open disk with its centre removed, while all bounded faces are homeomorphic to the open disk.

The *combinatorial description* of a pseudocircle arrangement is the associated labelled cell-complex. That is, the lists of faces, segments and intersection points together with a list of the incidences between them and a labelling indicating which is the infinite face and which segment belongs to which pseudocircle. The combinatorial description contains all the relevant combinatorial information of a pseudocircle arrangement, and in particular two pseudocircle arrangements are equivalent if and only if they have the same combinatorial description. Alternative notions of the combinatorial description of an arrangement of pseudocircles have been given by Linhart and Ortner in [15] and Goodman and Pollack in [12].

Recall that a *planar embedding* of a planar (multi)graph G is a map \mathcal{A} that assigns a point $\mathcal{A}(v) \in \mathbb{C}$ to each vertex $v \in V(G)$, and an arc $\mathcal{A}(e) \subseteq \mathbb{C}$ to each edge $e \in E(G)$ in such a way that **(i)** if $e \in E(G)$ has endpoints $u, v \in V(G)$, then $\mathcal{A}(u), \mathcal{A}(v)$ are the endpoints of $\mathcal{A}(e)$, and **(ii)** if $e, f \in E(G)$ are edges, then $\mathcal{A}(e)$ and $\mathcal{A}(f)$ do not meet except possibly at common endpoints. We speak of a *straight-line embedding* if each arc $\mathcal{A}(e)$ is simply the line segment $[\mathcal{A}(u), \mathcal{A}(v)]$ when $e = uv$. (Note that straight-line embeddings only make sense for simple graphs.) We will say that two embeddings \mathcal{A}, \mathcal{B} of a planar multigraph G are *equivalent* (up to the choice of outer face) if there exists a homeomorphism $\varphi : \overline{\mathbb{C}} \rightarrow \overline{\mathbb{C}}$ such that $\varphi(\mathcal{A}(v)) = \mathcal{B}(v)$ for all $v \in V(G)$ and $\varphi[\mathcal{A}(e)] = \mathcal{B}(e)$ for all $e \in E(G)$. (Note that our edges are labelled/named so that there is no confusion among parallel edges.)

A fact that will turn out to be quite useful to us is Whitney's unique embeddability theorem [26].

Theorem 2.1 (Whitney 1932) *If G is a simple, 3-connected planar graph then every two embeddings of G are equivalent.*

Let us remark that the condition that G be *simple* is essential in the theorem. The conclusion is clearly false if we allow multiple edges or loops.

The faces of a planar embedding are defined similarly to the faces of an arrangement of pseudocircles. Every face of a planar embedding is bounded by a closed walk in G , called a *face-walk*. Note that a face-walk need not be a cycle. In fact, a connected planar graph is 2-connected if and only if every face-walk is a cycle (see for instance Theorem 1.6.1 in [13]). It is well known and easy to see that the set of all face-walks determines the embedding of a connected graph up to equivalence.

Rotation systems provide another way to describe an embedding combinatorially. Given a multigraph G , a rotation system π assigns to each vertex of $v \in V(G)$ a cyclic permutation $\pi(v)$ of the edges that are incident with v . (Recall that a cyclic permutation of a set $S = \{s_1, \dots, s_n\}$ is one that can be written in cycle notation as $(s_{i_1} \dots s_{i_n})$, i.e. there is a single cycle containing each element. Put differently, the orbit of each element is the entire set S .) If \mathcal{A} is an embedding of a connected planar (multi)graph G , then the corresponding rotation system $\pi_{\mathcal{A}}$ assigns to each vertex $v \in V(G)$ the cyclic permutation corresponding to the order in which we encounter the edges incident with v if we go around the vertex in clockwise fashion. Observe that if we produce a new embedding \mathcal{B} by applying a reflection in a line to \mathcal{A} then $\pi_{\mathcal{B}}(v)$ will be exactly the opposite order of $\pi_{\mathcal{A}}(v)$ for every vertex $v \in V(G)$. This motivates the next definition. We will say that two rotation systems π, τ defined on a planar multigraph G are *equivalent* if either $\pi(v) = \tau(v)$ for every vertex $v \in V(G)$ or if $\pi(v)$ is the reverse order of $\tau(v)$ for every $v \in V(G)$.

From the rotation system $\pi_{\mathcal{A}}$, we can retrieve the set of all face-walks on \mathcal{A} . (Starting from some vertex, exit along some edge, and keep going “immediately clockwise” until traversal of the initial edge in the same direction. Then the edges traversed form a face-walk. Repeating the procedure for each vertex and edge clearly produces the set of all face-walks corresponding to the embedding.) Conversely, suppose we are given the set of face-walks of some embedding. Let us take two edges e, f sharing an endpoint and lying on a common face F of length at least three. If we declare that e is the successor of f on the “clockwise face-walk around F ”, then this determines all the cyclic orders $\pi(v)$ of a rotation system. And, if we had taken f as the successor of e on the clockwise walk around F , then we would get the reverse order everywhere. Since the face-walks determine the embedding up to equivalence as mentioned earlier, we have the following.

Lemma 2.2 *If G is a connected, planar graph and \mathcal{A}, \mathcal{B} are two embeddings of G , then \mathcal{A} is equivalent to \mathcal{B} if and only if the rotation systems $\pi_{\mathcal{A}}, \pi_{\mathcal{B}}$ they define are equivalent. ■*

A *line arrangement* is a system $\mathcal{L} := (\ell_1, \dots, \ell_n)$ of lines in the plane. If every two lines intersect and no point lies on more than two lines, then we call \mathcal{L} a *simple line arrangement*. A *pseudoline* is the image of a line under a homeomorphism of the plane. A *pseudoline arrangement* is a system of pseudolines satisfying the requirements that every two pseudolines intersect in at most one point, and that when two pseudolines meet in a point they must cross at that point. A *simple pseudoline arrangement* furthermore satisfies the requirement that every two pseudolines intersect and no point lies on more than two pseudolines. We will say that two pseudoline arrangements $\mathcal{L} = (\ell_1, \dots, \ell_n)$ and $\mathcal{M} = (m_1, \dots, m_n)$ are *equivalent* if there exists a homeomorphism φ from the plane onto itself with the property that $\varphi[\ell_i] = m_i$ for all $i \in \{1, \dots, n\}$. The faces, segments and intersection points and combinatorial description of a pseudoline arrangement are defined analogously as for pseudocircle arrangements.

Again, it can be seen that two pseudoline arrangements are equivalent if and only if their combinatorial descriptions coincide. We should mention that various alternative notions of a combinatorial description of a pseudoline arrangement are available in the literature such as local sequences, allowable sequences and oriented matroids of rank 3 (see for instance [9]).

A useful way to picture a pseudoline arrangement is as a *wiring diagram*. In a wiring diagram, each pseudoline is a finite union of line segments. Each pseudoline starts off and ends with half-infinite horizontal segments, all other segments are of finite length having slope either 0, 1 or -1 , and the endpoints of these line segments all lie on the integer grid \mathbb{Z}^2 . See Figure 4 for an example of a wiring diagram. It can be shown that every pseudoline arrangement is




Figure 4: A wiring diagram with three pseudolines.

equivalent to (moreover, can be “continuously deformed” into) a wiring diagram [10].

If a pseudoline arrangement is equivalent to a line arrangement, then we say it is *stretchable*. The name “stretchability” of course comes from imagining the pseudolines “stretching” into lines. STRETCHABILITY is the computational problem of deciding, given a combinatorial description of a pseudoline arrangement as input, whether it belongs to a stretchable arrangement. Analogously, SIMPLE STRETCHABILITY is the same computational problem when the input is restricted to combinatorial descriptions of *simple* pseudoline arrangements. Our proofs of Theorem 1.2 and Theorem 1.3 rely heavily on the following classical result.

Theorem 2.3 (Shor [24]) *SIMPLE STRETCHABILITY is NP-hard.*

This theorem is also a straightforward corollary of a deep topological result by Mnëv [20, 21]. Shor’s proof is more direct and uses Pappus’ and Desargues’ theorems to encode instances of a SAT variant as instances of the simple stretchability problem. (See also Chapter 8 of [3].)

3 Small pseudocircle arrangements are circleable

Let us say that an arrangement of pseudocircles $\mathcal{C} = (c_1, \dots, c_n)$ is *circleable in the extended complex plane* if there exists an arrangement of circles $\mathcal{D} = (d_1, \dots, d_n)$ and a homeomorphism φ of $\overline{\mathbb{C}}$ onto itself such that $\varphi[c_i] = d_i$. When dealing with an arrangement of pseudocircles in the plane there is always an outer face with a different topology from the other faces (i.e. the outer face is homeomorphic to a “punctured disk” while the other faces are homeomorphic to disks), but in the extended complex plane all faces behave the same. This might lead us to suspect that circleability in the plane and circleability in the extended complex plane are two different notions. On the contrary (and fortunately for us) circle inversions provide a short argument that the two notions coincide.

Lemma 3.1 *A pseudocircle arrangement is circleable in the plane if and only if it is circleable in the extended complex plane.*

Proof: Any self-homeomorphism of the plane also induces a self-homeomorphism of the extended complex plane — we just send ∞ to itself. In other words, if an arrangement is circleable in the ordinary plane it is also circleable in the extended complex plane.

To see the converse, suppose that $\mathcal{C} = (c_1, \dots, c_n)$ is circleable in the extended complex plane, and let $\mathcal{D} = (d_1, \dots, d_n)$ be an arrangement of circles such that there is a self-homeomorphism $\varphi : \overline{\mathbb{C}} \rightarrow \overline{\mathbb{C}}$ with $\varphi[c_i] = d_i$ for all $i \in \{1, \dots, n\}$. If it so happens that φ maps ∞ to itself, then we are done: the restriction of φ to \mathbb{C} is a self-homeomorphism of \mathbb{C} demonstrating that \mathcal{C} and \mathcal{D} are equivalent.

So suppose $\varphi(\infty) = p \in \mathbb{C}$. Observe that p does not lie on any of the circles d_1, \dots, d_n . Let $\sigma : \overline{\mathbb{C}} \rightarrow \overline{\mathbb{C}}$ be the circle inversion in a circle C of centre p and radius one, say (in fact any radius will do). Then σ maps each circle d_i to some other circle e_i and it is a self-homeomorphism of the extended complex plane. Observe that $(\sigma \circ \varphi)[c_i] = e_i$ for each $i \in \{1, \dots, n\}$, that $\sigma \circ \varphi$ is a self-homeomorphism of $\overline{\mathbb{C}}$ and that $(\sigma \circ \varphi)(\infty) = \infty$. These three properties show that the restriction of $\sigma \circ \varphi$ to \mathbb{C} demonstrates that \mathcal{C} is circleable in the (ordinary) plane. ■

For the rest of this section, all action will take place in the extended complex plane. We will often find it useful to apply a suitable circle inversion so that one or two circles of an arrangement are mapped to lines. This makes the ensuing case analysis conceptually simpler. We may always apply another circle inversion (in a circle whose centre does not lie on any pseudocircle) to convert an arrangement \mathcal{C} consisting of lines and circles into an equivalent arrangement consisting only of circles. We shall often be informal by referring to pictures and stating that two depicted arrangements are equivalent. It will always be easy to see that the “combinatorial structures” agree, keeping in mind that all lines meet in the point at ∞ . (It may also be helpful to think of arrangements as drawn on the sphere S^2 by taking the inverse stereographic projection and adding the north pole to pre-images of the lines.) Readers can check that equivalences can always be formally justified by repeated invocations of the Jordan-Schoenflies theorem. Before we proceed with the proof of Theorem 1.1, let us make one more fairly straightforward observation.

Lemma 3.2 *If $\mathcal{C} = (c_1, \dots, c_n)$ is such that (c_1, \dots, c_{n-1}) is circleable and c_n intersects at most one of c_1, \dots, c_{n-1} , then \mathcal{C} is circleable.*

Proof: We can assume that c_1, \dots, c_{n-1} are all circles. Let us furthermore assume that c_n intersects precisely one other pseudocircle. (The case when c_n intersects no other pseudocircle is similar and easier.) We can suppose without loss of generality that the unique pseudocircle that c_n intersects is c_{n-1} . By applying a suitable circle inversion (in a circle whose centre lies on c_{n-1} but not on any other c_i) we obtain the situation where c_1, \dots, c_{n-2} are all circles and c_{n-1} is a line. By applying a suitable isometry if needed, we can further assume that the origin lies inside c_n and that c_{n-1} is the x -axis (a.k.a. the real line). See Figure 5, left, for a depiction.




Figure 5: If (c_1, \dots, c_{n-1}) is circleable and c_n intersects only one pseudocircle

Let $A \subseteq [n-2]$ be such that c_i lies inside c_n for $i \in A$ and c_i lies outside c_n when $i \notin A$. If we apply a dilation by a small enough $\lambda_1 > 0$ to all $c_i : i \in A$ and a dilation by a large enough $\lambda_2 > 0$ to all $c_i : i \notin A$, then we can ensure that c_i is contained in a disk of radius $1/2$ around the origin for $i \in A$, while c_i does not intersect the disk of radius 2 for $i \in [n-2] \setminus A$. (We leave c_{n-1} and c_n intact.) This way we have an arrangement that is equivalent to the original such that c_1, \dots, c_{n-2} are circles and c_{n-1} is a line. See Figure 5, middle, for a depiction. It is now clear that we can replace c_n by the unit circle and still have an equivalent arrangement (see Figure 5, right). Applying a circle inversion in a circle with centre not on any c_i gives the sought arrangement of circles and concludes the proof. \blacksquare

Proof of Theorem 1.1: We first prove the theorem for simple pseudocircle arrangements, and later extend it to general pseudocircle arrangements.

Let $\mathcal{C} = (c_1, \dots, c_n)$ be a simple arrangement of pseudocircles on the extended complex plane with $n \leq 4$. By Lemma 3.2, \mathcal{C} is circleable if $n \leq 2$. Let us now consider the case $n = 3$. We may assume c_1 and c_2 meet, because otherwise each of them intersects at most one other pseudocircle, and we are done by Lemma 3.2. By the above, we can apply a homeomorphism of the extended complex plane, to arrive at an arrangement equivalent to \mathcal{C} where both c_1, c_2 are circles (and c_3 is a general Jordan curve). By applying an inversion in a circle with centre one of the intersection points of c_1, c_2 , we arrive at the situation where both c_1, c_2 are lines. Applying another homeomorphism if necessary, we arrive at the situation where c_1 coincides with the x -axis and c_2 coincides with the y -axis (and c_3 is still a general Jordan curve). If c_3 has the origin in its interior, then c_3 must intersect both the positive and the negative x -axis and both the positive and the negative y -axis. (Recall that two pseudocircles intersect either never or twice.) Clearly, the arrangement must then be equivalent to that in Figure 6, left.




Figure 6: Two pseudocircle arrangements with $n = 3$.

Otherwise, c_3 does not contain the origin and it intersects either the positive x -axis twice or the negative x -axis twice, and similarly for the y -axis. We can assume without loss of generality that c_3 intersects the positive x -axis and the positive y -axis twice. But now \mathcal{C} is equivalent to the arrangement in Figure 6, right.

This proves the statement for simple pseudocircle arrangements on at most three curves. Let us thus consider the case when $n = 4$. Again by Lemma 3.2, we can assume that each pseudocircle intersects at least two others.

Our first main case in considering simple \mathcal{C} is that each pseudocircle intersects exactly two others. Without loss of generality we can assume c_1 does not intersect c_3 and c_2 does not intersect c_4 . By the above, (c_1, c_2, c_3) is circleable. Arguing as before, we can apply a self-homeomorphism of $\bar{\mathbb{C}}$ that sends c_1 to the x -axis, c_2 to the y -axis and c_3 to a circle. Hence up to symmetries, the arrangement now looks as in Figure 7, left; and it is clearly circleable

(see Figure 7, right).




Figure 7: An arrangement of four pseudocircles in which each pseudocircle intersects exactly two others.

Our second main case is that at least one pseudocircle intersects all three others, while some pseudocircle intersects only two others. Note that there are exactly two pseudocircles intersecting all others and two intersecting exactly two others. Without loss of generality, we can assume c_3, c_4 do not intersect, while c_1, c_2 intersect each other as well as both of c_3, c_4 . Let us first assume that c_4 separates the two intersection points of c_1, c_2 (i.e. one of the two intersection points lies inside c_4 and the other outside c_4). By the previous cases, we can assume c_1, c_2, c_3 are circles. Applying a circle inversion if needed, we can further assume that c_3 is not contained inside c_4 . We can now replace c_4 by a small circle c'_4 that lies within c_4 and contains the relevant intersection point of c_1, c_2 , to obtain an arrangement of circles equivalent to \mathcal{C} (see Figure 8).




Figure 8: When c_4 separates the two intersection points of c_1, c_2 and c_4 does not intersect c_3 .

By symmetry, we are also done if c_3 separates the two intersection points of c_1, c_2 . Let us therefore consider the case that neither c_3 nor c_4 separate the two intersection points of c_1, c_2 . Note that the pseudocircle arrangement (c_1, c_2, c_3) must necessarily be equivalent (in the extended plane) to the arrangement on the right of Figure 6. By applying a suitable circle inversion, we see that this arrangement is equivalent to the one shown in Figure 9, left. (We have refrained from drawing it as a circle arrangement in this figure to avoid the figure getting too “cramped”.)




Figure 9: The situation where (c_1, c_2, c_3) is equivalent to Figure 6, right.

First we suppose that c_4 intersects some segment of this arrangement exactly once. Then this segment must be one of the segments labelled a, a', b, b', c, c' in Figure 9, left, because if c_4 hits one of the other segments exactly once, then it will separate the intersection points of c_1, c_2 or c_1, c_3 or c_2, c_3 (here we use that if c_4 crosses some segment exactly once then the two endpoints of the segment are in different components of $\overline{\mathcal{C}} \setminus c_4$). Hence, if c_4 intersects a segment other a, a', b, b', c, c' exactly once, then we will be done either by a previous case or by contradiction — if c_4 separates the intersection points of c_1, c_3 or c_2, c_3 , then it intersects c_3 . Let F be the 4-face whose boundary contains a, a' . Since c_4 intersects the boundary of F an even number of times, c_4 intersects a exactly once if and only if it intersects a' exactly once. Similarly c_4 intersects b exactly once if and only if it intersects b' exactly once; and c_4 intersects c exactly once if and only if it intersects c' exactly once. By considering the triangular faces abc and $a'b'c'$, we see that, similarly, if c_4 intersects one of a, b, c exactly once, then it intersects precisely two of them exactly once. And, if c_4 intersects one of a', b', c' exactly once, then it intersects precisely two of them exactly once. It follows that, if c_4 intersects some segment exactly once, then the situation must be as in Figure 9, middle. It is then easily seen that the arrangement \mathcal{C} is equivalent to an arrangement of circles (see Figure 9, right).

We can thus assume that c_4 intersects every segment of the arrangement (c_1, c_2, c_3) an even number of times. (Yet c_4 does not intersect c_3 , and (c_1, c_2, c_3) is equivalent to Figure 9, left.) Note that, since c_4 does not intersect c_3 and it intersects both c_1 and c_2 , we must have that it in fact intersects exactly two segments. There must be a face F of the arrangement (c_1, c_2, c_3) such that c_4 intersects two of its bounding segments. If these segments are consecutive on F , then we can replace c_4 by a sufficiently small circle, close to the point where the segments meet, similar to what we did in Figure 8. If c_4 intersects two segments of F that are not consecutive, then F must be a 4-face, and the situation must be as in Figure 10, left. It is easily seen that in this case we can find an equivalent arrangement of circles (see Figure 10, right).




Figure 10: The situation where (c_1, c_2, c_3) is as in Figure 6, right, and c_4 intersects two segments that bound a 4-face.

This concludes the analysis of the second main case. Our last main case for considering when \mathcal{C} is simple is that every pair of pseudocircles intersect. We first assume that (c_1, c_2, c_3) is as in Figure 6, left. If c_4 surrounds the origin (i.e. it separates the origin from the point ∞), then it must intersect both the positive and the negative x -axis, and both the positive and the negative y -axis. Figure 11 shows the possible positions of these intersection points with respect to c_3 (up to obvious symmetries).




Figure 11: The situation where (c_1, c_2, c_3) is as in Figure 6, left, and c_4 surrounds the origin.

In the first case, c_4 must intersect some segment of c_3 twice. By symmetry we can assume it is the north-east segment. In the second case, the position of the intersection points with c_1, c_2 forces c_4 to intersect c_3 in its north-west and north-east segments. Since c_4 can intersect c_3 at most twice, it is easily seen that the third option in Figure 11 in fact cannot occur. In the fourth case, c_4 necessarily intersects c_3 in its north-west and south-east segments. Observe that in the fifth case, if we apply a circle inversion in a circle with center the origin then we arrive at a situation equivalent to the second case (up to relabelling). Similarly, the sixth case is equivalent to the first case. Figure 12 shows that in the first, second and fourth cases the arrangement \mathcal{C} is circleable.




Figure 12: The arrangements of the first, second and fourth situations in Figure 11 correspond to circleable arrangements.

Let us thus assume that (c_1, c_2, c_3) is as in Figure 6, left, c_4 intersects all three other pseudocircles, and it does not separate the two intersection points of c_1, c_2 (which is the same as separating the origin from ∞ in the arrangement of Figure 6, left). Let us observe that if c_4 intersects some segment of the arrangement (c_1, c_2, c_3) exactly once, then it separates the two intersection points of a pair of pseudocircles. (Note that the arrangement (c_1, c_2, c_3) is in fact extremely symmetric.) Hence, if c_4 intersects some segment of the arrangement (c_1, c_2, c_3) exactly once, then we can relabel pseudocircles and apply a previous case. We can therefore assume c_4 intersects each segment either twice or not at all. There must be some face F such that two of its sides are intersected by c_4 . Now c_4 either intersects the third side of this face, or it intersects a segment of another face (which must share a side with F). Hence, up to symmetry, the situation must be one of those depicted in Figure 13.




Figure 13: The situations where (c_1, c_2, c_3) is equivalent to the arrangement in Figure 6, left, and c_4 does not separate any pair of intersection points.

Figure 14 shows that in each of these two cases the arrangement is circleable.




Figure 14: Arrangements of circles corresponding to those in Figure 13.

It thus remains to consider the case when (c_1, c_2, c_3) is as in Figure 6, right. If c_4 separates the intersection points of c_1, c_2 , then we see that (c_1, c_2, c_4) is as in Figure 6, left. We can then relabel the arrangement and apply a previous case. So we can assume this is not the case. Similarly, we can assume c_4 does not separate the intersection points of c_1, c_3 or of c_2, c_3 . Note that this in particular implies that each segment bounding a 2-face is intersected an even number of times by c_4 . Hence, if c_4 intersects some segment exactly once, then it is one of the segments a, a', b, b', c or c' shown in Figure 9, left. Moreover, by earlier arguments, if c_4 intersects such a segment exactly once, then there are precisely four segments that c_4 intersects exactly once; and up to relabelling these segments are a, a', b, b' in Figure 9. Hence we must be in one of the situations in Figure 15.




Figure 15: (c_1, c_2, c_3) is as in Figure 6, right, c_4 does not separate any pair of intersection points and c_4 intersects some segment exactly once.

By obvious symmetries, the second and fourth cases in Figure 15 are the same. Similarly, the first, third and fifth cases are also identical (by applying a suitable inversion to switch the inner triangular face and outer triangular face to see the third case is equivalent to the first case). It is thus again easy to see that in both cases there is an equivalent arrangement of circles (see Figure 16).




Figure 16: The arrangements of Figure 15 are circleable.

Hence we can assume that c_4 intersects each segment of (c_1, c_2, c_3) either zero or two times (and (c_1, c_2, c_3) is as in Figure 6, right, and every two pseudocircles intersect). There must be some face F such that c_4 intersects two of its sides. If F is a 2-face then, up to symmetry, the situation must be as in Figure 17, left.




Figure 17: The arrangement (c_1, c_2, c_3) is as in Figure 6, right, and c_4 intersects two sides of a 2-face.

If we perform a circle inversion in the point p labelled in the figure, then we get a situation as in Figure 17, middle, and this is again easily seen to be circleable (see Figure 17, right).

Let us now suppose F is a 3-face. If c_4 intersects all three sides of F , then the situation must be as in Figure 18, left. (Recall that by a suitable circle inversion we can make any face the outer face.)




Figure 18: The arrangement (c_1, c_2, c_3) is as in Figure 6, right, and c_4 intersects all three sides of a 3-face.

It is again easy to see \mathcal{C} is circleable (see Figure 18, right). If, on the other hand, F is a 3-face and c_4 intersects exactly two of its sides, then up to symmetry the situation is as in

Figure 19, left.




Figure 19: The arrangement (c_1, c_2, c_3) is as in Figure 6, right, and c_4 intersects two sides of a 3-face.

If we perform a circle inversion in a circle with centre the point labelled p in Figure 19, left, then we obtain a situation as in Figure 19, middle, and it again easy to see that \mathcal{C} is circleable (see Figure 19, right).

Let us then assume that F is a 4-face. If c_4 intersects the sides of a 2-face or 3-face twice, then we are done. Hence it remains to consider the case when c_4 intersects three sides of F . In that case, up to symmetry the situation is as in Figure 20, left, and the arrangement is easily seen to be circleable (see Figure 20, right).




Figure 20: The arrangement (c_1, c_2, c_3) is as in Figure 6, right, and c_4 intersects three sides of a 4-face.

This concludes our case analysis for the case when \mathcal{C} is simple. It remains to consider the case that \mathcal{C} is not simple. In this case, there must be a point p common to at least three pseudocircles. Let us first assume that p lies on all of the pseudocircles. If we perform an inversion in a circle with centre p , then we in fact get a pseudoline arrangement \mathcal{L} on at most four lines. By the result of Goodman and Pollack [11], \mathcal{L} is equivalent to a line arrangement \mathcal{L}' . Performing an inversion in a circle with centre not on any line of \mathcal{L}' yields a circle arrangement \mathcal{C}' that is equivalent to \mathcal{C} .

Let us thus assume that $n = 4$, that there is a point p common to three pseudocircles, but that there is no point common to all four pseudocircles. We can assume without loss of generality that p lies on c_1, c_2, c_3 . Applying a circle inversion to an arrangement in which c_1, c_2, c_3 are circles and c_4 is a general Jordan curve (such an arrangement exists by the last argument), we obtain a situation where c_1, c_2, c_3 are lines and c_4 is a Jordan curve.

Let us first assume that c_1, c_2, c_3 have another point q in common, which implies that the corresponding line arrangement is the unique non-simple line arrangement of three lines. Then, depending on whether c_4 separates p and q , whether it intersects only two other pseu-

dolines, or whether it does not separate p, q but intersects the other three pseudolines, the arrangement must be equivalent (up to relabelling and symmetries) to one of the arrangements in Figure 21, and hence is circeable.




Figure 21: The curves c_1, c_2, c_3 have two points in common.

Let us therefore assume that p is the only point common to c_1, c_2, c_3 , (so that the corresponding line arrangement is the unique simple line arrangement on three lines). If furthermore c_4 intersects only two other curves, then the arrangement must clearly be equivalent (up to relabeling and symmetries) to one of the arrangements in Figure 22.




Figure 22: The curves c_1, c_2, c_3 have exactly one point in common, and c_4 intersects two of them.

If p is the only point common to c_1, c_2, c_3 and c_4 has at least one of the three intersection points of the lines in its interior, then the arrangement must be equivalent (up to relabeling and symmetries) to one of the arrangements in Figure 23, and hence it is clearly circleable.




Figure 23: The curves c_1, c_2, c_3 have exactly one point in common, and c_4 intersects three of them and has at least one of the three intersection points in its interior.

Finally, let us assume p is the only point common to c_1, c_2, c_3 , that c_4 intersects all three other pseudolines and that c_4 does not have any of the three intersection points of the corresponding lines in its interior. Then, up to relabelling and symmetries, the situation must be as in one of the four cases in the top row of Figure 24. Hence the arrangement is circleable as well (Figure 24, bottom row).




Figure 24: The curves c_1, c_2, c_3 have exactly one point in common, and c_4 intersects three of them and has none of the three intersection points in its interior.

This concludes our (lengthy) case analysis, and proves that every pseudocircle arrangement on at most four pseudocircles is circleable. ■

4 A counterexample to the folklore conjecture

Before starting with the proof of Theorem 1.4, let us clarify the construction of the arrangement $\mathcal{C} = (c_1, \dots, c_{346})$ depicted in Figure 2. We start by taking the five curves c_1, \dots, c_5 from Figure 1 and overlay the curves of the arrangement $\mathcal{H} = (c_6, \dots, c_{346})$ shown in Figure 25, left. The curves from \mathcal{H} are placed in such a way that each of the initial five curves does not




Figure 25: The arrangement \mathcal{H} of convex pseudocircles and its intersection graph H .

intersect the unbounded face of \mathcal{H} , and for each of the 8-faces of \mathcal{H} , the corresponding eight curves intersect zero, one or two of c_1, \dots, c_5 in one of the ways drawn in Figure 26. This concludes our description of the construction of \mathcal{C} .




Figure 26: The different ways (up to rotation and reflection) in which the five initial pseudocircles may intersect the pseudocircles corresponding to an 8-face of \mathcal{H} .

The proof of Theorem 1.4 amounts to showing, under the assumption that \mathcal{C} can be realised as the intersection graph of a list of circles, that the pseudocircle arrangement of Figure 1 is circleable, a contradiction. The proof also relies on the following observation, the proof of which is a straightforward perturbation argument. Nevertheless, we include this argument for completeness.

Lemma 4.1 *Every graph that is an intersection graph of a list of circles is also the intersection graph of a simple arrangement of circles.*

Proof: For notational convenience, let us denote by $S(p, r)$ the circle with centre p and radius r . We will make use of the following three straightforward observations.

- (i) If $|S(p, r) \cap S(q, s)| = 0$, then there exists $\varepsilon > 0$ such that $|S(p', r') \cap S(q, s)| = 0$ for all $r' \in (r - \varepsilon, r + \varepsilon)$ and $p' \in B(p, \varepsilon)$.
- (ii) If $|S(p, r) \cap S(q, s)| = 2$, then there exists $\varepsilon > 0$ such that $|S(p', r') \cap S(q, s)| = 2$ for all $r' \in (r - \varepsilon, r + \varepsilon)$ and $p' \in B(p, \varepsilon)$.
- (iii) If $|S(p, r) \cap S(q, s)| = 1$ and $r \leq s$, then there exists $\varepsilon > 0$ such that $|S(p, r') \cap S(q, s)| = 2$ for all $r' \in (r, r + \varepsilon)$.

Let us fix a graph G that is an intersection graph of a list of circles. For every list $\mathcal{D} = (d_1, \dots, d_n)$ of circles that has G as its intersection graph, let us say that index i is *good* if d_i is not tangent to any other circle of \mathcal{D} , and no point of d_i lies on three or more circles of \mathcal{D} . If i is not good, then we call it *bad*. Let $\text{bad}(\mathcal{D})$ denote the number of bad indices in \mathcal{D} .

Now let $\mathcal{D} = (d_1, \dots, d_n)$ be an arbitrary list of circles with intersection graph G , and suppose $\text{bad}(\mathcal{D}) > 0$. Let p_i be the centre of d_i and r_i its radius. For now, let us assume that d_1, \dots, d_n are all distinct (i.e. $(p_i, r_i) \neq (p_j, r_j)$ if $i \neq j$). We can assume without loss of generality that the labelling is such that $r_1 \leq r_2 \leq \dots \leq r_n$.

Let $i \in \{1, \dots, n\}$ be the maximal index such that each of d_1, \dots, d_i are good. (So in particular d_{i+1} is bad.)

By observations (i)–(iii) above, there exists $\varepsilon > 0$ such that $S(p_{i+1}, r)$ intersects precisely the same subset of $d_1, \dots, d_i, d_{i+2}, \dots, d_n$ as d_{i+1} does for all $r \in [r_{i+1}, r_{i+1} + \varepsilon)$. We can therefore also choose $r' \in [r_{i+1}, r_{i+1} + \varepsilon)$ such that the circle $d'_{i+1} := S(p_{i+1}, r')$ intersects the same set of circles as d_{i+1} , is not tangent to any of the other circles, and does not pass through the intersection points of any pair of the other circles. Hence the list $\mathcal{D}' := (d_1, \dots, d_i, d'_{i+1}, d_{i+2}, \dots, d_n)$ of circles has intersection graph G and $\text{bad}(\mathcal{D}') < \text{bad}(\mathcal{D})$.

Clearly, after finitely many iterations of this procedure, we arrive at a list \mathcal{D}' of circles with intersection graph G and $\text{bad}(\mathcal{D}') = 0$. In other words, \mathcal{D}' is the sought arrangement of circles with intersection graph G .

This proves the lemma in the case when all circles in the initial list are distinct. Otherwise, we first remove duplicates, apply the above argument to the thinned list, and then use **(i)** and **(ii)** to replace the duplicates, but each with a centre p that is distinct from the centres of all the other circles. ■

We are now ready for the proof of Theorem 1.4.

Proof of Theorem 1.4: Let us denote the intersection graph of \mathcal{C} by G , with vertex set $\{1, \dots, 346\}$. Let H be the subgraph of G corresponding to \mathcal{H} . (Observe that H can also be obtained by taking the 11×11 -grid and subdividing each edge. A planar embedding of H is given in Figure 25, right.)

Aiming for a contradiction, suppose that $\mathcal{D} = (d_1, \dots, d_{346})$ is a list of circles whose intersection graph is G . By Lemma 4.1, we may assume without loss of generality that \mathcal{D} is also a simple arrangement of circles. For notational convenience, let p_i denote the centre of d_i and let r_i denote its radius, and let us set $\mathcal{H}' := (d_6, \dots, d_{346})$. Observe that H has the following property:

- for every three distinct vertices u, v, w of H , there is a uv -path P in $H \setminus w$ such that no internal vertex of P is adjacent to w .

This implies that there cannot be three distinct indices $i, j, k \in \{6, \dots, 346\}$ such that d_j is in the interior of d_i and d_k is outside of d_i . (Otherwise, if $i_1 \dots i_m$ is a jk -path in H that avoids i , then one of the circles d_{i_1}, \dots, d_{i_m} will intersect d_i , a contradiction.) Thus, if there is at least one circle of \mathcal{H}' in the interior of circle d_i , then every other circle of \mathcal{H}' either intersects d_i or is in its interior. It follows that, by applying a suitable circle inversion if needed, we can assume that no d_i has another d_j in its interior (with $i, j \in \{6, \dots, 346\}$).

For $i \in \{6, \dots, 346\}$, let $D_i := B(p_i, r_i)$ be the (open) disk with centre p_i and radius r_i (i.e. the boundary of D_i is d_i). Observe that

$$D_i \cap D_j \neq \emptyset \text{ if and only if } d_i \cap d_j \neq \emptyset \text{ for all pairs } i, j \in \{6, \dots, 346\}.$$

In other words, we may consider \mathcal{H}' as a realisation of H as an intersection graph of disks.

Since H is triangle-free and has minimum degree two, the points p_i together with the line segments $[p_i, p_j]$ for $ij \in E(H)$ constitute a straight-line embedding of H (see for instance Lemma 5.2 of [19]). Observe that, as H is a subdivision of a 3-connected planar graph, it also follows from Whitney's unique embeddability theorem that every two embeddings of H are equivalent. In fact, as we will now see, something even stronger is true.

Claim 4.2 *The arrangements of pseudocircles \mathcal{H}' and \mathcal{H} are equivalent in the extended complex plane.*

Proof of Claim 4.2: Consider the rotation system corresponding to the straight-line embedding of H we have obtained. By Lemma 2.2, we know that it is equivalent to the rotation system provided by the embedding in Figure 25, right. Applying a reflection if needed, we can assume that the two rotation systems are in fact identical. (And clearly \mathcal{D} still has all the properties we assumed and derived until now.)

Let us now fix $i \in \{6, \dots, 346\}$ and suppose that i is adjacent to four neighbours in H . (In other words, c_i is one of the square boxes not touching the outer face of \mathcal{H} .) Let i_1, i_2, i_3, i_4 denote the four neighbours of i . Without loss of generality, they appear in this order in the

cyclic permutation of i . Since the rotation systems agree, d_i intersects each of d_{i_1}, \dots, d_{i_4} twice, and H is triangle free, we see that there is also a straightforward correspondence between the intersection points and segments in our arrangement \mathcal{H}' and those in \mathcal{H} . See Figure 27. (For $j \in \{1, 2, 3, 4\}$, let $p_{(i,i_j)}^1, p_{(i,i_j)}^2$ be the two intersection points of c_i and c_{i_j} ,




Figure 27: The pseudocircles $c_i, c_{i_1}, c_{i_2}, c_{i_3}, c_{i_4}$ (left), the corresponding part of the “canonical” embedding of H from Figure 25 (middle) and the circles $d_i, d_{i_1}, d_{i_2}, d_{i_3}, d_{i_4}$ (right). On the right we have also drawn the relevant edges of the straight-line embedding of H that \mathcal{H}' defines.

and let $q_{(i,i_j)}^1, q_{(i,i_j)}^2$ be the intersection points of d_i and d_{i_j} . Without loss of generality, the labelling is such that $p_{(i,i_j)}^1$ comes before $p_{(i,i_j)}^2$ if we go around c_i in the clockwise direction; and similarly for $q_{(i,i_j)}^1, q_{(i,i_j)}^2$ and the clockwise order on d_i . Now observe that the clockwise order of the intersection points on c_i is $p_{(i,i_1)}^1, p_{(i,i_1)}^2, p_{(i,i_2)}^1, p_{(i,i_2)}^2, p_{(i,i_3)}^1, p_{(i,i_3)}^2, p_{(i,i_4)}^1, p_{(i,i_4)}^2$ and on d_i the clockwise order is $q_{(i,i_1)}^1, q_{(i,i_1)}^2, q_{(i,i_2)}^1, q_{(i,i_2)}^2, q_{(i,i_3)}^1, q_{(i,i_3)}^2, q_{(i,i_4)}^1, q_{(i,i_4)}^2$. A similar statement holds if c_i is one of the square boxes on the outer face, so we can indeed label the intersection points of \mathcal{H} and \mathcal{H}' in a consistent way. We can then also label the segments of \mathcal{H} and \mathcal{H}' such that a segment of \mathcal{H} and a segment of \mathcal{H}' are labelled the same if and only if their endpoints correspond and they are part of corresponding pseudocircles.

It is easy to see that in this way we have obtained two embeddings of the same planar graph (the “arrangement graph” of \mathcal{H} , whose vertices are the intersection points — not the intersection graph of the pseudocircles) and both embeddings define the same rotation system. It follows by Lemma 2.2 that there is a self-homeomorphism $\varphi : \overline{\mathbb{C}} \rightarrow \overline{\mathbb{C}}$ such that $\varphi[d_i] = c_i$ for all $i \in \{6, \dots, 346\}$. This proves the claim. \blacksquare

Let $\varphi : \overline{\mathbb{C}} \rightarrow \overline{\mathbb{C}}$ be the self-homeomorphism with $\varphi[d_i] = c_i$ for all $i \in \{6, \dots, 346\}$ provided by the last claim. Let us set $c'_i := \varphi[d_i]$ for $i \in \{1, \dots, 5\}$. Then $\mathcal{C}' := (c'_1, \dots, c'_5, c_6, \dots, c_{346})$ is an arrangement of pseudocircles equivalent to \mathcal{D} (in the extended complex plane), and

$$c'_i \cap c_j \neq \emptyset \text{ if and only if } c_i \cap c_j \neq \emptyset \text{ for all } i \in \{1, \dots, 5\} \text{ and all } j \in \{1, \dots, 346\}. \quad (1)$$

Our next aim in the proof will be to show that, for every $i \in \{1, \dots, 5\}$, the curve c'_i intersects precisely the same segments and faces of the arrangement \mathcal{H} as c_i . From this we will then be able to derive that the arrangements (c'_1, \dots, c'_5) and (c_1, \dots, c_5) are equivalent, which will give us a contradiction since (c_1, \dots, c_5) is not circleable by the result of Linhart and Ortner.

Let us thus pick $i \in \{1, \dots, 5\}$ and $j \in \{6, \dots, 346\}$ such that $c_i \cap c_j \neq \emptyset$. We first suppose that c_j is one of the non-square rectangles of \mathcal{H} . Let $j_1, j_2 \in \{6, \dots, 346\}$ be the two neighbours of j in H . Let b be the segment of c_j inside c_{j_1} , let b' be the segment of

c_j inside c_{j_2} , and let a, a' be the two other segments (see Figure 28). By the construction




Figure 28: The segments a, a', b, b' .

of \mathcal{C} , c_i intersects both of the segments a, a' exactly once and it does not intersect b, b' (see Figure 26). We aim to show the same is true for c'_i .

Let us first note that c'_i cannot intersect b , since it does not intersect c_{j_1} but it does intersect some c_k lying outside c_{j_1} . Similarly, c'_i does not intersect b' . Since \mathcal{C}' is an arrangement of pseudocircles, c'_i intersects c_j at exactly two points. Aiming for a contradiction, let us assume that both intersection points lie on a (and therefore not on a'). Let us consider the 8-face F of \mathcal{H} whose boundary contains a . Let us first assume that c_i intersects them as in the second (from left) situation in Figure 26. Let γ be the curve consisting of the segments bounding F , and let γ' be the curve bounding the union of F with the eight rectangles touching F . Let f, f' be the two other segments of these eight rectangles that c_i intersects. See Figure 29. Observe




Figure 29: The segments a, a', f, f' and the curves γ, γ' .

that c'_i cannot intersect any of the segments in Figure 29 besides a, f, f' (otherwise it would intersect a pseudocircle that c_i does not or it would intersect the inside of a pseudocircle that c_i does not intersect, both of which are impossible). Since c'_i intersects pseudocircles that lie outside of γ' , it follows that c'_i crosses both f and f' . Now recall that two Jordan curves cross an even number of times — this is a straightforward consequence of the Jordan curve theorem. Applying this fact to γ and γ' , it follows that c'_i intersects both f and f' at least twice. In particular, c'_i intersects some pseudocircle c_k at least four times, contradicting the fact that \mathcal{C}' is an arrangement of pseudocircles. For the other case, i.e. c_i intersects F as in the fifth situation of Figure 26, we obtain a contradiction in a completely analogous way. It thus follows that c'_i intersects each of a, a' precisely once.

Completely analogously, if c_j is one of the squares rather than one of the non-square rectangles, then c'_i intersects precisely the same segments of c_j in \mathcal{H} (precisely the same number of times) as does c_i . These are the third and fourth situations of Figure 26.

Hence c'_i intersects precisely the same faces and segments of \mathcal{H} as c_i . Let us fix $1 \leq i < j \leq 5$. Observe that if c_i, c_j do not intersect then they intersect distinct sets of edges and faces of \mathcal{H} . In that case c'_i, c'_j also do not intersect. Now suppose that c_i, c_j do intersect. Then there are two distinct faces of \mathcal{H} where they cross. Looking at the rightmost two cases of Figure 26, we see that if c_i, c_j (with $1 \leq i < j \leq 5$) cross inside a face of \mathcal{H} then c'_i, c'_j cross (at least once) inside the same face. Since c'_i, c'_j cross at most two times and cross at least once in two distinct faces, it follows that they cross exactly once inside each of these faces.

Similarly to what we did in the proof of Claim 4.2, we can make a correspondence between those intersection points and segments of \mathcal{C} and those of \mathcal{C}' and then apply Lemma 2.2 to see that \mathcal{C} and \mathcal{C}' are in fact equivalent. But this also proves that \mathcal{C} and \mathcal{D} are equivalent, and in particular (c_1, \dots, c_5) is equivalent to (d_1, \dots, d_5) . But this cannot be, because d_1, \dots, d_5 are circles and (c_1, \dots, c_5) is not equivalent to any arrangement of circles by the result of Linhart and Ortner [16]. This last contradiction proves that G is not the intersection graph of a list of circles, as required. \blacksquare

5 Hardness of circleability

In overview, the proof of Theorem 1.2 passes through two reductions, embodied in Lemmas 5.1 and 5.4, relying on an intermediate class of pseudocircle arrangements we now define. We will say that an arrangement of pseudocircles $\mathcal{C} = (c_1, \dots, c_n)$ is *pencil* if there is a point p common to all curves, and every point of $\mathcal{C} \setminus \{p\}$ is on at most two of the curves. See Figure 30 for a depiction of a pencil arrangement with $n = 3$.




Figure 30: A pencil pseudocircle arrangement.

Lemma 5.1 *Given a combinatorial description of a simple pseudoline arrangement \mathcal{L} , one can construct, in polynomial time, a combinatorial description of a pencil pseudocircle arrangement \mathcal{C} that is circleable if and only if \mathcal{L} is stretchable.*

Proof: Let \mathcal{L} be a simple pseudoline arrangement. Observe that if x is a point not on any line then circle inversion in a circle with centre x will produce a pencil pseudocircle arrangement \mathcal{C} . The face of \mathcal{L} containing x will get mapped to the outer face of \mathcal{C} , and x will be the unique point common to all pseudocircles of the pencil pseudocircle arrangement \mathcal{C} . More generally, there is a one-to-one correspondence between the faces, segments and intersection points of \mathcal{L} and those of \mathcal{C} , and the combinatorial description of \mathcal{C} can be easily read off from that of \mathcal{L} once we fix the face that the point x lies in.

It remains to see that \mathcal{L} is stretchable if and only if \mathcal{C} is circleable. Let us first assume that \mathcal{L} is stretchable. In that case, if we take an arrangement \mathcal{L}' of lines equivalent to \mathcal{L} , and we pick x' as the face corresponding to the face of \mathcal{L} that we took x in, then it is easy to see that circle inversion in a circle with centre x' will produce a circle arrangement \mathcal{C}' equivalent to \mathcal{C} . Similarly, suppose that \mathcal{C}' is an arrangement of circles equivalent to \mathcal{C} , and let p be the unique point contained in every circle of \mathcal{C}' . Then circle inversion in a circle with centre p will produce a line arrangement \mathcal{L}' that is equivalent to \mathcal{L} . ■

Lemma 5.2 *Every pseudoline arrangement is equivalent to a pseudoline arrangement in which each pseudoline is a convex curve consisting of finitely many line segments.*

Proof: Let \mathcal{L} be an arbitrary pseudoline arrangement. By a result of [11], we may assume \mathcal{L} is represented by a wiring diagram. Without loss of generality, we assume that all finite line segments of the wiring diagram (i.e. all except the beginning- and end-segment of each pseudoline) are contained within the square $[0, a]^2$. We first apply the self-homeomorphism of the plane $g(x, y) := (x, y \cdot 2^{-1000a})$, to obtain an equivalent pseudoline arrangement \mathcal{L}' . Next, we apply the self-homeomorphism

$$f(x, y) := \begin{cases} (x, y - x - 1) & \text{if } x < 0, \\ \left(x, y - \sum_{i=0}^j 2^{-i} - (x - j) \cdot 2^{-(j+1)} \right) & \text{if } j \leq x < j + 1 \text{ for some } j \in \{0, \dots, a - 1\}, \\ (x, y - \sum_{i=0}^a 2^{-i}) & \text{if } x \geq a. \end{cases}$$

to obtain a new pseudoline arrangement \mathcal{L}'' . See Figure 31 for a depiction.




Figure 31: Turning a wiring diagram into an arrangement of convex curves (not to scale).

By construction, \mathcal{L}'' is equivalent to \mathcal{L} and each of its pseudolines consists of at most $a + 2$ line segments, with their endpoints on the vertical lines $\{x = 0\}, \dots, \{x = a\}$. It remains to see that each pseudoline of \mathcal{L}'' is convex.

To see this, fix some pseudoline $\ell \in \mathcal{L}''$. Observe that $\ell \cap \{x \leq 0\}$ is a half-infinite line segment with slope -1 . For each $0 \leq j < a$, the segment $\ell \cap \{j \leq x < j + 1\}$ has slope between $-2^{-(j+1)} - 2^{-1000a}$ and $-2^{-(j+1)} + 2^{-1000a}$ (as it is the image under $f \circ g$ of a line segment with slope between -1 and $+1$). The last segment (i.e. $\ell \cap \{x \geq a\}$) has slope equal to zero. We conclude the slopes are strictly increasing, which amounts to the curve being convex. ■

This last lemma allows us to give a relatively easy proof of the fact that every pencil pseudocircle arrangement is equivalent to a convex pseudocircle arrangement. (Observe that this is not true for general pseudocircle arrangements by Theorem 1.3, which we will prove in the next section.)

Lemma 5.3 *Every pencil pseudocircle arrangement is equivalent to an arrangement of convex polygons.*

Proof: Let \mathcal{C} be an arbitrary pencil pseudocircle arrangement, and let \mathcal{L} be the simple pseudoline arrangement we get by an inversion in a circle whose centre is the unique point common to all pseudocircles. Let \mathcal{L}' be a wiring diagram equivalent to \mathcal{L} . Observe that the face of \mathcal{L}' corresponding to the outer face of \mathcal{C} is either the face “above” all the pseudolines or the face “below” all of them. Without loss of generality, let us assume it is below.

By Lemma 5.2, there is an equivalent pseudoline arrangement \mathcal{L}'' in which each line is a convex curve made up of finitely many line segments, with the face of \mathcal{L}'' below all pseudolines corresponding to the outer face of \mathcal{C} . If we apply an inversion in a circle completely contained within this face, then we get a pencil pseudocircle arrangement $\mathcal{C}' = (c'_1, \dots, c'_n)$ that is equivalent to \mathcal{C} . Moreover, since circle inversions are conformal and map lines outside the circle to circles inside, it follows that each c'_i is a closed convex curve through the origin, consisting of finitely many circular arcs. We can now replace the circular arcs of each c'_i by polygonal paths with the same endpoints (that are sufficiently fine approximations) so that each c'_i remains convex and the combinatorial description of \mathcal{C}' remains unchanged. In this way, we obtain the required arrangement \mathcal{C}' of convex polygons that is equivalent to \mathcal{C} . ■

The construction in the next lemma is inspired by a similar construction of Jaggi et al. [14] for oriented matroids.

Lemma 5.4 *Given a combinatorial description of a pencil pseudocircle arrangement $\mathcal{C} = (c_1, \dots, c_n)$, one can construct, in polynomial time, the combinatorial description of a simple, convex pseudocircle arrangement $\mathcal{C}' := (c_1^-, c_1^+, \dots, c_n^-, c_n^+)$ on twice as many curves that is circleable if and only if \mathcal{C} is.*

The proof of this last lemma relies on the following geometric fact, whose straightforward proof we leave to the reader.

Lemma 5.5 *Suppose that c, c' are two circles in the plane with c' inside c , and let the point p lie inside c but outside c' . Then there is a circle c'' through p such that c'' lies inside c and c' lies inside c'' . See Figure 32. ■*




Figure 32: An illustration of Lemma 5.5.

Proof of Lemma 5.4: Let $\mathcal{C} = (c_1, \dots, c_n)$ be an arbitrary pencil pseudocircle arrangement, and let p denote the unique point that is on all curves. By Lemma 5.3, we assume without

loss of generality that the curves c_1, \dots, c_n are convex polygons. We construct the curves c_i^- by tracing along c_i , very close to the original curve, yet completely inside c_i ; and construct c_i^+ similarly by tracing along and just outside of c_i (see Figure 33, left). By choosing the curves




Figure 33: Replacing each curve of a pencil pseudocircle arrangement by two curves.

c_i^-, c_i^+ sufficiently close to the original curves, we can ensure that each intersection point other than p lies in a unique 4-face of the arrangement $\mathcal{C}' := (c_1^-, c_1^+, \dots, c_n^-, c_n^+)$. The 4-face containing the intersection point of c_i, c_j other than p is then one of the two regions of the area between c_i^- and c_i^+ and between c_j^- and c_j^+ . Observe that, since \mathcal{C} was an arrangement of convex polygons, we can certainly ensure that the curves of \mathcal{C}' are convex too. We can also ensure that the arrangement \mathcal{C}' is simple. Moreover, while placing c_i^-, c_i^+ , we can inductively ensure that, for all $j, k < i$, the area containing p that is between c_j^-, c_j^+ and between c_k^-, c_k^+ is also completely contained between c_i^- and c_i^+ (see Figure 33, right).

This last step guarantees that this procedure uniquely determines the combinatorial description of a simple, convex pseudocircle arrangement. Having access to an explicit representation of \mathcal{C} in the plane is not necessary to determine \mathcal{C}' . Indeed, it is not hard to see that the combinatorial description of \mathcal{C}' can be produced in polynomial time directly from the combinatorial description of \mathcal{C} . (Notice for instance that for every face of \mathcal{C} we get a face in \mathcal{C}' , for every segment of \mathcal{C} we get a face and some segments bounding it, for every intersection point other than p we get a 4-face, etc., and the incidences between the faces, segments and intersection points of \mathcal{C}' are completely determined by the combinatorial description of \mathcal{C} .)

It remains to be seen that \mathcal{C} is circleable if and only if \mathcal{C}' is. First assume that \mathcal{C} is circleable, and let $\mathcal{D} = (d_1, \dots, d_n)$ be an equivalent arrangement of circles. Then it is not hard to see that we can pick circles d_i^-, d_i^+ , concentric with d_i for each i , such that the arrangement $\mathcal{D}' = (d_1^-, d_1^+, \dots, d_n^-, d_n^+)$ is equivalent to \mathcal{C}' .

Conversely, suppose that $\mathcal{D}' = (d_1^-, d_1^+, \dots, d_n^-, d_n^+)$ is an arrangement of circles equivalent to \mathcal{C}' . Let p' be an arbitrary point inside the face of \mathcal{D}' that corresponds to the face of \mathcal{C}' that contains p . By Lemma 5.5, there are circles d_1, \dots, d_n such that $p' \in d_i$ and d_i lies inside d_i^+ and outside d_i^- for each i . Clearly $\mathcal{D} = (d_1, \dots, d_n)$ is a pencil circle arrangement. To see that it is equivalent to \mathcal{C} , we can for instance make use of rotation systems again. Let φ be a self-homeomorphism of the plane satisfying $\varphi[d_i^-] = c_i^-, \varphi[d_i^+] = c_i^+$ for all i , and let the pseudocircle arrangement $\mathcal{E} = (e_1, \dots, e_n)$ be defined by $e_i := \varphi[d_i]$. Then every e_i contains

the point $\varphi(p')$, which lies in the same face of the arrangement \mathcal{C}' as p . And, for each pair $i \neq j$, the other intersection point of e_i and e_j lies in the same face of \mathcal{C}' as the intersection point of c_i and c_j not equal to p . Moreover, for each i , e_i intersects exactly the same faces and segments of \mathcal{C}' as c_i does. It follows that the planar multigraphs defined by the intersection points and segments of \mathcal{E} and of \mathcal{C} define the same rotation system, and so must be equivalent. It then also follows that \mathcal{D} and \mathcal{C} are equivalent, which shows \mathcal{C} to be circleable. ■

For completeness we collect our findings into an explicit proof.

Proof of Theorem 1.2: Given a combinatorial description of an arbitrary simple pseudoline arrangement \mathcal{L} , we can transform it, in polynomial time, into a combinatorial description of a pencil pseudocircle arrangement \mathcal{C} that is circleable if and only if \mathcal{L} is stretchable using Lemma 5.1. Using Lemma 5.4, we can then produce, in polynomial time, a combinatorial description of a simple, convex pseudocircle arrangement \mathcal{C}' that is circleable if and only if \mathcal{C} is circleable. Thus, if there is a polynomial-time algorithm to decide if an arbitrary simple arrangement of convex pseudocircles is circleable, then there must also be a polynomial-time algorithm to solve SIMPLE STRETCHABILITY. By Theorem 2.3, it follows that our computational problem is NP-hard. ■

6 Hardness of convexibility

Proof of Theorem 1.3: Let us consider an arbitrary simple arrangement $\mathcal{L} = (\ell_1, \dots, \ell_n)$ of n pseudolines. We may assume that \mathcal{L} is given as a wiring diagram. From \mathcal{L} , we shall construct an arrangement $\mathcal{C} = (c_1^-, c_1^+, \dots, c_n^-, c_n^+)$ of $2n$ pseudocircles with the property that \mathcal{C} is convexible if and only if \mathcal{L} is stretchable. For each of the “wires” ℓ_i , we construct a




Figure 34: The construction of the pseudocircle arrangement \mathcal{C} (right) from a simple pseudoline arrangement \mathcal{L} (left).

pair c_i^-, c_i^+ of pseudocircles that trace along ℓ_i , with c_i^- just below and c_i^+ just above it. We convert the curves c_i^-, c_i^+ into Jordan curves, using vertical and horizontal line segments, as illustrated in Figure 34, right. Observe that this can be done in such a way that c_i^+ and c_i^- each intersect both c_j^- and c_j^+ exactly twice if $j \neq i$; c_i^- and c_i^+ do not intersect; and the unique intersection point of ℓ_i, ℓ_j ($i \neq j$) lies in a unique 4-face of the arrangement \mathcal{C} . As in the last section, having access to an explicit representation of \mathcal{L} in the plane is not actually necessary to determine \mathcal{C} : the combinatorial description of \mathcal{C} can be produced in polynomial time directly from the combinatorial description of \mathcal{L} .

Let us now see how \mathcal{C} is convexible if and only if \mathcal{L} is stretchable. First assume that \mathcal{L} is stretchable. Then there is a representation of \mathcal{L} as an arrangement of lines in the plane. By performing an appropriate linear transformation if needed, we can assume without loss

of generality both that the lines have arbitrarily small slopes and that the initial vertical ordering of lines “agrees” with the corresponding initial ordering of pseudolines as given by the wiring diagram. Therefore, using a minor modification of the procedure that produced \mathcal{C} from the wiring diagram, we can produce a representation of \mathcal{C} as an arrangement of convex quadrilaterals (see Figure 35).




Figure 35: Repeating the construction when \mathcal{L} is stretchable.

This proves that \mathcal{C} is convexible if \mathcal{L} is stretchable. To complete the proof, we now show that \mathcal{L} is stretchable under the assumption that \mathcal{C} is convexible. Let $\mathcal{D} = (d_1^-, d_1^+, \dots, d_n^-, d_n^+)$ be a convex pseudocircle arrangement and φ be some self-homeomorphism of the plane such that $\varphi[d_i^-] = c_i^-$, $\varphi[d_i^+] = c_i^+$ for all $i \in \{1, \dots, n\}$. We then fix a line arrangement $\mathcal{M} = (m_1, \dots, m_n)$ where m_i is an arbitrary line separating d_i^- and d_i^+ . The existence of \mathcal{M} is guaranteed by multiple invocations of the hyperplane separation theorem. Let $\mathcal{L}' = (\ell'_1, \dots, \ell'_n)$ be defined by setting $\ell'_i := \varphi[m_i]$. See Figure 36. Then \mathcal{L}' will be a




Figure 36: A line m_i separating d_i^- and d_i^+ (left) and their images under φ (right).

pseudoline arrangement with the property that ℓ'_i separates c_i^- and c_i^+ for all i . It therefore follows that, for each $i \neq j$, the pseudolines ℓ'_i, ℓ'_j intersect inside exactly the same 4-face of \mathcal{C} where ℓ_i, ℓ_j intersect (and ℓ'_i and ℓ'_j intersect the same sides of this 4-face as ℓ_i and ℓ_j , respectively). Using rotation systems for example, it can now easily be seen that \mathcal{L} must in fact be equivalent to \mathcal{L}' . But then \mathcal{L} is also equivalent to \mathcal{M} , so \mathcal{L} is stretchable. This concludes the proof of the theorem. \blacksquare

7 Open problems

Theorem 1.3 proves the existence of an arrangement of pseudocircles that is not convexible. A natural question is for an analogue of Theorem 1.1 for convexibility.

Question 7.1 *What is the least n such that there exists a non-convexible pseudocircle arrangement on n pseudocircles?*

By Theorem 1.1, we must have $n \geq 5$ in this last question.

Theorem 1.4 proved the existence of a pseudocircle arrangement such that no circle arrangement has the same intersection graph. We made no attempt to minimize the number of pseudocircles (our example had no less than $n = 346$ pseudocircles), but the question of what is the smallest example seems very natural.

Question 7.2 *What is the smallest n for which there exists a graph on n vertices that is the intersection graph of a convex pseudocircle arrangement, but not of a circle arrangement?*

Another natural problem is to resolve the computational complexity of determining, given a combinatorial description of a pseudocircle arrangement as input, whether there is a circle arrangement with the same intersection graph.

Acknowledgement

We thank Artem Pyatkin for bringing the folklore conjecture to our attention and for helpful pointers to the literature. We thank the anonymous referees for careful proofreading and suggestions which have improved the presentation.

References

- [1] P. K. Agarwal, B. Aronov, and M. Sharir. On the complexity of many faces in arrangements of pseudo-segments and circles. In *Discrete and computational geometry*, volume 25 of *Algorithms Combin.*, pages 1–24. Springer, Berlin, 2003.
- [2] N. Alon, H. Last, R. Pinchasi, and M. Sharir. On the complexity of arrangements of circles in the plane. *Discrete Comput. Geom.*, 26(4):465–492, 2001.
- [3] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, and G. M. Ziegler. *Oriented matroids*, volume 46 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, second edition, 1999.
- [4] B. Bultena, B. Grünbaum, and F. Ruskey. Convex drawings of intersecting families of simple closed curves. In *11th Canadian Conference on Computational Geometry*, pages 18–21, 1998.
- [5] A. A. Dobrynin and L. S. Mel’nikov. Counterexamples to Grötzsch-Sachs-Koester’s conjecture. *Discrete Math.*, 306(6):591–594, 2006.
- [6] A. A. Dobrynin and L. S. Mel’nikov. Infinite families of 4-chromatic Grötzsch-Sachs graphs. *J. Graph Theory*, 59(4):279–292, 2008.
- [7] A. A. Dobrynin and L. S. Mel’nikov. 4-chromatic edge critical Grötzsch-Sachs graphs. *Discrete Math.*, 309(8):2564–2566, 2009.
- [8] H. Edelsbrunner and E. A. Ramos. Inclusion-exclusion complexes for pseudodisk collections. *Discrete Comput. Geom.*, 17(3):287–306, 1997.

- [9] S. Felsner. *Geometric graphs and arrangements*. Advanced Lectures in Mathematics. Friedr. Vieweg & Sohn, Wiesbaden, 2004. Some chapters from combinatorial geometry.
- [10] J. E. Goodman. Proof of a conjecture of Burr, Grünbaum, and Sloane. *Discrete Math.*, 32(1):27–35, 1980.
- [11] J. E. Goodman and R. Pollack. Proof of Grünbaum’s conjecture on the stretchability of certain arrangements of pseudolines. *J. Combin. Theory Ser. A*, 29(3):385–390, 1980.
- [12] J. E. Goodman and R. Pollack. The combinatorial encoding of disjoint convex sets in the plane. *Combinatorica*, 28(1):69–81, 2008.
- [13] J. L. Gross and T. W. Tucker. *Topological graph theory*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons Inc., New York, 1987. A Wiley-Interscience Publication.
- [14] B. Jaggi, P. Mani-Levitska, B. Sturmfels, and N. White. Uniform oriented matroids without the isotopy property. *Discrete & Computational Geometry*, 4:97–100, 1989.
- [15] J. Linhart and R. Ortner. On the combinatorial structure of arrangements of oriented pseudocircles. *Electron. J. Combin.*, 11(1):Research Paper 30, 13 pp. (electronic), 2004.
- [16] J. Linhart and R. Ortner. An arrangement of pseudocircles not realizable with circles. *Beiträge Algebra Geom.*, 46(2):351–356, 2005.
- [17] J. Linhart and R. Ortner. A note on convex realizability of arrangements of pseudocircles. *Geombinatorics*, 18(2):66–71, 2008.
- [18] J. Linhart and Y. Yang. Arrangements of arcs and pseudocircles. *Beiträge Algebra Geom.*, 37(2):391–398, 1996.
- [19] C. J. H. McDiarmid and T. Müller. Integer realizations of disk and segment graphs. *J. Combin. Theory Ser. B*, 103(1):114–143, 2013.
- [20] N. E. Mnëv. Realizability of combinatorial types of convex polyhedra over fields. *J. Soviet Mathematics (in Russian)*, 28:606–609, 1985.
- [21] N. E. Mnëv. The universality theorems on the classification problem of configuration varieties and convex polytopes varieties. In *Topology and geometry—Rohlin Seminar*, volume 1346 of *Lecture Notes in Math.*, pages 527–543. Springer, Berlin, 1988.
- [22] R. Ortner. Arrangements of pseudocircles. Ph.D. Thesis, University of Salzburg, 2001.
- [23] A. V. Pyatkin. Personal communication, December 2011.
- [24] P. W. Shor. Stretchability of pseudolines is NP-hard. In *Applied geometry and discrete mathematics*, volume 4 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 531–554. Amer. Math. Soc., Providence, RI, 1991.
- [25] J. Stillwell. *Geometry of surfaces*. Universitext. Springer-Verlag, New York, 1992.
- [26] H. Whitney. Congruent Graphs and the Connectivity of Graphs. *Amer. J. Math.*, 54(1):150–168, 1932.