Stabbing and Separation

by

Rephael Wenger

School of Computer Science
McGill University

February, 1988

A dissertation submitted to the Faculty of Graduate Studies and Research
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

ISBN 0-315-45941-7

Stabbing and Separation

Abstract

Computational geometry is the study of the algorithmic properties of geometric objects. Some of the favourite questions in computational geometry are about the intersection properties of objects: determine if some objects intersect, report all the objects which do intersect or find all the intersections. This thesis addresses a different type of intersection problem. Given a family of convex polytopes, find a k-flat, an affine subspace of dimension k, which intersects all members of the family. Examples of k-flats are points, lines, planes and hyperplanes. Mathematicians are interested in a similar question. What are the necessary and sufficient conditions for the existence of a k-flat which intersects all members of a given family of convex sets? Such a k-flat is known as a common transversal or stabber.

In this thesis, conditions are given for the existence of line and hyperplane transversals. Theorems are presented about the order in which a line transversal intersects a given family of convex sets. Algorithms are developed for finding line transversals in 3-space.

Hyperplanes which separate convex sets play an important role in conditions for the existence of hyperplane transversals. They also determine the order in which line transversals intersect a family of convex sets. A chapter in this thesis is devoted to studying some algorithmic questions concerning such separating hyperplanes.

Résumé

La géométrie informatique est l'étude des propriétés algorithmiques d'objects géométriques. Certaines des questions préferées en géométrie informatique concernent les propriétés d'intersection d'objects: déterminer si des objects se croisent, rapporter tous les objects qui se croisent ou trouver toutes les intersections. Cetter thèse aborde un genre différent de problème d'intersection. Etant donné une famille d'ensembles polyédraux convexes, trouvez un sous-espace affin de dimension k qui rencontre tous les membres de la famille. Des exemples de sous-espace affin sont les points, droites, plans et hyperplans. Les mathématiciens sont intéressés par une question similaire. Quelles sont les conditions nécessaires et suffisantes à l'existence d'un sous-espace affin de dimension k qui rencontre tous les membres d'une famille donnée d'ensembles convexes? Un tel sous-espace affin est appelé sous-espace affin transversal.

Dans cette thèse, des conditions sont données pour l'existence d'une droite transversale et d'un hyperplan transversal. Des théorèmes sont présentés sur l'ordre dans lequel une droite transversale rencontre une famille donnée d'ensembles convexes. Des algorithmes sont développés pour trouver des droites transversales dans un espace tri-dimensionel.

Les hyperplans qui séparent des ensembles convexes jouent un rôle important dans les conditions pour l'existence d'un hyperplans transversal. Ils déterminent aussi l'ordre dans lequel les droites transversales rencontrent une famille d'ensembles convexes. Un chapitre de cette thèse est consacré à l'étude de certaines questions algorithmiques concernant ces hyperplans séparants.

Acknowledgements

Any thesis is a long, arduous task, which owes itself to many although its authorship is one. Without the help of my advisor, David Avis, I would never have completed this thesis. I thank him for introducing me to computational geometry, for working with me on many of the problems which compose this thesis, and for showing me how to present my results in a style which is concise and coherent (his, not mine). Most of all, I thank him for telling me to finish my thesis. For others who decide to pursue a doctorate, I can only quote the following advice. "Begin at the beginning and go on till you come to the end: then stop."[10]

I owe much to my parents, who sparked my initial interest in mathematics and computer science. I am also indebted to my undergraduate advisor, Kenneth Steiglitz, for guiding me as an undergraduate at Princeton into computer science. During my years at McGill I was supported by research grants from NSERC and CRIM. Many of the students and faculty I encountered at McGill contributed to my education in computer science. Hossam ElGindy, Ryan Hayward, Bill Lenhart, David Rappaport, Tom Shermer and Godfried Toussaint increased my understanding through many, various discussions. Gilles Pesant translated the thesis abtract into French. Meir Katchalski from the Technion introduced me to stabbing problems which form the basis of this thesis. Richard Pollack from N.Y.U. has been a continual inspiration and encouragement, showing great enthusiasm for my research. Some of the work in this thesis was done with him. Finally I would like to thank my wife, Shifra, without whom it wouldn't have been worth doing.

To my parents,

for vast and sundry things.

Table of Contents

Chapter 1

Introduction

Mathematicians have long been interested in the necessary and sufficient conditions for the existence of a k-flat, an affine subspace of dimension k, which intersects all members of a given family of convex sets in E^d, d dimensional Euclidean space. Recently computer scientists have asked their version of this problem. Given a finite family of convex polytopes in E^d and a fixed k, find some k-flat which intersects all members of the family. Mathematicians refer to k-flats which intersect all members of a family of sets as **common transversals** while these objects are known as **stabbers** in the computer science literature.

When k equals 0, the problem becomes one of finding a point which lies at the intersection of a family of compact convex sets or convex polytopes. Helly found a necessary and sufficient condition for the existence of such a point[34]. Helly's Theorem and related others by Radon and Carathéodory spawned a whole area of research into similar types of geometric theorems.

The problem of finding a point stabber for a finite family of convex polytopes can be one way of viewing problems in linear programming. If the convex polytopes are described as the intersection of half-spaces, the problem becomes one of finding a point which lies at the intersection of a finite family of half-spaces. Obviously this problem can be solved using linear programming algorithms. Equally true, any linear programming problem can be reformulated as a problem of finding a point which lies at the intersection of a finite family of half-spaces.

The importance of Helly's Theorem to convexity theory and of linear programming to operations research and computer science motivated me to consider k-flat transversals for values of k greater than 0. By investigating the mathematical properties of these transversals, I hoped to devise efficient algorithms to find them. By asking algorithmic questions about these transversals, I tried to probe their structure and so derive their mathematical properties.

Most research in k-transversals for $k > 0$ has centered on line transversals in the plane and hyperplane transversals in any dimension. Hadwiger found necessary and sufficient conditions for the existence of line transversals for families of pairwise disjoint convex sets[31] in the plane. Katchalski[37] and Goodman and Pollack[29] generalized Hadwiger's Theorem in different ways to theorems about hyperplane transversals. Hadwiger's Theorem relies upon the order in which a directed line intersects pairwise disjoint convex sets. Katchalski and others studied this order, particularly for families of translates[38, 39, 40].

An important property of convex sets concerns separating hyperplanes. Two convex sets are disjoint if and only if they can be strictly separated by a hyperplane. In this thesis I show that separating hyperplanes play an important role in hyperplane transversals as well. Whether there exists a hyperplane transversal depends upon whether the normals of separating hyperplanes cover the unit sphere. Additionally, the order in which a line intersects a family of convex sets is completely determined by the arrangement of lines which separate the convex sets.

Algorithms for finding k-transversals have been recently proposed for line transversals in $\mathbf{E}^2$ and hyperplane transversals in any dimension[2, 4, 18, 19]. Little has been accomplished on k-stabbing in $\mathbf{E}^d$ for values of k other than 0 and $d-1$. I provide an algorithm which D. Avis and I devised for finding line transversals in $\mathbf{E}^3$.

The relationship of separation to stabbing led me to investigate separation properties of convex sets. For instance, how many hyperplanes are needed to separate every pair in a family of n convex sets? How can these hyperplanes be found? Deciding whether m hyperplanes suffice to separate every pair is NP-complete but other questions admit efficient algorithms.

Chapter 2 contains background information on convex sets, polytopes, linear algebra, and some elementary topology and graph theory. It also contains the definitions of the standard terms and the notations I will use. Any new terms which I create or which have only been used recently are also defined within the relevant chapter, so this chapter can be used solely for reference.

Chapter 3 has results on necessary and sufficient conditions for hyperplane transversals. Hadwiger's Theorem is generalized to families of compact convex sets which are not necessarily pairwise disjoint. A sufficient condition is given for the existence of hyperplane transversals for many different families of compact convex sets. This condition includes theorems of Katchalski and Goodman and Pollack on hyperplane transversals. The results in Section 3.2 are contained in a paper "A generalization of Hadwiger's Theorem to intersecting sets" [58], which has been submitted for publication. The results in Section 3.3 are by Richard Pollack and myself and will soon be submitted for publication.

Chapter 4 examines the number of different orders in which a line or hyperplane can intersect a family of convex sets. As noted before, these orders are related to the arrangement of hyperplanes which separate the convex sets. One of the lemmas in this chapter is interesting in its own right. A family of compact convex sets can be embedded in a family of convex polygons using 'few' edges. Section 4.2 is contained in a paper "Upper bounds on geometric permutations"[59] which will appear in "Discrete and Computational Geometry".

In Chapter 5 a theory for line stabbers in E^3 is developed. A Helly-type theorem is given for line stabbing of lines in E^3. The theory for line stabbing is used in algorithms for line stabbing in E^3. The work in Section 5.2 was done by David Avis and myself and was presented at the 3rd ACM Conference on Computational Geometry under the title "Algorithms for line stabbers in space"[5]. A modified version of this paper, entitled "Polyhedral line transversals in space"[6], will appear in a special issue of "Discrete and Computational Geometry".

Chapter 6 investigates some algorithmic questions about hyperplane separators for families of convex sets. Finding a minimum set of hyperplanes which separate a family of convex sets is NP-complete. However, a polynomial time algorithm is given for finding the minimum number of slopes needed to separate the family. An efficient algorithm is also given for finding the minimum size point set needed to cover a set of arcs on the circle.

Finally, Chapter 7 concludes with some open problems and areas of research.

Chapter 2

Background

The computational problems considered in this thesis fall under the area in computer science called computational geometry. Computational geometry encompasses a wide range of problems concerning the algorithmic properties of geometric objects. Two texts are particularly relevant for this thesis. "Computational Geometry" by Preparata and Shamos[48] is an introductory text to many of the basic problems and algorithms studied in computational geometry. "Algorithms in Combinatorial Geometry" by Herbert Edelsbrunner[17] is a more advanced text which covers many of the computational problems related to stabbing. A whole chapter in Edelsbrunner's book is devoted to line stabbing in the plane.

An invaluable reference for geometry, both Euclidean and projective, is "Analytic Geometry," by K. Borsuk[8]. Borsuk gives an excellent treatment of homogeneous coordinates and duality. The classical work on convex polytopes is "Convex Polytopes" by Grunbaum[30], who also gives a good introduction to convex sets. A new book, "Convex Sets and their Applications" by S. Lay[42], also covers convex sets quite well. Many theorems on convex sets are included in Hadwiger, Debrunner and Klee's book, "Combinatorial Geometry in the Plane"[32], and in the article "Helly's Theorem and its Relatives" by Danzer, Grunbaum and Klee[13]. For topology there are many texts available. I suggest "Principles of Mathematical Analysis" by Rudin[50]. Standard texts on graph theory are "Graph Theory with Applications" by Bondy and Murty [7] and "Graph Theory" by Harary[33].

The rest of this chapter introduces the mathematical definitions and properties of many of the mathematical and geometric objects in this thesis. Cartesian and homogeneous coordinates, k-flats, convex sets, polytopes, orientation, projective space, and some elementary topology and graph theory are all discussed here. All terms which are new to this thesis or of recent origin will be defined as necessary within the main text, so this chapter can be merely used as reference by those familiar with the subject matter.

The objects which are studied in this thesis lie in d-dimensional **Euclidean space, E^d**. Each point in E^d is represented by its **Cartesian coordinates**, $(\alpha_1, \ldots, \alpha_d)$, a d-dimensional vector of real numbers. If $x = (\alpha_1, \ldots, \alpha_d)$ and $y = (\beta_1, \ldots, \beta_d)$ are two points in E^d, the **inner product** or **dot product** of x and y, $x \cdot y$, is $\sum_{i=1}^{d} \alpha_i \beta_i$. The **norm** of a vector v, denoted $||v||$, is $\sqrt{v \cdot v}$. The **distance** between two points x and y is $||x-y||$. **R** denotes the set of real numbers and $\mathbf{R}^d$ is the set of d-dimensional real vectors. 0 is used both for the real number 0 and the vector $(0, \ldots, 0)$. The point $(0, \ldots, 0)$ is called the **origin**.

Two sets of points are **isometric** if there is a 1-1 onto mapping between the two which preserves distance. Such a mapping is called an **isometric mapping** or **isometry**. A set of points in E^d which is isometric to E^k is known as an **affine subspace** of dimension k or a **k-flat**. **Points** are 0-flats, **lines** are 1-flats and **planes** are 2-flats. $d-1$-flats in E^d are known as **hyperplanes**. By convention, the empty set is a flat of dimension -1. If X is a set of points, the **affine hull** of X is the smallest k-flat containing X. The **dimension** of X is the dimension of the affine hull of X.

A hypersphere is the set of all points a fixed positive distance from some point x in E^d. The unit hypersphere in E^d, denote S^{d-1}, is the set of all points-distance one from the origin. A circle is a hypersphere in E^2 and a sphere is a hypersphere in E^3.

Often the homogeneous coordinates of points in E^d are useful. Homogeneous coordinates are $d+1$ dimensional vectors, $(\alpha_0, \alpha_1, \ldots, \alpha_d)$, $\alpha_i \in \mathbf{R}$, where $\alpha_0 \neq 0$. The point x with Cartesian coordinates $(\alpha_1, \alpha_2, \ldots, \alpha_d)$ is assigned to the set of vectors $(\lambda, \lambda\alpha_1, \lambda\alpha_2, \ldots, \lambda\alpha_d)$ over all $\lambda \in \mathbf{R}$, $\lambda \neq 0$. Any one of these $d+1$-vectors represents the point x in homogeneous coordinates. Given the homogeneous coordinates $(\alpha_0, \alpha_1, \alpha_2, \ldots, \alpha_d)$, $\alpha_0 \neq 0$, of a point $x \in E^d$, the Cartesian coordinates of the points are $(\frac{\alpha_1}{\alpha_0}, \frac{\alpha_2}{\alpha_0}, \ldots, \frac{\alpha_d}{\alpha_0})$. Sometimes the homogeneous coordinates of points in E^d will be restricted to $(\alpha_0, \alpha_1, \ldots, \alpha_d)$ where $\alpha_0 > 0$.

Both Cartesian and homogeneous coordinates are useful for describing points in E^d. Thus we will define many objects and properties in both Cartesian and homogeneous coordinates. It is left to the reader to show that these definitions are equivalent.

Lines and hyperplanes have many different coordinate parameterizations. If $x = (\alpha_1, \ldots, \alpha_d)$ and $y = (\beta_1, \ldots, \beta_d)$ are two points in E^d represented in Cartesian coordinates, the line through x and y is the set of all points of the form $\lambda x + \mu y$ where $\lambda + \mu = 1$. The closed line segment or just line segment $\overline{xy}$ joining x and y is the set of all points of the form $\lambda x + \mu y$ where $\lambda \geq 0$, $\mu \geq 0$ and $\lambda + \mu = 1$. The open line segment joining x and y is $\overline{xy}$ minus its endpoints, x and y. Lines can be parametrized by $u\gamma + x$, where $u \in \mathbf{R}^d$, $x \in E^d$, and γ varies over all the reals. In E^2 the parameterization of a line is $(\alpha_1, \alpha_2)\gamma + (\beta_1, \beta_2)$. $\frac{\alpha_2}{\alpha_1}$ is known as the slope.

If $x = (\alpha_0, \alpha_1, \ldots, \alpha_d)$ and $y = (\beta_0, \beta_1, \ldots, \beta_d)$ are two points in E^d represented in homogeneous coordinates, then the **line** through x and y is the set of all points of the form $\lambda x + \mu y$, $\lambda \neq 0$ or $\mu \neq 0$. With the additional restriction that $\alpha_0 > 0$ and $\beta_0 > 0$, the **line segment** $\overline{xy}$ is the set of points of the form $\lambda x + \mu y$ where $\lambda \geq 0$, $\mu \geq 0$ and $\lambda \neq 0$ or $\mu \neq 0$.

In Cartesian coordinates a hyperplane h is the set $\{x : u \cdot x = c, u \in R^d, c \in R, \textit{and } u \neq 0\}$. The vector u is said to be a **normal** to the hyperplane. In homogeneous coordinates a hyperplane h is the set $\{x : u \cdot x = 0, u \in R^{d+1}$ and $u \neq 0\}$. For convenience, we sometimes omit the conditions on u and c.

A hyperplane h divides E^d into two parts known as **half-spaces.** h **bounds** these two half-spaces. These half-spaces or **closed half-spaces** are parametrized as $\{x : u \cdot x \geq c\}$ and $\{x : u \cdot x \leq c\}$. In homogeneous coordinates these half-spaces are $\{x : u \cdot x \geq 0\}$, and $\{x : u \cdot x \leq 0\}$. A half-space in E^2 is called a **half-plane.** An **open half-space** is a closed half-space minus its bounding hyperplane h. A set H of hyperplanes divides Euclidean space into connected regions called **cells.** Such a division is called an **arrangement of hyperplanes.**

If X and Y are two sets of points in E^d, the hyperplane h **separates** X from Y if X lies in one half-space bounded by h while Y lies in the other. If, in addition, X and Y do not intersect h, then h **strictly separates** X from Y. If X lies in one half-space bounded by h and h intersects X, then h is a **supporting hyperplane** for X. Two sets of points X and Y can be **separated** or **strictly separated** if there exists a hyperplane which separates or strictly separates X from Y.

$u \cdot x$ is a polynomial of degree one in x. Hyperplanes are sets of points whose Cartesian coordinates solve the equation $f(x) = 0$ where f is a polynomial of degree one.

The set of points $\{x : f(x) = 0, f$ is a polynomial of degree two $\}$ is called a **quadric surface**. In homogeneous coordinates a **quadric surface** is $\{x : f(x) = 0,$ all terms of f have degree two $\}$. A quadric surface is **ruled** if it can be decomposed into a set of lines. A quadric surface is **doubly ruled** if there are two such decompositions.

Let X be a set of points in E^d. The **upper envelope** of X is the set of points $\{(\beta_1, \ldots, \beta_{d-1}, \beta_d) : (\beta_1, \ldots, \beta_{d-1}, \beta_d) \in X$ and $\beta_d \geq \gamma$ for all $(\beta_1, \ldots, \beta_{d-1}, \gamma) \in X\}$. **Ackerman's function** is defined inductively as follows:

$$A_1(n) = 2n,$$

$$A_m(1) = 2,$$

$$A_m(n) = A_{m-1}(A_m(n-1)), \quad m \geq 2, n \geq 2.$$

The **inverse of Ackerman's function** is $\alpha(n) = min(i : A_i(i) \geq n)$. Sharir and Hart proved that the upper envelope of n line segments is composed of at most $O(n\,\alpha(n))$ line segments[52].

X is a **convex** set if for each pair of points x and y in X, the line segment $\overline{xy}$ is a subset of X. Any two disjoint convex sets can be separated by a hyperplane. If X is a set of points, the **convex hull** of X, $conv(X)$, is the smallest convex set containing X. For a family A of sets of points, $conv(A)$ will represent the convex hull of the union of all the sets in A.

There are three related theorems by Helly, Radon and Carathéodory on convex sets.

Helly's Theorem. There exists a point which intersects every member of a family of compact convex sets in E^d if and only if the intersection of every $d+1$ sets is non-empty.

Carathéodory's Theorem. Let X be a set of points in $\mathbf{E}^d$. Each point in $conv(X)$ is the convex combination of $d+1$ or fewer points of X.

Radon's Theorem. Each set of $d+2$ or more points in $\mathbf{E}^d$ can be expressed as the union of two disjoint sets whose convex hulls have a common point.

A decomposition of a set X of points in $\mathbf{E}^d$ into two disjoint sets, $X = Y \cup Z$, $Y \cap Z = \emptyset$, such that $conv(Y) \cap conv(Z) \neq \emptyset$ is called a **Radon partition**.

The convex hull of a finite set of points is called a **convex polytope**. Since all polytopes in this thesis are convex, we will simply refer to convex polytopes as polytopes. The **dimension of a polytope** is the dimension of the smallest k-flat which contains the polytope. A **convex polygon** is a polytope of dimension two. Sometimes line segments and points will also be considered convex polygons, albeit degenerate ones. A **polyhedral set** is the intersection of a finite number of half-spaces. All polytopes are polyhedral sets although not all polyhedral sets are polytopes.

If p is a polytope, there is a smallest set of points whose convex hull is p. A point in this set is known as a **vertex** of p. A k**-dimensional simplex** is a polytope with $k+1$ vertices and dimension k. A **line segment** is a 1-dimensional simplex, a **triangle** is a 2-dimensional simplex and a **tetrahedron** is a 3-dimensional simplex.

A **face** of p is a set of points which lies in the intersection of p and some supporting hyperplane of p. Each face has a dimension which is the dimension of the smallest k-flat which contains the face. A face of dimension k is known as a k-face. 0-faces are just the vertices of p. 1-faces are known as **edges**. If p is a d-dimensional polytope, the **facets** of p are the $d-1$-faces of p.

A hyperplane h is defined as $\{x : u \cdot x = 0\}$ in homogeneous coordinates. If $u = \{1, \alpha_1, \ldots, \alpha_d\}$, then replacing u by $\{\lambda, \lambda\alpha_1, \ldots, \lambda\alpha_d\}$ for any $\lambda \in \mathbf{R}$, $\lambda \neq 0$, will define the same hyperplane. Thus h can be associated with the set of vectors $\{\lambda, \lambda\alpha_1, \ldots, \lambda\alpha_d\}$. This is quite similar to the definition of points in homogeneous coordinates and suggests a mapping between the two. The **dual map** D maps points in $\mathbf{E}^d$ to hyperplanes and hyperplanes to points. For a point x in homogeneous coordinates in $\mathbf{E}^d$, $D(x) = \{u : x \cdot u = 0\}$. For a hyperplane $h = \{x : u \cdot x = 0\}$, $D(h) = u$, the point with homogeneous coordinates u. $D(D(x)) = x$ and $D(D(h)) = h$, so D is its own inverse. Conveniently, D preserves incidence relations between points and hyperplanes. If x lies on h, then $D(h)$ lies on $D(x)$. D is undefined on the origin and on hyperplanes through the origin.

A set of points is **collinear** if all of the points lie on a line. A set of points is **coplanar** if all the points lie on a plane. A set of points in $\mathbf{E}^d$ is in **general position** if no k of the points are contained in a $k-2$-flat, $2 \le k \le d+1$. A set of hyperplanes in $\mathbf{E}^d$ is in general position if the intersection of any k of the hyperplanes is contained in a $d-k$-flat, $2 \le k \le d+1$.

Let $(x_1, \ldots, x_{d+1})$ be an ordered set of $d+1$ points in $\mathbf{E}^d$ represented by their Cartesian coordinates where $x_i = (\alpha_{i,1}, \ldots, \alpha_{i,d})$, $\alpha_{i,j} \in \mathbf{R}$. Let M be the $d+1 \times d+1$ matrix

$$\begin{bmatrix} 1 & \alpha_{1,1} & \alpha_{1,2} & \cdots & \alpha_{1,d} \\ 1 & \alpha_{2,1} & \alpha_{2,2} & \cdots & \alpha_{2,d} \\ . & . & . & \cdots & . \\ 1 & \alpha_{d+1,1} & \alpha_{d+2,2} & \cdots & \alpha_{d+1,d} \end{bmatrix}.$$

***det*(*M*)** is the determinant of M. The **orientation** of $(x_1, \ldots, x_{d+1})$ is the sign of ***det*(*M*)** or ***sgn*(*det*(*M*))**. Note that the orientation depends upon the order of the points in

$(x_1, \ldots, x_{d+1})$. In $\mathbf{E}^2$ the orientation of (x_1, x_2, x_3) is equivalent to whether someone traveling from x_1 to x_2 and then to x_3 makes a left turn, a right turn, or continues straight. If $u_1, \ldots, u_n$ are n vectors of length n, then $det(u_1, \ldots, u_n)$ is the determinant of the $n \times n$ matrix with rows u_1 through u_n. The **orientation** of $(x_1, \ldots, x_{d+1})$ can then be described as $sgn(det(x_2 - x_1, x_3 - x_1, \ldots, x_{d+1} - x_1))$. If $x_1, \ldots, x_{d+1}$ are given in homogeneous coordinates, $x_i = (\alpha_{i,0}, \alpha_{i,1}, \ldots, \alpha_{i,d})$, $\alpha_{i,0} > 0$, then the **orientation** of $(x_1, \ldots, x_{d+1})$ is $sgn(det(x_1, \ldots, x_{d+1}))$.

k-flats can also be given an orientation. An isometry Γ from a k-flat to $\mathbf{E}^k$ gives the k-flat an **orientation** by specifying the orientation of every $k+1$ points in general position in the k-flat. The orientation of the $k+1$ points is their orientation under Γ. The isometries from the k-flat to $\mathbf{E}^k$ can be divided into two equivalence classes where each equivalence class gives the k-flat the same orientation. If two isometries Γ and Γ' are in the same equivalence class, then the orientation of $k+1$ points under Γ is the same as the orientation of $k+1$ points under Γ'. If two isometries Γ and Γ' are in different equivalence classes, then the orientation of $k+1$ points under Γ is the opposite of the orientation of $k+1$ points under Γ'.

Lines can also be given orientations or directions based upon their parametrization. A line parametrized by $u\gamma + x$ can be given a **direction** u and is then referred to as a **directed line**. A point $y = u\gamma + x$ comes before a point $y' = u\gamma' + x$ if $\gamma < \gamma'$. Under a suitable isometry of the line to $\mathbf{R}$, this is equivalent to saying that the orientation of (y, y') is positive.

Hyperplanes can be given an orientation based upon their normals. A hyperplane $h = \{x : u \cdot x = c\}$ in Cartesian coordinates has a normal $u = (\beta_1, \ldots, \beta_d)$. If $x_1, \ldots, x_d$

are d points in h, $x_i = (\alpha_{i,1}, \ldots, \alpha_{i,d})$, then u gives them the **orientation** $sgn(det(M))$ where

$$M = \begin{bmatrix} 1 & \beta_1 & \beta_2 & \cdots & \beta_d \\ 1 & \alpha_{1,1} & \alpha_{1,2} & \cdots & \alpha_{1,d} \\ 1 & \alpha_{2,1} & \alpha_{2,2} & \cdots & \alpha_{2,d} \\ . & . & . & \cdots & . \\ 1 & \alpha_{d,1} & \alpha_{d,2} & \cdots & \alpha_{d,d} \end{bmatrix}.$$

Let a set of points X be represented in Cartesian coordinates in $\mathbf{E}^d$. X is **centrally symmetric** if $-x \in X$ whenever $x \in X$. Two isometric mappings of $\mathbf{E}^d$ to itself are of particular interest. X' is a **reflection** of X about the hyperplane h if $X' = \{y-\lambda u : y+\lambda u \in X, y \in h, \lambda \in \mathbf{R}, u \text{ is a normal to } h\}$. X' is a **translation** of X if $X' = \{x+v : x \in X\}$ for some fixed $v \in \mathbf{R}^d$. Two k-flats are **parallel** if one is a translation of the other.

Parallel k-flats must often be handled as special cases for theorems and algorithms in Euclidean space. Projective space avoids this annoying property of parallelism. Divide the lines in $\mathbf{E}^d$ into disjoint sets of parallel lines. **Projective space** is constructed from Euclidean space by adding a point for each such set of parallel lines. These new points are called **improper points**. All the lines in a given set of parallel lines are extended to include the improper point corresponding to that set. These extended lines are called **projective lines**.

There is no notion of distance in $\mathbf{P}^d$, so isometric mappings which preserve distance make no sense in $\mathbf{P}^d$. Instead a **projective mappings** is a 1-1 onto mapping of a set of points to a set of points which preserves lines. A set of points in $\mathbf{P}^d$ which is some projective map of $\mathbf{E}^k$ is known as a k**-flat** in projective space. $d-1$-flats in $\mathbf{P}^d$ are called **hyperplanes**. Whether a k-flat or hyperplane is Euclidean or projective will be understood from the context. Every two lines in $\mathbf{P}^2$ intersect in a point and every two planes in $\mathbf{P}^3$ intersect in a line. This generalizes to every two hyperplanes in $\mathbf{P}^d$ intersect in a $d-2$-flat.

The points in $\mathbf{P}^d$, d-dimensional projective space, are coordinatized using **homogeneous coordinates**. Each point x in $\mathbf{P}^d$ corresponds to the set of vectors $(\lambda\alpha_0, \lambda\alpha_1, \ldots, \lambda\alpha_d)$ over all $\lambda \in \mathbf{R}$, $\lambda \neq 0$, $(\alpha_0, \alpha_1, \ldots, \alpha_d) \neq 0$. The homogeneous coordinates $(\alpha_0, \alpha_1, \ldots, \alpha_d)$ where $\alpha_0 \neq 0$ correspond to both a point in Euclidean space and a point in projective space. This isomorphism defines a natural embedding of Euclidean space into projective space. The homogeneous coordinates $(\alpha_0, \alpha_1, \ldots, \alpha_d)$ where $\alpha_0 = 0$ correspond to the improper points in projective space.

If $x = (\alpha_0, \ldots, \alpha_d)$ and $y = (\alpha_0, \ldots, \alpha_d)$ are points in $\mathbf{P}^d$, the **projective line** through x and y is parameterized by $\lambda x + \mu y$, $\lambda \neq 0$ or $\mu \neq 0$. A **projective hyperplane** is $\{x : u \cdot x = 0, u \in \mathbf{R}^{d+1}$ and $u \neq 0\}$. A **quadric surface** is $\{x : f(x) = 0,$ all terms of f have degree two $\}$. The points x and y divide the projective line through x and y into two connected pieces. These pieces are **projective line segments** with endpoints x and y.

A little topology is necessary to describe some properties of sets of points. A ball around a point $x \in \mathbf{E}^d$ is the set of all points whose distance from x is less than some fixed radius $r > 0$. An **open set** in $\mathbf{E}^d$ is the union of some collection of balls. A **closed set** Y in $\mathbf{E}^d$ is the complement of some open set X, $Y = \mathbf{E}^d - X$. Closed line segments are closed sets in $\mathbf{E}^1$ while open line segments are open sets in $\mathbf{E}^1$. Closed half-spaces are closed sets while open half-spaces are open sets.

If X is a subset of $\mathbf{E}^d$, an **open set** in X is a set $X \cap Y$, where Y is an open set in $\mathbf{E}^d$. The **open sets** in $\mathbf{P}^d$ are unions of projective transformations of open sets in $\mathbf{E}^d$. If X is a subset of $\mathbf{P}^d$, an **open set** in X is a set $X \cap Y$, where Y is an open set in $\mathbf{P}^d$.

For a function f, a point x and a set of points X, $f^{-1}(x) = \{y : f(y) = x\}$ and $f^{-1}(X) = \{y : f(y) \in X\}$. A function f is **continuous** if $f^{-1}(X)$ is an open set whenever X is an open set. In $\mathbf{E}^d$ a function f is continuous if and only if $\lim_{x \to a} f(x) = f(a)$.

The **interior** of a set X, $int(X)$, is the largest open set contained in X. The **closure** of a set X, $cl(X)$, is the smallest closed set containing X. The **boundary** of X, $bd(X)$, is $cl(X) - int(X)$. The **relative interior** of a set $X \subseteq \mathbf{E}^d$, $relint(X)$, where X has dimension k, is the interior of X when X is embedded in $\mathbf{E}^k$.

A set is **bounded** if there is some ball which contains the set. A **compact set** is a set which is closed and bounded. A **covering** of a set X is a collection of sets whose union contains X. This collection is said to **cover** X. An **open covering** of a set X is a collection of open sets which cover X. A **subcovering** of a set X is some subset of a covering of X which also covers X. The subcovering is finite if it has a finite number of elements. The **Heine-Borel Theorem** states that any open covering of a compact set has a finite subcovering. The proof is available in any standard topology text.

Two theorems which use topological properties are worth mentioning. A polytope is a bounded polyhedral set. Any two compact convex sets can be strictly separated by a hyperplane. See any standard text on convex sets for proofs.

A set of points is called **disconnected** if it can be split into two disjoint open sets Y, Y', neither of which is empty. A set of points which is not disconnected is **connected**. A hypersphere is connected. **Arcs** are connected subsets of the circle. **Closed arcs** are arcs which contain their endpoints while **open arcs** are arcs which do not contain their endpoints.

Finally, some terminology from graph theory is useful. A **graph** is a collection of **vertices** and **edges**, where each edge is **incident** with two vertices. Two vertices are **adjacent** or **neighbours** if there is an edge between them, i.e. they are incident with the same edge. A graph where every edge is incident with two distinct vertices and no two edges are incident with the same two vertices is known as a **simple graph**. A **clique** is a subset of the vertices in a graph where every two vertices in the clique are adjacent. An **embedding** of a graph is a drawing of the graph on a surface such that edges only intersect at their endpoints. A **straight line embedding** is an embedding where all edges are drawn as line segments. A **planar graph** is a graph which can be embedded in the plane. A simple planar graph always has a straight line embedding[23]. A straight line embedding of a planar graph is also called a **planar subdivision**.

An embedding of a graph in the plane divides the plane into connected regions called **faces**. One of these connected regions is not bounded and is called the **external face**. All other faces are **internal faces**. Each face is surrounded by vertices and edges which form its boundary. A planar subdivision is a **triangulation** if each internal face is a triangle. A **complete triangulation** is a triangulation where every face, including the external one, is bounded by three edges. Any face of a planar subdivision can be broken into triangles, or **triangulated**, by the addition of line segments between vertices on the boundary of the face. Similarly a planar subdivision can be **triangulated** by triangulating each of its faces. **Euler's formula** states that if n is the number of vertices in a planar graph, m is the number of edges, and f is the number of faces, then $n+f-m=2$. In a complete triangulation, $2m=3f$ and $m=3n-6$.

This thesis will loosely adhere to some notational conventions. **R** is the set of reals and $\mathbf{R}^d$ is the set of d-dimensional real vectors. Euclidean d-space and projective d-space are denoted $\mathbf{E}^d$ and $\mathbf{P}^d$, respectively. $\mathbf{S}^{d-1}$ is the unit hypersphere in $\mathbf{E}^d$. Lower case greek letters will be used for real numbers and u,v,w for real vectors. d is the dimension of Euclidean or projective space. Points in Euclidean or projective space will be labelled x,y or z. Lines will be called l, line segments s, hyperplanes h, flats g, and polytopes p. Half-spaces and half-planes will be labelled h^+ and h^-. Sets of objects will generally be in capital letters. X,Y,Z will be sets of points, L,S,K,H,P will be sets of lines, line segments, arcs, hyperplanes and polytopes, respectively. E will be used for sets of edges and V for sets of vertices. Graphs are labelled G. One exception to sets being capitalized is convex sets which will be labelled a,b,c. a,b,c will also be used for connected sets of points. Families of convex sets will be A,B,C. All objects can be subscripted or primed to refer to many different objects of the same type. $(x_1, \ldots, x_n)$ is the set of n points from x_1 to x_n. $(x_1, \ldots, \hat{x}_i, \ldots, x_n)$ is the set of $n-1$ points from x_1 to x_n excluding x_i. D is the dual map which takes points to hyperplanes and hyperplanes to points. By standard convention, $A_m(n)$ is Ackerman's function and $\alpha(n)$ is the inverse of Ackerman's function. $|S|$ is used for the size of set S.

Chapter 3

Hadwiger's Theorem and Generalizations

3.1 Introduction

Let A be a family of compact convex sets in E^d. A k-**transversal** or k-**stabber** of A is a k-flat which intersects every member of A. 1-transversals, 2-transversals and $d-1$-transversals in E^d will also be called **line transversals, plane transversals** and **hyperplane transversals**, respectively. If the line intersecting every member of A is directed, then we will refer to it as a **directed line transversal**. Similarly, if the k-flat intersecting every member of A is oriented, we will refer to it as an **oriented** k-**transversal**. What are necessary and sufficient conditions for the existence of a k-transversal of A?

In 1913, Helly proved his famous theorem on convex sets which is restated here using the terminology of transversals[34].

Helly's Theorem. A family of compact convex sets in E^d has a 0-transversal if and only if every $d+1$ of the sets have a 0-transversal.

Helly's Theorem makes no restriction that the family of compact convex sets be finite. The theorem also applies to finite families of convex sets which are not necessarily compact.

Helly's Theorem is so useful because a property of the entire family can be found by only looking at small subsets of the family. In fact, if there exists a polynomial time algorithm for determining whether every $d+1$ sets in some family of convex sets have a 0-transversal, then Helly's Theorem gives a polynomial time algorithm for determining whether the entire family has a 0-transversal. Simply examine every $d+1$ subsets of the family to determine if they have a 0-transversal. This algorithm runs in polynomial time for fixed d.

After Helly published his theorem, mathematicians started looking for theorems with a similar form. They wanted to show that some property was true for a set A if and only it was true for every m elements of A. For line transversals this produces the conjecture that a family A of compact convex sets in E^2 have a line transversal if and only if every m of them have one, for some fixed m. Unfortunately, this conjecture is false. Figure 3.1, consisting of a family of four line segments and a point, is a counterexample for $m = 4$. It is a modification of an example by Lewis[43]. Any line transversal must go through the point but any line through the point misses one of the lines The four line segments have a line transversal and every three line segments and the point have a line transversal. It is easy to see how this example may be extended for any m.

A directed line transversal for a family of pairwise disjoint compact convex sets generates an ordering on those sets. A directed line transversal is **consistent** with some ordering of the family if the order generated by the transversal is the same as the given ordering. In 1957, Hadwiger provided the following necessary and sufficient condition for line transversals in E^2[31].

Hadwiger's Theorem. A family of pairwise disjoint compact convex sets in E^2 has a line transversal if and only if there exists some ordering of the family such that every three of the sets have a directed line transversal consistent with the ordering.

Like Helly's Theorem, Hadwiger's Theorem is not restricted to finite families. However, the condition that the sets are compact and pairwise disjoint implies that the families must be

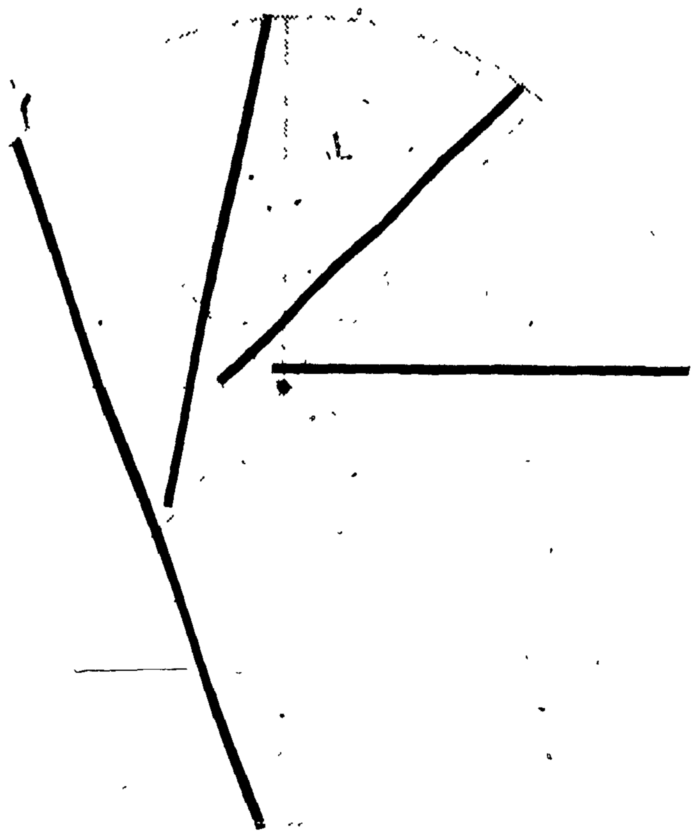

Figure 3.1. Five Convex Sets with No Line Transversal

countable.

Hadwiger's Theorem includes an ordering condition. This ordering condition makes it difficult to turn Hadwiger's Theorem into a polynomial algorithm for finding a line transversal on a family of size n because a-priori there could be up to $n!$ orderings to check. As we shall see in Chapter 4, there are only a polynomial number of orderings in which a line can intersect a family of n compact convex sets.

Hadwiger's Theorem only gives necessary and sufficient conditions for families of pairwise disjoint compact convex sets. What of families in which some sets intersect? Section 3.2 gives a generalization of Hadwiger's Theorem to necessary and sufficient conditions for the existence of a line transversal for any family of compact convex sets in E^2. In fact, the theorem applies to the even more general category of families of compact connected sets in E^2. The theorem applies to infinite families of compact connected sets, even those that are not countable.

Hadwiger's Theorem does not give any information about the order in which the line transversal intersects the entire family. One conjecture is that an ordered family of pairwise disjoint compact convex sets in E^2 has a directed line transversal consistent with the ordering if and only if every three of the sets have a directed line transversal consistent with the ordering. This conjecture is false as can be seen from Figure 3.2, where the ordering is (a, b, c, d). However, if we replace the condition that every three sets have a line transversal consistent with the ordering with the condition that every four sets have such a transversal, then the conjecture is true. Section 3.2 provides a proof of this theorem and its generalization to families of compact convex sets which are not necessarily pairwise disjoint.

Lines are hyperplanes in E^2. In 1980, Katchalski gave a sufficient condition for the existence of a hyperplane transversal for a family of pairwise disjoint compact convex sets in E^d [37].

Katchalski's Theorem. A family of pairwise disjoint compact convex sets in E^d has a hyperplane transversal if there exists some ordering of the family such that every three of the sets have a directed line transversal which intersects them consistent with the ordering.

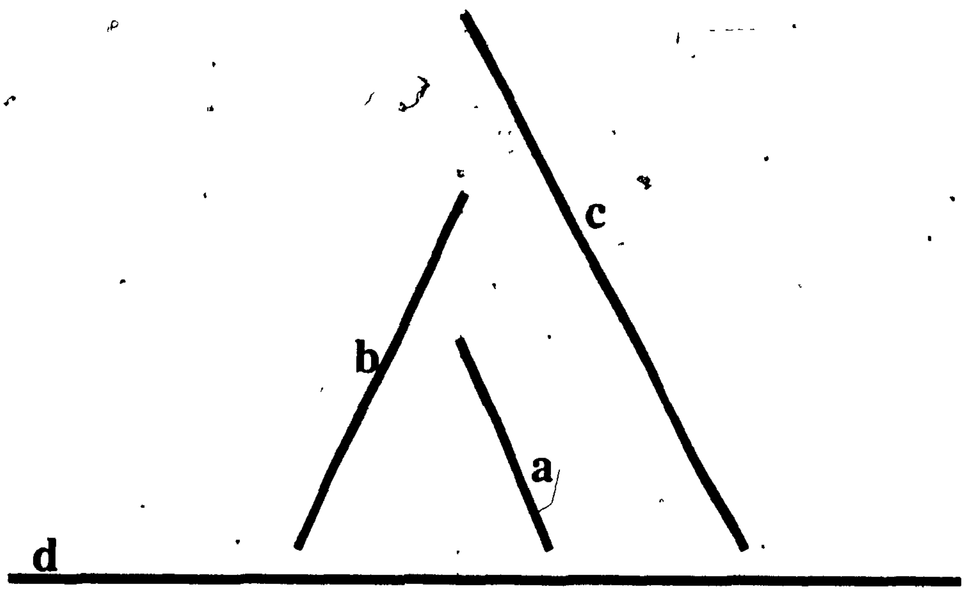

Figure 3.2. Four Line Segments with No Line Transversal in Order "abcd"

The family in Katchalski's Theorem may be infinite but must be countable. Katchalski's Theorem is only a sufficient, not a necessary, condition for the existence of a hyperplane transversal. It also includes the pairwise disjointness condition of Hadwiger's Theorem.

In 1986, Goodman and Pollack were able to provide necessary and sufficient conditions for the existence of hyperplane transversals[29]. To do so, they needed to generalize the ideas of ordering and pairwise disjointness. The **order type** of a set of points in E^d is the family of orientations of its $d+1$ tuples[27]. The order type of points in E^1 is the family of relative orders in which any two points lie on E^1. A **k-ordering** of a set is the order type produced by the association of each element of the set with some point in E^k.

A family A of sets of points is **k-separable** if every j sets can be strictly separated by a hyperplane from every other $k+2-j$ sets in A, $1 \le j \le k+1$. A family is 0-separable if every set can be strictly separated from every other set and a family is 1-separable if every two sets can be strictly separated from every other set. **Separable** may be used as a synonym for 0-separable. If a family is k-separable, then there exists no k-transversal for any $k+2$ sets in the family. For a family of convex sets, this is a necessary and sufficient condition. A family of convex sets is k-separable if and only if there exists no k-transversal for any $k+2$ sets in the family. Proofs follow the argument in Lemma 3.7 and are left to the reader. Associating each set a in a $d-1$-separable family in E^d with some point in a generates a d-ordering on the family. Since the family is $d-1$-separable, this d-ordering is independent of the choice of points and is unique.

An oriented k-flat g which intersects a $k-1$-separable family generates a k-ordering by associating each set a with some point in $a \cap g$. As with d-orderings above, this k-ordering is independent of the choice of points. An oriented k-transversal is **consistent** with a k-ordering of a $k-1$-separable family of convex sets if the k-ordering generated by the transversal is the same as the given k-ordering.

Goodman and Pollack's Theorem. A $d-2$-separable family of compact convex sets in E^d has a hyperplane transversal if and only if there exists a $d-1$-ordering of the family such that every $d+1$ sets have an oriented hyperplane transversal which is consistent with the $d-1$-ordering.

This theorem is also a generalization of a result by Valentine[56]. Again the theorem is only true for certain families in E^d, those which are $d-1$-separable. These families may be infinite but must be countable.

Both Katchalski's result and Goodman and Pollack's result are generalizations of Hadwiger's Theorem to hyperplane transversals. However, neither of these generalizations contains the other. In fact, these two theorems apply to different families of convex sets. Families in Katchalski's Theorem have the property that every three have a line transversal whereas $d-2$-separability for families in Goodman and Pollack's Theorem implies that no three have a line transversal when $d \geq 3$. Hadwiger's Theorem with suitable modifications is true for all families of compact convex sets, not only families which are pairwise disjoint or 0-separable. This suggests that Goodman and Pollack's Theorem could be generalized to all families also, not only $d-1$-separable ones. Unfortunately, I was unable to prove such a generalization. Instead R. Pollack and I proved a theorem about hyperplane transversals for k-separable families where k is not restricted to 0 or $d-2$. This result, presented in Section 3.3, includes Katchalski's result where $k = 0$ and Goodman and Pollack's result where $k = d-2$. The proof uses many of the techniques presented in their papers.

3.2 Permitting Intersections

Hadwiger's Theorem contains the condition that every three convex sets have a directed line transversal consistent with some ordering. If convex sets are permitted to intersect, then it is no longer clear when a transversal intersects the sets "consistent" with a given ordering. Furthermore, if compact convex sets are replaced by compact connected sets, then a transversal may intersect these sets more than once. Again it is unclear in what order a transversal intersects these sets. To generalize Hadwiger's Theorem, we must generalize this notion of "consistency". We do this by ignoring pairs of sets which are not separable.

Let A be a family of compact connected sets in $\mathbf{E}^2$. A directed line l is **consistent** with an ordering if, for every separable pair of sets a and b intersected by l, a precedes b in the ordering if and only if l intersects a before b. Note that if a and b intersect, l may intersect b before a and still be consistent with the ordering ab. The following theorem generalizes Hadwiger's Theorem to families of connected sets in $\mathbf{E}^2$.

Theorem 3.1. A family of compact connected sets in $\mathbf{E}^2$ has a line transversal if and only if there exists some ordering of the family such that every three of the sets have a directed line transversal consistent with the ordering.

A few lemmas precede the proof of Theorem 3.1.

Let A be a family of compact connected sets in $\mathbf{E}^d$. A normal to a hyperplane transversal for A is called a **stabbing normal** for A.

Lemma 3.1. A family of compact connected sets in $\mathbf{E}^d$ has a hyperplane transversal if and only if there exists some non-zero vector which is a stabbing normal for every pair of sets in the family.

Proof: If a family has a hyperplane transversal, then a normal to that hyperplane transversal is a stabbing normal for every pair of sets in the family. Assume there exists some non-zero vector v which is a stabbing normal for every two sets. Project the sets onto a line with direction v. The sets are compact and connected, so the projection of each set is a closed line segment. Every pair of these line segments intersect. By Helly's Theorem there exists some point in the intersection of all the line segments. The hyperplane through this point with normal v is a hyperplane transversal for the family. ■

If an oriented hyperplane with normal v strictly separates two sets, a, b, then v is a separation normal for a and b.

Lemma 3.2. Any non-zero vector in $\mathbf{E}^d$ is either a stabbing normal or a separation normal for a pair of compact connected sets but not both.

Proof: Let a, b be two compact connected sets and let v be a non-zero vector in $\mathbf{E}^d$. Project a and b onto a line with direction v. The two projections form two closed line segments. If the line segments intersect at some point, the hyperplane through this point with normal v intersects a and b. If the line segments do not intersect, there is some point on the line which separates the two. The hyperplane through this point with normal v separates a from b. ■

Lemma 3.3. There exists a hyperplane transversal for a family of compact connected sets in $\mathbf{E}^d$ if and only if the union of the sets of separation normals over all pairs of sets in the family does not cover the unit hypersphere $\mathbf{S}^{d-1}$.

Proof: By Lemma 3.1 there exists a hyperplane transversal if and only if there exists some non-zero vector which is a stabbing normal for every pair of sets. By Lemma 3.2 there exists a non-zero vector which is a stabbing normal for every pair of sets in the family if and only if that normal is not a separation normal for any pair of sets. Some normal is not a separation normal for any pair of sets if and only if the separation normals over all pairs of sets do not cover $\mathbf{S}^{d-1}$. ■

Let A be a family of compact convex sets. If every pair of sets in A intersect then no non-zero vector is a separation normal for any pair. By Lemma 3.1 every non-zero vector is a stabbing normal for A. If every non-zero vector is a stabbing normal for A, then no non-zero vector separates any pair of sets in A by Lemma 3.2. Thus every pair of sets in a family A of compact convex sets intersect if and only if every non-zero vector is a stabbing normal for the pair.

Lemmas 3.1, 3.2 and 3.3 make no assumption that the families are finite. The following lemma reduces the problem of hyperplane transversals for infinite families to the problem of hyperplane transversals for finite families. It is a simple application of the Heine-Borel Theorem, any open covering of a compact set has a finite subcovering.

Lemma 3.4. There exists a hyperplane transversal for a family of compact connected sets in E^d if and only if there exists a hyperplane transversal for every finite subfamily of the family.

Proof: Let A be a family of compact connected sets in E^d. If A has a hyperplane transversal, then every subfamily of A has the same hyperplane transversal. Assume there exists no hyperplane transversal for A. By Lemma 3.3 the separation normals for all pairs of sets in A cover S^{d-1}. For each pair of sets $a, b \in A$, let c_{ab} be the set of all separation normals for the pair a, b. c_{ab} is an open set in E^d. Let C be a collection of all such sets of separation normals. C is an open covering of S^{d-1} and S^{d-1} is a compact set. By the Heine-Borel Theorem, C has a finite subcovering C'. Let $A' = \{a : c_{ab} \in C'$ for some $b \in A\}$. Since C' is finite, A' is finite. C' covers S^{d-1}, so by Lemma 3.3 $A' \subseteq A$ has no hyperplane transversal. ∎

Lemmas 3.1, 3.2, and 3.3, have special formulations in E^2. Let A be a family of compact connected sets in E^2. The direction of a directed line transversal for A is called a **stabbing direction** for A. If a directed line l strictly separates two sets a, b, then the direction of l is a **separation direction** for a and b. Stabbing directions and separation directions are merely stabbing normals and separation normals in E^2 rotated ninety degrees. It is more convenient to use stabbing directions and separation directions in the proof of Theorem 3.1. In this terminology, Lemma 3.1 states that a family of compact connected sets in E^2 has a line transversal if and only if there exists some direction which is a stabbing direction for every pair of sets. By Lemma 3.2 a direction is either a stabbing direction or a separation direction for a pair of compact convex sets but not both. Lemma 3.3 is there exists a line transversal for a family of compact connected sets in E^2 if and only if the separation directions over all pairs of sets in the family do not cover the unit circle.

In E^2 all lines with the same direction intersect a separable pair of connected sets in the same order. Thus a stabbing direction induces an ordering on separable pairs, namely the order in which any directed line with the given direction stabs the pair.

Lemma 3.5. An ordered family of compact connected sets in E^2 has a directed line transversal consistent with the ordering if and only if there exists some direction which is a stabbing direction for every two sets and the induced ordering on separable pairs of sets is consistent with the given ordering.

Proof: If an ordered family has a directed line transversal consistent with the ordering, then the direction of the line transversal is the desired stabbing direction for every two separable sets. Assume that there exists some direction which is a stabbing direction for every two sets and the induced ordering on separable pairs is consistent with the given ordering. By Lemma 3.1, there exists a directed line with the given stabbing direction which stabs the family. Every two separable sets are stabbed by some translate of this directed line in the given order so the directed line must stab every two separable sets consistent with the given ordering. ■

Directions in E^2 can be mapped to points on the unit circle. The stabbing directions of two compact connected sets which are not separable map to the entire circle. The stabbing directions of two separable sets, a, b, map to two disjoint closed arcs on the circle. Each of these arcs can be associated with a different ordering of the sets, either ab or ba. The separation directions also map to two disjoint open arcs on the circle. Each of these arcs can be associated with a different ordering of the sets, either ab or ba, representing a to the left of b or b to the left of a. By Lemma 3.2 the circle is covered by the four arcs representing stabbing and separation directions and the intersection of any two of these arcs is disjoint. We are now ready for the proof of Theorem 3.1.

Proof of Theorem 3.1: By Lemma 3.4 if the theorem is true for any finite family then it is also true for infinite families. A family that has a line transversal has a directed line transversal. This directed line transversal generates an ordering on the family and intersects every three sets consistent with that ordering.

Let A be a finite family of compact connected sets. Assume that there exists some ordering of A such that every three sets in A have a directed line transversal consistent with the ordering. Let K be the finite set of all arcs corresponding to separation directions for every pair of separable sets in A. Associate with each arc in K a unique label ab or ba depending on whether a or b is to the left of separating lines with the given direction. Let

K_1 be the set of all arcs labelled ab where a precedes b in the given ordering and let K_2 be the set of all other arcs.

We wish to prove that the intersection of any arc in K_1 with any arc in K_2 is empty. Certainly the intersection of arcs ab and ba is empty. If arcs $ab \in K_1$ and $ab' \in K_2$ intersect, then there must be some line l separating a from b and b'. Since $ab \in K_1$ and $ab' \in K_2$, a precedes b and b' precedes a in the ordering. By assumption there must be some directed line which intersects b', a, b in that order. This directed line would have to cross line l twice, an impossibility. Therefore the intersection of arcs ab and ab' must be empty.

If arcs $ab \in K_1$ and $b'a \in K_2$ intersect, then there must be some line l separating b' from a and b and a parallel line l' separating b from a and b'. There must also be some directed line which intersects a and then b and b'. This line would have to intersect l or l' twice. Therefore the intersection of arcs ab and $b'a$ is empty. Similarly, the intersection of arcs ab and $a'b$ and the intersection of arcs ab and ba' are empty.

The case where arcs $ab \in K_1$ and $a'b' \in K_2$ intersect, a, b, a', b' distinct, reduces to the previous cases. Let l be the line separating a from b with direction represented by the point of intersection. Since there is a line parallel to l which separates a' from b', l also separates a from b' or b from a'. Assume l separates a from b'. Then arc ab' intersects ab. Since we proved that $ab \in K_1$ and $ax \in K_2$ cannot intersect, ab' must lie in K_1. Similarly, ab' intersects $a'b'$ and so ab' must lie in K_2, a contradiction. Therefore l does not separate a from b'. By the same reasoning, l does not separate b from a'. Therefore ab and $a'b'$ do not intersect and the intersection of any two arcs in K_1 and K_2 is empty.

Assume all the arcs in K cover the circle. K must not be empty, so K_1 and K_2 must not be empty. Let X_1 be all the points covered by arcs in K_1 and let X_2 be all the points covered by arcs in K_2. X_1 and X_2 are open sets in the circle, and $X_1 \neq \emptyset$, $X_2 \neq \emptyset$, $X_1 \cap X_2 = \emptyset$. Since the circle is connected, $X_1 \cup X_2$ must not cover the circle. We conclude that K does not cover the circle. By Lemma 3.3 there exists a line transversal for A. ■

The importance of ordering in Hadwiger's Theorem suggests the problem of finding necessary and sufficient conditions for the existence of a line transversal which intersects a family consistent with a given order. The following theorem gives such conditions for families of compact connected sets.

Theorem 3.2. An ordered family of compact connected sets in the plane has a directed line transversal consistent with the ordering if and only if every six sets have a directed line transversal consistent with the ordering.

Proof: If an ordered family has a directed line transversal consistent with the ordering, then that directed line transversal intersects every six sets consistent with the ordering. Let K be the set of all arcs corresponding to stabbing directions for pairs of disjoint sets which are consistent with the given ordering. Since every six sets can be intersected by some directed line consistent with the given ordering, every three arcs intersect. Since each arc has measure less than 180 degrees, the intersection of two arcs is still an arc. Choose one arc arbitrarily, say ab, and intersect each of the other arcs with this arc. Let K' be the set of new arcs formed. All arcs in K' lie on arc ab and every two arcs in K' intersect. Applying Helly's Theorem, there exists some point at the intersection of all the arcs. By Lemma 3.5 there exists a line transversal consistent with the given ordering. ■

The number six in Theorem 3.2 cannot be reduced as can be seen from Figure 3.3. Every five sets can be intersected consistent with the ordering *abcdef* but there is no line transversal for all six consistent with that ordering. The number six in Theorem 3.2 is replaced by four for families of compact convex sets which are pairwise disjoint.

Theorem 3.3. An ordered family of pairwise disjoint compact convex sets in the plane has a directed line transversal consistent with the ordering if and only if every four sets can be intersected by some directed line consistent with the ordering.

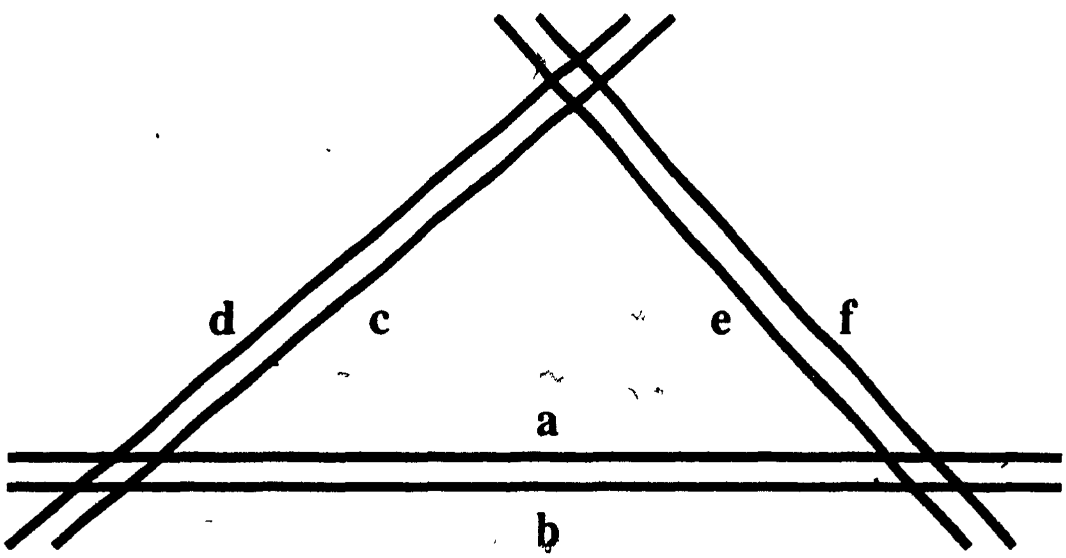

Figure 3.3. Six Line segments with No Line Transversal in Order "abcdef"

Proof: If an ordered family has a directed line transversal consistent with the ordering, then that directed line transversal intersects every four sets consistent with the ordering. We first show that Theorem 3.3 is true for a family of five sets. Let $\{a, b, c, d, e\}$ be a family of pairwise disjoint compact convex sets in the plane with the alphabetic ordering. For every four sets there is a directed line stabber consistent with the alphabetic order. Let x_{abcd} be the point on the circle corresponding to the stabber of *abcd*. In the same manner define points x_{abce}, x_{abde}, x_{acde} and x_{bcde}. Any three of these points are covered by some arc corresponding to a stabbing direction for a pair of ordered sets. For instance x_{abcd}, x_{abce} and x_{abde} are covered by the arc corresponding to the stabbing directions for ab. Thus every three of these points must lie within some half-circle. If the center of the circle lies within the convex hull of these five points, then it lies within the convex hull of three of the points by Carathéodory's Theorem, and these three points would not lie in a half-circle. It follows that all the points must lie in some half-circle θ.

Let K be the set of all arcs corresponding to stabbing directions for pairs of disjoint sets which are consistent with the alphabetic ordering. Intersect each of the arcs in K with the half-circle θ to form the set of arcs K'. The intersection of each pair of arcs in K contains one of the five points, x_{abcd}, x_{abce}, x_{abde}, x_{acde} or x_{bcde}. These points line in θ, so the intersection of each pair of arcs in K' is non-empty. By Helly's Theorem and Lemma 3.5, the intersection of all the arcs is non-empty and there exists a line transversal consistent with the alphabetic ordering.

We now show that Theorem 3.3 is true for a family of six sets, $\{a, b, c, d, e, f\}$ where every four sets have a stabber consistent with the alphabetic ordering. By the argument above, every five of the sets have a stabber consistent with the alphabetic ordering.

Choose the points x_{abcde}, x_{abcdf}, x_{abcef}, x_{abdef}, x_{acdef} and x_{bcdef} corresponding to the six stabbers. By the same argument as before, all these points lie in some half-circle θ, the pairs of arcs corresponding to stabbing directions intersect in θ and there is a stabber consistent with the alphabetic ordering.

Given any ordered family of pairwise disjoint compact convex sets in the plane we just showed that if every four sets can be intersected by some directed line consistent with the given ordering, then every six sets can be so intersected. By Theorem 3.2, there exists a line transversal for the entire family consistent with the given ordering. ∎

As shown before in Figure 3.2, the number four in Theorem 3.3 cannot be reduced.

Hadwiger, Debrunner and Klee used the intersection of arcs corresponding to stabbing directions for proving the following theorem, Proposition 27 in their book "Combinatorial Geometry in the Plane"[32]. Let A be a family of compact convex sets in the plane where every pair of sets in A can be strictly separated by a vertical line. A has a line transversal if and only if three sets in A have a line transversal. The theorem was originally proposed by P. Vincensini with the condition that every four sets in A have a line transversal[57]. V.L. Klee, Jr. improved the condition to every three sets have a line transversal[41].

The same arguments for generalizing Hadwiger's Theorem eliminate the pairwise disjointness condition in Katchalski's Theorem. Jacob E. Goodman and Richard Pollack observed that the arguments would allow an even more general statement of Katchalski's Theorem, namely:

Theorem 3.4. Let Θ be a connected, centrally symmetric region on the hypersphere in E^d. A family of compact connected sets in E^d has a hyperplane transversal with normal in Θ if there exists some ordering of the sets such that every three of the sets have a directed line transversal consistent with the ordering.

The proof follows the proof of Theorem 3.1 and is left to the reader.

3.3 Generalizations to Hyperplane Transversals

A **Radon partition** of a set X of points in E^d is a decomposition of X into two disjoint sets, $X = Y \cup Z$ and $Y \cap Z = \varnothing$ such that $conv(Y) \cap conv(Z) \neq \varnothing$. Radon's Theorem is that any set of d+2 or more points in E^d has a Radon partition. If a set has exactly d+2 points in general position in E^d, Goodman and Pollack proved that the set has a unique Radon partition which is determined by the order type of the set[29]. Their lemma is restated here to include the converse that the unique Radon partition determines the order type up to reflection. The reflection of an order type is the order type obtained by reflecting the points in E^d. The notation $x_1, ..., \hat{x}_i, ..., x_{d+1}$ refers to all elements from x_1 to x_{d+1} excluding x_i.

Lemma 3.6. There exists a unique Radon partition for each order type of d+2 points in general position such that any set with the given order type has the given Radon partition. Conversely, there exists an order type, unique up to reflection, for each Radon partition on a set of d+2 points in general position, such that every set with the given Radon partition has the given order type.

Proof: Let $X = \{x_1, \ldots, x_{d+2}\}$ be a set of $d+2$ points represented in homogeneous coordinates where $x_i = \{\alpha_{i,0}, \alpha_{i,1}, \ldots, \alpha_{i,d}\}$, $\alpha_{i,0} > 0$. A point x is a convex combination of the x_i if $x = \sum_{i=1}^{d+2} \sigma_i x_i$, for some $\sigma_i \geq 0$. Given some Radon partition $X = Y \cup Z$, there exist σ_i not all equal to 0, such that

$$\sum_{i=1}^{d+2} \sigma_i x_i = 0, \tag{3.1}$$

where $\sigma_i > 0$ when $x_i \in Y$ and $\sigma_i \leq 0$ when $x_i \in Z$. Without loss of generality, assume x_1 is an element of Y, i.e. $\sigma_1 > 0$.

$$det(\sum_{i=1}^{d+2} \sigma_i x_i, x_2, x_3, \ldots, \hat{x}_j, \ldots, x_{d+2}) = \sum_{i=1}^{d+2} \sigma_i \; det(x_i, x_2, x_3, \ldots, \hat{x}_j, \ldots, x_{d+2})$$

$$= \sigma_1 \; det(x_1, \ldots, \hat{x}_j, \ldots, x_{d+2}) + \sigma_j \; (-1)^j \; det(x_2, \ldots, x_{d+2}) . \tag{3.2}$$

From equation (3.1),

$$det(\sum_{i=1}^{d+2} \sigma_i x_i, x_2, x_3, \ldots, \hat{x}_j, \ldots, x_{d+2}) = det(0, x_2, x_3, \ldots, \hat{x}_j, \ldots, x_{d+2})$$

$$= 0 . \tag{3.3}$$

Equating equations (3.2) and (3.3) gives

$$\sigma_j = (-1)^{j+1} \sigma_1 \frac{det(x_1, x_2, \ldots, \hat{x}_j, \ldots, x_{d+2})}{det(x_2, \ldots, x_{d+2})} , \quad \text{and} \tag{3.4}$$

$$det(x_1, x_2, \ldots, \hat{x}_j, \ldots, x_{d+2}) = (-1)^{j+1} \frac{\sigma_j}{\sigma_1} det(x_2, \ldots, x_{d+2}) . \tag{3.5}$$

Since $\sigma_1 > 0$, $sgn(\sigma_j)$ is dependent upon the signs of the determinants in (3.4) which are determined by the order type of X. x_j belongs to Y or Z depending upon $sgn(\sigma_j)$ and the Radon partition of X is determined by the order type of X. Similarly, by reflection of X we may assume that $det(x_2, \ldots, x_{d+2}) > 0$. Then $sgn(det(x_1, \ldots, \hat{x}_j, \ldots, x_{d+2}))$ in (3.5) is completely determined by $sgn(\frac{\sigma_j}{\sigma_1})$ which depends upon the Radon partition of X. ∎

Let A be a $k-1$-separable family of $k+2$ compact convex sets in E^d with a given k-ordering. This k-ordering corresponds to an order type on $k+2$ points in E^k. By Lemma 3.6 there is a unique Radon partition of these $k+2$ points which depends solely upon the order type. By transforming these points into the corresponding elements of A, this partition becomes a partition $A = B \cup C$ of the convex sets in A.

Lemma 3.7. Let $A = \{a_1, \ldots, a_{k+2}\}$ be a $k-1$-separable family of $k+2$ compact convex sets in E^d with a given k-ordering. Let $A = B \cup C$ be the partition of A corresponding to the unique Radon partition of points with the given order type. There exists a k-flat which stabs A consistent with the given k-ordering if and only if

$$conv(B) \cap conv(C) \neq \emptyset .$$

Proof: Without loss of generality, assume $B = \{a_1, \ldots, a_j\}$ and $C = \{a_{j+1}, \ldots, a_{k+2}\}$. Assume there exists a k-flat which stabs A consistent with the given k-ordering. There exist points $x_i \in a_i$, $i = 1, \ldots, k+2$, all of which lie on the k-flat. Furthermore, there exists a Radon partition Y and Z of the x_i such that some point x lies in $conv(Y)$ and $conv(Z)$. The partition of A into B and C corresponds to this Radon partition into Y and Z. Thus, all the points in Y lie in $conv(B)$ and all the points in Z lie in $conv(C)$. It follows that $x \in conv(B) \cap conv(C)$.

For the converse assume that there exists some $x \in conv(B) \cap conv(C)$. Let Y be a set of j points each from a different convex set in B such that $x \in conv(Y)$. Similarly, let Z be a set of $k-j+2$ points each from a different convex set in C such that $x \in conv(Z)$.

There exists some $j-1$-flat containing all the points in Y and some $k-j+1$-flat containing all the points in Z. Since both the $j-1$-flat and the $k-j+1$-flat contain x, they are contained in some k-flat. Furthermore, Y and Z are the Radon partition of the points x_i in the k-flat. By Lemma 3.6, there is an order type, unique up to reflection, associated with this Radon partition. By properly choosing the orientation of the k-flat, this order type is the same as the given k-ordering. ■

Goodman and Pollack introduced the following crucial lemma[29]. The proof is repeated here for completeness.

Lemma 3.8. Let Y, Z be the Radon partition of a set $X = \{x_1, \ldots, x_{d+2}\}$ of $d+2$ points in general position in $\mathbf{E}^d$. If x_{d+3} is any point of $\mathbf{E}^d$ in general position with respect to the points in X, then there exists a point $x_{i*} \in X$ such that Y', Z' is the Radon partition of $X-\{x_{i*}\}+\{x_{d+3}\}$ where $x_{d+3} \in Y'$, $Y'-\{x_{d+3}\} \subseteq Y$ and $Z' \subseteq Z$.

Proof: Represent the points in homogeneous coordinates where $x_i = \{\alpha_{i,0}, \alpha_{i,1}, \ldots, \alpha_{i,d}\}$, $\alpha_{i,0} > 0$. Since Y, Z is a Radon partition of X,

$$\sum_{i=1}^{d+2} \sigma_i x_i = 0 \tag{3.6}$$

for some σ_i, $i=1, \ldots, d+2$ where $\sigma_i > 0$ if $x_i \in Y$ and $\sigma_i < 0$ if $x_i \in Z$. $d+1$ points in general position in $\mathbf{E}^d$ do not have a Radon partition, so $\sigma_i \neq 0$, $i=1, \ldots, d+2$.

$$x_{d+3} + \sum_{i=1}^{d+2} \phi_i x_i = 0 \tag{3.7}$$

for some ϕ_i, $i=1, \ldots, d+2$. Let $\mu = \min \left[\dfrac{\phi_i}{\sigma_i} \right]$. By subtracting equation (3.6) multiplied by μ from equation (3.7), we get

$$x_{d+3} + \sum_{i=1}^{d+2} (\phi_i - \mu\sigma_i) x_i = 0 .$$

Let $\psi_i = \phi_i - \mu\sigma_i$, $i=1, \ldots, d+2$, and $\psi_{d+3} = 1.$ $\frac{\phi_i}{\sigma_i} \geq \mu$, $i=1, \ldots, d+2$. If $\sigma_i > 0$, $\psi_i = \phi_i - \mu\sigma_i \geq 0$. If $\sigma_i < 0$, $\psi_i = \phi_i - \mu\sigma_i \leq 0$. For some $i*$, $\psi_{i*} = 0$. Since x_{d+3} is in general position with respect to the point in X, $\psi_i \neq 0$ for all $i \neq i*$. Thus, $Y' = \{ x_i : \psi_i > 0 \}$, $Z' = \{ x_i : \psi_i < 0 \}$ form the desired Radon partition of $X - \{x_{i*}\} + \{x_{d+3}\}$. ■

Theorem 3.5 is the generalization of Katchalski's Theorem and Goodman and Pollack's Theorem from 0-separable and $d-1$-separable families to k-separable families.

Theorem 3.5. A family of $k-1$-separable compact convex sets in E^d has a $d-1$-transversal if there exists some k-ordering of the family such that every $k+2$ sets are intersected by some oriented k-flat consistent with the k-ordering.

Proof: By Lemma 3.4 if the theorem is true for any finite family, then it is also true for infinite families. Let $A = \{a_1, a_2, \ldots, a_n\}$ be a finite family of n $k-1$-separable compact convex sets in E^d. For each $a_i \in A$, choose a point $x_i \in a_i$. For $\gamma \in [0,1]$, let $a_i(\gamma)$ be the contraction of a_i by a factor of γ about x_i where:

$$a_i(\gamma) = \{ x_i + \gamma(y - x_i) \mid y \in a_i \} .$$

Here x_i and y are represented in Cartesian coordinates. Let τ be the largest number such that for any $\gamma < \tau$ some $k+2$ sets of $A(\gamma) = \{a_i(\gamma)\}$ have no oriented k-flat which intersects them consistent with the k-ordering. For all $\gamma > \tau$, every subfamily of $k+2$ convex sets, $\{a_{i_1}(\gamma), a_{i_2}(\gamma), \ldots, a_{i_{k+2}}(\gamma)\}$, has a k-transversal consistent with the k-ordering. By Lemma 3.7 there exists a unique partition of this subfamily corresponding to the Radon partition of their k-ordering. Let $B(\gamma) \cup C(\gamma) = \{a_{i_1}(\gamma), a_{i_2}(\gamma), \ldots, a_{i_{k+2}}(\gamma)\}$, be this partition. All the sets in $B(\gamma)$ and $C(\gamma)$ are compact, so $conv(B(\gamma))$ and $conv(C(\gamma))$ are compact sets and $conv(B(\gamma)) \cap conv(C(\gamma))$ is a compact set. For all $\gamma > \tau$, $conv(B(\gamma)) \cap conv(C(\gamma)) \neq \emptyset$, so $conv(B(\tau)) \cap conv(C(\tau)) \neq \emptyset$. By Lemma 3.7 every k+2 sets in $A(\tau)$ have some k-flat which intersects them consistent with the k-ordering.

If $\tau = 0$, each of the sets is contracted to a point. These points all must lie on some k-flat and be contained in some hyperplane, so the theorem is trivial. Hence we may assume $\tau > 0$.

Without loss of generality, assume that for $\gamma < \tau$ the convex sets $a_1(\gamma), \ldots, a_{k+2}(\gamma)$ have no oriented k-flat which intersects them consistent with the k-ordering. Let B, C be the unique partition of $a_1(\tau), \ldots, a_{k+2}(\tau)$ corresponding to the Radon partition of the k-ordering. $relint(conv(B)) \cap relint(conv(C)) = \emptyset$, or else the convex sets in B and C could still be shrunk and by Lemma 3.7, there would still be a k-flat which intersects them consistent with the k-ordering. Thus there is a hyperplane h which separates $conv(B)$ from $conv(C)$. We claim that h is the desired $d-1$-transversal.

Let h^+ be a closed half-space bounded by h and containing $conv(B)$ and let h^- be the other closed half-space which contains $conv(C)$. We first show that $h^+ \cap a_i \neq \emptyset$, for any a_i. By Lemma 3.8 there exists some $a_j \in B \cup C$ such that B', C' is the the Radon partition of $B \cup C - \{a_j\} + \{a_i\}$ where $B' \subseteq B$, $a_i \in C'$ and $C' - \{a_i\} \subseteq C$. There exists some point $y \in conv(B') \cap conv(C')$. Since $B' \subseteq B$ and $conv(B)$ lies in h^+, y must lie in h^+. Since $A(\gamma)$ is $k-1$-separable, $C' - \{a_i\}$ and B' can be strictly separated by some hyperplane. y lies in $conv(B')$ so y cannot lie in $conv(C' - \{a_i\})$. $conv(C' - \{a_i\})$ lies in h^-. For y to lie in h^+ some point of a_i must lie in h^+. Thus $h^+ \cap a_i \neq \emptyset$. By the same reasoning, $h^- \cap a_i \neq \emptyset$.

Therefore $h \cap a_i \neq \emptyset$ and h is a $d-1$-transversal which intersects all the elements of A. ■

Chapter 4

Geometric Permutations

4.1 Introduction

Let A be a finite family of n disjoint compact convex sets in E^2. A directed line transversal for A generates an ordering on the elements of A. This ordering plays an important role in Hadwiger's Theorem. Katchalski, Lewis, Zaks and Liu called this ordering and its 'reversal' a **geometric permutation** of the family [39, 40].

Orderings can also be examined for lines which intersect only some of the elements of A. If l is a directed line which intersects $A' \subseteq A$, l generates an ordering on A'. Any two directed lines with the same direction which intersect $A' \subseteq A$ generate the same ordering on A'.

In how many different orders can a directed line intersect A? Restricted to lines which intersect all the elements of A, i.e. line transversals for A, the question is equivalent to asking for an upper bound on the number of different geometric permutations on A. Katchalski, Lewis, Zaks and Liu gave an upper bound of $\binom{n}{2}$.

In Section 4.2 I show that the directed lines in E^2 can be partitioned into at most $12n$ sets where any two lines in the same set which intersect any $A' \subseteq A$ generate the same ordering on A'. This bounds the number of geometric permutations of A by $6n$. In my bound I prove that n disjoint compact convex sets in the plane can be embedded in n disjoint convex polygons with a total of at most $12n$ edges. These edges have at most $6n$ distinct slopes. Edelsbrunner, Robison, and Shen improved upon this theorem by showing that n convex sets can be embedded in n convex polygons with a total of at most $6n-9$ edges and at most $3n-6$ slopes[20]. Their theorem implies that the directed lines in $\mathbf{E}^2$ can be partitioned into at most $6n-12$ sets where any two lines in the same set which intersect any $A' \subseteq A$ generate the same ordering on A'. This decreases the bound on the number of geometric permutations to $3n-6$. Edelsbrunner and Sharir independently showed that that the maximum number of geometric permutations for A is $2n-2$[21]. However, their argument does not extend to lines which intersect only a subset of A.

The number of geometric permutations can be studied for line transversals in higher dimensions. Let H be a finite set of hyperplanes in $\mathbf{E}^d$. Define $\Psi(H)$ to be the number of cells in the arrangement obtained by translating the hyperplanes in H to the origin. Let $\Psi^d(m)$ be the maximum number of cells in any arrangement of m hyperplanes through the origin in $\mathbf{E}^d$. R. Winder proved that $\Psi^d(m) = 2\sum_{i=0}^{d-1}\binom{m-1}{i}$[60]. The directed lines in $\mathbf{E}^d$ can be partitioned into $\Psi^d(\binom{n}{2})$ sets, where any two lines in the same set which intersect $A' \subseteq A$ generate the same order on A'. The number of geometric permutations is bounded by $\frac{1}{2}\Psi^d(\binom{n}{2})$. This is the only known bound.

Geometric permutations can also be defined on hyperplane transversals or even on k-transversals for any k. A **geometric permutation** of k-transversals for a family of $k-2$-separable convex sets in $\mathbf{E}^d$ is the k-ordering produced by an oriented hyperplane transversal and the reverse k-ordering produced by the same hyperplane with reverse orientation.

Section 4.3 presents a result on geometric permutations of plane transversals in $\mathbf{E}^3$. For hyperplane transversals in higher dimensions, I only have a conjecture which is included at the end of the section. Unfortunately, even the theorem proven in section 4.3 does not give good bounds on the number of different 2-orderings generated by plane transversals for a family in E^3.

4.2 Upper Bounds in the Plane

A **separation set** for A is a set H of hyperplanes such that every pair of elements, $a, b \in A$, is separated by some hyperplane in H. A **strict separation set** is a separation set H where every pair of elements is strictly separated by some hyperplane in H.

Theorem 4.1. Let A be a family of pairwise disjoint compact convex sets in E^d and let $H = \{h_1, \ldots, h_m\}$ be some strict separation set for A. The directed lines in E^d can be partitioned into $\Psi(H)$ sets such that any two lines in the same set which intersect any $A' \subseteq A$ generate the same order on A'.

Proof: Let $\{u_1, \ldots, u_m\}$ be a set of normal vectors for H where $h_i = \{x : u_i \cdot x = c_i\}$. Let v be a vector in E^d and let L be the set of all lines with direction v. Let a, b be two sets in A represented in Cartesian coordinates and let h_i be the hyperplane separating a from b, Assume $u_i \cdot (y-x) > 0$, for all $x \in a$, $y \in b$. If $v \cdot u_i > 0$, then any line in L must intersect a before b. If $v \cdot u_i < 0$, then any line in L must intersect b before a. Thus $sgn(v \cdot u_i)$, $i = 1, \ldots, m$, determines the relative ordering of a and b generated by lines in L.

Consider the arrangement created by hyperplanes through the origin with normals u_i, $i = 1,...,m$. There are at most $\Psi(H)$ cells in this arrangement. Partition the directed lines in E^d into at most $\Psi(H)$ sets, assigning directed lines which point into the same cell to the same sets. The values of $sgn(v \cdot u_i)$, $k = 1, \ldots, m$, are determined by the cell into which v points. Any two directed lines which intersect $a, b \in A$ and point to the same cell generate the same relative ordering on a and b. For any $A' \subseteq A$ the order in which a line intersects A' is completely determined by the relative order in which a line intersects $a, b \in A'$. Thus any two directed lines which intersect $A' \subseteq A$ and lie in the same set generate the same ordering on A'. ■

We next show how to find a strict separation set for a family of convex polygons in the plane. Here polygons are non-degenerate, i.e. not line segments or points.

Theorem 4.2. If P is a family of pairwise disjoint convex polygons in E^2, then there exists a strict separation set L for A where each line in L is parallel to some edge of a polygon in A.

Proof: Let a and b be any two polygons in P and let $l*$ be some line separating a from b and tangent to a at some vertex. Let l_1 and l_2 be the two lines containing the edges of a which meet at that vertex. l_1 and l_2 divide the plane into four cells or quadrants, q_1, q_2, q_3, q_4, with a lying wholly in one quadrant, say q_1. (See Figure 4.1.) The boundary of each quadrant is included in that quadrant.

Since $l*$ separates a from b, b does not intersect quadrant q_1. If b does not intersect quadrant q_2, then l_2 separates a from b. If b does not intersect quadrant q_4, then l_1 separates a from b. Translating either of these separators slightly toward b, produces a line strictly separating a from b and parallel to a line through an edge of a.

Assume b contains a point from quadrant q_2 and from quadrant q_4. Some edge of b must intersect l_1. Let l_3 be the line containing that edge. l_3 separates a from b so by

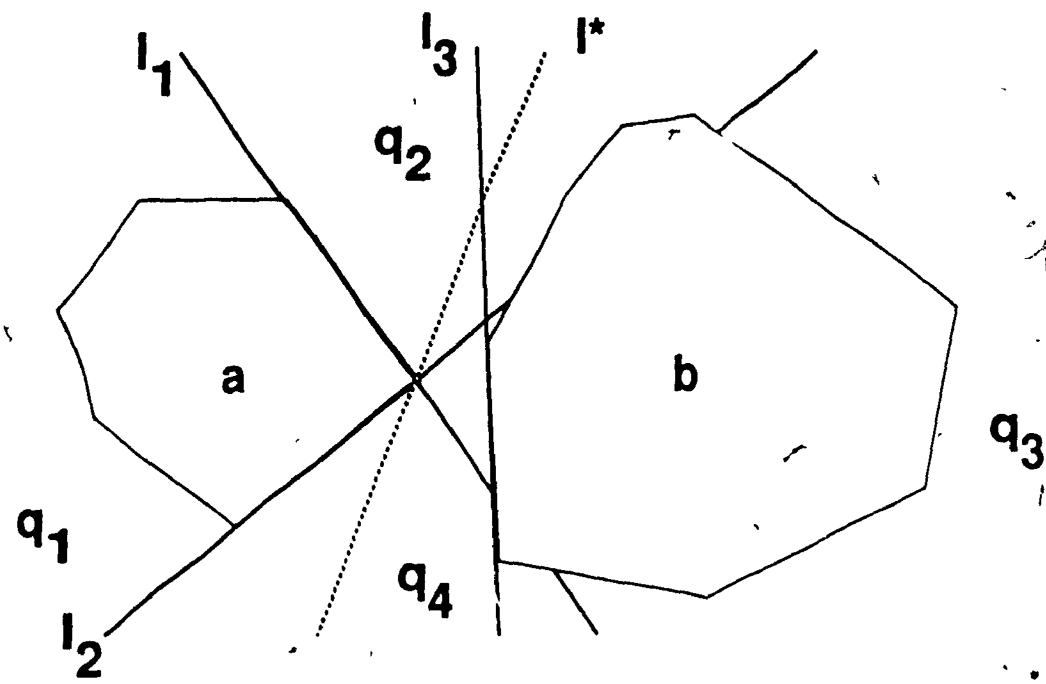

Figure 4.1. Convex Polygons a and b.

translating it slightly toward a we have a line strictly separating a from b and parallel to an edge of b.

For every pair of polygons $a, b \in P$, add to L a line strictly separating a from b and parallel to some edge in a or b. L is a separation set for P such that every line in L is parallel to some edge of a polygon in P. ■

Let P be a finite family of pairwise disjoint convex polygons. P can be considered as a planar subdivision. A **complete triangulation** of P is a planar subdivision which contains

P and some additional line segments between vertices of distinct polygons in P such that each face which is not a polygon is bounded by three additional line segments and three polygons. (See Figure 4.2.) A **complete triangulation** of P defines a graph whose vertices are the polygons in P and whose edges are the additional line segments placed between polygons. This graph may contain multiple edges, i.e. more than one edge between any two vertices. Any straight edge embedding of this graph in the plane is a triangulation.

The following theorem shows that convex sets can be embedded in polygons with few edges. Using such an embedding, we can then find a separation set for the polygons and hence for the convex sets. Bill Lenhart provided the inspiration for this theorem.

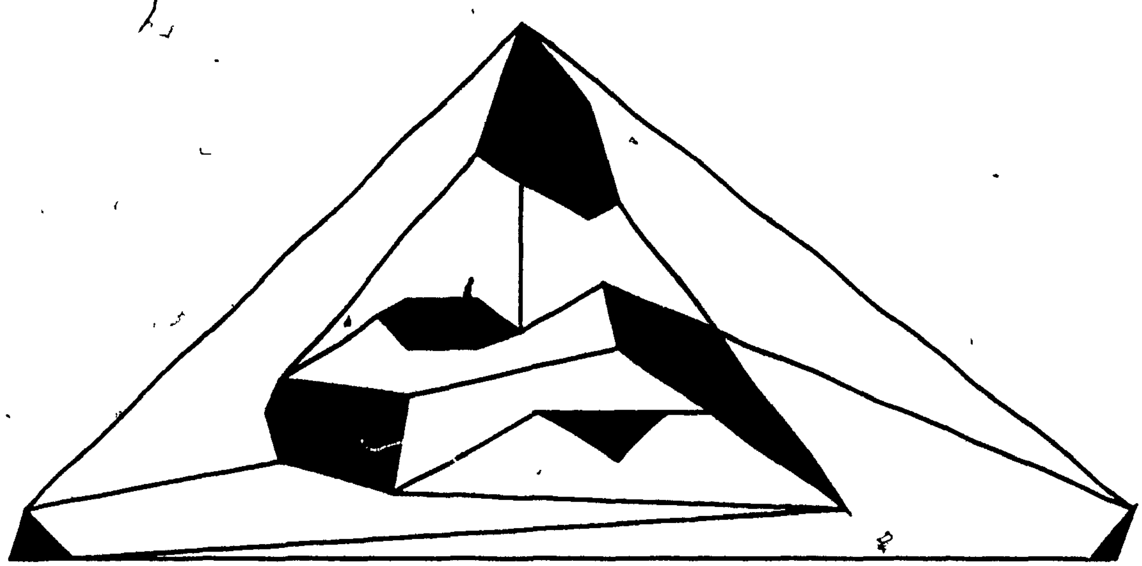

Figure 4.2. A Complete Triangulation of Convex Polygons

Theorem 4.3. **Let A be a finite family of n pairwise disjoint compact convex sets in E^2. There exists a family P of n pairwise disjoint compact convex polygons, such that:**

i) each convex set in A is entirely contained in a unique polygon in P,

ii) the total number of edges in all the polygons in P is at most $12n$ and

iii) if L is the set of lines containing the edges of the polygons in P, then $\Psi^2(L) \leq 12n$.

Proof: We first embed the convex sets in A in a family Q of n pairwise disjoint convex polygons. Choose some convex set $a \in A$ and find $n-1$ lines separating a from the $n-1$ other convex sets. These $n-1$ lines bound $n-1$ half-planes containing a. If the intersection of these half-planes is unbounded, we can add three suitable half-planes containing a such that the intersection of all the half-planes is bounded. The intersection of these half-planes forms a polygon which contains a. Repeat the procedure n times to get the family Q of n pairwise disjoint convex polygons containing the n convex sets.

Let c be some polygon in Q which lies on the boundary of $conv(Q)$. Add to Q two small triangles, d and d', to form Q' such that the boundary of $conv(Q')$ is three line segments between c, d and d'. We construct a complete triangulation of Q'. Q' forms a planar subdivision. Triangulate the external face of this planar subdivision by adding triangulation line segments between vertices of polygons. The triangulation faces are bounded by two or three triangulation line segments. If a face is bounded by only two triangulation line segments, one of the bounding line segments is redundant. Remove the redundant line segments until all faces are bounded by three triangulation line segments and three polygons of Q'. The resultant subdivision is a complete triangulation of Q'. Let G be the planar graph whose vertices are polygons in Q' and whose edges are the triangulation line segments.

For each edge e of G, let l_e be the line containing e. For each pair of polygons $a, b \in Q'$, let $l_{a,b}$ be some line strictly separating a from b. For each polygon $a \in Q'$, let $N(a)$ be a list of all the neighbours of a in the graph G. Let $F(a)$ be a list of the faces which are bounded by a. Let $L_1(a) = \{l_{a,b} : b \in N(a)\}$ and let $L_2(a) = \{l_e : e$ lies on a face of $F(a)$ and connects $b, b' \in N(a)\}$. Let $L_0(a)$ be the union of the lines in $L_1(a)$ and $L_2(a)$ slightly translated towards a. Let $P(a)$ be the intersection of the half-planes containing a and bounded by the lines in $L_0(a)$. For any $a \in Q$ other than c, $P(a)$ is bounded by the cycle of the neighbours of a and the triangulation line segments between them. (See Figure 4.3.) Thus $P(a)$ is a convex polygon containing a. By a judicious choice of $l_{c,d}$ and $l_{c,d'}$, $P(c)$ is also a convex polygon containing c. We claim that $P = \{P(a) : a \in Q\}$ is a family of n pairwise disjoint compact convex polygons such that each $a \in Q$ is contained in $P(a)$ and the total number of edges in the polygons in P is $12n$.

Assume $a, b \in Q$ are neighbours in G. $P(a)$ and $P(b)$ are both bounded by a translation of the same separator, $l_{a,b}$ of a and b. Since $P(a)$ and $P(b)$ are in different halfplanes bounded by $l_{a,b}$, $l_{a,b}$ must separate $P(a)$ from $P(b)$ and $P(a)$ and $P(b)$ are pairwise disjoint.

Now, assume $a, b \in A$ are not neighbours in G. $P(a)$ is entirely contained in a cycle of the neighbours of a while $P(b)$ is entirely contained in a cycle of the neighbours of b. Since b is not a neighbour of a, these two cycles must contain different regions in the plane and so $P(a)$ and $P(b)$ must be pairwise disjoint.

Let m be the number of edges in G. G is a complete triangulation and there are $n + 2$ vertices in G, so $m = 3n$. Each edge connects two neighbours, so the total number of neighbours over all vertices in the graph is $6n$. The total number of edges in the polygons in

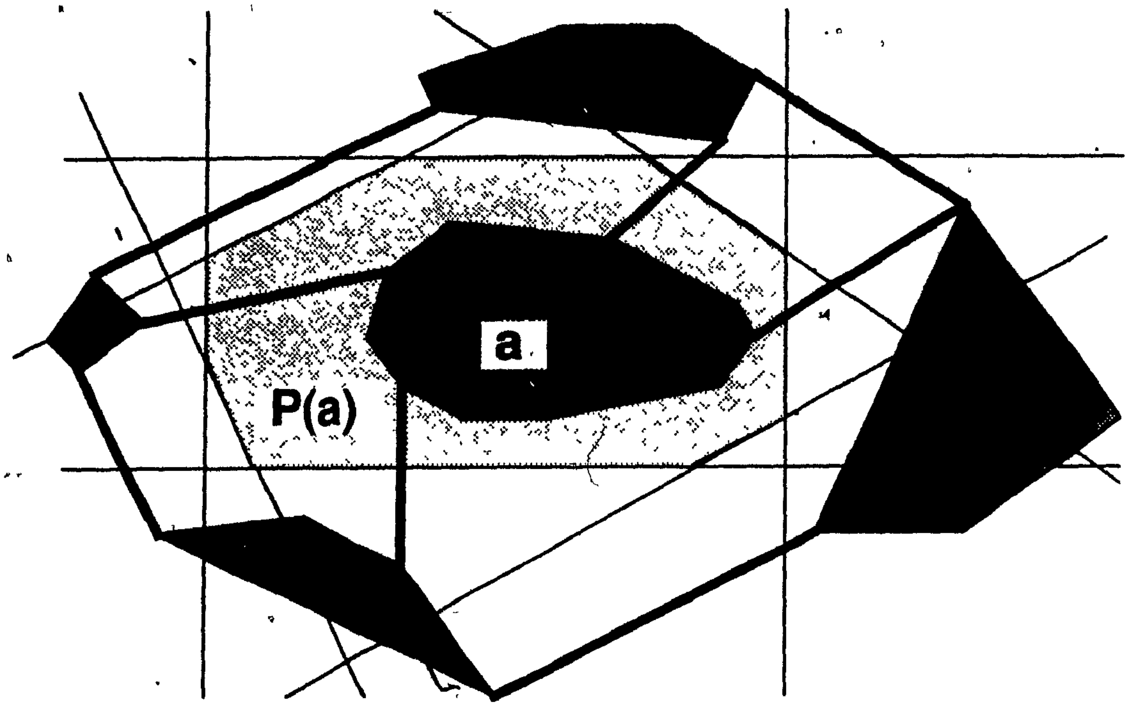

Figure 4.3. Cycle of Neighbours of a.

P is at most the total number of lines in $L_0(a)$ over all $a \in Q$.

$$\sum_{a \in Q} |L_0(a)| \le \sum_{a \in Q'} |L_0(a)| = \sum_{a \in Q'} |L_1(a)| + \sum_{a \in Q'} |L_2(a)| \le 6n + 6n = 12n$$

Therefore, the total number of edges in the polygons in P is at most $12n$.

Let L be the set of lines containing the edges of the polygons in P. L is a subset of $L' = \bigcup_{a \in Q'} L_0(a)$. Each line in $L_0(a)$ is parallel to some line in some $L_0(b)$, $a, b \in Q'$. Since the total number of lines in L' is at most $12n$, $\Psi(L') \le 12n$ and so $\Psi(L) \le 12n$. ■

For every pair of disjoint compact convex sets, there exists some hyperplane which strictly separates the pair. Construct a separation set H for a family of n compact convex sets by choosing a separating hyperplane for each pair. The size of H, $|H|$, is at most $\binom{n}{2}$. Since $\Psi(H) \leq \Psi^d(|H|)$, the following is a corollary to Theorem 4.1:

Corollary 4.1. Let A be a family of n pairwise disjoint compact convex sets in $\mathbf{E}^d$. The directed lines in $\mathbf{E}^d$ can be partitioned into $\Psi^d(\binom{n}{2})$ sets such that any two lines in the same set which intersect some $A' \subseteq A$ generate the same order on A'.

Finally, by Theorems 4.2 and 4.3, there exists a separation set L for n convex sets in $\mathbf{E}^2$ such that $\Psi^2(L) \leq 12n$. Applying Theorem 4.1, we have the following corollary:

Corollary 4.2. Let A be a family of n pairwise disjoint compact convex sets in $\mathbf{E}^2$. The directed lines in $\mathbf{E}^2$ can be partitioned into $12n$ sets such that any two lines in the same set which intersect some $A' \subseteq A$ generate the same order on A'.

Each geometric permutation corresponds to two sets of directed lines in the partition of directed lines in $\mathbf{E}^d$. Thus the number of geometric permutations on A is bounded by half the size of the partition. Corollaries 4.1 and 4.2 imply that there are at most $\frac{1}{2}\Psi^d(\binom{n}{2})$ geometric permutations of A in $\mathbf{E}^d$ and at most $6n$ geometric permutations of A in $\mathbf{E}^2$.

Katchalski, Lewis and Zaks asserted that for every d, there exists a constant β_d and a family A of n pairwise disjoint compact convex sets in $\mathbf{E}^d$ such that there are at least $\beta_d n^{d-1}$ geometric permutations of A [40]. Corollary 4.1 implies an upper bound of $O(n^{2d-2})$ for the number of geometric permutations of A, leaving a wide gap for improvement. Villanger showed that for any n there exist families of n line segments in E^3 in which any hyperplane separates at most one line segment from one other line segment[55]. Thus, there exists a family A of n compact convex sets in E^d, for any $d \geq 3$, such that any separation set of A must have $\binom{n}{2}$ elements. In fact, by embedding Villanger's line segments in rectangular prisms, we see that Theorem 4.2 does not generalize to two dimensions. Reduction of the upper bounds must come from other directions.

4.3 Geometric Permutations in Higher Dimensions

Let A be a $k-1$-separable family of n compact convex sets in E^d. How many different k-orderings are generated by oriented k-flat stabbers of A? If $k = d-1$, any hyperplane transversal can be continuously transformed until it is tangent to d convex sets. For any d convex sets which are $d-2$ separable, there are at most 2^{d+1} tangent oriented hyperplanes, depending upon which half-space bounded by the hyperplane contains which convex set. This gives a trivial bound of $2^{d+1} \binom{n}{d}$.

If A is $d-2$-separable in E^d, any two hyperplane transversals for A with the same normal must generate the same $d-1$-ordering on A. This motivates the following problem: Given a $d-2$-separable family A of compact convex sets in E^d, partition the space of all normals such that any two hyperplane stabbers with normals in the same partition generate the same $d-1$-ordering on A. Such a partition is constructed in the proof of Theorem 4.1 for line transversals in E^2. Theorem 4.4 provides a partition for normals in E^3.

The definition of separation sets must be extended to k-separable families. A **k-separation set** for a k-separable family of sets is a set of hyperplanes such that every j sets are strictly separated from every $k+2-j$ sets by some hyperplane, $1 \le j \le k+1$. The previously defined "separation set" becomes a 0-separation set under this notation.

Theorem 4.4. Let A be a 1-separable family of compact convex sets in E^3, and let H be a 1-separation set for A. For every set $a \in A$ and every other pair of sets $b, c \in A$, let $u_{a|bc}$ be a normal to the plane in H separating a from b and c. The 2-ordering generated by any plane stabber with normal u is determined by $sgn(det(u, u_{a|bc}, u_{c|ab}))$ over all $a, b, c \in A$.

Theorem 4.4 means that E^3 can be partitioned by the set of planes containing the origin and vectors $u_{a|bc}$ and $u_{c|ab}$ over all $a, b, c \in A$. All the plane stabbers with normals pointing into a given cone in the partition generate the same 2-ordering of A.

Before proving Theorem 4.4, we need the following lemma relating the orientation of three points in the plane to the orientation of three vectors separating them. All points and vectors are represented in Cartesian coordinates. The order type of three points in Cartesian coordinates is $sgn(x_1-x_0, x_2-x_0)$. The cone formed by two vectors $v_0, v_1 \in E^2$ is the set of vectors $\{\sigma_0 v_0 + \sigma_1 v_1 : \sigma_0 \ge 0, \sigma_1 \ge 0\}$. If u lies in the cone formed by v_0, v_1, $u = \sigma_0 v_0 + \sigma_1 v_1$, then

$$sgn(det(v_0, u)) = sgn(det(v_0, \sigma_1 v_1)) = sgn(det(v_0, v_1)) \quad \text{and}$$

$$sgn(det(u, v_1)) = sgn(det(\sigma_0 v_0, v_1)) = sgn(det(v_0, v_1)).$$

If $v = (\alpha_1, \alpha_2)$, then $v^P = (-\alpha_2, \alpha_1)$, a vector perpendicular and to the left of v. $v^P \cdot v' = det(v, v')$. Subscripts are computed *mod* 3 where appropriate.

Lemma 4.1. Let x_0, x_1, x_2 be three points in $\mathbf{E}^2$ and let u_0, u_1, u_2 be three vectors in $\mathbf{E}^2$. If $u_i \cdot (x_i - x_j) > 0,\ j \neq i$, then

$$sgn(det(x_1 - x_0,\ x_2 - x_0)) = sgn(sgn(det(u_0, u_1)) + sgn(det(u_1, u_2)) + sgn(det(u_2, u_0))) .$$

Proof: If x_0, x_1, x_2 were collinear, $u_i \cdot (x_i - x_j)$ would equal 0 for some $i,\ j,\ i \neq j$, so $x_0,\ x_1, x_2$ must form a triangle. Let $v_0,\ v_1,\ v_2$ be outward pointing normals to the sides of this triangle, where

$$v_i \cdot\ (x_{i+1} - x_{i+2}) = 0\ , \quad \text{and} \quad v_i \cdot\ (x_i - x_j) < 0\ , \quad j \neq i\ .$$

(See Figure 4.4.) It is easy to see that $\mu_0 v_0 + \mu_1 v_1 + \mu_2 v_2 = 0$, for some $\mu_0,\ \mu_1,\ \mu_2 > 0$, and that u_{i+2} lies in the cone formed by v_i and v_{i+1}. Furthermore,

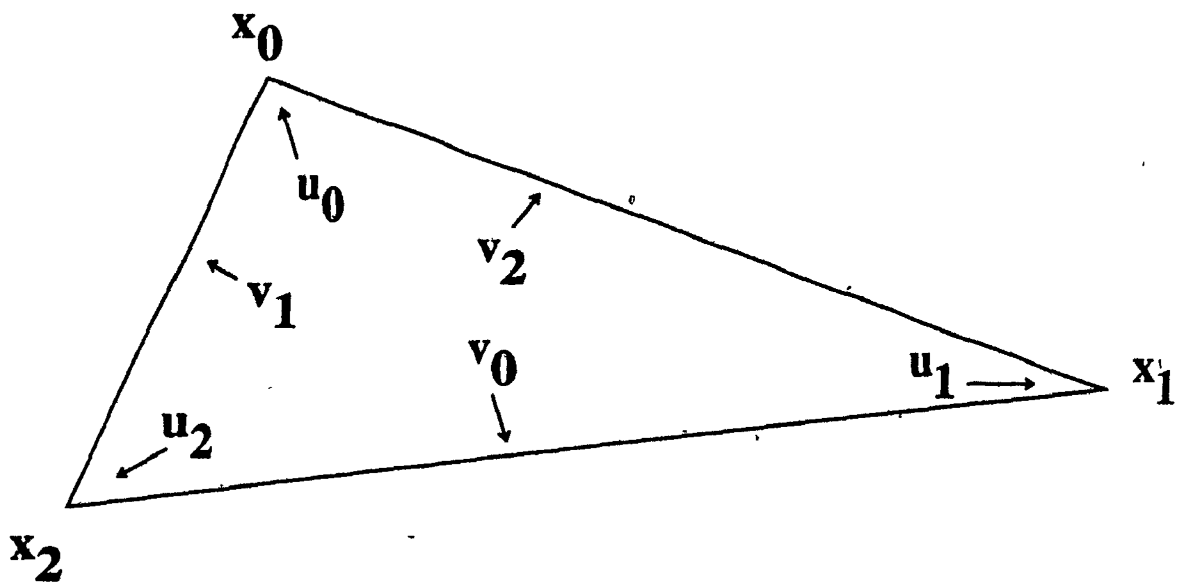

Figure 4.4. Triangle x_1, x_2, x_3.

$$sgn(det(v_i,v_{i+1})) = sgn(det(v_i,u_{i+2})) = sgn(det(u_{i+2},v_{i+1})).$$

We show that the orientation of any two of the outward pointing normals, v_i, v_{i+1}, is the same as the orientation of x_0, x_1, x_2. For some $\alpha, \beta \in \mathbf{R}$, $x_1 - x_0 = \alpha v_2^{\perp}$ and $x_2 - x_0 = \beta v_1^{\perp}$. By taking the dot product with v_1 and v_2, respectively, we find

$$\alpha det(v_1,v_2) = -(x_1-x_0)\cdot v_1 > 0\,, \qquad \text{and} \qquad \beta det(v_1,v_2) = (x_2-x_0)\cdot v_2 < 0\,.$$

Therefore, $sgn(\alpha) = -sgn(\beta)$. It follows that

$$\begin{aligned} sgn(det(x_1-x_0,\ x_2-x_0)) &= sgn(det(\alpha v_2^{\perp},\ \beta v_1^{\perp})) \\ &= sgn(\alpha\beta det(v_2,\ v_1)) \\ &= sgn(det(v_1,v_2)). \end{aligned}$$

Since $\mu_0 v_0 + \mu_1 v_1 + \mu_2 v_2 = 0$,

$$sgn(det(v_0,v_1)) = sgn(det(-\mu_1 v_1 - \mu_2 v_2,\ v_1)) = sgn(det(v_1,v_2)) \qquad \text{and}$$

$$sgn(det(v_2,v_0)) = sgn(det(v_2,\ -\mu_1 v_1 - \mu_2 v_2)) = sgn(det(v_1,v_2))\,.$$

Therefore,

$$sgn(x_1-x_0,\ x_2-x_0) = sgn(det(v_0,v_1)) = sgn(det(v_1,v_2)) = sgn(det(v_2,v_0))\,.$$

It remains to show that

$$sgn(det(v_1,v_2)) = sgn(sgn(det(u_0,u_1)) + sgn(det(u_1,u_2)) + sgn(det(u_2,u_0))).$$

If $sgn(det(u_j,\ u_{j+1})) = sgn(det(v_1,v_2))$ for all values of j, then the claim is obviously true. Assume $sgn(det(u_0,u_1)) = -sgn(det(v_1,v_2))$. $u_0 = \alpha v_1 + \beta v_2$ and $u_2 = \alpha' v_1 + \beta' v_0$, for some $\alpha, \beta, \alpha', \beta' > 0$. If $sgn(det(v_1,u_1)) = sgn(det(v_1,v_2))$, then

$$sgn(det(u_0,u_1)) = sgn(\alpha det(v_1,u_1) + \beta det(v_2,u_1)) = sgn(det(v_1,v_2))\,.$$

Therefore, $sgn(det(v_1,u_1)) = -sgn(det(v_1,v_2))$.

$$sgn(det(u_1,u_2)) = sgn(\alpha' det(u_1,v_1) + \beta' det(u_1,v_0)) = sgn(det(v_1,v_2))\,.$$

By a similar argument, $sgn(det(u_1,u_2)) = sgn(det(v_1,v_2))$. It follows that

$$sgn(sgn(det(u_0, u_1)) + sgn(det(u_1, u_2)) + sgn(det(u_2, u_0))) = sgn(det(v_1, v_2))$$

$$= sgn(det(x_1 - x_0, x_2 - x_0)) \ . \blacksquare$$

If two vectors v, w in lie on a plane with normal u in E^3, this plane can be mapped to E^2 under an isometry Γ to preserve orientation, i.e. $sgn(det(u, v, w)) = sgn(det(\Gamma(v), \Gamma(w))$. Thus Lemma 4.1 also applies to three vectors and three points which lie on an oriented plane in E^3. Theorem 4.4 follows immediately.

Proof of Theorem 4.4: Let u be the normal to any plane stabber h of A. Let a_0, a_1, a_2 be three sets in A, and let h intersect a_i at point x_i. The orientation of x_0, x_1, x_2 with respect to h is $sgn(det(u, x_1 - x_0, x_2 - x_0))$. For $i = 0, 1, 2$, let u_i be the normal to the plane in H separating a_i from a_{i+1} and a_{i+2}. Since $sgn(det(u, v, w))$ is determined by $sgn(det(u, v, -w))$, we may assume $u_i \cdot (x_i - x_j) > 0$, $j \neq i$. Let u_i' be the projection of u_i onto h.

$$det(u, u_i', u_{i+1}') = det(u, u_i, u_{i+1}) \ .$$

By Lemma 4.1,

$$\begin{aligned} sgn(det(u, x_1 - x_0, x_2 - x_0)) &= sgn(sgn(det(u, u_0', u_1')) + sgn(det(u, u_1', u_2')) \\ &\qquad + sgn(det(u, u_2', u_0'))) \\ &= sgn(sgn(det(u, u_0, u_1)) + sgn(det(u, u_1, u_2)) \\ &\qquad + sgn(det(u, u_2, u_0))) \ . \blacksquare \end{aligned}$$

Unfortunately, Theorem 4.4 by itself does not generate a good bound on the number of k-orderings generated by plane stabbers of A. If H is the set of all planes containing the origin and the vectors $u_{a|bc}$ and $u_{c|ab}$ over all $a, b, c \in A$, then H has size $3 \binom{n}{3}$ or $O(n^3)$. Applying Winder's Theorem[60], H partitions E^3 into $O(n^6)$ cones, a worse bound than the trivial one. Many of the planes in H contain the same lines so it is possible that there is a tighter bound for H. Also open is a tight bound on the size of the 1-separation set for A. The trivial bound is $O(n^3)$ but this may not be optimal.

While I was unable to extend Theorem 4.4 to dimensions greater than three, I do have the following conjecture.

Conjecture 4.1. Let A be a $d-2$-separable family of compact convex sets in $\mathbf{E}^d$ and let H be a $d-2$-separation set for A. For every subset B of d elements of A, let $H_{\{B\}}$ be a subset of H which forms a $d-2$-separation set for B and let $U_{\{B\}}$ be the set of all normals to the hyperplanes in $H_{\{B\}}$. The $d-1$-ordering generated by any hyperplane stabber with normal v is completely determined by $sgn(det(v, u_{i_1}, u_{i_2}, \ldots, u_{i_{d-1}}))$, $u_{i_j} \in U_{\{B\}}$, over all subsets of $d-1$ elements of $U_{\{B\}}$ for all $B \subseteq A$, $|B| = d$.

Chapter 5

Stabbing Algorithms

5.1 Introduction

Given a finite set of m polytopes in $\mathbf{E}^d$ and an integer k, find a k-transversal or k-stabber for the set. When $k = 0$, the problem reduces to a linear programming problem. For a fixed dimension d, Megiddo's algorithm solves this problem in time proportional to the number of hyperplanes bounding the polytopes[44, 45].

When $k = 1$, the problem becomes one of finding a line stabber for a family of polytopes. Line stabbing has applications for hidden line problems[19], set partitioning[3] and updating triangulations[22]. Edelsbrunner, Maurer, Preparata, Rosenberg, Welzl and Wood found an $O(n \log n)$ algorithm for line stabbing n line segments in $\mathbf{E}^2$[19]. This algorithm was generalized by Atallah and Bajaj to line stabbing of convex polygons in $\mathbf{E}^2$[2]. The algorithm runs in $O(n \log n\, \alpha(n))$ time, where n is the total number of edges over all polygons and $\alpha(n)$ is the inverse of Ackerman's function.

Edelsbrunner, Guibas and Sharir extended the algorithm for line stabbing convex polygons in the plane to an $O(n^2 \alpha(n))$ algorithm for plane stabbing of convex polytopes with a total of n edges in $\mathbf{E}^3$ and an $O(n^2)$ algorithm for plane stabbing of n line segments in $\mathbf{E}^3$[18]. In higher dimensions, Avis and Doskas gave an $O(n^{d-1} m)$ algorithm for hyperplane stabbing of m convex polyhedra with a total of n edges in d-space[4].

Lemma 3.3 states that a family of compact connected sets has a hyperplane transversal in $\mathbf{E}^d$ if and only if the set of separation normals for all pairs of sets in the family do not cover S^{d-1}. All the algorithms above can be considered as algorithms for determining whether these normals cover this hypersphere.

The algorithm of Atallah and Bajaj for line stabbing in the plane constructs a representation of all the line stabbers in the dual space. From this representation one can determine which directions are the stabbing directions. As observed in Section 3.2, every pair of convex sets in a family of compact convex sets intersect if and only if every direction is a stabbing direction. Thus Atallah and Bajaj's algorithm also determines in $O(n \log n \, \alpha(n))$ time whether every pair of convex polygons in a family of convex polygons intersect.

Little is known about algorithms for k-stabbing when k is not 0 or $d-1$, $d \geq 3$. It is noteworthy that the values of k for which computer scientists have found polynomial algorithms are the same values of k for which mathematicians have devised good necessary and sufficient conditions.

D. Avis and I developed an $O(n^4 \log n)$ time algorithm for line stabbing of polytopes with a total of n edges in E^3. Section 5.2 presents some theoretical results about how lines intersect lines, line segments and polytopes in E^3. These results are used to develop line stabbing algorithms in Section 5.3. Jaromczyk and Kowaluk subsequently improved upon these results with an algorithm for line stabbing of polytopes which runs in $O(n^3 2^{\alpha(n)} \log n)$ time[36].

5.2 Theory for Line Stabbing in Three Dimensions

In the study of line intersections parallel lines must be handled as special cases. To avoid these extra cases it is convenient to study line intersections in projective space, where every two lines which lie on some projective plane must intersect.

Two lines which do not lie in the same plane in E^3 or P^3 are called **skewed**. In P^3 this is equivalent to the two lines not intersecting. A set of lines is **skewed** if every two lines in the set are skewed. A set of line segments is **collinear** if one line contains all the segments. A set of line segments is **co-planar** if one plane contains all the line segments. A set of line segments is **skewed** if the set of lines containing the line segments is skewed.

Let x and y be the coordinates of two distinct points in P^3 parametrized in homogeneous coordinates. The unique line through x and y is parametrized by $\lambda x + \mu y$ where λ and μ vary over R and either $\lambda \neq 0$ or $\mu \neq 0$. If two lines l_1 and l_2 parametrized by $\lambda_1 x_1 + \mu_1 y_1$ and $\lambda_2 x_2 + \mu_2 y_2$ intersect, then there exists $\lambda_1, \mu_1, \lambda_2, \mu_2$, not all zero, such that $\lambda_1 x_1 + \mu_1 y_1 = \lambda_2 x_2 + \mu_2 y_2$. Equivalently, l_1 and l_2 intersect if and only if $det(x_1, y_1, x_2, y_2) = 0$.

Lemma 5.1. Two skew lines l_1 and l_2 and a point z which is not on l_1 or l_2 admit one stabber. If l_1 and l_2 are parametrized by $\lambda_1 x_1 + \mu_1 y_1$ and $\lambda_2 x_2 + \mu_2 y_2$, then this stabber intersects l_1 at

$$z* = det(x_2, y_2, z, y_1)\, x_1 - det(x_2, y_2, z, x_1)\, y_1$$

Proof: There is a unique hyperplane h containing l_2 and z. Since l_1 and l_2 are skew, l_1 intersects h at exactly one point. The line $l*$ through this point and z is the unique stabber of l_1, l_2 and z.

The hyperplane h is given by the following equation:

$$h = \{x : det(x_2, y_2, z, x) = 0,\ x \in P^3\}.$$

Let $z* = det(x_2, y_2, z, y_1)\, x_1 - det(x_2, y_2, z, x_1)\, y_1$. Since x_1 and y_1 cannot both lie on h, either $det(x_2, y_2, z, y1) \neq 0$ or $det(x_2, y_2, z, x_1) \neq 0$ and so $z*$ is a point in $\mathbf{P}^3$ lying on l_1.

Now $z*$ lies on h since

$$det(x_2, y_2, z, z*) = det(x_2, y_2, z, det(x_2, y_2, z, y_1)\, x_1 - det(x_2, y_2, z, x_1)\, y_1)$$

$$= det(x_2, x_2, z, x_1)\, det(x_2, y_2, z, y_1) - det(x_2, y_2, z, y_1)\, det(x_2, y_2, z, x_1)$$

$$= 0.$$

Thus the stabber of l_1, l_2 and z intersects l_1 at $z*$.■

Lemma 5.2. The stabbers of three skew lines in $\mathbf{P}^3$ form a quadric surface Q.

Proof: We first show that the line stabbers of three skew lines lie on a quadric surface. Let $L = \{l_1, l_2, l_3\}$ be a set of three skew lines in $\mathbf{P}^3$ where l_i is parametrized by $\lambda x_i + \mu y_i$, $x_i, y_i \in \mathbf{P}^3$, $\lambda, \mu \in \mathbf{R}$. Let l' be a line stabber of L and let z be a point on l' which does not lie on l_1, l_2 or l_3. By Lemma 5.1, l' intersects l_1 at

$$z* = det(x_2, y_2, z, y_1)\, x_1 - det(x_2, y_2, z, x_1)\, y_1$$

and hence

$$det(x_3, y_3, z, z*) = det(x_3, y_3, z, det(x_2, y_2, z, y_1) x_1 - det(x_2, y_2, z, x_1) y_1)$$

$$= det(x_3, y_3, z, x_1)\, det(x_2, y_2, z, y_1) -$$

$$det(x_3, y_3, z, y_1)\, det(x_2, y_2, z, x_1).$$

Since l' intersects l_3, $det(x_3, y_3, z, z*) = 0$, and

$$det(x_3, y_3, z, x_1)\, det(x_2, y_2, z, y_1) - det(x_3, y_3, z, y_1)\, det(x_2, y_2, z, x_1) = 0. \tag{5.1}$$

If z lies on l_1 or l_2 or l_3, z also satisfies equation 5.1. Equation 5.1 has degree two in the coordinates of z and defines a quadric surface Q; so if z lies on a stabber of L then z must lie on this quadric surface.

We now show that every point on the quadric surface Q lies on some stabber. Let z be any point on Q. If z lies on l_1, then by Lemma 5.1 there is a stabber of L and z. If z does not lie on l_1, then by Lemma 5.1 there is some point $z*$ on l_1 such that the line through z and $z*$ stabs l_1 and l_2. Since z lies on the quadric surface defined above, $det(x_3, y_3, z, z*) = 0$ and this stabber of l_1, l_2 and z also stabs l_3. ■

Any line which does not lie on a quadric surface intersects the quadric surface in at most two points. Thus any four skew lines have at most two stabbers or an infinite number of stabbers. It is easy but tedious to check that this statement is true for any four lines which are not skew.

A set of skew lines that admit an infinite number of stabbers is called **ruled**. The terminology derives from the fact that the set of stabbing lines forms a doubly ruled surface. A set of three skewed lines is trivially ruled. A doubly ruled quadric surface can be partitioned into two sets of lines: every pair of lines from the same set is skewed; every pair of lines from different sets is intersecting. There is thus an obvious duality between the stabbing lines and the lines to be stabbed.

Lemma 5.3. For $m \geq 3$, let $L = \{l_1, \ldots, l_m\}$ be a ruled set of lines lying on the quadric surface Q formed by the stabbers of l_1, l_2, l_3. For every point $z \in Q$ there is a unique stabbing line through z that intersects each line in L.

Proof: We prove the lemma by induction. If $m = 3$, the lemma is true by Lemma 5.2. For $m \geq 4$, assume the lemma is true for $m-1$. Let z be a point on Q. By Lemma 5.1, there is a unique stabber $l*$ of l_1, l_2 and z. By the inductive assumption, there is a unique stabber l' of $L' = \{l_1, l_2, \ldots, l_{m-1}\}$ and z. Since l' stabs l_1, l_2, l_3 and z and $l*$ is the unique stabber of l_1, l_2, l_3 and z, $l* = l'$. Similarly, there is a unique stabber l'' of $L'' = \{l_1, l_2, \ldots, l_{m-2}, l_m\}$ and ~~z,~~ and $l* = l''$. Thus $l*$ is the unique stabber of $L = L' \cup L''$ and z. ■

Let $L = \{l_1, l_2, l_3\}$ and let l_i be parametrized by $\lambda x_i + \mu y_i$. Given the quadric surface Q formed by the stabbers of L, define a function Φ which maps every point z on Q to the unique point $z*$ on l_1 such that the line through z and $z*$ is a stabber of L. Formally,

$$\Phi(z) = \begin{cases} \det(x_2, y_2, z, y_1)\, x_1 - \det(x_2, y_2, z, x_1) y_1 & z \in Q - l_2 \quad (5.2) \\ \det(x_3, y_3, z, y_1)\, x_1 - \det(x_3, y_3, z, x_1) y_1 & z \in l_2 \quad (5.3) \end{cases}$$

It follows from Lemmas 5.1 and 5.2 that Φ performs the function described above. In the sequel we will need the following fact about $\Phi(z)$.

Lemma 5.4. Φ is a continuous function from Q to l_1.

Proof: For every point $z \in Q - l_2$, we can define a neighbourhood $N(z)$ which does not intersect l_2. For every point $z' \in N(z)$,

$$\Phi(z') = det(x_2, y_2, z', y_1)\, x_1 - \ det(x_2, y_2, z', x_1)\, y_1 .$$

Note that if $z' \in l_1$, this formula sets $\Phi(z') = z'$. Thus Φ is continuous at z. Now suppose $z \in l_2$, and consider a neighbourhood $N(z)$ small enough to be disjoint from l_3. If we first apply Lemma 5.1 to l_1 and l_2 and then reapply it to l_1 and l_3, we see that equations (5.2) and (5.3) agree up to a non-zero multiple for any point not on either l_2 or l_3. Therefore, for any $z* \in N(z)$, $\Phi(z*)$ is given by equation (5.3) and is therefore continuous in this region. Again Φ is continuous at z. ■

Applying Lemma 5.4, if s is a line segment on Q with endpoints x and y, then $\Phi(s)$ is a line segment on l_1 with endpoints $\Phi(z)$ and $\Phi(y)$. Note that a line segment in projective space may correspond to either a segment or two half-lines in affine space.

If p is a polyhedron which intersects Q, then $\Phi(Q \cap p)$ is a closed set in l_1 and is composed of the union of line segments in l_1. If N is an open set inside p, then $\Phi(Q \cap N)$ is an open set in l_1. Therefore, the endpoints of the line segments forming $\Phi(Q \cap p)$ must correspond to stabbers which do not intersect the interior of p. These stabbers must pass through some edge of p.

As a consequence of Lemmas 5.1 and 5.2, we have the following Helly-like theorem.

Theorem 5.1. A set of $m \geq 6$ lines in P^3 have a stabbing line if and only if every six of the lines has a stabbing line.

Proof: Let $L = \{l_1, \ldots, l_m\}$ denote the set of lines. Assume at first that they are skewed. If some set of four lines admits a unique stabber, then the conclusion is immediate. Suppose next that some set of four lines, say l_1, l_2, l_3, l_4, admits exactly two stabbing lines l' and l''. If neither is a stabber for L, then there is some line l_i missed by l' and some line l_j missed by l''. But this is impossible, since then there would be no stabber of $l_1, l_2, l_3, l_4, l_i, l_j$, a contradiction. There remains the case that each set of four lines admits an infinity of stabbers. But in this case, it follows from Lemma 5.2 that all of the lines must lie in a quadric surface Q. Any line in this surface that intersects three of the lines must intersect all of them.

Now suppose that the lines are not skewed. Two lines, say l_1 and l_2, intersect at point z. If all the lines contained point z, the conclusion is immediate, so assume some line l_3 does not contain z. Let h denote the plane containing l_1 and l_2, let h' denote the plane containing l_3 and z and let l be the line $h \cap h'$. If all the lines lie in plane h or h' or intersect l, then l is stabbing line for L. Otherwise some line l_4 does not lie in h or h' and does not intersect l. l_1, l_2, l_3, l_4 have at most two stabbers so we can proceed as above. The theorem follows. ∎

The theorem generalizes to all dimensions. A skewed set of line segments is ruled if the set of lines containing the line segments is ruled. Theorem 5.1 has the following corollary.

Corollary 5.1. A set of $m \geq 6$ skewed segments in E^3 that are not ruled have a stabbing line if and only if every six segments has a stabbing line.

We now turn to the problem of finding stabbing lines for convex polyhedra in E^3. Let $P = \{p_1, \ldots, p_m\}$ be a set of disjoint polyhedra in E^3. Using the previous lemmas, we can prove an 'extremal' theorem for polyhedra. We first need a lemma about 'extremal' stabbing lines in the plane.

Lemma 5.5. Let $P = \{p_1, \ldots, p_m\}$ be a set of $m \geq 2$ disjoint convex polygons, in the plane that admit a stabbing line. There exist two distinct polygons p_i and p_j and vertices $x \in p_i$ and $y \in p_j$, such that the line through x and y is a stabbing line for P.

Proof: Let l be a stabber for P. Translate l until it goes through some vertex x of some polygon p_i. Rotate l about x, until it passes through a vertex y of some polygon $p_j \neq p_i$. l is still a stabber for P and l is the line through x and y. ∎

Theorem 5.2. $P = \{p_1, \ldots, p_m\}$, $m \geq 2$, has a stabbing line if and only if there exists a stabbing line through:

(a) Two vertices in two distinct p_i; or

(b) One vertex and two skewed edges in three distinct p_i; or

(c) Two co-planar non-collinear edges in two distinct p_i; or

(d) Three skew edges in three distinct p_i.

Proof: Assume there is a stabbing line of P. There are two cases.

Case 1. There exists a stabbing line l of P which passes through a vertex of some polyhedron in P.

Assume l passes through the vertex x of polyhedron p_i. Rotate l around x in any direction, until it passes through an edge e of some polyhedron $p_j \neq p_i$. Let h be the plane containing x and e, let $p_i' = x$, $p_j' = e$, and $p_k' = h \cap p_k$, $k \neq i, j$. By Lemma 5.5, there exists a stabbing line l' through two vertices of two distinct polygons. $p_i' = x$ is a point, so l' must contain x. l' must also contain some other vertex y belonging to some $p_k' \neq p_i'$. If y is a vertex of p_k, then l' is a stabbing line through two vertices, x, y, satisfying condition a). If y is not a vertex of p_k, then it must lie on some edge e' which is not contained in h, and l' is a stabbing line through x, e and e', satisfying condition b).

Case 2. No stabbing line of P passes through a vertex of any polyhedron in P.

Let l be a stabbing line of P. Translate l in any direction, until it intersects an edge e of some polyhedron p_i. Let h be the plane containing l and e, let $p_i' = e$ and let $p_k' = h \cap p_k$, $k \neq i$. By Lemma 5.5, there exists a line stabber l' through vertices $x \in p_j'$ and $y \in p_k'$. Line l' must not go through an endpoint of e or else l' would be a line stabber through a vertex of p_i. Therefore p_i', p_j' and p_k' are distinct polygons.

Let x and y lie on edges e' and e'' of p_j and p_k, respectively. By assumption l' does not intersect the endpoints of e' or e'', so e' and e'' do not lie in h and are not collinear. If e' and e'' are co-planar, then condition c) is satisfied. Otherwise, e, e', e'' form three skew edges and condition d) is satisfied. ■

5.3 Algorithms for Line Stabbing in Three Dimensions

The algorithms follow quite naturally from the theory. Given a set $S = \{s_1, \ldots, s_n\}$ of n pairwise disjoint line segments in E^3, we can find a stabber of S in $O(n \log n)$ time.

Choose three line segments s_1, s_2, s_3. If two of these line segments are collinear, then the line containing the line segments is the only candidate stabber. Check whether this line stabs S in $O(n)$ time. If two of these line segments are co-planar, then a stabber of S must lie in the plane containing the line segments. Apply the algorithm of Edelsbrunner et al. for stabbing line segments in the plane in $O(n \log n)$ time[19].

If the three line segments, s_1, s_2, s_3, are skew embed E^3 in P^3 and let l_1, l_2, l_3 be the projective lines containing s_1, s_2, s_3, respectively. The stabbers of these lines form a quadric surface Q. Check in constant time whether each s_i lies on Q. If some s_i does not lie on Q, then it intersects Q in at most two points. By lemma 5.3, there are at most two stabbers of s_1, s_2, s_3 and s_i. Check whether these stabbers stab S in $O(n)$ time.

Finally, if all s_i lie on Q, define the function Φ as in Lemma 5.4 which maps every point of Q onto l_1 such that the line from z to $\Phi(z)$ is a stabber of l_1, l_2 and l_3. By sorting the endpoints of the line segments $\Phi(s_i)$, we can find the intersection of all the line segments

$\Phi(s_i)$ in $O(n \log n)$ time. To each point in that intersection there corresponds a line stabber of S.

Let $P = \{p_1, \ldots, p_m\}$, $m \geq 2$, be a set of disjoint polyhedra in E^3. Let n be the total number of vertices and edges in P. We will show how to find a line stabber for P in $O(n^4 \log n)$ time.

Let p be any polyhedra with t edges which has been preprocessed using the techniques given by Dobkin and Kirpatrick in [14] for fast reporting of polyhedral intersections. Let Q be a quadric surface formed by the stabbers of three skew lines, l_1, l_2, l_3, and let Φ be the function defined in Lemma 5.4 mapping Q to l_1. We can construct $\Phi(Q \cap p)$ in $O(t \log t)$ time. Let $X = \{\Phi(Q \cap e) : e$ is an edge of p which does not lie in $Q\}$. X is composed of at most $2t$ points on l_1. Sort the points of X. These points divide l_1 into at most $2t$ line segments, each of whose interiors is either contained in $\Phi(Q \cap p)$ or in $l_1 - \Phi(Q \cap p)$. For every such line segment s_i, determine if the interior of s_i is in $\Phi(Q \cap p)$ by choosing some point x from the interior of s_i and querying whether the unique stabber of x, l_2 and l_3 also stabs p. $\Phi(Q \cap p)$ is the union of all the line segments s_i whose interior lies in $\Phi(Q \cap p)$ and the points in X.

Since p has been preprocessed using the techniques in [14], it takes $O(\log t)$ time to find if a line intersects p. There are at most $2t$ such queries. Sorting takes $O(t \log t)$ time so the total complexity for this algorithm is $O(t \log t)$.

To find a line stabber for the polyhedra in P we first preprocess the polyhedra as in[14]. We then test for the four possible cases in Theorem 5.2:

For every two vertices in distinct polyhedra p_i, p_j:

Check if the line through the two vertices stabs all the polyhedra.

For every vertex and two skew edges in distinct polyhedra p_i, p_j, p_k:

Find the stabber of the vertex and two edges if it exists and check if it stabs all the polyhedra.

For every two co-planar edges in distinct polyhedra p_j, p_k:

Let h be the plane containing the two edges. Intersect each of the polyhedra with h to form m polygons and use the algorithm of Edelsbrunner et al[19]. to find any stabbers of P which lie in h.

For every three skew edges in distinct polyhedra p_i, p_j, p_k:

Embed $\mathbf{E}^3$ in $\mathbf{P}^3$ and let l_1, l_2, l_3 be the projective lines containing the three skew edges. Let Q be the quadric surface formed by the stabbers of these lines. Define Φ as in Lemma 5.4. Apply the algorithm above to construct $\Phi(Q \cap p)$ for each $p \in P$. Let X be the set of all the endpoints of segments in $\Phi(Q \cap p), p \in P$. X is a collection of at most $2n$ points. Sort the points in X. By scanning the points in X in order, intersect all the $\Phi(Q \cap p), p \in P$. A point in this intersection has a unique stabbing line which stabs all the elements in P.

Preprocessing all the polyhedra takes $O(n^2)$ time. The total time of this algorithm is dominated by the last step. This step takes a total of $O(n \log n)$ time to construct $F(Q \cap p)$ for each polyhedron p. It also takes $O(n \log n)$ time to sort X. The last step is executed $O(n^3)$ times for a total of $O(n^4 \log n)$ time complexity.

Chapter 6

Separation Algorithms

6.1 Introduction

Let A be a family of n pairwise disjoint convex polygons in the plane. Line segments and points are also considered convex polygons, albeit degenerate ones. The upper bounds on the number of orders in which directed lines can intersect A is a function of the size of the minimum strict separation set for A. If a family A has a small strict separation set, then by Theorem 4.1 there are few geometric permutations on A. This property of separation sets suggests the problem of finding small ones.

Theorems 4.2 and 4.3 ensure that there exists a strict separation set for a family of n pairwise disjoint convex polygons with size $12n$. In fact, the proof of Theorem 4.3 is a construction which can be turned into an algorithm for finding a strict separation set of size $12n$. As in Theorem 4.3, triangulate the family of polygons and find the neighbours of each polygon. Strictly separate each of the polygons from each of its neighbours and from each of the triangulation edges between neighbours. These separation lines strictly separate a polygon not only from its neighbours but from all other polygons, as argued in Theorem 4.3. There are at most $12n$ separation lines and they form a strict separation set of size $12n$. The algorithm is dominated by the triangulation time. Using standard techniques[48], triangulation takes $O(m \log m)$ time where m is the number of edges.

Given a family of convex polygons in the plane, decide if there exists a separation set of size k. Given a family of convex polygons in the plane, decide if there exists a strict separation set of size k. These problems we call the the **separation set problem** and the **strict separation set problem**. In Section 6.2 it will be shown that both these problems are NP-complete.

Theorem 4.1 does not really depend upon the size of a strict separation set, but on the number of different slopes of lines in the set. Each strict separation set has an associated set of slopes which are the slopes of all the lines in the set. Given a family A of n convex polygons with a total of m edges in the plane, find the minimum size set of slopes associated with any strict separation set for A. This problem is called the **separation slope problem.** By mapping slopes to points on a circle, this problem can be transformed to one of finding a minimum point cover for a set of k open arcs on the circle, called the **point cover of arcs problem**. Section 5.2 solves the point cover of arcs problem in $O(k \log k)$ time. Finding the minimum set of slopes takes $O(n^2 \log n + n^2 \log m)$ time.

6.2 The Separation Set Problem

The proof that the separation set problem is NP-complete is a reduction from vertex cover for planar graphs. A **vertex cover** is a set of vertices such that every edge is incident with some vertex. The **vertex cover problem** for planar graphs is decide if there exists a vertex cover of size k for a given planar graph. The vertex cover problem for planar graphs is NP-complete[24]. D is a dual transform which maps points to lines and lines to points. For a point x in homogeneous coordinates, $D(x) = \{u : x \cdot u = 0\}$. For a hyperplane $h = \{x : u \cdot x = 0\}$, $D(h) = u$.

Theorem 6.1. The separation set problem is NP-complete.

Proof: Let A be a family of convex polygons in the plane. The lines in any separation set can be perturbed until they go through two vertices of two distinct polygons. The possible separation sets for A need only be chosen from the polynomial number of lines which pass through two vertices of polygons in A. In polynomial time a non-deterministic Turing machine could guess a separation set of size k and check whether it did in fact separate all the pairs of polygons in A. Thus the separation set problem is in NP.

Let s_1 and s_2 be two line segments which intersect at their endpoints. All the lines which separate s_1 from s_2 lie in the cone formed by s_1 and s_2. (See Figure 6.1.) If no line through the origin separates s_1 and s_2, then D maps the separators of s_1 and s_2 to a line segment in the dual space. (See Figure 6.2.)

The process can be reversed, mapping a line segment in the dual space to the separators of a pair of line segments in the primal. Given some line segment in the dual space with endpoints x, y, let z be the point $D(x) \cap D(y)$ in the primal space and let l be the line through z and the origin. Let s_1 and s_2 be two line segments lying on $D(x)$ and $D(y)$, respectively, which share a common endpoint z and lie to one side of l. The line segment in the dual space with endpoints x, y maps under D to the set of line separators of s_1 and s_2. Note that s_1 and s_2 may have any positive length.

We are now ready to show how an instance of the planar vertex cover problem transforms to an instance of the weak separation problem. Let G be any planar graph. G has a straight line embedding in the plane. (See Figure 6.3.) Perturb the vertices of this planar embedding so that no two graph edges lie on the same line and no edge lies on a line

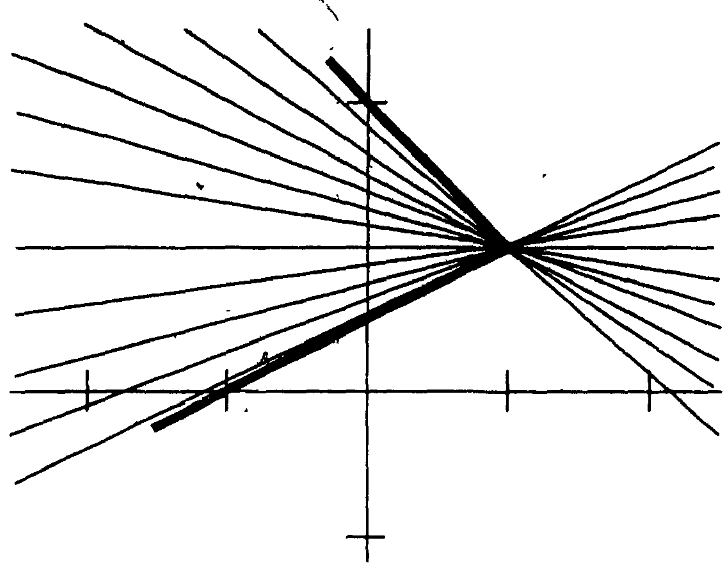

Figure 6.1. Separators of Two Line Segments

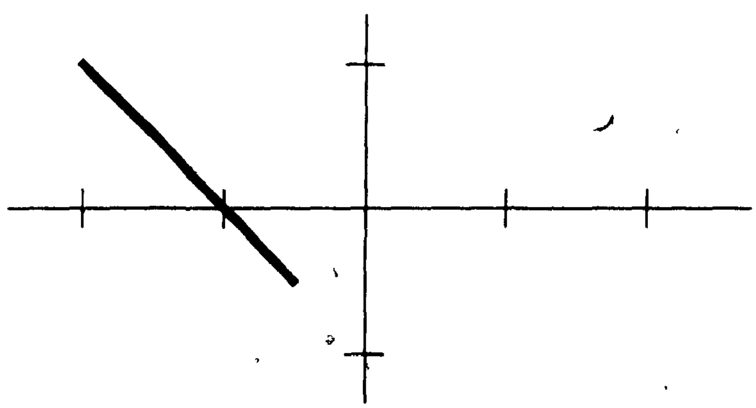

Figure 6.2. Separators of Two Line Segments in Dual.

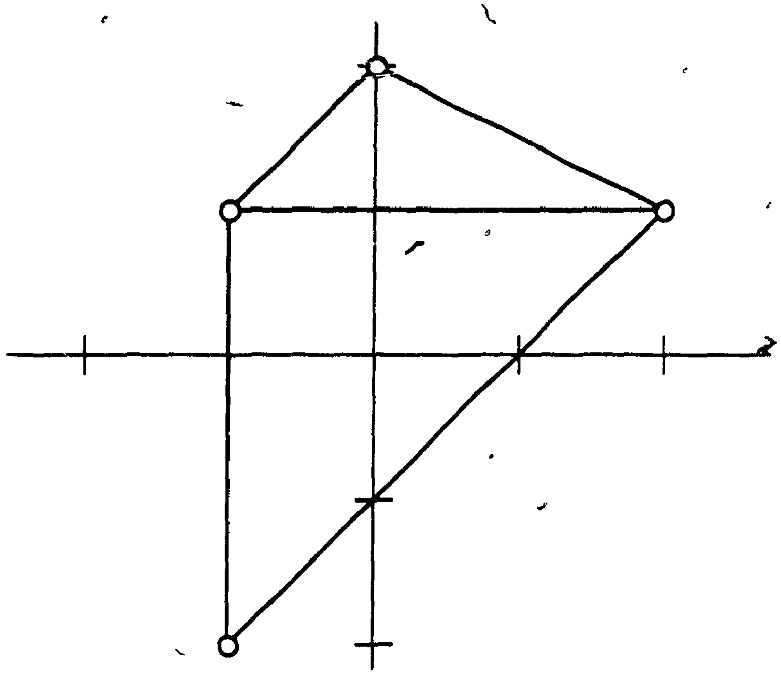

Figure 6.3. Straight Edge Embedding of Planar Graph.

through the origin. Transform the lines containing graph edges to points under the mapping D and let X be the set of these points. Each edge in the planar embedding now corresponds to a unique point in X. Let $\varepsilon > 0$ be a lower bound on the vertical and horizontal distance between any two points in X. For each edge in the planar embedding construct a pair of line segments such that the points on the edge are mapped to the separators of these line segments under the transform D. (See Figure 6.4.) Let T be a set of these pairs of line segments. Each edge in G corresponds to an element, a pair of line segments, in T. Choosing each line segment to have a positive length less than $\frac{\varepsilon}{4}$, ensures that every two elements of T, i.e. every two pairs of line segments, can be separated by a vertical or horizontal line.

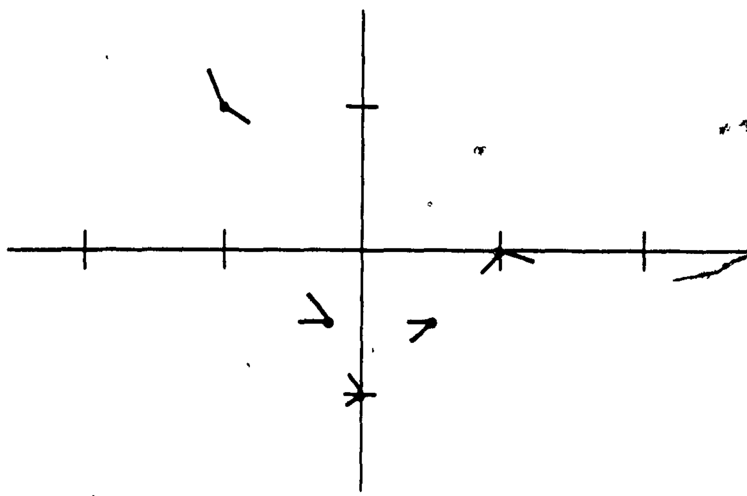

Figure 6.4. Pairs of Line Segments Corresponding to Edges.

Consider the problem of finding a set of k points such that each edge in the planar embedding of G contains at least one point. Such a set of points will be called a point cover for G. Since the relative interiors of no two edges intersect, each point can be moved to an edge endpoint to form a vertex cover. Thus, a point cover for G of size k gives a vertex cover of G of the same size. The mapping D transforms a point covering of edges in G to a set of lines such that each pair of line segments which are an element of T is separated by some line. Similarly, D transforms a set of lines which separate every element of T to a point cover of G. Therefore, there exists a set of k lines which separate every element of T if and only if there exists a vertex cover of size k for G.

Let L'' be a set of vertical and horizontal line segments which separate every two pairs of line segments in T but do not separate any two line segments which are paired. Let l_1 be a translate of the x-axis which lies below all the pairs of line segments in T. Let l_2 be a translate of the y-axis which lies to the left of all the pairs of line segments in T. Place abutting rectangles along l_1 and l_2 so that the each of the lines in L'' is the unique separator for two adjacent rectangles as in Figure 6.5. Let L' be the minimum separation set for these rectangles. L' contains L''.

Let A be the set of rectangles plus the line segments which form the pairs in T. L' is a subset of every separation set for A. The two lines separating the rectangle in the left, bottom corner from its neighbours also separate all the rectangles from all the line segments. Let L be a separation set for A of size $k + |L'|$. Removing L' from L leaves a set of k lines which separate every element of T. Transforming the lines in $L - L'$ to points in the

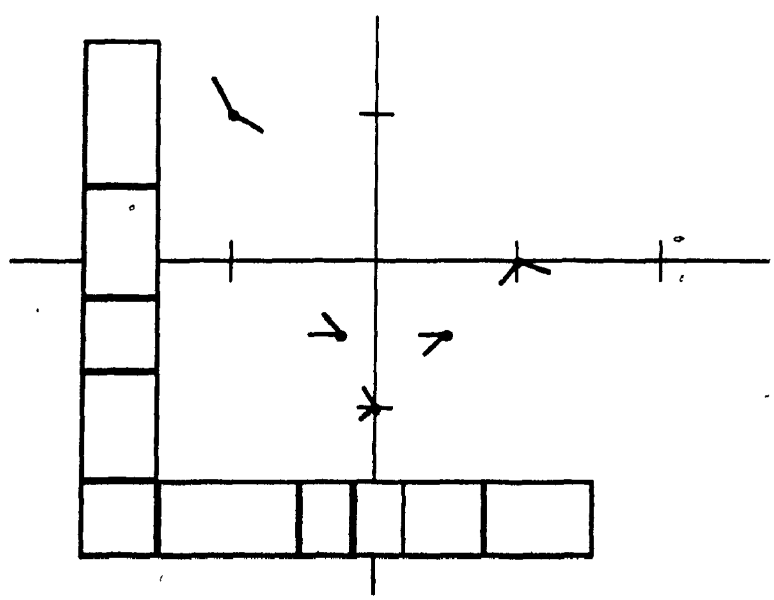

Figure 6.5. Abutting Rectangles.

dual space and moving the points to vertices, produces a vertex cover of size k. A vertex cover of size k gives a set of k lines which separate every element of T. Adding L' to this set, would result in a separation set of size $k + |L'|$. Thus there exists a vertex cover of G of size k if and only if there exists a separation set of size $k + |L'|$ for A.

Finding a straight line embedding of a planar graph takes polynomial time[54]. In this straight line embedding vertices are mapped to points with integer coordinates. By scaling up all the coordinates by a suitable factor, polynomial in the number of vertices and the coordinates of the vertices, the vertices in the embedding can be perturbed to new integer coordinates so that no two edges lie on the same line and no edge lies on a line through the origin. The transformation of edges to points, the creation of line segments and of rectangles, and determining the size of L' all use polynomial time. Therefore, in polynomial time an instance of the vertex cover problem can be reduced to an instance of the separation set problem and the separation set problem is NP-complete.∎

Given a set X of points in $\mathbf{E}^d$ and a closed half-space h^+, $X \cap h^+$ is called a **k-set** where k is the number of points in $X \cap h^+$. The number of different k-sets over all k is $O(n^d)$ and these k-sets can be constructed in $O(n^d)$ time[17].

Theorem 6.2. The strict separation set problem is NP-complete.

Proof: Let A be a family of convex polygons in the plane and let X be the set of all vertices of polygons in A. For all k, construct the k-sets for X. For each k-set choose a line which strictly separates the k-set from the rest of the points in X. Let L be the set of all such chosen lines. Lines in a strict separation set can always be perturbed so that the strict separation set is a subset of L. In polynomial time a non-deterministic Turing machine could guess a strict separation set of size k from the lines in L and check whether it did in fact strictly separate all the pairs of polygons in A. Thus the strict separation set problem is in NP.

To prove NP-completeness a special version of the separation set problem is reduced to the strict separation set problem. Let A be a set of convex polygons in the plane whose relative interiors are pairwise disjoint. In the proof of Theorem 6.1, the vertex cover problem is reduced to determining whether there exists a separation set of a given size for this restricted family of convex polygons. Thus determining whether there exists a separation set of a given size for a family of convex polygons whose relative interiors are pairwise disjoint is also NP-complete. We show how to reduce this problem to the strict separation set problem.

If the polygons in A are pairwise disjoint, then the lines in a separation set for A can be perturbed in polynomial time to form a strict separation set for A. A strict separation set for A is also a separation set for A. Thus finding a separation set or a strict separation set of size k is equivalent for a family of pairwise disjoint polygons.

However, A may be a family of polygons whose relative interiors are pairwise disjoint but which are not themselves pairwise disjoint. In this case, A has a separation set but no strict separation set. Each polygon in A will be shrunk to form a new family of polygons A' such that the A' has a separation set of size k if and only if A has one.

For every pair of polygons A which intersect, there must be some vertex x in one of the polygons a which does not lie in the other polygon, b. For $\gamma \in [0,1]$, let $a(\gamma)$ be the contraction of a by a factor of γ about x where:

$$a(\gamma) = \{ x + \gamma(y - x) \mid y \in a \} .$$

x and y are represented in Cartesian coordinates. For each line through two vertices in A, there is a set of vertices not contained in that line. Let γ' be the maximum value of γ such that a line through two vertices v and v' in $A - \{a\} + \{a(\gamma)\}$ contains a point which is not contained in the line through v and v' in A. γ' is a rational number whose numerator and denominator are bounded by polynomial functions of the coordinates of the vertices. Shrink a by some $\gamma* = \frac{\gamma' + 1}{2}$ and scale up all the coordinates by a suitable factor so that the vertices of $a(\gamma*)$ have integer coordinates. $\gamma*$ is less than one, the relative interiors of a and b are disjoint, and x does not lie in b, so $a(\gamma*)$ does not intersect b.

Let $A' = A - \{a\} + \{a(\gamma*)\}$. By the choice of $\gamma*$, any line through two vertices v and v' in A' separates the same sets as the line through the two vertices v and v' in A and vice versa. A separation set for A is obviously a separation set for A'. A separation set for A' can be transformed to a separation set for A of the same size. Perturb the lines in the separation set in A' so that every line goes through two vertices. Transform the line which goes through vertices v and v' in A' to the line which goes through vertices v and v' in A. Since each line separates the same polygons before and after the transformation, the transformation maps a separation set for A' to a separation set for A.

Repeating the shrinking procedure for each pair of intersecting sets, produces a family A' of pairwise disjoint convex sets which has a separation set of size k if and only if A has a separation set of size k. Since A' is pairwise disjoint, A' has a separation set of size k if and only if A' has a strict separation set of size k. Thus the separation set problem reduces to the strict separation set problem in polynomial time. ■

6.3 The Separation Slope Problem and the Point Cover of Arcs Problem

Let A be a family of n pairwise disjoint convex polygons with a total of m edges. Assume that the relative interiors of the polygons in A are pairwise disjoint so A has a strict separation set. The slopes of the strict separators of a pair of polygons in A map to an open arc b on the unit circle. This pair of polygons has two critical separators which are tangent to both polygons in the pair. The slopes of the critical separators map to the endpoints of arc b. The critical separators can be found in $O(\log m)$ time[49], so each arc can be constructed in $O(\log m)$ time. The entire set of $\binom{n}{2}$ arcs can be constructed in $O(n^2 \log m)$ time.

Each point which lies in an open arc generated from a pair of polygons corresponds to the slope of some strict separator of the pair. Finding a minimum size set of separation slopes corresponds to finding a minimum set of points on the circle such that each open arc contains some point.

Problems concerning arcs on the circle have been studied in graph theory where arcs are transformed to circular-arc graphs[9, 25, 26]. **Circular arc graphs** are created by mapping the arcs on a circle to vertices of a graph. Two vertices are joined by an edge if their corresponding arcs intersect.

Each point on the unit circle corresponds to a clique in the circular-arc graph. **A clique cover of a graph** is a set of cliques such that each vertex is in some clique. Finding a minimum point cover seems quite like finding a minimum clique cover of the circular-arc graph. Unfortunately, each clique in a circular arc graph does not correspond to one point on the unit circle. (See Figure 6.6.) I saw no way to transform the point cover of arcs problem to a problem on circular arc graphs and so was forced to develop my own algorithm.

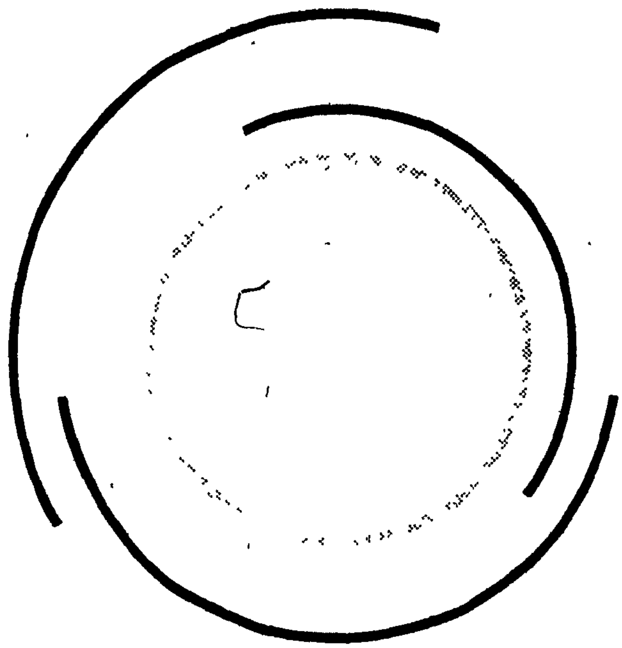

Figure 6.6. Three Arcs in a Clique.

Let K be a set of open arcs on a circle. If the endpoint of some arc is not the endpoint of any other arc, the arc can be lengthened and the size of the minimum cover will not change. If some arc contains some other arc, then the first arc can be removed and the size of the minimum point cover will again not change. A **normalized set of arcs** is a set of arcs where every endpoint of an arc is the endpoint of some other arc and no arc contains any other arc in the set. If an arc extends in the clockwise direction from an endpoint, then the endpoint is called a left endpoint. Otherwise, the endpoint is called a right endpoint. In a normalized set of arcs, no point may be the left endpoint or the right endpoint of two arcs.

Given a set K of k open arcs, a normalized set of arcs K' with the same size minimum cover can be constructed from K in $O(k \log k)$ time. Initialize K' to be the set of arcs K. Sort the endpoints of the arcs in K' on the circle. By scanning the endpoints first in clockwise and then in counter-clockwise order, lengthen each arc until both its left endpoint lies on the right endpoint of some other arc and its right endpoint lies on the left endpoint of some other arc.

Scan the circle in a clockwise direction. At each left endpoint encountered in the scan, add the new arc to the beginning of a list $LARC$. At each right endpoint belonging to some arc a, check whether a is contained in list $LARC$. If a is in the list, all the arcs behind it in the list contain a. Mark these arcs for deletion, push them onto a stack $ADEL$ and remove them and arc a from list $LARC$.

For each endpoint position x on the circle, there is some smallest arc a whose left endpoint lies at x and some smallest arc b whose right endpoint lies at x. Set pointers from the left endpoint of a to b and from the right endpoint of b to a.

Pop an arc a from the top of stack $ADEL$ and delete a from K'. If the left endpoint of a contains a pointer to the right endpoint of some arc b, move the right endpoint of b clockwise around the circle to the position y of the closest left endpoint of some arc. Let c be the smallest arc with left endpoint at y. The left endpoint of c points to some arc d. If b is smaller than d, set the right endpoint of b to point to c and the left endpoint of c to point to b. Mark d for deletion and push d on stack $ADEL$, if it is not already there. If d is smaller than b, mark b for deletion and push b on stack $ADEL$, if it is not already there. Repeat this procedure until stack $ADEL$ is empty.

K' was produced by moving the endpoints of arcs not marked for deletion in stack *ADEL* to meet other arcs and by deleting redundant arcs, so K and K' have the same size minimum cover. The endpoints of each arc not marked for deletion are always moved so that they always coincide with the endpoints of some other arc. The deletions ensure that no arc may be contained in any other arc. Thus K' is a normalized set of arcs. The time to construct a normalized set of arcs is dominated by the initial sorting which takes $O(k \log k)$ time.

Given a normalized set K' of k' arcs, a minimum cover can be found for K' in $O(k' \log k')$ time. Sort the endpoints of arcs in K'. An arc a in K' **connects** to an arc b on its right if the right endpoint of a meets the left endpoint of b. a **connects** to b on its left if the left endpoint of a meets the right endpoint of b. A **chain** of arcs from a_1 to a_j is a set of arcs $\{a_1, \ldots, a_j\}$ where a_i connects to a_{i+1} on its right.

Let x be a point on the circle which is not an endpoint of some arc and let m be the number of arcs containing x. Let y be any other point which is not an endpoint. The m arcs containing x begin m chains which lead to m arcs containing y. If $m' > m$ arcs contained y, these m' arcs would begin m' chains leading to m' arcs containing x. Thus every point in the circle which is not an endpoint is contained in exactly m arcs.

Let ε be the minimum distance between any two left endpoints of arcs in K'. Choose some arc a from K'. There is a chain C of $\left\lceil \frac{k'}{m} \right\rceil$ arcs starting at a whose closure covers the circle. Let X be the set of $\left\lceil \frac{k'}{m} \right\rceil$ points distance $\frac{\varepsilon}{2}$ from the left endpoints of arcs in C. Since C covers the circle and no arc may be contained in an arc of C, X is a point cover for all the arcs in K'.

Assume there was a point cover Y of size less than $\left\lceil \frac{k'}{m} \right\rceil$. If a point in Y lies on the left endpoint of some arc, the point can be moved slightly and Y will remain a point cover of the arcs. Assume no point in Y lies on a left endpoint. Sort the points in Y in order around the circle. Between some two adjacent points $x, y \in Y$, there must lie at least $\left\lceil \frac{k'}{\left\lceil \frac{k'}{m} \right\rceil - 1} \right\rceil \geq m+1$ left endpoints of arcs. There are only m chains of arcs extending from x to y. One of the chains must contain two left endpoints. This means that some arc lies between x and y. This arc is not covered by Y and Y is not a point cover for K'. It follows that X is a minimum point cover for K'.

Sorting the points in K' takes $O(k' \log k')$ time. Finding a set Y takes $O(k')$ time. A minimum point cover for K' can be found by reversing the process by which K' was derived from K. A point cover of arcs for K can be found in $O(k \log k)$ time.

The separation slope problem on n polygons with m edges can be reduced to a point cover of $O(n^2)$ arcs problem in $O(n^2 \log m)$ time. Applying the algorithm above for the arc cover problem gives an $O(n^2 \log m + n^2 \log n)$ algorithm for solving the separation slope problem.

Chapter 7

Conclusion

This thesis grew out of a number of questions, both mathematical and algorithmic, which turned out to be interrelated. Questions about ordering, separation, the necessity of certain conditions in Hadwiger's Theorem, all pointed to the relationship between stabbing and separation.

Some progress was made in generalizing Hadwiger's Theorem. The pairwise disjointness condition was eliminated from the theorem. Theorems by Katchalski and Goodman and Pollack, themselves generalizations of Hadwiger's Theorem, were generalized to new families of compact convex sets. However, a necessary and sufficient condition for the existence of hyperplane transversals for unrestricted families of compact convex sets is still unknown. Just as Hadwiger's Theorem was true without the pairwise disjointness condition, I conjecture that Goodman and Pollack's Theorem is true without the $d-2$-separability condition. Necessary and sufficient conditions for the existence of k-transversals for k-flats other than points or hyperplanes are even more of a mystery.

The number of ways in which a line intersects a family of convex sets was shown to relate to the arrangement of hyperplane separators. The number of ways in which a plane in E^3 intersects a family also relates to such an arrangement. For higher dimensions, I conjecture but cannot prove that there is also such a relation.

Much work has been done by computer scientists and mathematicians on point and hyperplane stabbing. The algorithms in Chapter 5 for line stabbing in E^3 are a first attempt at k-flat stabbing for k-flats other than points and hyperplanes. It is highly unlikely that the algorithm for stabbing polyhedra is optimal. With the tools of algebraic geometry, it may be possible to generalize these algorithms to k-flat stabbing in higher dimensions.

Chapter 6 contains some problems about the construction of separation sets. Other problems on separation sets are still unexplored. How does one construct separation sets of small size in higher dimensions? How does one construct small k-separation sets for k-separable families?

Index

Bibliography

1. Aho, A.V., Hopcroft, J.E., and Ullman, J.D., *The Design and Analysis of Computer Algorithms*, Addison-Wesley, Reading, Mass. (1974).

2. Atallah, M. and Bajaj, C., "Efficient algorithms for common transversals," *Information Processing Letters* **25** pp. 87-91 (May 1987).

3. Avis, D., "Diameter partitioning," *Discrete and Computational Geometry* **1** pp. 265-276 (1986).

4. Avis, D. and Doskas, M., "Algorithms for high dimensional stabbing problems," Technical Report SOCS 87.2, McGill University, Montreal (January 1987).

5. Avis, D. and Wenger, R., "Algorithms for line stabbers in space," *Proceedings of the 3rd ACM Conference on Computational Geometry*, pp. 300-307 (1987).

6. Avis, D. and Wenger, R., "Polyhedral line transversals in space," *Discrete and Computational Geometry*, (to appear).

7. Bondy, A. and Murty, U.S.R., *Graph Theory with Applications*, American Elsevier (1976).

8. Borsuk, K., *Multidimensional Analytic Geometry*, Polish Scientific Publishers, Warsaw (1969).

9. Buckingham, M. A., "Circle graphs," Report No. NSO-21, Courant Institute of Mathematical Sciences (October 1980).

10. Carroll, L., *Alice's Adventures in Wonderland*, Octopus Books Limited, London (1982).

11. Chazelle, B. and Dobkin, D., "Detection is easier than computation," *Proceedings of the 12th Annual ACM Symposium on Theory of Computing*, pp. 146-153 (1980).

12. Chvátal, V., *Linear Programming*, Freeman, New York (1983).

13. Danzer, L., Grünbaum, B., and Klee, V., "Helly's theorem and its relatives," *Convexity, Proc. Symp. Pure Math.* 7 pp. 100-181 Amer. Math. Soc., (1963).

14. Dobkin, D. and Kirpatrick, D., "Fast detection of polyhedral intersection," *Theoretical Computer Science* 27 pp. 241-253 (1983).

15. Edelsbrunner, H., "Intersection problems in computational geometry," PhD. Thesis, Tech. Univ. Graz (1982).

16. Edelsbrunner, H., "Finding transversals for sets of simple geometric figures," *Theoretical Computer Science* 35 pp. 55-69 (1985).

17. Edelsbrunner, H., *Algorithms in Combinatorial Geometry*, Springer-Verlag (1987).

18. Edelsbrunner, H., Guibas, L. J., and Sharir, M., "The upper envelope of piecewise linear functions: algorithms and applications," Report UIUCDCS-R-87-1390, University of Illinois (November 1987).

19. Edelsbrunner, H., Maurer, H. A., Preparata, F. P., Rosenberg, A.L., Welzl, E., and Wood, D., "Stabbing line segments," *BIT* 22 pp. 274-281 (1982).

20. Edelsbrunner, H., Robison, A., and Shen, X., "Covering convex sets with non-overlapping polygons," *Manuscript*, (1987).

21. Edelsbrunner, H. and Sharir, M., "The maximum number of ways to stab n convex non-intersecting objects in the plane is 2n-2," *Discrete and Computational Geometry*, (to appear).

22. ElGindy, H. and Toussaint, G. T., "Efficient algorithms for inserting and deleting edges from triangulations," in *Proc. Int. Conference on Foundations of Data Organization*, Kyoto (May 22-24, 1985).

23. Fary, I., "On straight line representations of planar graphs," *Acta. Sci. Math. Szeged.* 11 pp. 229-233 (1948).

24. Garey, M.R. and Johnson, D.S., *Computers and Intractability*, Freeman, New York (1979).

25. Gavril, F., "Algorithms on circular-arc graphs," *Networks* 4 pp. 357-369 (1974).

26. Golumbic, M. C., *Algorithmic Graph Theory*, Academic Press, New York (1980).

27. Goodman, J. E. and Pollack, R., "Multidimensional sorting," *SIAM J. Computing* 12 pp. 484-507 (1983).

28. Goodman, J. E. and Pollack, R., "Semispaces of configurations, cell complexes of arrangements," J. Combin. Theory Series A 37 pp. 257-293 (1984).

29. Goodman, J. E. and Pollack, R., "Hadwiger's transversal theorem in higher dimensions," *Journal of Am. Math. Soc.* 1(2)(to appear).

30. Grünbaum, B., *Convex Polytopes*, Wiley (1967).

31. Hadwiger, H., "Über Eibereiche mit gemeinsamer Treffgeraden," *Portugal Math.* **16** pp. 23-29 (1957).

32. Hadwiger, H., Debrunner, H., and Klee, V., *Combinatorial Geometry in the Plane*, Holt, Rinehart and Winston (1964).

33. Harary, F., *Graph Theory*, Addison-Welsley, Reading, Massachusetts (1969).

34. Helly, E., "Über Mengen konvexer Körper mit gemeinschaftlichen Punkten," *Jahrb. Deut. Math. Verein* 32 pp. 175-176 (1923).

35. Horn, A. and Valentine, F., "Some properties of L sets in the plane," *Duke Mathematical Journal* 16 pp. 131-140 (1949).

36. Jaromczyk, Jerzy W. and Kowaluk, Miroslaw, "Skewed projections with an application to line stabbing in R^3," *Proceedings of the 4th ACM conference on computational geometry*, (1988).

37. Katchalski, M., "Thin sets and common transversals," *J. of Geometry* 14 pp. 103-107 (1980).

38. Katchalski, M., Lewis, T., and Liu, A., "Geometric permutations and common transversals," *Discrete and Computational Geometry* 1 pp. 371-377 (1986).

39. Katchalski, M., Lewis, T., and Liu, A., "Geometric permutations of disjoint translates of convex sets," *Discrete Mathematics*, (to appear).

40. Katchalski, M., Lewis, T., and Zaks, J., "Geometric permutations for convex sets," *Discrete Mathematics* 54 pp. 271-284 (1985).

41. Klee, V. L., Jr., "Common secants for plane convex sets," *Proc. Am. Math. Soc.* **5** pp. 639-641 (1954).

42. Lay, S. R., *Convex Sets and Their Applications*, John Wiley & Sons, New York (1982).

43. Lewis, T., "Two counterexamples concerning transversals for convex subsets of the plane," *Geometriae Dedicata* 9 pp. 461-465 (1980).

44. Megiddo, N., "Linear-time algorithms for linear programming in R^3 and related problems," *SIAM J. Computing* 12 pp. 759-776 (1983).

45. Megiddo, N., "Linear programming in linear time when the dimension is fixed," *J. ACM* 31 pp. 114-127 (1984).

46. Pach, J. and Sharir, M., "The upper envelope of piecewise linear functions: combinatorial analysis," *Discrete and Computational Geometry*, (to appear).

47. Papadimitriou, C. H. and Steiglitz, K., *Combinatorial Optimization*, Prentice-Hall, Englewood Cliffs (1982).

48. Preparata, F. P. and Shamos, M. I., *Computational Geometry*, Springer-Verlag (1985).

49. Rohnert, H., "Shortest paths in the plane with convex polygonal obstacles," Tech. Rept. A 85/06, University of Saarbrucken (1985).

50. Rudin, W., *Principles of Mathematical Analysis*, McGraw-Hill (1964).

51. Sharir, M., "Almost linear bounds on the length of general Davenport-Schinzel sequences," Technical Report 29/1985, Tel Aviv U., Tel Aviv (1985).

52. Sharir, M. and Hart, S., "Nonlinearity of Davenport-Schinzel sequences and of generalized path compression schemes," *Combinatorica* 6 pp. 151-177 (1986).

53. Sharir, M. and Livne, R., "On minima of functions, intersection patterns of curves, and Davenport-Schinzel sequences," *Proc. 26th Ann. IEEE Symp. on Theory of Computing*, pp. 312-320 (1986).

54. Tutte, W. T., "How to draw a graph," *Proc. London Math. Soc.* 13 pp. 743-768 (1963).

55. Tverberg, H., "A separation property of plane convex sets," *Mathematica Scandinavica* 45 pp. 255-260 (1979).

56. Valentine, F.A., "The dual cone and Helly type theorems," *Convexity, Proc. Symp. Pure Math.* 7 pp. 473-493 Amer. Math. Soc., (1963).

57. Vincensini, P., "Sur une extension d'un théorème de M. J. Radon sur les ensembles de corps convexes," *Bull. Soc. Math. France* **67** pp. 115-119 (1939).

58. Wenger, R., "A generalization of Hadwiger's Theorem to intersecting sets," Technical Report, TR-SOCS-87.14 (October 1987).

59. Wenger, R., "Upper bounds on geometric permutations," *Discrete and Computational Geometry*, (to appear).

60. Winder, R. O., "Partitions of n-space by hyperplanes," *SIAM Journal of Applied Math.* 14 pp. 811-818 (1966).