Beiträge zur Algebra und Geometrie
Contributions to Algebra and Geometry
Volume 45 (2004), No. 1, 133-153.

Unbounded Regions in an Arrangement of Lines in the Plane

Alan West

School of Mathematics, University of Leeds, Leeds LS2 9JT, England

Abstract. We take a set Ω of n points and an arrangement Σ of m lines in the plane which avoid these points but separate any two of them. We suppose these satisfy the following unboundedness property: for each point $x \in \Omega$ there is a homotopy from Σ to Σ' avoiding Ω so that x is in an unbounded component of the complement of Σ'. It is proved that then $n \leq 2m$. This result is required to partially solve a problem in differential geometry which is described briefly.

1. Introduction

The object of this paper is to prove a result about arrangements of lines in the plane that is required to deal with a problem in differential geometry. We first give a very brief discussion of the differential geometry background.

We prove the result by a rather roundabout route using ideas about ordering. On the way we describe a number of structures and prove results about them that are interesting in themselves.

We concentrate on dimension 2, that is, to arrangements of lines in the plane. However, the differential geometry background does not have this restriction and the required result can be stated in terms of abstract arrangements of hyperplanes without such a restriction. Indeed it seems very likely that it is true in any dimension. Therefore we state this as a conjecture and discuss whether it is possible to generalise the methods we use in the 2-dimensional case to prove the general case.

The author would like to thank Ilda da Silva for many useful discussions and for introducing him to the work done on oriented matroids and, in particular, for showing that the

axioms given in this paper are equivalent to axioms used by workers in that field. He would also like to thank Kevin Buzzard for suggesting the method used to prove Theorem 5.5. The resulting proof is much shorter than the original.

2. Differential Geometry

Let $f : M \to \mathbb{R}^{m+d}$ be a smooth immersion of a compact connected m-dimensional manifold with flat, trivial, normal bundle $N(f) \subset M \times \mathbb{R}^{m+d}$. Let $\eta : N(f) \to \mathbb{R}^{m+d}$ be the end-point map given by $\eta(p,x) = f(p) + x$. The set of critical points of η is $\Sigma(f) \subset N(f)$. Thus $(p,x) \in \Sigma(f)$ means that $\eta(p,x) = f(p) + x$ is a focal point of the immersed manifold (also called a centre of curvature).

For any $p \in M$, $N_p(f) = \{x \in \mathbb{R}^{m+d} : (p,x) \in N(f)\}$ is a d-dimensional subspace of $\mathbb{R}^{m+d}$ and $\Sigma_p(f) = \Sigma(f) \cap N_p(f)$ is an algebraic variety in $N_p(f)$. In our case the normal bundle is flat and trivial. In other words, the normal holonomy group is trivial. This implies two things.

First, the normal bundle is isomorphic to $M \times \mathbb{R}^d$ and we can define the projection $\Phi : N(f) \cong M \times \mathbb{R}^d \to \mathbb{R}^d$ using parallel transport. This means that if $x \in \mathbb{R}^d$ then $\Phi^{-1}(x) \subset N(f)$ is a cross-section representing a parallel normal field (see [4]).

Secondly, for any $p \in M$ the variety $\Sigma_p(f)$ is the union of a set of m hyperplanes arranged in the d-dimensional vector space $N_p(f)$ (allowing hyperplanes to coincide and including the hyperplane at infinity).

Thus for each $p \in M$ there is a corresponding arrangement Σ_p of m hyperplanes in $\mathbb{R}^d$ such that $\Phi(\Sigma_p(f)) = \Sigma_p$. The sets Σ_p clearly vary smoothly with $p \in M$.

Now a parallel immersion to $f : M \to \mathbb{R}^{m+d}$ is given by a parallel normal field $\nu : M \to N(f)$ such that $\eta \circ \nu$ is also an immersion. We think of $\eta \circ \nu(M)$ as being obtained by pushing out each point p of M to the end of the normal $\nu(p)$. The condition that $\eta \circ \nu$ is an immersion is that the cross-section $\nu(M) \subset N(f)$ does not intersect $\Sigma(f)$.

Thus, in our case, the immersions parallel to f correspond precisely to the set of vectors $x \in \mathbb{R}^d$ for which $\Phi^{-1}(x) \cap \Sigma(f) = \emptyset$. That is, to the set $\Omega(f) = \mathbb{R}^d \setminus \bigcup\{\Sigma_p : p \in M\}$. Notice that $0 \in \Omega$. We call $\Omega(f)$ the *push-out space* for f and are interested in its topology and, in particular, how many connected components it can have (see [2] and [3]).

Let $\gamma(m,d) = 1 + \binom{m}{1} + \binom{m}{2} + \cdots + \binom{m}{d} - \binom{m-1}{d}$ be the number of unbounded connected components in $\mathbb{R}^d \setminus \Sigma$ where Σ is any arrangement of m hyperplanes in general position in $\mathbb{R}^d$.

Conjecture 2.1. *Given an embedded manifold $f : M \to \mathbb{R}^{m+d}$ with trivial normal holonomy group the number of connected components of the push-out space Ω is at most equal to $\gamma(m.d)$.*

The reason this seems likely is because any $x \in \Omega(f)$ must have an "unboundedness property" which we shall now briefly describe.

For any $x \in \Omega(f)$ let us write f_x for the corresponding parallel immersion described above. That is, $f_x = \eta \circ \nu$ where ν is given by $\Phi \circ \nu(p) = x$ for all $p \in M$. The normal bundle $N(f_x)$ is again flat and trivial so $\Phi_x : N(f_x) \to \mathbb{R}^d$ is defined in analogy to Φ. The map $\Phi_x \circ \Phi^{-1} : \mathbb{R}^d \to \mathbb{R}^d$ is the simple translation $u \mapsto u - x$ and $\Phi(\Sigma_p(f)) = \Phi_x(\Sigma_p(f_x))$. Hence $\Omega(f_x)$ is just the translation of $\Omega(f)$.

Since M is compact any non-degenerate height function for the immersion f_x is bounded and attains its bounds. So there must be a height function L_z for f_x given by some $z \in \mathbb{S}^{d-1}$ that attains its maximum at some $q \in M$. This is equivalent to saying that $(q, z) \in N(f_x)$ and if $t > 0$ then $(q, tz) \notin \Sigma(f_x)$. Thus, from the relation between Φ and Φ_x and between $\Sigma_p(f)$ and $\Sigma_p(f_x)$ above we deduce that x lies in an unbounded component of $\mathbb{R}^d \setminus \Sigma_q$.

Thus for $x \in \Omega$ there is some homotopy $\Sigma_{p(t)}$ of arrangements of hyperplanes with $p(0) = p$ and $p(1) = q$ for which $\Omega \cap \Sigma_{p(t)} = \emptyset$ and x lies in an unbounded component of $\mathbb{R}^d \setminus \Sigma_q$. This is the "unboundedness property".

The ideas of a homotopy and of unboundedness arise here from a geometrical approach but we need to formulate this in a more abstract way. For one thing Σ_p may contain hyperplanes at infinity and hyperplanes that coincide and these are easier to make precise in a more algebraic context. Also we are not really interested in Σ_p but only the connected components of $\mathbb{R}^d \setminus \Sigma_p$ and we are not interested in the manifold M and the way the homotopy of $\Sigma_p(t)$ arises from parallel transport. We only need to keep track of the components of $\mathbb{R}^d \setminus \Sigma_p(t)$ during the homotopy.

Thus we are given an open set $\Omega \subset \mathbb{R}^d$ with $0 \in \Omega$ and we choose a finite set $S \subset \Omega$ such there is exactly one point of S in each connected component of Ω. We may choose S so that $0 \in S$ and, since Ω is open, no point in S lies on the hyperplane at infinity. We are also given a collection of m hyperplanes Σ with $\Sigma \cap S = \emptyset$ with the "unboundedness property" that for each $x \in S$ there is some collection of hyperplanes Σ' which is homotopic to Σ in $\mathbb{R}^d \setminus S$ for which x lies in an unbounded connected component of $\mathbb{R}^d \setminus \Sigma'$.

Then Conjecture 2.1 will follow if we prove the following.

Conjecture 2.2. *Let Σ be any arrangement of m hyperplanes in $\mathbb{R}^d$ and let $S \subset \mathbb{R}^d$ be a finite set such that no connected component of $\mathbb{R}^d \setminus \Sigma$ contains two points of S. Suppose that S has the "unboundedness property" that for each $x \in S$ there is an arrangement of m hyperplanes, Σ', which is homotopic to Σ in $\mathbb{R}^d \setminus S$, for which x lies in an unbounded component of $\mathbb{R}^d \setminus \Sigma'$. Then the number of points in S cannot exceed $\gamma(m, d)$.*

A hyperplane on $\mathbb{R}^d$ which does not go through the origin 0 is completely determined by an element $\lambda \in \mathbb{R}^d$. The hyperplane at infinity corresponds to $\lambda = 0$. A point $v \in \mathbb{R}^d$ lies on the positive or negative side of the hyperplane according to whether $< \lambda, v > +1$ is positive or negative. Thus, given a set of m hyperplanes $E = \{\lambda_1, \lambda_2, \ldots, \lambda_m\}$ a connected component of the complement of these hyperplanes is completely determined by a map $\omega : E \to \{+1, -1\}$. Not every such map represents a connected component but given an arrangement of hyperplanes Σ we have an associated set of such maps W.

We can still simplify the problem because the set of arrangements of m hyperplanes which do not intersect S is open in the set of all such arrangements. This means that we can always assume that that the arrangements of hyperplanes, Σ, Σ' are in general position and do not include the hyperplane at infinity. In a homotopy Σ_t between them we can assume that Σ_t never contains the line at infinity. So we can assume that the sets of maps, W, is obtained from an arrangement of hyperplanes Σ in general position.

We can abstract axioms for a set of such maps, W, that must be satisfied when the maps really do represent the connected components of the complement of such an arrangement of hyperplanes, Σ, and these probably contain sufficient information to prove the result we require.

Given such a W, the axioms are not enough to reconstruct an arrangement of hyperplanes which will give rise to the given maps. So we would definitely be throwing away information in using the axioms only.

We can also assume that the homotopy Σ_t with $\Sigma_0 = \Sigma$ and $\Sigma_1 = \Sigma'$ is such that for some $0 = t_0 < t_1 < t_2 \cdots < t_{r+1} = 1$ all the arrangements Σ_t for $t \in (t_i, t_{i+1})$ are in general position and have the same set of maps, W_i, when $0 \leq i \leq r$, and that the change from W_i to W_{i+1} is either the result of one hyperplane crossing the vertex formed by the intersection of d other hyperplanes or the result of such a vertex going off to infinity.

Further, no hyperplane of Σ_t crosses a point in S so the points of S, which determine a subset of $\Omega \subset W$, always lie on the same side of each hyperplane in Σ_t. In other words $\Omega \subset W_i \cap W_{i+1}$.

Such a change in the arrangement is called an "elementary modificaton" in the set of maps W_i.

A homotopy of an abstract arrangement W modulo S is then a finite sequence of such elementary modifications.

Finally, it is easy to see that a map $\omega \in W$ represents an *unbounded* connected component of $\mathbb{R}^d \setminus \Sigma$ if and only if the map $-\omega$ also represents a connected component of $\mathbb{R}^d \setminus \Sigma$, that is, $-\omega \in W$ also.

In this way the original problem from differential geometry depends on a problem about the maps W associated with arrangements of hyperplanes in general position and elementary modifications of them.

We are going to prove Conjecture 2.2 in the particular case when $d = 2$. We shall not use an axiomatic approach but will point out where the axioms come from.

Let us give a sketch of the proof here. The basic tool is to consider the way the lines of E split the set S. This is obviously invariant under the homotopies we consider. Using this we pick out certain subsets that we call $\mathcal{C}$-sets that have a natural cyclic ordering on them. If C is a $\mathcal{C}$-set then any line in E either does not split C or splits it into two connected arcs. $E(C) \subset E$ is the set of lines that split the edges of C. Since each line in $E(C)$ splits precisely two edges of C, it is easy to see that the number of elements in C is at most equal to twice the number of lines in $E(C)$. We pick out a collection of disjoint maximal $\mathcal{C}$-sets and show that at the same time the corresponding $E(C)$ are disjoint.

We also pick out subsets of S which contain no $\mathcal{C}$-set. These we call $\mathcal{T}$-sets. If T is a $\mathcal{T}$-set we again define $E(T) \subset E$ as the set of lines that split T. Such a set can be given the natural structure of a tree such that any line in $E(T)$ splits only one edge. So again the number of elements of T is at most equal to twice the number of lines in $E(T)$. Again we can choose a collection of disjoint $\mathcal{T}$-sets such that the corresponding $E(T)$ are also disjoint and so that, in fact, the $\mathcal{T}$-sets are disjoint from the $\mathcal{C}$-sets and the corresponding $E(T)$ and $E(C)$ are also disjoint.

Since every element of S lies either in a $\mathcal{C}$-set or in a $\mathcal{T}$-set the result follows.

3. Arrangements of lines in the plane

Let E be a set of m lines arranged in the plane and let their union be Σ. To each connected component of $\mathbb{R}^2 \setminus \Sigma$ there corresponds a map $\omega : E \to \{-1, +1\}$ that will be called the corresponding *region* of the arrangement. The set of regions is W.

Definition 3.1. *Let E and E' be sets of lines in the plane and with corresponding sets of regions W and W'. Let $S \subset W \cap W'$. If there is a homotopy E_t from E to E' in $\mathbb{R}^2 \setminus X$ with corresponding regions W_t for which $S \subset W_t$ we say that W is homotopic to W' modulo S and write $W \sim_S W'$.*

The set of regions W contains all the information we require about E. Indeed we could give axioms for W and define ideas like "homotopy" entirely in terms of regions. However in our case it is easier to use the geometry of the set of lines E to justify intuitive statements that require a rather tedious derivation from these axioms alone.

Also we may assume in general that the lines E are in general position. In this case we can characterise the unbounded regions as those $\omega \in W$ for which $-\omega \in W$ also.

This will not be true throughout a homotopy but we can think of a homotopy as a modification of W to W' in which some regions which are bounded in W could become unbounded in W'.

Then Conjectures 2.1 and 2.2 in the case when $d = 2$ will follow from the following theorem.

Theorem 3.2. *Let W be the regions for a set of lines E in dimension 2 and let $S \subset W$ satisfy the following "unboundedness property": for any $\omega \in S$ there exists W' such that $W \sim_S W'$ and $-\omega, \omega \in W'$. Then $\sharp S \leq 2 \times \sharp E$.*

This theorem is obviously true if $\sharp E \leq 2$.

We shall assume throughout this paper that S has this "unboundedness property". Notice that if S has the "unboundedness property" then so does any subset of S and so also does the "projection" $P_\chi(S)$ of S in the "projection" $P_\chi(S)$ of W given by simply restricting all maps $\omega \in W$ to a subset $\chi \subseteq E$. This idea of projection is so useful that we incorporate it in a definition for reference.

Definition 3.3. *Let W be the set of regions associated with the lines E in general position in dimension 2. Let $\chi \subset E$ then if $\bar{W}$ is the set of regions associated with χ we define the projection $P_\chi : W \to \bar{W}$ by letting $P_\chi \omega \in \bar{W}$ be the restriction of ω to χ.*

If $\lambda \in E$ we also talk about the restriction of W to the line λ. This is an arrangement of hyperplanes in dimension 1 where the "hyperplanes" are the points of intersection of λ with the lines in $E \setminus \{\lambda\}$ and the regions represent the intervals between them. These intervals are linearly ordered and correspond to the regions of $u \in W$ with the property that $u' \in W$ also where $v(\lambda) = -u(\lambda)$ and $v(i) = u(i), \forall i \in E \setminus \{\lambda\}$.

In the abstract setting in higher dimensions the ideas of projection and restriction are the main tools. The fact that the regions of the restriction to a line are linearly ordered is the basis of the betweenness axiom. In our geometrical setting this is obvious.

The restriction to a line is mainly used in the application of the following lemma which gives another characterisation of unboundedness.

Lemma 3.4. *If $x \in W$ is unbounded then there exists a line $\lambda \in E$ such that x restricts to an unbounded interval on λ.*

Proof. The conditon for there to be an interval representing the restriction of x to λ is that an interval on λ is part of the boundary of the connected component of $\mathbb{R}^2 \setminus \Sigma$ that x represents. When expressed geometrically in this way the result is obvious. □

This result is not obvious in an axiomatic approach but needs a fairly long proof. This point will be discussed later.

Definition 3.5. *If A and B are two subsets of W we shall write $(A \mid B)$ as an abbreviation for the statement "there is a line in E that separates A from B". That is, for some $\lambda \in E$ we have $\forall a \in A, b \in B$, $a(\lambda) = -b(\lambda)$. A similar meaning is given to $(a, b \mid c, d)$ where $a, b, c, d \in W$.*

Similarly, if $\lambda \in E$ we shall say that λ splits *a subset $C \subset W$ if C can be written as the disjoint union of two proper subsets which are separated by λ.*

Note that if $(A \mid B)$ with respect to W and $A \cap B \subset S$ then $W \sim_S W'$ implies that $(A \mid B)$ with respect to W' also. Thus it is the separation properties of S which are important to us.

For instance an elementary but important result is contained in the following lemma which forms the basis for the "simplex axiom" satisfied by an abstract arrangement.

Lemma 3.6. *Let $\sigma \in E$ consist of just three lines (σ is a simplex). For any region $x \in W$ there is a line $\ell \in \sigma$ which does not separate x from the vertex defined by the other two lines in σ.*

An immediate consequence of this will be one of the the basic tools in using the unboundedness property.

Lemma 3.7. *Let $\sigma \in E$ consist of just three lines (σ is a simplex). Suppose that $x \in W$ is a region such that no line $\ell \in \sigma$ separates x from the vertex defined by the other two lines in σ. Then $-x \notin W$.*

This says that a simplex formed by three lines from the set E (which are in general position) has an inside and any region inside this simplex is bounded. We do not give a proof as in our situation it is obvious from the geometry. A vertex is the intersection of two lines so we can identify a vertex with a subset $\{\lambda, \mu\} \subset W$.

The lemma implies that four-element subsets in S must have certain separation properties as given in the next lemma.

Lemma 3.8. *Let $S \subset W$ have the unboundedness property and let $\{a, b, c, d\} \subseteq S$. Then at least one of the statements $(a, b \mid c, d)$, $(a, c \mid b, d)$ and $(a, d \mid c, b)$ must be false.*

Proof. It will be sufficient to prove the following: Let $\{a,b,c,d\} \subseteq S \subset W$ satisfy $(a,b \mid c,d)$, $(a,c \mid b,d)$ and $(a,d \mid c,b)$. Then there is one element $x \in \{a,b,c,d\}$ such $-x \notin W'$ for any W' with $W \sim_S W'$.

So we assume that there exists a line λ that separates a,b from c,d, a line ν that separates a,c from b,d and a line μ that separates a,d from c,b.

Then the line ν separates the vertex $\{\lambda,\mu\}$ either from a,c or from b,d. Let us suppose it separates this vertex from a,c.

Similarly, the line μ separates the vertex $\{\lambda,\nu\}$ either from a,d or from b,c. Let us suppose it separates this vertex from b,c.

We now apply Lemma 3.6 to the simplex σ and the region d we see that it must be the line λ that does not separate d from the vertex $\{\nu,\mu\}$. Then d satisfies the requirement for x in Lemma 3.7 and so $-x \notin W$. Further, since under any homotopy modulo S the same lines separate the same points in S in the same way, this implies that $-x \notin W'$ for any $W' \sim_S W$. □

Note that the result of Theorem 3.2 follows immediately from this lemma if $\sharp E = 3$.

4. $\mathfrak{T}$-sets

In this section we shall discuss $\mathfrak{T}$-sets. These cannot be very big and we do not need the unboundedness property to prove this; a $\mathfrak{T}$-set is not necessarily a subset of S.

Definition 4.1. *A subset $T \subseteq W$ is called a $\mathfrak{T}$-set if T does not contain four points a,b,c,d such that $(a,b \mid c,d)$ and $(b,c \mid a,d)$.*

Let us start by proving the basic lemma about $\mathfrak{T}$-sets.

Lemma 4.2. *Let $E' \subseteq E$ split T into the subsets $T_1, T_2, \ldots, T_k$ and let $\lambda \in E \setminus E'$ split T_i. Then λ does not split T_j if $i \neq j$.*

Proof. Suppose $\lambda \in E \setminus E'$ splits T_i into T_i^0 and T_i^1 and also splits T_j into T_j^0 and T_j^1. There must be some $\mu \in E'$ that separates T_i from T_j. So we can take $a \in T_i^0$, $b \in T_i^1$, $c \in T_j^0$ and $d \in T_j^1$ such that λ separates b,c from a,d then μ separates a,b from c,d and this contradicts the hypothesis. □

Proposition 4.3. *Let W be the regions for an arrangement of a set E of lines in dimension 2 and let $T \subset W$ be a $\mathfrak{T}$-set. Then $\sharp T \leq \sharp E(T) + 1$ where $E(T)$ is the set of lines in E that split T.*

Proof. We use Lemma 4.2 and the following inductive hypothesis: if $E' \subset E$ splits T into n non-empty subsets then $n \leq \sharp E' + 1$. This is obviously true if E' contains only one line. But then Lemma 4.2 says that if we take another line $\lambda \in E \setminus E'$ this line splits at most one of the sets T_i. In other words, when $\sharp E'$ is increased by one the number of subsets of T is increased by at most one. This proves the inductive step. So the hypothesis is true for all subsets E' of E.

If we take $E' = E(T)$ we see that $E(T)$ splits T into those subsets which consist of a single element of T. The hypothesis we have just proved then says that $\sharp T \leq \sharp E(T) + 1$. □

5. $\mathcal{C}$-sets

The idea of a cyclic ordering on a subset of W that is compatible with the way E splits W is crucial to our proof. Note that a cyclic ordering is inherited by subsets and that deleting an edge of a cyclic ordering gives a linear ordering.

Definition 5.1. *Let $C \subset W$ have a cyclic ordering then a proper subset $A \subset C$ is an* arc *of C if it inherits a linear ordering from C that has the property that an edge in A is an edge in C.*

We shall often denote the cyclic ordering by writing $C = \langle s_1, s_2, \ldots, s_p \rangle$. An edge will be denoted by $\langle a, b \rangle$.

Definition 5.2. *We define a subset $C \subseteq W$ with $\sharp C \geq 4$ to be a $\mathcal{C}$-set if it has a cyclic ordering that satisfies the conditions:*

(1) *any line in E either does not split C or splits C into two arcs or, equivalently; any line in E splits at most two edges of C;*
(2) *any two edges of C with no common end-points are separated by some line in E.*

The reason that we require a $\mathcal{C}$-set to have at least three elements is because any set with fewer elements is clearly a $\mathcal{T}$-set. If $D \subset C$ then D clearly has an inherited cyclic ordering but D may not be a $\mathcal{C}$-set.

Let us consider the simple case when $\sharp E = 3$ and $\sharp S = 4$.

Lemma 5.3. *Let $\{a, b, c, d\} \subset W$. Then $\langle a, b, c, d \rangle$ is a $\mathcal{C}$-set if and only if $(a, b \mid c, d)$ and $(b, c \mid a, d)$ but not $(a, c \mid b, d)$.*

Proof. The proof is just a matter of checking with the definition of a $\mathcal{C}$-set. □

Lemma 5.4. *Let $C \subseteq W$ be any $\mathcal{C}$-set and let $a, b \in C$. Then one can find $c, d \in C$ such that a, b, c, d in some order, form a $\mathcal{C}$-set.*

Proof. There is some line $\lambda \in E$ that separates a from b. The line λ splits precisely two edges of C . Let them be $\langle u, u' \rangle$ and $\langle v, v' \rangle$. We may suppose λ separates u, a, v from u', b, v'. Since C is a $\mathcal{C}$-set there is some $\mu \in E$ that separates u, u' from v, v'. If μ separates a from b we may choose the notation so that μ separates u, u', a from v, v', b. We then take $c = u'$ and $d = v$ and see that by Lemmas 3.8 and 5.3 $\langle a, d, b, c \rangle$ is a $\mathcal{C}$-set.

On the other hand if μ does not separate a from b we can choose the notation so that μ separates a, b, u, u' from v, v' and take $c = v'$ and $d = v$. Again, Lemmas 3.8 and 5.3 show that $\langle a, b, c, d \rangle$ is a $\mathcal{C}$-set. □

The point of this lemma is that it shows that any $T \subseteq W$ that does not contain a $\mathcal{C}$-set must be a $\mathcal{T}$-set.

The next proposition shows that the cyclic ordering on a $\mathcal{C}$-set is unique because it is really determined solely by the way the lines of E split up the set C.

Theorem 5.5. *Let $C \subseteq W$ be a $\mathcal{C}$-set. Then the compatible cyclic ordering is unique.*

Proof. Suppose that C is a $\mathcal{C}$-set. Take any $x \in W$ and for each $\lambda \in E$ define $A(\lambda) \subset C$ to be all those elements of C that are separated from x by λ.

We now define a collection of subsets of C that we shall call (temporarily) $\mathcal{A}$-sets by the following rules.

(a) $A(\lambda)$ is an $\mathcal{A}$-set for all $\lambda \in E$,

(b) if A is an $\mathcal{A}$-set then $C \setminus A$ is an $\mathcal{A}$-set,

(c) if A and B are $\mathcal{A}$-sets and $A \cap B \neq \emptyset$ then $A \cup B$ is an $\mathcal{A}$-set.

The collection of $\mathcal{A}$-sets is defined without using the cyclic ordering on C. However we claim that the edges in this cylic ordering are precisely those $\mathcal{A}$-sets that contain just two elements.

First observe that all proper $\mathcal{A}$-sets are arcs of C. Then take any edge $\langle u, v\rangle$ of C. We must show that this is an $\mathcal{A}$-set. Let A be the largest $\mathcal{A}$-set that contains neither u nor v. Since by Lemma 5.4 C must contain at least four elements we may suppose that $\langle u', u, v, v'\rangle$ is an arc of C. If $u', v' \in A$ then $\langle u, v\rangle = C \setminus A$ is an $\mathcal{A}$-set. So suppose either u' or v' is not in A then we can find an edge $\langle a, b\rangle$ of C so that $a \in A$ and $b \in C \setminus A$ with $b \neq u$, $b \neq v$.

Since C is a $\mathcal{C}$-set there exists some line $\lambda \in E$ that separates the edges $\langle a, b\rangle$ and $\langle u, v\rangle$. Suppose $\langle u, v\rangle$ lies in $A(\lambda)$ and $\langle a, b\rangle$ lies in $C \setminus A(\lambda)$. In this case we define $A' = A \cup (C \setminus A(\lambda))$ and observe that since $b \in A \cap (C \setminus A(\lambda))$ this is an $\mathcal{A}$-set. Yet neither u nor v are in A', $A \subset A'$ and $b \in A' \setminus A$. So A is not maximal.

On the other hand if $\langle u, v\rangle$ lies in $C \setminus A(\lambda)$ and $\langle a, b\rangle$ lies in $A(\lambda)$ we define $A' = A \cup A(\lambda)$ and observe that since $b \in A \cap A(\lambda)$ this is an $\mathcal{A}$-set. Yet again neither u nor v are in A', $A \subset A'$ and $b \in A' \setminus A$. So A is not maximal.

We deduce that the edge $\langle u, v\rangle = C \setminus A$ and this edge is an $\mathcal{A}$-set. □

This theorem is useful when we are considering whether a $\mathcal{C}$-set can be extended to a larger $\mathcal{C}$-set. We know that the ordering on the larger $\mathcal{C}$-set must always be such that the ordering inherited by the smaller $\mathcal{C}$-set is the one given originally.

It also shows us that the structure of a $\mathcal{C}$-set is completely determined by the way it is separated by the lines in E that separate the edges of C. This leads us to the following definition.

Definition 5.6. *Let $C \subset W$ be a $\mathcal{C}$-set. We define $E(C)$ to be the set of lines in E that split C into two arcs each of which contains an edge of C.*

Lemma 5.7. *Let $C \subset W$ be any $\mathcal{C}$-set. Then $\sharp C \leq 2 \times \sharp E(C)$.*

Proof. Clearly since any line in $E(C)$ splits C into two arcs it splits exactly two edges of C. Also any edge of C is split by a line in $E(C)$. So the number of edges in C is less than twice the number of lines in $E(C)$. Since the number of edges in C is equal to $\sharp C$ the result follows. □

The basic idea in defining a $\mathcal{C}$-set is that it has an inside and any region inside a $\mathcal{C}$-set is bounded. Hence, if the $\mathcal{C}$-set belongs to S and we take another region in S, this region must lie outside the $\mathcal{C}$-set. Of course this is an intuitive idea but it can be expressed by the following central proposition which depends crucially on the assumption that all elements in S have the unboundedness property.

Proposition 5.8. *Let $C \subset S$ be a $\mathcal{C}$-set and let $x \in S \setminus C$ where S has the unboundedness property. Then there is some edge of C that is separated from x by no line of E that splits C.*

Proof. Suppose on the contrary that for every edge $\langle u, u' \rangle$ in C there is some $\lambda \in E$ such that λ splits C and separates u, u' from x. This property and the property for C to be a $\mathcal{C}$-set are invariant under homotopies modulo S. So we can suppose that x is unbounded and use Lemma 3.4.

Each line $\lambda \in E$ will split C into two arcs; we define $A(\lambda)$, to be the arc that is separated from x by λ. For an edge $\langle u, u' \rangle$ in C there will, in general, be several lines that split C and separate the edge $\langle u, u' \rangle$ from x. So we will start by reducing E to a subset that has no superfluous lines.

We choose $E' \subset E$ so that
(i) every line in E' splits C into two arcs,
(ii) for every edge $\langle u, u' \rangle$ in C there is some $\lambda \in E'$ such that $\langle u, u' \rangle \subseteq A(\lambda)$,
(iii) the set E' is minimal (under inclusion) among such sets.

The first condition throws away all lines that play no part in the definition of C as a $\mathcal{C}$-set. The minimality condition is equivalent to the following condition.

(iiia) For every $\lambda \in E'$ there is some edge that lies in $A(\lambda)$ but lies in no other arc $A(\mu)$, $\mu \in E'$.

The proof is divided into two parts. In the first part we consider the case when $\sharp E' \geq 3$ and in the second part we consider the case when any such E' contains only two lines. We will use the first part to prove the second part.

Part 1: Suppose that $\sharp E' \geq 3$. In this case E' has some useful properties described in the following lemma.

Lemma 5.9. *Let E' be chosen as above. Then if $\lambda \in E'$ and $\mu \in E'$ either $A(\lambda) \cap A(\mu) = \emptyset$ or their intersection is an arc that is split by no line in E'. Each $u \in C$ lies in at most two arcs $A(\lambda)$, $\lambda \in E'$. Further for any $\lambda \in E'$ there are precisely two lines $\mu, \nu \in E'$ such that $A(\lambda)$ has non-empty intersection with $A(\mu)$ and $A(\nu)$.*

Proof. Suppose $\lambda, \lambda' \in E'$ and $A(\lambda) \cap A(\lambda') = D$ is an arc with at least two elements. (Note that the condition $\sharp E' \geq 3$ implies that D must be an arc.) Suppose there is some $\mu \in E'$ that splits D. The arc $A(\mu)$ cannot lie in either $A(\lambda)$ or $A(\lambda')$ else μ would not satisfy (ii). We can assume then that $A(\lambda) \subset D \cup A(\mu)$. Now consider the edge $\langle u, u' \rangle$ of C given for λ by condition (ii). This cannot lie in $A(\mu)$ or in D but it must lie in $D \cup A(\mu)$. So, say, $u \in D$, $u' \in A(\mu)$ but $u \notin A(\mu)$, $u' \notin D$ which is impossible since μ splits D. This proves the first result.

Now suppose that there are three lines $\lambda, \mu, \nu \in E'$ such that $A(\lambda) \cap A(\mu) \cap A(\nu) = D$ is non-empty. Since the intersection of any two of these arcs cannot be split by the line determining the other one we conclude that D has an end-point in common with two of these arcs. But this means that one of these arcs must be a sub-arc of one of the others. This is impossible as it contradicts condition (iiia). This proves the second result.

The last statement follows easily from the first two results and condition (ii). □

Now consider the projection $P_{E'}$. If x is unbounded in W then $P_{E'}x$ is unbounded in $P_{E'}W$. Take any line $\lambda \in E'$. The results of Lemma 5.9 show that there are $\mu, \nu \in E'$ so that $A(\lambda)$ is the union of two or three arcs $D = A(\lambda) \cap A(\mu)$, $D' = A(\lambda) \cap A(\nu)$ and D^* (which could be empty). The only line in E' that separates D^* from x is λ so $P_{E'}D^*$ and $P_{E'}x$ have the same restriction to the line λ; call this interval I. Further the only lines in E' that separate D from x are λ and μ and so $P_{E'}D$ has a restriction to the line λ which is an interval distinct from I, but with the vertex $\{\lambda, \mu\}$ an end-point in common with I. The same applies to $P_{E'}D'$ where the common end-point is now the vertex $\{\lambda, \nu\}$. Thus I has two distinct end-points and cannot be unbounded. Since this applies to all $\lambda \in E'$ we can apply Lemma 3.4 and deduce that $P_{E'}x$ is bounded and this in turn implies that $x \in S$ is always bounded thus contradicting the unboundedness property.

Part 2: Now let us suppose that the only sets E' with properties (i), (ii), and (iii) have just two lines. Let $E' = \{\lambda, \mu\}$ be such a set.

First consider the case when $\sharp C = 4$. We can write $C = \langle a, b, c, d\rangle$ where $A(\lambda) = \langle a, b, c\rangle$ and $A(\mu) = \langle c, d, a\rangle$.

There must exist $\nu \in E$ that separates a, b from c, d and we can choose the notation so that $\langle a, b\rangle \subset A(\nu)$.

There must also exist $\ell \in E$ that separates b, c from a, d. If $\langle b, c\rangle \subset A(\ell)$ then we would have $(b, c \mid x, a)$, $(x, c \mid a, b)$ and $(b, x \mid a, c)$ from ℓ, ν and μ which is impossible by Lemma 3.8. So $\langle a, d\rangle \subset A(\ell)$.

It is easy to check then that $C^* = \langle a, d, x, b\rangle$ is a $\mathfrak{C}$-set. Further $c \in S \setminus C^*$ is separated from every edge of C^* by one of the lines λ, μ, ν, ℓ and these lines all split C^*. Furthermore these lines form a set E' (with respect to C^*) that satisfies the above conditions (i), (ii), and (iii) so we can apply the argument in Part 1 and derive a contradiction to the unboundedness property.

Thus the required result is true if $\sharp C = 4$ in both cases so we are in a position to use induction on $\sharp C$. Suppose then that the result is true for all C^* with $\sharp C^* < \sharp C$ and we are given $E' = \{\lambda, \mu\}$ as above.

We let $A(\lambda) \cap A(\mu) = A \cup B$ where A and B are disjoint arcs of C. There are unique elements $u, v, w, z \in C$ such that neither u nor v is in $A(\lambda)$, neither w nor z is in $A(\mu)$ but $\langle w, A, u\rangle$ and $\langle v, B, z\rangle$ are arcs of C. It is possible that $u = v$ or $w = z$ but this is taken into account in the argument.

Let us note first that no line $\nu \in E$ can separate $\langle w, A, u\rangle$ from $\langle z, B, v\rangle$. To see this, we may suppose that ν is such a line and $\langle w, A, u\rangle \subset A(\nu)$. Then we would have $(x, u \mid w, b)$, $(x, w \mid b, u)$ and $(w, u \mid b, x)$ from λ, μ and ν which is impossible by Lemma 3.8. A similar argument applies if $\langle z, B, v\rangle \subset A(\nu)$.

Let $C^* = A \cup \{u, x, w\}$ have the cyclic ordering defined by specifying $\langle w, A, u\rangle$ and $\langle u, x, w\rangle$ to be arcs of C^*. We claim that C^* is a $\mathfrak{C}$-set.

Let $\ell \in E$ split C^*. If ℓ does not split C it must split C^* into the arcs $\{x\}$ and $\langle w, A, u\rangle$. If $A(\ell) \cap C^*$ is not an arc of C^*it must contain both u and w but not the whole of A. Then for any $a \in A \setminus A(\ell)$ we would have $(a, x \mid w, u)$, $(a, w \mid x, u)$ and $(a, u \mid w, x)$ from ℓ, λ and μ, which is impossible by Lemma 3.8. Thus $A(\ell) \cap C^*$ is an arc and ℓ splits C^* into the arcs $A(\ell) \cap C^*$ and $C^* \setminus A(\ell)$. This proves condition (1) of Definition 5.2. Condition (2) follows easily from the fact that λ separates $\langle u, x\rangle$ from $A \cup \{w\}$ and μ separates $\langle w, x\rangle$ from $A \cup \{u\}$. This shows that C^* is a $\mathcal{C}$-set. Similarly $C^{**} = B \cup \{z, x, v\}$ is a $\mathcal{C}$-set.

If $\sharp C^* < \sharp C$ or if $\sharp C^* = 4$ we can apply the inductive hypothesis to C^* and any $b \in B$. On the other hand if $\sharp C^* = \sharp C$ we must have $\sharp C^{**} = 4$ so we can apply the inductive hypothesis to C^{**} and any $a \in A$. By changing the labelling if necessary we may suppose that we can apply the inductive hypothesis to C^* and deduce that for any $b \in B$ there is some edge $\langle a, a'\rangle$ of C^* that is not separated from b by any line in E that splits C^*. Since λ and μ both split C^* and since $b \in A(\lambda) \cap A(\mu)$ means that λ separates $\langle u, x\rangle$ from b and μ separates $\langle w, x\rangle$ from b we see that the edge $\langle a, a'\rangle$ must be an edge of $\langle w, A, u\rangle$.

Further, if $\langle b, b'\rangle$ is any edge of $\langle z, B, v\rangle$ that has no end-point in common with $\langle a, a'\rangle$, then there is some $\ell \in E$ that separates $\langle a, a'\rangle$ from $\langle b, b'\rangle$ and so by the above cannot split C^*. This means that ℓ separates C^* from $\langle b, b'\rangle$. However ℓ cannot separate $\langle w, A, u\rangle$ from $\langle z, B, v\rangle$ so ℓ must split $\langle z, B, v\rangle$.

Take some $a \in A$ and let $E^* \subset E$ be defined by $\ell \in E^*$ whenever $A(\ell) \cap C^* = \emptyset$. Clearly all the lines in this set split C^{**}. Then the above says that every edge of $C^{**} \setminus C^*$ is separated from a by a line in E^*. Note that the edges $\langle v, x\rangle$ and $\langle x, w\rangle$ are separated from a by λ and μ respectively. So if $C^* \cap C^{**} = \emptyset$ this means that every edge of C^{**} is separated from a by a line in $E^* \cup \{\lambda, \mu\}$. Since all the lines in this set clearly split C^{**} and any subset with the same property must include λ, μ and a line in E^* we can apply Part 1 and obtain a contradiction.

Suppose however that $w = z$ then there must exist an edge $\langle a, a'\rangle$ of $\langle A, u\rangle$ and a line $\nu \in E$ that separates the edge $\langle a, a'\rangle$ from the edge $\langle w, b\rangle$ of $\langle w, B\rangle$ since both these edges are also edges of C. Similarly if $u = v$ there must be a line $\bar{\nu} \in E$ that separates the edge $\langle \bar{b}, u\rangle$ of $\langle B, u\rangle$ from a. If $C^* \cap C^{**} \neq \emptyset$ all we have to do is to repeat the above argument after increasing the set of lines $E^* \cup \{\lambda, \mu\}$ by adding ν or $\bar{\nu}$ or both as necessary.

We obtain a contradiction in every case and this proves the proposition. □

We need to consider maximal $\mathcal{C}$-sets and so we need to consider when we can add new elements to a given $\mathcal{C}$-set to make another $\mathcal{C}$-set. That is the purpose of the following propositions.

Proposition 5.10. *Let $C \subset W$ be a $\mathcal{C}$-set and let $x \in S \setminus C$. Suppose there are at least two edges which are not separated from x by any line in E that splits C. Then there are are just two such edges with a unique common end-point $v \in C$.*

Further, the point $v \in C$ is completely determined by the property that no line in $E(C)$ separates x from v.

Proof. We define $\mathcal{A}(\lambda)$ as the set of edges of C that are separated from x by λ. By hypothesis there are at least two edges that lie in no $\mathcal{A}(\lambda)$, $\lambda \in E$. However C is a $\mathcal{C}$-set so if these edges are distinct there is some $\lambda \in E$ that separates them and hence one of them must lie in $\mathcal{A}(\lambda)$. We conclude that there can only be two edges and they have a common end point v.

To prove the last statement let v have the given property and let $\langle u, v, w\rangle$ be an arc of C. Consider a line that separates u, v from x. It cannot split $C \setminus \{v\}$ else it would belong to $E(C)$. So since $u \in C \setminus \{v\}$ it must separate C from x and hence cannot split C. We can apply the same argument to v, w so both the edges $\langle u, v\rangle$ and $\langle v, w\rangle$ are such that no line in E that separates the edge from x also splits C.

Conversely, suppose the two edges have this property and a line separates x from v. If it separates either u, v from x or v, w from x it cannot split C and so cannot belong to $E(C)$. On the other hand if such a line separates v from u, x, w it must separate C into the two arcs v and $C \setminus \{v\}$ and so cannot belong to $E(C)$. Thus no line in $E(C)$ separates x from v and the proposition is proved. □

Lemma 5.11. *Let $C \subset W$ be a $\mathcal{C}$-set and let $x \in W$. Then the following two statements are equivalent.*

There is no line in $E(C)$ that separates x from $v \in C$.

There is no line that separates x from $v \in C$ that also splits $C \setminus \{v\}$.

Proof. Suppose the first statement holds. Then a line that separates x from v cannot belong to $E(C)$ and so clearly cannot split $C \setminus \{v\}$. So the second statement holds.

Suppose the second statement holds then if a line belongs to $E(C)$ it splits $C \setminus \{v\}$ and so cannot separate x from v. So the first statement holds. □

There are either just two edges as described by Proposition 5.8 or such an edge is unique. In the first case $(C \setminus \{v\}) \cup \{x\}$ is a $\mathcal{C}$-set and in the second case $C \cup \{x\}$ is a $\mathcal{C}$-set. This is proved in the following two propositions. Of course if C is not a subset of S or $x \notin S$ there may be no such edges even though C is a $\mathcal{C}$-set.

Proposition 5.12. *Let $C \subset W$ be a $\mathcal{C}$-set and let $x \in W \setminus C$. Then $C \cup \{x\}$ is a $\mathcal{C}$-set if and only if there is a unique edge that is separated from x by no line of E that splits C.*

Proof. Suppose that $C \cup \{x\}$ is a $\mathcal{C}$-set. If there is more than one edge that is separated from x by no line of E that splits C Proposition 5.10 tells us that we can find an arc of C, $\langle u, v, w\rangle$ such that there is no $z \in C$ such that $(u, v \mid x, z)$ or $(v, w \mid x, z)$.

Now there must be some edge $\langle a, b\rangle$ of C such that $\langle a, x, b\rangle$ is an arc of $C \cup \{x\}$. But then any $\lambda \in E$ that splits $C \cup \{x\}$ cannot separate a, b from x. This means that $\langle a, b\rangle$ must be one of the edges $\langle u, v\rangle$ or $\langle v, w\rangle$.

Suppose that $\langle u, x, v, w\rangle$ is an arc of $C \cup \{x\}$. Then since $C \cup \{x\}$ is a $\mathcal{C}$-set there is some $\lambda \in E$ that separates v, w from x, u, which is impossible. The same argument applies if $\langle u, v, x, w\rangle$ is an arc of $C \cup \{x\}$. We deduce that the edge given by Proposition 5.8 is unique.

Conversely, suppose that $\langle u, v\rangle$ is the unique edge with the property that no line in E that splits C separates u, v from x. We claim that $C \cup \{x\}$ is a $\mathcal{C}$-set in which $\langle u, x, v\rangle$ is an arc.

Let $\lambda \in E$. If λ splits C but does not split $\langle u, v\rangle$ it splits $C \cup \{x\}$ into two arcs, one of which contains the arc $\langle u, x, v\rangle$. On the other hand if it does split $\langle u, v\rangle$ it splits $C \cup \{x\}$ into two arcs one of which has x as an end-point. If it does not split C then it either does not split $C \cup \{x\}$ or splits it into the arcs $\{x\}$ and C.

It is clear that we only have to check that the two "new" edges $\langle u, x\rangle$ and $\langle x, v\rangle$ are separated from edges of C.

Consider first those edges which do not have u or v as an end-point. There is some line E that separates such an edge from $\langle u, v\rangle$ since C is a $\mathcal{C}$-set. Since this line splits C it cannot separate u, v from x so it separates the edge from the arc $\langle u, x, v\rangle$.

Consider now the edge $\langle v, w\rangle$ in C that has v but not u as an end-point. We have to show that $(u, x \mid v, w)$. By the uniqueness of the edge $\langle u, v\rangle$ there must exist some $\lambda \in E$ that splits C and separates v, w from x. Now λ cannot separate u, v from x so it must separate x, u from v, w. We show similarly that $(\bar{w}, u \mid x, v)$ if $\langle \bar{w}, u\rangle$ is the edge of C that has u but not v as an end-point.

This is enough to prove that $C \cup \{x\}$ is a $\mathcal{C}$-set with arc $\langle u, x, v\rangle$. □

Proposition 5.13. *Let $C \subset W$ be a $\mathcal{C}$-set and let $x \in W \setminus C$. Suppose that $v \in C$ is such that no line in $E(C)$ separates x from v. Then $(C \setminus \{v\}) \cup \{x\}$ is a $\mathcal{C}$-set and the ordering must be that in which x replaces v.*

Proof. Write $C' = (C \setminus \{v\}) \cup \{x\}$ and consider the cyclic ordering on C' obtained by just replacing $v \in C$ by $x \in C'$ in the ordering on C. Let $\lambda \in E$ split C into two arcs A and B where $v \in B$. Then by definition A is an arc of C' and $(B \setminus \{v\}) \cup \{x\}$ is an arc of C'. If $\lambda \in E$ does not split C then either it does not split C' or it separates x from C in which case λ splits C' into $\{x\}$ and $C \setminus \{v\}$ which are arcs of C'. So the first condition for C' to be a $\mathcal{C}$-set is satisfied.

We now have to show that C' satisfies the second condition. Let $\langle u, v, w\rangle$ be an arc of C so that $\langle u, x, w\rangle$ is an arc of C'. Then by Proposition 5.10 the condition that no line in $E(C)$ separates x from v means that there are just the two edges $\langle u, v\rangle$ and $\langle v, w\rangle$ that are not separated from x by a line of E that splits C; that is, there is no $z \in C$ such that $(u, v \mid x, z)$ or $(v, w \mid x, z)$.

Clearly we only have to check that there are lines in E separating the "new" edges in C'; that is, those edges with x as an end-point. So let $\langle a, a'\rangle$ be an edge of C' for which x is not an end-point then it is an edge of C.

If u is not an end-point of a, a' we know, since C is a $\mathcal{C}$-set, that $(u, v \mid a, a')$ and we cannot have $(u, v \mid x, a)$ so $(u, v, x \mid a, a')$. Similarly, if w is not an end-point of a, a' we can show $(x, v, w \mid a, a')$. Thus we have shown $(u, x \mid a, a')$ and $(x, w \mid a, a')$ and the fact that C' is a $\mathcal{C}$-set follows easily.

Theorem 5.5 shows that the ordering must be the one obtained by replacing v by x. □

In the next propositions we show that $\mathcal{C}$-sets are not necessarily obtained by adding on one point at a time.

Proposition 5.14. *Let C be a $\mathcal{C}$-set and let $x, y \in W$. Suppose that $u, v \in C$ are such that no line in $E(C)$ separates x from u or separates y from v. Suppose also that $(x, y \mid C)$. Let B be an arc of C with end-points u and v such that the complementary arc $A = C \setminus B$ contains at least one edge. Let $C' = A \cup \langle u, x, y, v\rangle$. Then C' is a $\mathcal{C}$-set.*

Proof. The edges of C' are the edges of A, the edges of the arc $\langle u, x, y, v\rangle$ together with the edges of C $\langle \bar{u}, u\rangle$ and $\langle v, \bar{v}\rangle$ where $\bar{u}, \bar{v}$ are the end-points of A. We will use Lemma 5.11 without comment.

Let $\lambda \in E$ split C. Suppose λ splits A then it splits both $C \setminus \{u\}$ and $C \setminus \{v\}$ so it cannot separate u from x or v from y. If it also splits B it separates u from v and hence u, x from v, y thus it must split one edge in A and the edge $\langle x, y\rangle$. So it splits C' into two arcs. If it does not split B then it does not separate u from v so it does not separate u, x from v, y and hence does not split $\langle u, x, y, v\rangle$ and so splits just two edges of C' ; they are in A. So it splits C' into two arcs.

Suppose λ does not split A but splits both $C \setminus \{u\}$ and $C \setminus \{v\}$. If it does not separate u from v then it does not separate u, x from v, y so it could split C' into the two arcs A and $\langle u, x, y, v\rangle$. If it separates u from v it separates u, x from v, y and splits just two edges of C'; $\langle x, y\rangle$ and either $\langle \bar{u}, u\rangle$ or $\langle v, \bar{v}\rangle$. So it splits C' into two arcs.

The only other possibilities are that λ separates u from $C \setminus \{u\}$ so it splits $C \setminus \{v\}$ and must split off one of the arcs u or $\langle u, x\rangle$ from C'. There is a similar case with u and v interchanged. Finally λ may not split C at all. So either it does not split C' or it splits off one of the arcs $\langle x, y\rangle$ or $\{x\}$ or $\{y\}$ from C'.

We now have to show that we can separate any two edges of C'. By hypothesis we do not have to worry about $\langle x, y\rangle$. Also $(\bar{u}, u \mid v, \bar{v})$ implies $(\bar{u}, u, x \mid y, v, \bar{v})$. Again, if $\langle a, b\rangle$ is any edge of A, $(\bar{u}, u \mid a, b)$ implies $(\bar{u}, u, x \mid a, b)$ and $(a, b \mid v, \bar{v})$ implies $(a, b \mid y, v, \bar{v})$. This proves the result. □

The elements of B apart from u, v, if there are any, all lie "inside" the new $\mathcal{C}$-set C' and are bounded. So this proposition is a little misleading because we are going to be concerned only with $C \subset S$ and $x, y \in S$ when this would contradict the "unboundedness" property. This is partly proved in the following lemma.

Lemma 5.15. *In the situation of Proposition* 5.14 *if* $B \neq \langle u, v\rangle$ *there is some* $z \in B$, $z \neq u$, $z \neq v$ *such that each edge of* C' *is separated from* z *by a line of* E *that splits* C'.

Proof. We take $z \in B$ so that $\langle u, z\rangle$ is an edge of B. For any edge $\langle a, b\rangle$ of C we have $(a, b \mid u, z)$ so we only have to consider the "new edges" $\langle x, y\rangle$, $\langle y, v\rangle$ and $\langle u, x\rangle$. The first follows from $(x, y \mid C)$. The second follows because $(u, z \mid v, \bar{v})$ implies $(u, z \mid y, v, \bar{v})$ For the last we take the edge $\langle z, \bar{z}\rangle$ of B with $\bar{z} \neq u$ and apply the same argument. This completes the proof. □

6. Maximal $\mathcal{C}$-sets

We are going to consider here $\mathcal{C}$-sets that are maximal under inclusion. If C is a maximal $\mathcal{C}$-set and x is any element of $S \setminus C$ then clearly $C \cup \{x\}$ cannot be a $\mathcal{C}$-set. Hence by Propositions 5.10 and 5.12 there is a unique $v \in C$ such that no line in $E(C)$ separates x from v.

Definition 6.1. *Let* C *be a maximal* $\mathcal{C}$*-set. We define* $F_C : S \to C$ *by putting, for any* $x \in D$, $F_C(x) = v \in C$ *where* v *is the unique element such that no line in* $E(C)$ *separates* x *from* v.

Given $x, y \in S$, *we say that* $x \sim y$ *if for each maximal* $\mathcal{C}$*-set* C *we have* $F_C(x) = F_C(y)$.

Given two maximal $\mathcal{C}$*-sets* C *and* C' *we say that* $C \cong C'$ *if* $F_C|'_C$, *that is* F_C *restricted to* C', *is a monomorphism.*

By Proposition 5.13 if $F_C(x) = v \in C$ then $C' = (C \setminus \{v\}) \cup \{x\}$ is a $\mathcal{C}$-set. We are going to use this often so we shall call it simply the "replacement property". In the case of maximal

$\mathcal{C}$-sets we show that C' is again maximal and so the replacement process can be repeated and the result is always a maximal $\mathcal{C}$-set.

Proposition 6.2. *The relation $\sim$ is an equivalence relation and each equivalence class is a $\mathcal{T}$-set.*

Proof. The fact that $\sim$ is an equivalence relation follows immediately from the way it is defined.

Let T be an equivalence class. Each element of a maximal $\mathcal{C}$-set must belong to a different equivalence class, since the cyclic ordering on a $\mathcal{C}$-set is unique. Thus T contains no $\mathcal{C}$-set and thus by Lemma 5.4 and Definition 4.1 T is a $\mathcal{T}$-set. □

We have already defined $E(C)$ when C is a $\mathcal{C}$-set but we repeat it here to emphasise that $E(T)$ where T is a $\mathcal{T}$-set is defined differently.

Definition 6.3. *For each $\mathcal{C}$-set $C \in S$ we define $E(C) \subseteq E$ to be those lines that split C into two arcs each of which contains an edge of C.*

For each $\mathcal{T}$-set we define $E(T) \subset E$ to be those lines that split T.

Note that $E(C)$ is the set of lines that determines uniquely the structure of a $\mathcal{C}$-set on C. The projection $P_{E(C)}$ defined by restricting the elements of W to $E(C)$ (recall that these elements are defined as maps from E to $\{-1, +1\}$) is an isomorphism when restricted to C and preserves all the separation properties of C.

Proposition 6.4. *Let C be a $\mathcal{C}$-set and let $x \in S$. Then $E(C)$ and $E([x])$ are disjoint.*

Proof. Suppose that $\lambda \in E([x])$ so λ separates x from y where $F_C(x) = F_C(y) = v \in C$. We may suppose that λ separates x from y, v. By definition of F_C this means that λ cannot split $C \setminus \{v\}$.

Now consider $C' = (C \setminus \{v\}) \cup \{y\}$ which is a $\mathcal{C}$-set by the replacement property. Since $C \setminus \{v\} = C' \setminus \{y\}$ the defining property for $F_C(y) = v$ that no line in E can separate y from v and also split $C \setminus \{v\}$ tells us that $F_{C'}(v) = y$.

Further since any line in $E(C)$ splits $C \setminus \{v\}$ and hence cannot separate x from v it also belongs to $E(C')$. By the symmetry we have just pointed out we also deduce that $E(C') \subset E(C)$. Hence $E(C) = E(C')$. So we may assume $y = v$. However in this case we know that λ does not split $C \setminus \{v\}$ and so cannot belong to $E(C)$ which is what we wanted to prove. □

Proposition 5.14 and Lemma 5.11 when applied to maximal $\mathcal{C}$-sets have very strong consequences. For example, they imply that $\cong$ is an equivalence relation. They also imply that replacement can be carried out repeatedly and in any order and the result is still a maximal $\mathcal{C}$-set. The crucial result is contained in the next proposition.

Proposition 6.5. *Let $C \subset S$ be a maximal $\mathcal{C}$-set and suppose that $x, y \in S$ where $F_C(x) = u$, $F_C(y) = v$ and $u \neq v$. Then no line in E separates x, y from C.*

Proof. Let $C \subset S$ be a maximal $\mathcal{C}$-set and suppose $(x, y \mid C)$.

If $\langle u, v\rangle$ is an edge of C we appeal to Proposition 5.14 which says that $C \cup \{x, y\}$ is a $\mathcal{C}$-set. This would contradict the maximality of C.

If $\langle u, v\rangle$ is not an edge of C we appeal to the same proposition and Lemma 5.11 which would contradict the unboundedness property. We deduce that such a line does not exist. Hence the result. □

Proposition 6.6. *Let $x, y \in S$ and suppose $[x] \neq [y]$. Then $E([x])$ and $E([y])$ are disjoint.*

Proof. Since $[x] \neq [y]$ there is some maximal $\mathcal{C}$-set C so that $F_C(x) = u \neq v = F_C(y)$. Now suppose $\lambda \in E$ splits both $[x]$ and $[y]$. Then we may suppose that x, y are chosen so that λ separates x, y from u, v. Proposition 6.5 says that this is impossible. This proves the result. □

Proposition 6.7. *Let C be a maximal $\mathcal{C}$-set and let $D^* \subset S$ be such that F_C restricted to D^* is a monomorphism with $F_C(D^*) = D$. Define $C^* = (C \setminus D) \cup D^*$. Then C^* is a maximal $\mathcal{C}$-set, $E(C) = E(C^*)$ and $F_C \circ F_{C^*} = F_C$.*

Proof. Let $\lambda \in E(C)$. Then λ splits $C \setminus \{v\}$ for every $v \in C$. Hence it cannot separate x from $F_C(x)$ for any $x \in S$. Hence λ separates $x \in C^*$ from $y \in C^*$ if and only if it separates $F_C(x) \in C$ from $F_C(y) \in C$. Since the property of being a $\mathcal{C}$-set (and, by Theorem 5.5, the cyclic ordering) is completely determined by the way it is separated we conclude that C^* is also a $\mathcal{C}$-set and $E(C) = E(C^*)$.

To show that C^* is maximal we suppose on the contrary that C' is another $\mathcal{C}$-set with C^* a proper subset of C'. Then $E(C) = E(C^*) \subset E(C')$ and the same argument shows that we can replace C^* by C in C' and obtain another $\mathcal{C}$-set which contains C as a proper subset. Since this is impossible we conclude that C^* is also maximal.

Finally we observe that, given $z \in S$, the fact that there is no line in $E(C)$ that separates z from v defines $F_C(z) = v$. But if $x \in C^*$ is given by $F_C(x) = v$ no line that belongs to $E(C)$ can separate x from v. So no line in $E(C) = E(C^*)$ separates z from x which means $F_{C^*}(z) = x$. Hence $F_C \circ F_{C^*} = F_C$ as required. □

Corollary 6.8. *Let $C \subset S$ and $C' \subset S$ be $\mathcal{C}$-sets such that $F_C|_{C'}$ is a monomorphism. Then $F_C|_{C'}$ is an isomorphism with inverse $F_{C'}|_C$' and $E(C) = E(C')$.*

Proof. Apply the above proposition with $D^* = C'$. Since $C' \subset C^*$ and C' is maximal we deduce that $C' = C^*$. So $F_C|_{C'}$ is an isomorphism. The relation $F_C \circ F_{C'} = F_C$ when restricted to C gives $F_C|_{C'} \circ F_{C'}|_C = \mathrm{id}_C$ so $F_{C'}|_C$ is the inverse of $F_C|_{C'}$.

The fact that $E(C) = E(C')$ is also given by the above proposition since $C^* = C'$. □

Proposition 6.9. *The relation $\cong$ is an equivalence relation between maximal $\mathcal{C}$-sets. Further $C \cong C'$ if and only if $E(C) = E(C')$.*

Proof. We see from the above Corollary 6.8 that $C \cong C'$ if and only if $F_C|_{C'}$ is an isomorphism and then $E(C) = E(C')$.

It follows from Proposition 6.7 that if C^* is also a maximal $\mathcal{C}$-set then $F_C|_{C^*} \circ F_{C^*}|_{C'} = F_C|_{C'}$. It follows immediately that $\cong$ is an equivalence relation.

Suppose that $E(C) = E(C')$ but $F_C|_{C'}$ is not a monomorphism. Then there is some $x, y \in C'$ with $F_C(x) = F_C(y) = v$, but $x \neq y$. However there must be some line in $E(C')$ that separates x from y and so separates one of them, say x, from v. Since this line is in $E(C)$ also this is impossible. So $F_C|_{C'}$ is a monomorphism and $C \cong C'$. □

In fact the condition $E(C) = E(C')$ could have been taken as the definition of $C \cong C'$ which perhaps throws some light on the meaning of this equivalence relation.

On the other hand the next result shows that if $C \ncong C'$ then $E(C)$ and $E(C')$ are disjoint. We could also have proved that, in this case, $F_C|_{C'}$ is a constant map but this is not necessary for our main theorem.

Proposition 6.10. *Let $C, C' \subset S$ be maximal $\mathcal{C}$-sets where $C \ncong C'$. Then $E(C)$ and $E(C')$ are disjoint.*

Proof. Since $C \ncong C'$ neither F_C nor $F_{C'}$ is a monomorphism by Corollary 6.8. There must be some $v \in C$ such that $D = F_C^{-1}(v)$ contains at least two elements. Suppose that $C' \setminus D$ also has more than one element. Then we can find edges $\langle x, x' \rangle$ and $\langle y, y' \rangle$ of C' which have no common end-point such that $x, y \in D$ and $x', y' \notin D$. There exists $\lambda \in E(C')$ that separates these edges and we may suppose that λ separates x, x', v from y, y'. Then λ cannot split $C \setminus \{v\}$.

Suppose it separates x, x', v from $y, y', C \setminus \{v\}$. Then we can put $F_C(x') = u \neq v$ and λ separates x' from u so it cannot split $C \setminus \{u\}$ either. This is impossible so we conclude that λ does not split C but separates x, x', C from y, y'. We can then apply Proposition 6.5 and get a contradiction. We conclude that $C' \setminus D$ has at most one element.

Now any line in $E(C) \cap E(C')$ would have to split D and at the same time it must not separate any $x \in D$ from $F_C(x) = v$. Clearly this is impossible so $E(C) \cap E(C') = \emptyset$. □

Theorem 6.11. *Let W be the regions for an arrangement of a set E of lines in dimension 2 and let $S \subset W$ have the "unboundedness property". Then $\sharp S \leq 2 \times \sharp E$.*

Proof. For each $x \in S$ we know from Proposition 6.2 that $[x]$ is a $\mathfrak{T}$-set and has the structure of a connected tree. We take in each equivalence class a "base point" $a \in [x]$ which is just any element which is the end-point of only one edge. If the equivalence class only contains one point we let a be that point. Thus $T = [x] \setminus \{a\}$ is either empty or is a $\mathfrak{T}$-set. Let $\mathcal{T}$ be the collection of all these non-empty $\mathfrak{T}$-sets and let D be the set of all the chosen base points. Associated with $T \in \mathcal{T}$ is the set $E'(T) = E([x]) \subset E$. By Proposition 6.6, if $T, T' \in \mathcal{T}$ and $T \neq T'$ then $E'(T)$ and $E'(T')$ are disjoint. Further Proposition 4.3 tells us that $\sharp T \leq \sharp E'(T)$

Now consider an equivalence class $[C]$ of maximal $\mathcal{C}$-sets. By the replacement property given by Proposition 6.7 we can choose a representative of each equivalence class that lies in D. Let C be the set of these representatives. Then if $C \in \mathcal{C}$ and $T \in \mathcal{T}$ we know from Proposition 6.4 that $E(C)$ and $E'(T)$ are disjoint. Also if $C, C' \in \mathcal{C}$ and $C \neq C'$ Proposition 6.10 tells us that $E(C)$ and $E(C')$ are disjoint. Further Lemma 5.7 says that $\sharp C \leq 2 \times \sharp E(C)$.

Every element of S occurs either as an element of some $T \in \mathcal{T}$ or as an element of D and hence as an element of some $C \in \mathcal{C}$ so clearly

$$\sharp S \leq \sum_{S \in \mathcal{C}} \sharp C + \sum_{T \in \mathcal{T}} \sharp T.$$

Also

$$\bigcup_{C \in \mathcal{C}} E(C) \cup \bigcup_{T \in \mathcal{T}} E(T) \subseteq E.$$

We noted above that the sets $E(C)$ for $C \in \mathcal{C}$ and the sets $E(T)$ for $T \in \mathcal{T}$ are all disjoint, consequently

$$\sum_{C \in \mathcal{C}} \sharp E(C) + \sum_{T \in \mathcal{T}} \sharp E(T) \leq \sharp E.$$

So we obtain

$$\sharp S \leq \sum_{S \in \mathcal{C}} \sharp C + \sum_{T \in \mathcal{T}} \sharp T \leq 2 \times \Big(\sum_{C \in \mathcal{C}} \sharp E(C) \Big) + \sum_{T \in \mathcal{T}} \sharp E(T) \leq 2 \times \sharp E.$$

This completes the proof of the theorem. □

7. Higher Codimensions

It is clear that the majority of the results we have obtained do not really depend on the fact that E is a set of lines in general position in the plane and that the elements of W correspond to the region in the complement of the union Σ of these planes. They depend only on the fact that each element of E corresponds to a subdivision of W into two subsets and on the corresponding subdivisions of $S \subset W$.

This can be axiomatised and extended to higher codimensions as follows. We let $\mathbb{K} = \{-1, +1\}$ and let $\mathcal{P}_r(E)$ be the collection of all subsets of E with r elements. For any $\chi \subset E$ we let $P_\chi : \mathbb{K}^E \to \mathbb{K}^\chi$ denote the map $\varphi \mapsto \varphi|_\chi$.

Definition 7.1. *Given a set E, an abstract arrangement of E in dimension d is a subset $W \subseteq \mathbb{K}^E$ satisfying the following axioms if $d \leq \sharp E$:*

(1) *for all $\omega \in W$, there exists $v \in \mathcal{P}_d(E)$ such that $P_v\omega' = P_v\omega$ implies that $\omega' \in W$;*
(2) *for all $v \in \mathcal{P}_d(E)$ there exists $\theta(v) \in \mathbb{K}^{E \setminus v}$ such that $P_{E\setminus v}\omega = \theta(u)$ implies that $\omega \in W$;*
(3) (the simplex axiom) *for all $\sigma \in \mathcal{P}_{d+1}(E)$ and all $\omega \in W$ there exists $i \in \sigma$ such that $P_i\omega = P_i\theta(\sigma \setminus \{i\})$;*
(4) (the betweenness axiom) *for any $\lambda \in \mathcal{P}_{d-1}(E)$ any three distinct elements in $E \setminus \lambda$ can be labelled i, j, k so that if we write $v_i = \lambda \cup \{i\}$, $v_j = \lambda \cup \{j\}$ and $v_k = \lambda \cup \{k\}$ then*

$$\begin{aligned} P_i\theta(v_j) &= P_i\theta(v_k), \\ P_j\theta(v_i) &= -P_j\theta(v_k), \\ P_k\theta(v_i) &= P_k\theta(v_j). \end{aligned}$$

If $\sharp E \leq d$ we put $W = \mathbb{K}^E$.

The idea is that E represents a set of hyperplanes in general position in $\mathbb{R}^d$ so that $v \in \mathcal{P}_d(E)$ represents a vertex that is the intersection of d hyperplanes. Similarly $\lambda \in \mathcal{P}_{d-1}(E)$ represents a line and so on. The definition can be expressed entirely in terms of the correspondence $v \in \mathcal{P}_m(E) \mapsto \theta(v) \in \mathbb{K}^{E \setminus v}$ that we shall call the vertex map.

We define an element $\omega \in W$ to be unbounded if $-\omega \in W$ also. The concept of a homotopy is represented by a finite sequence of elementary modifications which are defined as follows.

Definition 7.2. *Let W and W' be two abstract arrangements of E in d dimensions with vertex maps θ and θ'. Then W' is an elementary modification of W in either of the following two situations.*

(1) *There is some $\sigma \in \mathcal{P}_{d+1}(E)$ such that a unique $\omega \in W$ satisfies $P_i\omega = \theta((E \setminus \sigma) \cup \{i\})$ for all $i \in \sigma$ and $W \setminus \{\omega\} = W' \setminus \{\omega'\}$ where $\omega' \in W'$ is defined by $P_j\omega' = P_j\omega$ if $j \notin \sigma$ and $P_i\omega' = -P_i\omega$ if $i \in \sigma$.*

(2) *There is some $\omega \in W$ such that a unique $v \in \mathcal{P}_d(E)$ satisfies $P_{E\setminus\{v\}}\omega = \theta(v)$ and $W \setminus \{\omega\} = W' \setminus \{\omega'\}$ where $\omega' \in W'$ is defined by $P_j\omega' = -P_j\omega$ if $j \notin v$ and $P_i\omega' = P_i\omega$ if $i \in v$.*

Definition 7.3. *Given $S \subset W$ the equivalence relation $\sim_S$ is generated by $W \sim_S W'$ when W' is an elementary modification of W and $S \subset W \cap W'$.*

Then Conjecture 2.2 has an abstract version which should be easier to prove.

Conjecture 7.4. *Let E be a set with m elements and let W be an abstract arrangement of E in d dimensions. Suppose $S \subset W$ has the unboundedness property: for all $\omega \in S$ there exists $W' \sim_S W$ such that $\omega, -\omega \in W'$. Then $\sharp S \leq \gamma(m, d)$.*

The author is grateful to Ilda da Silva for drawing his attention to the fact that the ideas of Definitions 7.1 and 7.2 are equivalent to ideas in [1] and results proved there could well be useful in proving this conjecture.

It is not too difficult to prove directly from the axioms in Definition 7.1 that $\gamma(m, d) = 1 + \binom{m}{1} + \binom{m}{2} + \cdots + \binom{m}{d} - \binom{m-1}{d}$ is the number of elements in $W \cap (-W)$.

These axioms enable one to generalise many of the results of this paper to the abstract situation in d dimensions. For instance Lemma 3.4 generalises. It is possible to generalise Proposition 5.8. We need to replace the idea of a $\mathcal{C}$-set by a set which has a structure representing a piece-wise linear representation of $\mathbb{S}^{d-1}$ compatible in some way with the structure of the abstract arrangement. It does seem that Proposition 5.8 contains the essential information about the effect of the unboundedness property so some sort of generalisation of this proposition is going to be necessary.

It is not clear how this should be done but it is clear that some aspect of homology will be involved. We need to do it in such a way that we obtain a generalisation of Propositions 6.5, 6.6 and 6.10.

Note however that, since $\gamma(r, d) + \gamma(m - r, d) \leq \gamma(m, d)$, we only have to consider the situation where the problem cannot be reduced to the case when $S = S_1 \cup S_2$ and $E_1 \cap E_2 = \emptyset$ where there are enough lines in $E_1 \subset E$ to separate any two elements of S_1 and enough lines in $E_2 \subset E$ to separate any two elements of S_2.

References

[1] Björner, A.; Las Vergnas, M.; Sturmfels, B.; White, N.; Ziegler, G. M.: *Oriented Matroids*. Encyclopedia of Math. and its Appl. **46**, Cambridge University Press 1993. Zbl 0773.52001

[2] Carter, S.; Şentürk, Z.: *The space of immersions parallel to a given immersion*. J. of London Math. Soc. II. Ser. **50** (1994), 404–416. Zbl 0807.53042

[3] Carter, S.; West, A.: *The push-out region of an immersed manifold*. J. of Geom. **74** (2002), 44–60. Zbl pre01890530

[4] Palais, R. S.; Terng, C-L.: *Critical point theory and submanifold geometry*. Lecture Notes in Math. **1353**, Springer, Berlin 1988. Zbl 0658.49001

Received January 25, 2003