

Embeddability of Arrangements of Pseudocircles into the Sphere

Ronald Ortner

*Department Mathematik und Informationstechnologie, Montanuniversität Leoben,
Franz-Josef-Straße 18, 8700-Leoben, Austria*

Abstract

An arrangement of pseudocircles is a finite set of oriented closed Jordan curves each two of which cross each other in exactly two points. To describe the combinatorial structure of arrangements on closed orientable surfaces, in [6] so-called *intersection schemes* were introduced. Building up on results about the latter, we first clarify the notion of *embedding* of an arrangement. Once this is done, it is shown how the embeddability of an arrangement depends on the embeddability of its subarrangements. The main result presented is that an arrangement of pseudocircles can be embedded into the sphere if and only if all of its subarrangements of four pseudocircles are embeddable into the sphere as well.

1 Introduction

In [2], Bokowski used so-called hyperline sequences to introduce an alternative axiomatisation of oriented matroids that resulted in a new direct proof of the Folkman-Lawrence topological representation theorem for rank 3 [3] and for arbitrary rank [4]. Using the hyperline sequences approach, Linhart and Ortner [6] introduced a generalisation of uniform oriented matroids of rank 3, so-called *intersection schemes*, to describe the combinatorial properties of arrangements of pseudocircles and proved an analogue of the Folkman-Lawrence topological representation theorem. This paper continues this work by dealing with arrangements of pseudocircles where each pair of curves intersects in exactly two crossing points. Intersection schemes for these particular arrangements have a special form and are introduced as *intersection matrices*. With the aid of the latter and building up on results of [6], we first introduce the notion of *embedding of an arrangement in a closed orientable surface*. Then

Email address: ronald.ortner@unileoben.ac.at (Ronald Ortner).

we proceed showing two minor results about isomorphy and embeddability of arrangements in terms of their subarrangements. Finally, we prove that an arrangement of pseudocircles can be embedded into the sphere if and only if all of its subarrangements of four pseudocircles are embeddable into the sphere as well.

2 Preliminaries

A *pseudocircle* is an oriented closed Jordan curve on some closed orientable surface. We call a pseudocircle γ *separating*, if its complement consists of two connected components. With regard to one of these γ is oriented counterclockwise. This component is called the *interior* of γ , denoted by $\text{int}(\gamma)$.

Definition 1 *An arrangement of pseudocircles is a finite set of oriented closed Jordan curves on some closed orientable surface such that*

- (i) *no three curves meet each other at the same point,*
- (ii) *if two pseudocircles have a point in common, they cross each other in that point,*
- (iii) *each pair of curves intersects exactly two times.*

An arrangement is said to be strict if all its pseudocircles are separating.

Given an arrangement Γ of two or more pseudocircles on an orientable closed surface \mathcal{S} , we may consider the intersection points of the pseudocircles as vertices and the curves between the intersections as edges. Thus we obtain an embedding of a graph in \mathcal{S} which we call the *arrangement graph*. If this induced embedding is cellular, i.e. all its faces are homeomorphic to an open disc, we say that Γ is a *cellular* arrangement.

We may describe an arrangement of (labelled) pseudocircles $\{\gamma_1, \gamma_2, \dots, \gamma_n\}$ as follows. Consider a walk on each pseudocircle following its orientation beginning in an arbitrary vertex. Whenever we meet a vertex, we note the label i of the curve γ_i we cross provided with a sign that indicates whether γ_i comes from the left (+) or from the right (-). Thus we obtain for each pseudocircle a cyclic list of length $2(n-1)$. Ordering these lists according to the labels of the corresponding pseudocircles, one obtains an $n \times 2(n-1)$ -matrix which we call the *intersection matrix* of the arrangement. Generally, (abstract) intersection matrices may be defined as follows.

Definition 2 *Let $\mathcal{L} \subset \mathbb{N}$ be a set of n labels. An abstract intersection matrix is an $n \times 2(n-1)$ -matrix whose rows are labelled with distinct elements of \mathcal{L} in ascending order, such that the row (with label) i consists of a permutation*

of the elements $\{+k \mid k \in \mathcal{L}, k \neq i\} \cup \{-k \mid k \in \mathcal{L}, k \neq i\}$.

We call an intersection matrix A representable if there is an arrangement of pseudocircles with intersection matrix A . If A is representable with separating pseudocircles we say that A is strictly representable.

Definition 3 Let $A = (a_{ik})$ be an intersection matrix and a_{ij} , a_{ik} , $a_{i\ell}$ three distinct entries in the same row of A . Then we say that a_{ij} is between a_{ik} and $a_{i\ell}$, if either $k < j < \ell$, or $\ell < k < j$, or $j < \ell < k$.

An intersection matrix A is consistent if for all pairwise disjoint i, j, k the entry $\pm i$ in row k is placed between $+j$ and $-j$ if and only if the entry $\mp k$ in row i is placed between $+j$ and $-j$.

Linhart and Ortner [6] considered arrangements of pseudocircles of a more general nature. Condition (iii) of Definition 1 is relaxed such that an arbitrary finite number of intersection points is allowed. These generalised arrangements can be described by so-called *intersection schemes*, a generalisation of our intersection matrices. Thus, we may apply the following two main results of [6]:

The Face Algorithm. First, all intersection matrices are representable. For a given intersection matrix A there is an algorithm (called the *face algorithm*) that derives from A a unique (cellular) embedding of an arrangement with intersection matrix A in a closed orientable surface \mathcal{S}_g of minimal genus g . Moreover, there can be no cellular embedding of an arrangement with intersection matrix A in a surface $\mathcal{S}_{g'}$ of genus $g' \neq g$. Thus the combinatorial information of an arrangement is encoded in its intersection scheme, so that we call two arrangements of pseudocircles (*combinatorially isomorphic*) if they can be described by the same intersection matrix.

An arrangement Γ is (*cellularly, strictly*) embeddable into \mathcal{S} if there is a (cellular, strict) arrangement in \mathcal{S} that is isomorphic to Γ . The relation between an intersection matrix A and an arrangement described by A is analogous to that of a graph and its embedding. However, whereas a graph may be cellularly embeddable into several closed orientable surfaces (cf. [5], p. 132ff), for an arrangement of pseudocircles there is always a unique surface it can be cellularly embedded into.

The face algorithm can be used to enumerate all arrangements of pseudocircles. Thus Figure 1 shows all arrangements of three (counterclockwise oriented) pseudocircles in the plane (we call them $\alpha, \beta, \gamma, \delta$). On closed orientable surfaces \mathcal{S}_g there is an additional arrangement (called ε) resulting from an α -arrangement when turning its pseudocircles inside-out, i.e., reversing their orientation. Actually, on the sphere and generally on \mathcal{S}_g the arrangements α, β, γ and ε can be distinguished from each other only by their orientation. That

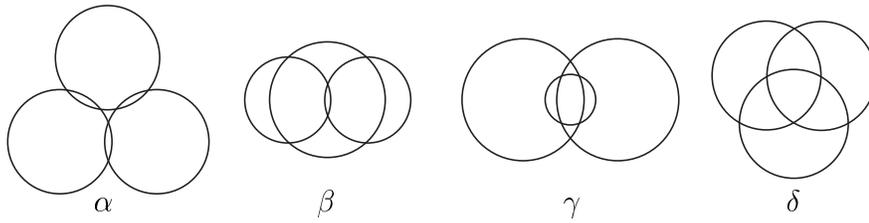


Fig. 1. Arrangements of three pseudocircles in the plane.

is, one can obtain each of these arrangements from each other by reversing the orientation of suitable pseudocircles. Unlike that, reversing any number of pseudocircles in a δ -arrangement always results in another δ -arrangement. We remark that no arrangement of three pseudocircles can be cellularly and strictly embedded into a surface of genus > 0 .

Characterisation of Strict Representability. Secondly, the following theorem establishes a one-to-one correspondence between consistent intersection matrices and strict arrangements of pseudocircles.

Theorem 4 *An arrangement is strictly embeddable into some closed orientable surface if and only if its intersection matrix is consistent.*

Intersection matrices can be considered as a generalisation of oriented matroids of rank 3 as defined via hyperline sequences by Bokowski (cf. [2], p. 576). More precisely, let A be an $n \times 2(n-1)$ -intersection matrix. Then A is a uniform oriented matroid of rank 3, if for all k : $\pm k$ occurs in position j ($1 \leq j \leq n-1$) of a row in A if and only if $\mp k$ occurs in position $j+n-1$ of the same row. Given an arrangement that can be described by a uniform oriented matroid, all its subarrangements of three pseudocircles are of type δ (cf. [1], p. 247ff).

3 Embeddability and Isomorphy via Subarrangements

In this section we present some results concerning embeddability and isomorphy of arrangements in terms of their subarrangements.

Definition 5 *Let A be an intersection matrix. Then submatrices of A are defined recursively as follows.*

- (i) *The matrix arising from A when deleting row j and all entries $\pm j$ from the other rows is a submatrix of A .*
- (ii) *If B is a submatrix of A , then any submatrix of B is also a submatrix of A .*

We will shortly say m -submatrices for $m \times 2(m-1)$ -submatrices of A . Furthermore, A_j^ denotes the submatrix obtained from A by applying rule (i). Obviously, an m -submatrix of A is the intersection matrix of an m -subarrangement*

(subarrangement of m pseudocircles) of an arrangement with intersection matrix A .

Proposition 6 *Let Γ, Γ' be two strict arrangements of $n \geq 4$ pseudocircles on an arbitrary closed orientable surface, and assume that the pseudocircles in Γ and Γ' are labelled using the same set of labels. Then Γ, Γ' are isomorphic if and only if after a suitable permutation of the labels they have the same set of labelled 4-subarrangements.*

Proof. We show that an intersection matrix with $n \geq 4$ rows is uniquely determined by the set of its 4-submatrices. This is trivial for $n = 4$. Proceeding by induction, let $A = (a_{ik}), B = (b_{ik})$ be two intersection matrices with $n + 1$ rows, such that the matrices A_i^*, B_i^* are identical for $i \in \{1, 2, \dots, n + 1\}$. Now suppose that A and B have different rows j . Since the rows are read cyclically, we may assume that $a_{j1} = b_{j1} = j + 1$ (modulo $n + 1$). Let a_{jk} be the first entry in row j of A that is different from the corresponding entry b_{jk} in B ($k > 1$). Now, because $n \geq 4$, there is an $m \in \{1, 2, \dots, n + 1\} \setminus \{j, |a_{j1}|, |a_{jk}|, |b_{jk}|\}$ (remember that $a_{j1} = b_{j1}$). By assumption $A_m^* = B_m^*$ so that especially the j -th rows of A_m^* and B_m^* are identical. But this leads to a contradiction:

- ▷ If entries $\pm m$ neither are placed between a_{j1} and a_{jk} nor between b_{j1} and b_{jk} , then A_m^* and B_m^* still differ in entries $a_{jk}^* = a_{jk} \neq b_{jk} = b_{jk}^*$.
- ▷ On the other hand, if there are entries $\pm m$ between a_{j1} and a_{jk} (or between b_{j1} and b_{jk} , respectively), then they have to appear on the same position in A and B . Otherwise the k -th entry would not be – as has been assumed – the first one where A differs from B . Hence, when deleting the entries $\pm m$ from row j , A_m^* and B_m^* still have different rows j .

□

Remark 7 *Proposition 6 is not true for 3-subarrangements. Figure 2 shows two non-isomorphic arrangements (with all curves oriented counterclockwise) all of whose 3-subarrangements are of type β .*



Fig. 2. An example illustrating Remark 7.

As a corollary to Theorem 4 one obtains the following proposition.

Proposition 8 *An arrangement Γ is strictly embeddable into some closed orientable surface if and only if all 3-subarrangements of Γ are strictly embeddable into some closed orientable surface.*

Proof. By Theorem 4, it is sufficient to prove that an intersection matrix A is consistent if and only if all 3-submatrices of A are consistent. Clearly, if there is an inconsistent submatrix of A , A itself is inconsistent. On the other hand, if A is inconsistent, then there are indices i, j, k , such that the entry $\pm j$ in row k is placed between $+i$ and $-i$, but the entry $\mp k$ in row j is not. Hence, the 3-submatrix of A consisting of the rows i, j, k is inconsistent, too. \square

Remark 9 *It is not sufficient that all 3-subarrangements of an arrangement Γ are embeddable into some surface \mathcal{S} to guarantee that Γ is embeddable into \mathcal{S} as well. Thus the arrangement Γ in Figure 3 can only be embedded into the torus (or a surface of higher genus), while all arrangements of three pseudocircles and hence all 3-subarrangements of Γ are embeddable into the sphere.*

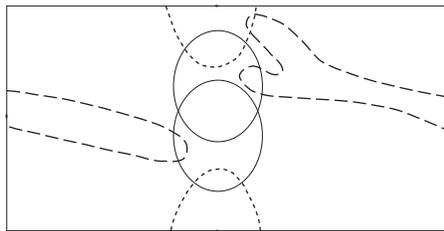


Fig. 3. An example illustrating Remark 9.

However, for 4-subarrangements the following theorem holds.

Theorem 10 *An arrangement Γ is embeddable into the sphere if and only if all of its 4-subarrangements are embeddable into the sphere.*

4 Proof of Theorem 10

First, we collect some observations about embeddings of graphs and arrangements in surfaces of genus > 0 .

Lemma 11 *Let $G = (V, E)$ be a graph that is cellularly embedded into \mathcal{S}_g with genus $g > 0$. Then the embedding contains a non-separating cycle.*

Proof. We apply repeatedly one of the following two operations to the embedding of G , until this is no longer possible:

- (i) Remove an edge that is incident with two different faces.
- (ii) Remove a vertex of degree 1 together with the single incident edge.

Obviously, neither of the two operations has an effect on the Euler characteristic of the embedding. Thus, when it is not further possible to ap-

ply (i) or (ii), for the remaining embedded graph $G' = (V', E')$ we have $\chi(G') = \chi(G) = 2 - 2g$. Since we assumed that $g > 0$, G' cannot be a tree. Furthermore, there are no separating cycles in G' , because otherwise we could apply (i). Hence, there is at least one non-separating cycle in the embedding of G' and consequently in the embedding of G . \square

Lemma 12 *Let $G = (V, E)$ be a graph (not necessarily cellularly) embedded into \mathcal{S}_g ($g > 0$) and $\mathcal{C} = (v_1, v_2, \dots, v_m)$ a non-separating cycle in G ($v_i \in V$, $v_1 = v_m$). Furthermore, let $\mathcal{P} = (v_j, v'_1, v'_2, \dots, v'_\ell, v_k)$ with $k > j$ be a path in G , such that each $v'_i \in V \setminus \{v_1, \dots, v_m\}$ and the cycle $\mathcal{C}_2 = (v_j, v_{j+1}, \dots, v_k, v'_\ell, v'_{\ell-1}, \dots, v'_1, v_j)$ is separating (cf. Figure 4).*

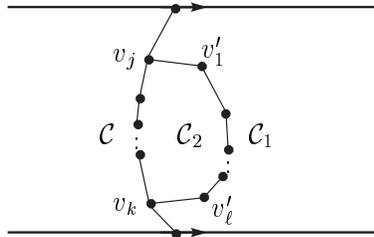


Fig. 4. Illustration of Lemma 12.

Then the cycle $\mathcal{C}_1 = (v_1, v_2, \dots, v_j, v'_1, v'_2, \dots, v'_\ell, v_k, \dots, v_m)$ is non-separating.

Proof. Assume that \mathcal{C}_1 is separating. Then $\mathcal{C}_1, \mathcal{C}_2$ are the boundary walks of two sub-surfaces F_1 and F_2 whose common border corresponds to the path \mathcal{P} . It follows that \mathcal{C} is the boundary walk of a sub-surface consisting of the union of F_1 and F_2 , which contradicts our assumption that \mathcal{C} is non-separating. \square

Lemma 13 *Let Γ be a strict arrangement of pseudocircles cellularly embedded into \mathcal{S}_g ($g > 0$). Then (possibly after reorientation of some pseudocircles) Γ contains two α -arrangements Γ_1, Γ_2 such that:*

- (i) *Each Γ_i has two non-separating boundary curves, i.e. cycles consisting of edges not contained in the interior of any pseudocircle of Γ_i .*
- (ii) *$\Gamma_1 \cup \Gamma_2$ can be cellularly embedded into some surface of genus > 0 .*
- (iii) *$\Gamma_1 \cup \Gamma_2$ consists of four or five pseudocircles.*

Proof. According to Lemma 11, the arrangement graph of $\Gamma = \{\gamma_1, \dots, \gamma_n\}$ contains a non-separating cycle \mathcal{C} . First we are going to show that we may assume that this cycle consists of edges of only three pseudocircles, which will give us the first α -arrangement. Afterwards we show how to obtain the second one.

Given an arbitrary non-separating cycle, we may assign to it a (cyclic) sequence of numbers $\in \{1, \dots, n\}$, each number i corresponding to a path con-

sisting of consecutive edges lying on the same pseudocircle γ_i . We show how to obtain from an arbitrary non-separating cycle one whose sequence consists of only three values. This is done in two steps.

(1) First we delete occurring multiple entries j in the sequence of \mathcal{C} . We are looking for pairs of vertices v', v'' of \mathcal{C} on γ_j such that on an oriented path \mathcal{P} from v' to v'' on \mathcal{C} there are neither further intersection points with γ_j nor any edges on γ_j . For each pair v', v'' we consider the cycle consisting of \mathcal{P} together with one of the two curves connecting v' and v'' on γ_j . If all these cycles were separating for all possible pairs v', v'' , then some of their interiors could be composed to give a well-defined interior of \mathcal{C} , contradicting our assumption. Thus for a suitable pair v', v'' we find a non-separating cycle \mathcal{C}' whose sequence contains only a single entry of j . Furthermore all edges of \mathcal{C}' not on γ_j were already contained in \mathcal{C} (cf. Figure 5).

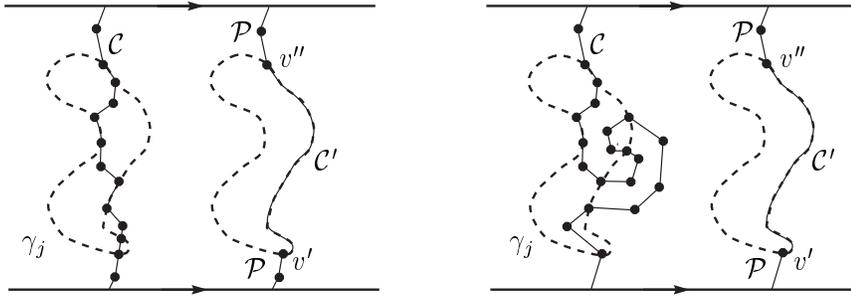


Fig. 5. Two examples for deleting multiple entries.

(2) Having removed multiple occurrences in the sequence of \mathcal{C} , we may reduce it further due to the fact that each two pseudocircles intersect. Given a cycle \mathcal{C} with sequence $\langle \dots i, j, k, \dots \rangle$ of length ≥ 4 , the basic idea is to walk on γ_i ignoring the intersection with γ_j until one arrives at an intersection with γ_k . Then continuing the way on γ_k one gets back to \mathcal{C} (see left picture of Figure 6). There are two difficulties to consider:

First, it may happen that we cannot apply Lemma 12, because the triangle Δ consisting of the detour via γ_i, γ_k together with the edges on $\mathcal{C} \cap \gamma_j$ (i.e. the cycle corresponding to \mathcal{C}_2 in Lemma 12) is not separating. However, in this case we are done, since we have found a non-separating cycle consisting only of edges on the three pseudocircles $\gamma_i, \gamma_j, \gamma_k$.

Secondly, our detour may cross some part of \mathcal{C} (see right picture of Figure 6). Assuming that the aforementioned triangle Δ is separating, there are at least three cycles consisting of edges of Δ and of \mathcal{C} , one of which is evidently separating. However, it is an easy consequence of Lemma 12 that one of the other cycles has to be non-separating. This cycle \mathcal{C}' consists only of edges of pseudocircles that also participated in \mathcal{C} . Moreover, the number of pseudocircles participating in \mathcal{C}' is obviously smaller than in \mathcal{C} . Hence, repeated application of this strategy finally yields a non-separating cycle with sequence $\langle i, j, k \rangle$.

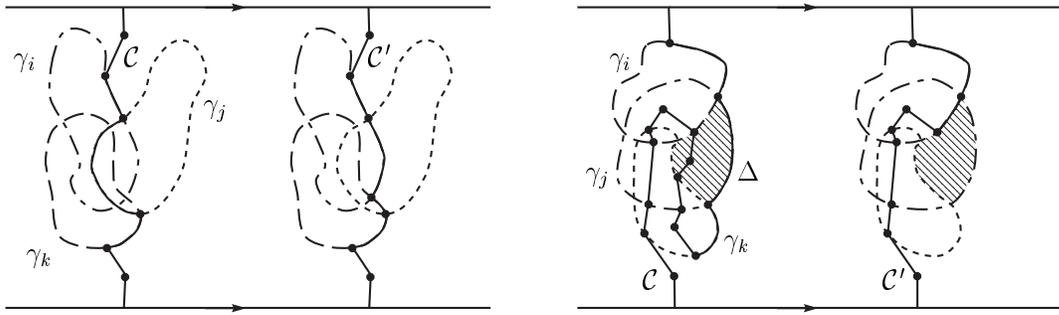


Fig. 6. Two examples for reducing the sequence.

Now, all types of arrangements of three pseudocircles (cf. Section 2) except δ can be strictly embedded into surfaces of genus > 0 so that non-separating boundary cycles arise. In case $\Gamma_1 = \{\gamma_i, \gamma_j, \gamma_k\}$ is an α -arrangement it is easy to see that this arrangement has two non-separating boundary curves $\mathcal{C}_1, \mathcal{C}_2$ (cf. Figure 7). Otherwise, we have already mentioned in Section 2 that such an α -arrangement can be obtained by a suitable reorientation of some of the three pseudocircles $\gamma_i, \gamma_j, \gamma_k$. Thus we may assume without loss of generality that Γ_1 is an α -arrangement.

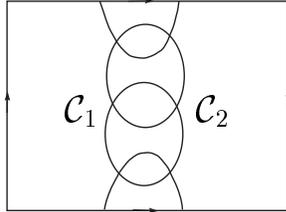


Fig. 7. An arrangement of three pseudocircles embedded into some \mathcal{S}_g ($g > 0$) with non-separating cycles.

Obviously, in order to obtain a cellular embedding of the arrangement graph in \mathcal{S}_g , the cycles \mathcal{C}_1 and \mathcal{C}_2 must be connected by a simple path whose edges are not contained in $int(\Gamma_1) := \bigcup_{\gamma \in \Gamma_1} int(\gamma)$. If there is a pseudocircle $\gamma_\ell \in \Gamma \setminus \Gamma_1$ that connects \mathcal{C}_1 and \mathcal{C}_2 with edges not contained in $int(\Gamma_1)$, we have found (again maybe after reorientation of γ_ℓ) another α -arrangement Γ_2 consisting of γ_ℓ together with two pseudocircles of Γ_1 .

Otherwise, we add the pseudocircles in $\Gamma \setminus \Gamma_1$ one by one to Γ_1 until a path as described above is established. Let γ_ℓ be the last pseudocircle added. If γ_ℓ intersects only one of the boundary curves, say \mathcal{C}_1 , it is clear that there must be some γ_m that cuts \mathcal{C}_2 so that at least one intersection point of $\gamma_\ell \cap \gamma_m$ is $\notin int(\Gamma_1)$. Then (possibly after reorientation) γ_ℓ, γ_m together with an arbitrary pseudocircle in Γ_1 form an α -arrangement Γ_2 embedded into \mathcal{S}_g ($g > 0$) with two non-separating boundary curves. The case where γ_ℓ intersects both \mathcal{C}_1 and \mathcal{C}_2 is similar. As before there is some γ_m that cuts either \mathcal{C}_1 or \mathcal{C}_2 such that γ_ℓ, γ_m together with an arbitrary pseudocircle $\in \Gamma_1$ form an α -arrangement Γ_2 .

Finally, note that Γ_1 and Γ_2 were chosen such that $\Gamma_1 \cup \Gamma_2$ can be cellularly embedded into some surface of genus > 0 . \square

Lemma 14 *Let Γ be a strict arrangement of four pseudocircles on a closed orientable surface \mathcal{S}_g of genus $g > 0$ such that*

- (i) Γ contains an α -arrangement Γ' with two non-separating boundary curves.
- (ii) Γ can be embedded into the sphere.

Then the single pseudocircle $\gamma \in \Gamma \setminus \Gamma'$ together with two pseudocircles $\in \Gamma$ forms an α -arrangement that has two non-separating boundary curves.

Proof. We have to conduct an extensive case distinction. We shall see that in each single case γ (the dashed pseudocircle in the figures) forms together with two curves $\in \Gamma'$ (the continuous ones in the figures) an α -arrangement with the asked property. The figures are simplified so that the interiors of all appearing pseudocircles are homeomorphic to an open disc. However, the argumentation applies to the general case, too. Let $\mathcal{C}_1, \mathcal{C}_2$ be the two non-separating boundary curves of Γ' .

Case 1. γ cuts \mathcal{C}_1 and \mathcal{C}_2 such that there are edges of γ not contained in $\text{int}(\Gamma')$ that connect a vertex $\in \mathcal{C}_1$ with another vertex $\in \mathcal{C}_2$.

This can only happen if Γ is cellularly embeddable into some surface of genus > 0 , contradicting assumption (ii).

Case 2. γ cuts \mathcal{C}_1 and \mathcal{C}_2 such that all edges of γ connecting a vertex $\in \mathcal{C}_1$ with another vertex $\in \mathcal{C}_2$ are contained in $\text{int}(\Gamma')$.

Note that either \mathcal{C}_i can only contain an even number of the six intersection points of γ with the pseudocircles $\in \Gamma'$ so that only the following two cases may occur.

- (A) *There are two vertices of γ on each \mathcal{C}_i :* First note that the two vertices on the same \mathcal{C}_i cannot be placed on the same pseudocircle γ_j . Otherwise it could not be avoided that γ has more than two intersection points with γ_j . Furthermore, since four vertices are distributed over three pseudocircles, there has to be a pair of vertices placed on the same pseudocircle (but on

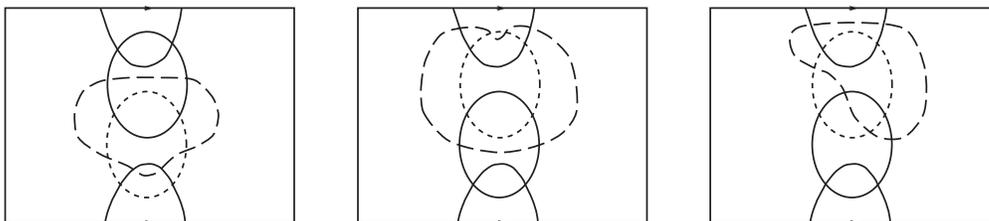


Fig. 8. Case (2.A)

different boundary curves). If all four intersection points of γ with $\mathcal{C}_1 \cup \mathcal{C}_2$ lie on two pseudocircles, then the situation is as shown in one of the two left pictures in Figure 8 (in the following we do not distinguish between symmetric cases).

Otherwise, if each pseudocircle $\in \Gamma'$ contains at least one intersection point of γ with $\mathcal{C}_1 \cup \mathcal{C}_2$, then Γ looks as in the right picture of Figure 8.

- (B) *There are two vertices of γ on \mathcal{C}_1 and four on \mathcal{C}_2 (or vice versa):* As argued in case (A), the two vertices in $\gamma \cap \mathcal{C}_1$ have to be placed on different pseudocircles γ_i, γ_j . Furthermore, two of the vertices of γ on \mathcal{C}_2 have to lie on the same pseudocircle γ_k ($k \neq i, j$). Hence, the other two intersection points of $\gamma \cap \mathcal{C}_2$ are placed on γ_i, γ_j , one on each. Thus, the situation is as shown in Figure 9.

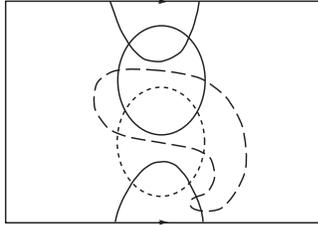


Fig. 9. Case (2.B).

Case 3. γ has intersection points with \mathcal{C}_1 , but not with \mathcal{C}_2 (or vice versa).

- (A) γ has two intersection points with \mathcal{C}_1 .
 (A.1) *Both of these vertices lie on the same pseudocircle $\gamma_i \in \Gamma'$:* The other four vertices of γ have to be placed inside γ_i so that the situation is as shown in Figure 10.

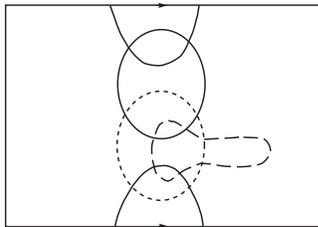


Fig. 10. Case (3.A.1).

- (A.2) *The two vertices in $\gamma \cap \mathcal{C}_1$ lie on different pseudocircles $\in \Gamma'$:* There are two possibilities dependent on how these two vertices on \mathcal{C}_1 are connected. Both are shown in Figure 11.
 (B) γ has four intersection points with \mathcal{C}_1 .
 (B.1) *These four vertices are placed on two curves $\gamma_i, \gamma_j \in \Gamma'$:* There are essentially two ways Γ may look like, both shown in Figure 12.

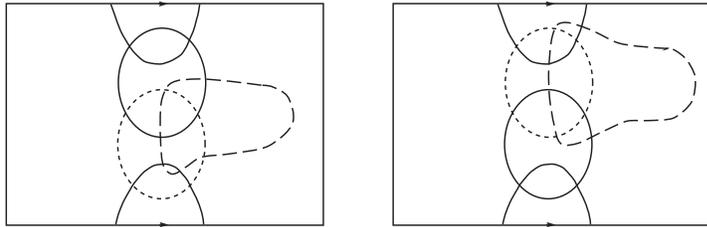


Fig. 11. Case (3.A.2).

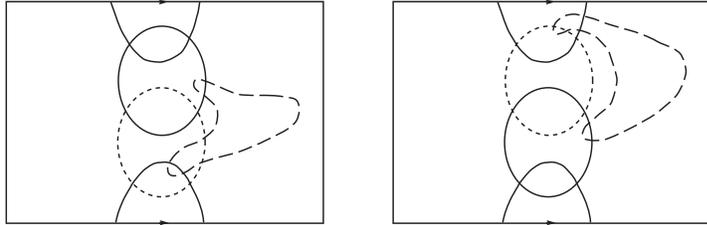


Fig. 12. Case (3.B.1).

(B.2) *The four vertices of γ on \mathcal{C}_1 are placed on all three curves of Γ' : This case allows only one type of arrangement that can be seen in Figure 13.*

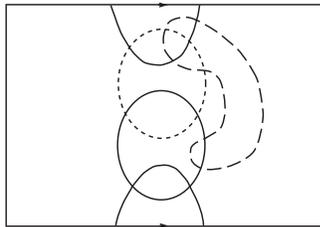


Fig. 13. Case (3.B.2).

(C) *γ has six intersection points with \mathcal{C}_1 : Figure 14 shows the only possible type of arrangement satisfying this condition.*

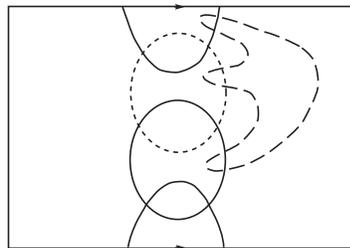


Fig. 14. Case (3.C).

Case 4. γ has no intersection points with either C_i .

In this final case, the vertices of γ on pseudocircles $\in \Gamma'$ have to be placed in pairs on three edges as shown in Figure 15.

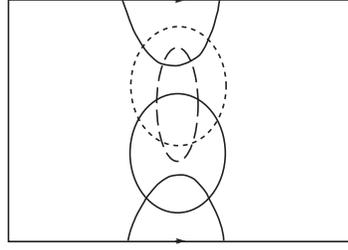


Fig. 15. Case (4).

In all cases we may remove the dotted pseudocircle in the figures from Γ obtaining an α -arrangement with two non-separating boundary curves. \square

Proof of Theorem 10. Obviously, an arrangement can only be embeddable into the sphere if all of its 4-subarrangements are embeddable into the sphere as well. To see that this condition is also sufficient, let A be an intersection matrix not representable on the sphere. We show that A has a 4-submatrix that is not representable on the sphere either. If A is inconsistent, we have already seen in Proposition 8 that there must be an inconsistent 3-submatrix and hence an inconsistent 4-submatrix of A , which is not representable on the sphere: Since pseudocircles on the sphere are always separating, in this case representability and strict representability coincide. Therefore, by Theorem 4 any inconsistent intersection matrix is not representable on the sphere.

Thus let us assume that A is consistent. By Theorem 4, A is strictly representable. We may assume that the embedding of the corresponding arrangement Γ in a closed orientable surface \mathcal{S} is cellular. By hypothesis, A is not representable on the sphere, so that \mathcal{S} is of genus > 0 . By Lemma 13, possibly after reorientation of some pseudocircles in Γ there are two distinct α -subarrangements $\Gamma_1, \Gamma_2 \subseteq \Gamma$ with non-separating boundary curves such that $\Gamma_1 \cup \Gamma_2$ consists of four or five pseudocircles and can be cellularly embedded into \mathcal{S}_g ($g > 0$). Now if $|\Gamma_1 \cup \Gamma_2| = 4$, we have obviously found a 4-subarrangement of Γ that can be cellularly embedded into \mathcal{S}_g and hence not into the sphere (cf. Section 2). Thus let us assume that $|\Gamma_1 \cup \Gamma_2| = 5$, and let $\Gamma_1 = \{\gamma_1, \gamma_2, \gamma_3\}$, $\Gamma_2 = \{\gamma'_1, \gamma'_2, \gamma'_3\}$ such that $\gamma_1 = \gamma'_1$. Applying Lemma 14, we are going to show that one can always remove one of the pseudocircles from $\Gamma_1 \cup \Gamma_2$ so that the remaining arrangement is still cellularly embeddable into a surface of genus > 0 . Thus, if either $\Gamma_1 \cup \{\gamma'_2\}$ or $\Gamma_1 \cup \{\gamma'_3\}$ is cellularly embeddable into some surface of genus > 0 , we have found what we are looking for. Otherwise, $\Gamma_1 \cup \{\gamma'_2\}$ satisfies the conditions of Lemma 14 and we may infer that there are

two pseudocircles $\gamma_i, \gamma_j \in \Gamma_1$ such that $\Gamma'_1 = \{\gamma_i, \gamma_j, \gamma'_2\}$ is an α -arrangement with two non-separating boundary curves. For the remaining $\gamma_k \in \Gamma_1$ ($k \neq i, j$) there are two possibilities:

- (i) If $\gamma_k \neq \gamma_1$, we may remove the pseudocircle γ_k from $\Gamma_1 \cup \Gamma_2$. This leaves the α -arrangement Γ_2 untouched while Γ_1 is replaced with Γ'_1 , so that $\Gamma'_1 \cup \Gamma_2$ still is cellularly embeddable into a surface of genus > 0 and hence not embeddable into the sphere:

Indeed, consider a cellular embedding of $\Gamma_1 \cup \Gamma_2$ in \mathcal{S}_g ($g > 0$) as given by Lemma 13, so that both, Γ_1 and Γ_2 have two non-separating boundary curves. Let $\mathcal{C}_1, \mathcal{C}_2$ be the non-separating boundary curves of Γ_1 . Then, as the embedding is assumed to be cellular, the pseudocircles γ'_2, γ'_3 in Γ_2 must connect \mathcal{C}_1 and \mathcal{C}_2 with edges not contained in the interior of Γ_1 (cf. also the end of the proof of Lemma 13). Let us assume that such a path starting at \mathcal{C}_1 first follows γ'_2 , then switches to γ'_3 until finally reaching \mathcal{C}_2 . Now let us consider what happens if Γ_1 is replaced with Γ'_1 by removing γ_k . The new boundary curves \mathcal{C}'_1 and \mathcal{C}'_2 of Γ'_1 will be such that γ'_3 still intersects \mathcal{C}'_2 : This is clear, if γ'_3 intersected any pseudocircle $\neq \gamma_k$ on \mathcal{C}_2 . If on the other hand, γ'_3 had only intersection points with γ_k on \mathcal{C}_2 , the other intersection points with pseudocircles in Γ_1 were in the interior of γ_k (remember that we assumed that $\Gamma_1 \cup \{\gamma'_3\}$ is not cellularly embeddable into a surface of genus > 0 , so that this corresponds to case (3.A.1) in the proof of Lemma 14). However, after removal of γ_k two of these intersection points lie on \mathcal{C}'_2 . Moreover, the segment of γ'_2 which intersects γ'_3 will be part of \mathcal{C}'_1 (this can be easily verified by considering the single cases in the proof of Lemma 14). Thus, γ'_3 connects \mathcal{C}'_1 and \mathcal{C}'_2 with edges not contained in the interior of Γ'_1 . It follows that $\Gamma'_1 \cup \{\gamma'_3\} = \Gamma'_1 \cup \Gamma_2$ can only be cellularly embedded into a surface of genus > 0 .

- (ii) If $\gamma_k = \gamma_1$, we consider the arrangements $\Gamma_2 \cup \{\gamma_2\}$ and $\Gamma_2 \cup \{\gamma_3\}$. If one of these is cellularly embeddable into a surface of genus > 0 , we are done. Otherwise, we find analogously as described above two pseudocircles $\gamma'_\ell, \gamma'_m \in \Gamma_2$ such that $\Gamma'_2 = \{\gamma'_\ell, \gamma'_m, \gamma_2\}$ is an α -arrangement whose two boundary curves are non-separating. If the remaining pseudocircle $\in \Gamma_2$ is not γ'_1 , the situation is as in case (i) with Γ_1, Γ_2 interchanged. Otherwise, a similar argument as given in (i) shows that removing $\gamma_1 = \gamma'_1$ from $\Gamma_1 \cup \Gamma_2$ (i.e. Γ_1, Γ_2 are replaced with $\Gamma'_1 = \{\gamma_2, \gamma_3, \gamma'_2\}$, $\Gamma'_2 = \{\gamma'_2, \gamma'_3, \gamma_2\}$) leaves an arrangement of four pseudocircles that still is cellularly embeddable into some surface of genus > 0 . \square

5 Final Remarks

Proposition 8 can be extended to the case of connected arrangements with relaxed condition (iii) (cf. the paragraph after Definition 3). Unlike that, The-

orem 10 cannot be transformed into a valid version for these generalised arrangements. Quite to the contrary, for each $n \geq 5$ one can give a cellularly embedded generalised arrangement Γ of n pseudocircles on the torus, such that each $(n - 1)$ -subarrangement of Γ is embeddable into the sphere (see Figure 16).

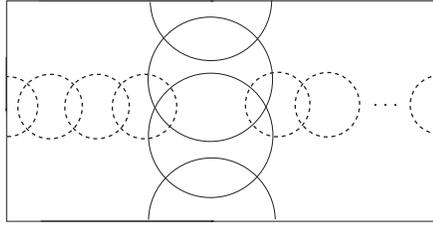


Fig. 16. A cellular generalised arrangement on the torus with all subarrangements embeddable into the sphere.

It is an interesting question whether it is possible to obtain similar results concerning embeddability into surfaces of genus > 0 . One could e.g. conjecture that embeddability into a surface \mathcal{S}_g of genus g holds if and only if all $(4 + g)$ -subarrangements are embeddable into \mathcal{S}_g . However, at the moment we neither have any evidence for nor against this conjecture, and it may be that a generalisation of Theorem 10 looks totally different.

An enumeration of all 72 arrangements of four pseudocircles that can be embedded into the sphere can be found in [7]. By the way, each of these arrangements can be realised with proper circles (cf. Appendix B of [7]).

References

- [1] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, and G. Ziegler. *Oriented Matroids*, volume 46 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, second edition, 1999.
- [2] J. Bokowski, Oriented matroids, in: P. M. Gruber and J. M. Wills, *Handbook of Convex Geometry*, pp. 555–602, Elsevier Science Publishers B.V., Amsterdam, 1993.
- [3] J. Bokowski, S. Mock, and I. Streinu, On the Folkman-Lawrence Topological Representation Theorem for Oriented Matroids of Rank 3, *Europ. J. Combin.* **22** (2001), 601–615, doi:10.1006/eujc.2000.0482.
- [4] J. Bokowski, S. King, S. Mock, and I. Streinu, The Topological Representation of Oriented Matroids, *Discrete Comput. Geom.* **33**(4) (2005), 645–668, doi:10.1007/s00454-005-1164-4.
- [5] J. L. Gross, T. W. Tucker, *Topological Graph Theory*, Wiley, New York, 1987.

- [6] J. Linhart, R. Ortner, On the combinatorial structure of arrangements of oriented pseudocircles, *Electron. J. Combin.* **11** (2004), Research Paper 30, 13 pp. (electronic).
- [7] R. Ortner, *Arrangements of Pseudocircles*, PhD Thesis, University of Salzburg, 2001.