THE COMPLEXITY OF COMPUTING MINIMUM CONVEX COVERS FOR POLYGONS

JOSEPH O'ROURKE
Department of Electrical Engineering and Computer Science
The Johns Hopkins University
Baltimore, MD 21218

ABSTRACT

The computational complexity of the problem of finding a minimum convex cover for a polygon is reviewed. Although an NP-hard lower bound has been proven, no upper bound on the complexity is known. The lack of known restrictions on the Steiner points that may be necessary prevents naive searching algorithms from establishing the solvability of the problem. In this paper the problem is proven to be decidable by using Tarski's decision procedure for geometry.

1. Introduction

Given an arbitrary polygonal region of the plane, perhaps with polygonal holes, it is frequently useful to decompose the region into convex polygons. Such decompositions find application in pattern recognition [17-20] [23], artificial intelligence and graphics [15], and VLSI design [6]. If the convex polygons used in the decomposition are permitted to overlap, the decomposition is called a *convex cover*. The problem discussed in this paper is the computational complexity of finding the *minimum* convex cover, minimum in the sense of having the fewest number of pieces.

The problem formulation in the style of Garey and Johnson [5] is as follows [6]:

COVERING BY CONVEX POLYGONS (CCP)

INSTANCE: A polygon P (not necessarily simply-connected), described by a sequence of n pairs of integer coordinate points in the plane, each pair representing the endpoints of a line segment on the boundary of P, and a positive integer K.

QUESTION: Are there convex polygons C_i, $1 \leqslant i \leqslant K$, whose union (as planar figures) is P ?

Decomposing a polygonal region into *non*-overlapping pieces (finding a *partition*) is a closely related but distinct problem. Minimum partitions of both simply-connected [3] [2] and multiply-connected regions can be found in polynomial time. These results unfortunately do not bear directly on minimum covers.

The use of Steiner points is discussed in the next section, and Section 3 shows that naive algorithms could miss finding the true minimum cover. Section 4 summarizes a proof of NP-hardness for the problem, establishing a lower bound, and the section following proves that the problem is decidable (a new result), establishing an upper bound. Open problems are formulated in the final section.

2. Steiner Points and Naive Algorithms

It is instructive to consider brute-force searching algorithms for finding minimum convex covers: such algorithms at least establish that the problem they solve is *decidable*. In general, *Steiner points*, points that are not vertices of the polygon being decomposed, may be required as vertices of the convex pieces of a minimum cover. If the problem is restricted to find minimum covers without the use of Steiner points, then a searching algorithm is possible: con-

* This research was supported by the National Science Foundation under Grant MCS 81-17424.

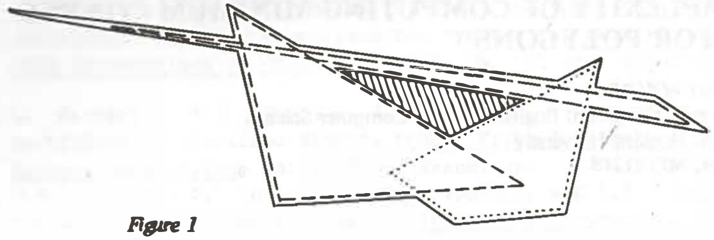

Figure 1

struct all possible convex pieces from the vertices, and then check subsets of these pieces for complete coverage. If the search examines smaller cardinality subsets first, the first cover found will be a minimum cover. Clearly this algorithm requires exponential time.

If Steiner points are allowed, it seems natural to restrict the convex pieces to be bounded by extensions of the polgon's edges (see Figure 1). Pavlidis observed that only *maximal* convex subsets need to be considered [17]: those which are not enclosed within any other internal convex subset. It is then clear that there are only a finite number of possible covers, and a searching algorithm is possible. Pavlidis investigated such algorithms for pattern recognition applications [17-20], although he was not interested in minimum covers. In [11] it is shown that searching for minimum covers under the edge-extension restriction can be accomplished by minimizing a Boolean expression, a known NP-complete problem. This does not, however, establish that the restricted minimum cover problem is NP-complete, and the complexity of this problem remains unknown.

A special case where the restriction to edge extensions is especially appropriate is the class of *rectilinear* polygons: those whose edges are parallel to one of two orthogonal directions. The restriction in this case implies that each piece must be a rectangle. Again it is clear that the problem is decidable by a searching algorithm. Masek showed that finding a minimum rectangle cover of a multiply-connected rectlinear polygon is NP-complete [9]. Recently, Chaiken *et al* have shown that if the class of rectlinear polygons is further restricted to be "convex" in the sense that every horizontal and vertical line meets the region in an interval, then a minimum rectangle cover can be found in polynomial time [1]. The problem for the important class of simply-connected rectilinear polygons remains open.

3. Counterexamples to Steiner Point Restrictions

Although the restriction that the pieces be bounded by edge extensions may seem natural for pattern recognition applications, it is a true restriction, in that there are polygons whose minimum covers require the use of pieces that are not bounded by such edge extensions. In fact there are rectilinear polygons that need edges in their minimum covers that are neither horizontal or vertical. Figure 2 shows a simply-connected example. As 9 mutually invisible points can be located in the interior of the polygon, any convex cover must use at least 9 pieces. The eight horizontal and vertical rectangles shown in Figure 3 can be proven to be necessary, leaving the ninth piece to use diagonal edges, as illustrated [12].

Although the diamond region in Figure 3 is not bounded by extensions of the polygon's edges, its vertices can be chosen to lie at intersections of edge extensions. It is natural to ask, then, whether restricting the Steiner points in this way is a true restriction. In fact it is, as can be verified by merging two copies of the polygon of Figure 2; the result is shown in Figures 4 and 5. The left diamond region requires a vertex somewhere along the segment labeled λ in Figure 5, but no two edge extensions intersect there.

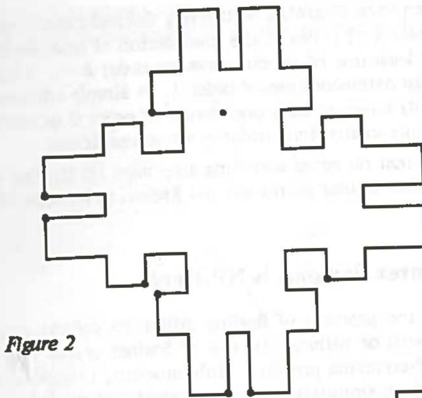

Figure 2

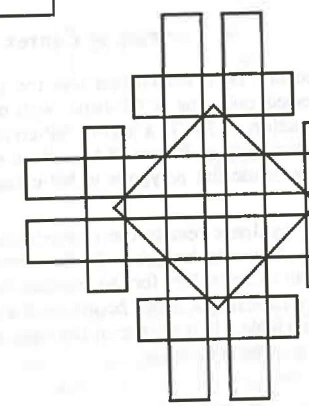

Figure 3

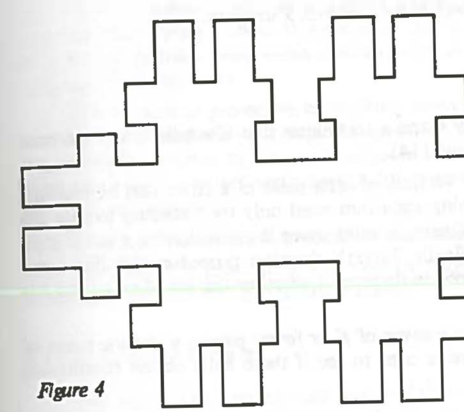

Figure 4

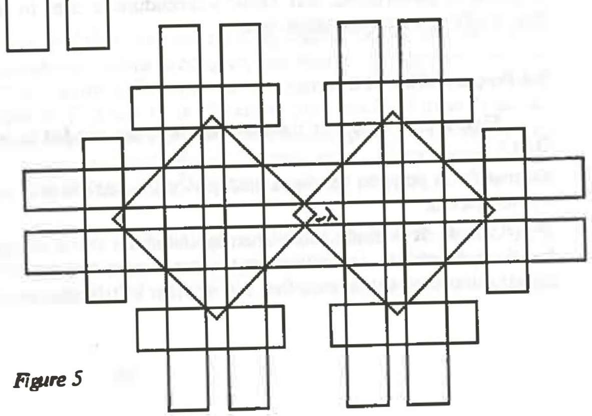

Figure 5

One could also restrict Steiner points to the following recursively defined class: a vertex of the polygon is of order 0; a point is of order k iff it lies at the intersection of lines determined by points of no more than order $k-1$, at least one of which is exactly order $k-1$. Under this definition, points at the intersection of edge extensions are of order 1. A simply-connected polygon can be constructed that requires in its cover at least one vertex of order 2 or more [13], and I conjecture that restricting Steiner points to any finite order is a true restriction.

The implication of these examples is that no naive searching algorithm for finding general minimum covers is known, since the possible Steiner points are not known to be restrictable to a finite class.

4. Covering by Convex Polygons is NP-Hard

It has recently been established that the problem of finding minimum convex covers of multiply-connected polygons is NP-hard, with or without the use of Steiner points [15]. The proof is by reduction of 3SAT, a known NP-complete problem. Subsequently, Lingas simplified by proof by reducing from Planar 3SAT, which simultaneously established that minimum partition of multiply-connected polygons is NP-complete, again with or without the use of Steiner points [8].

These NP-hardness results use multiply-connected polygons in their proofs, and so they do not establish any complexity results for simply-connected polygons. Also, the cover problem is not known to be in NP, for the reasons discussed in the previous two sections. Thus the NP-hard result establishes a lower bound on the complexity; it is possible for an NP-hard problem to be undecidable. It is shown in the next section that, unlike several infinite tiling problems [7], CCP is in fact decidable.

5. Decidability of Covering by Convex Polygons

5.1 Proof Sketch

We now show that CCP is decidable by using a technique that Chazelle employed on a similar problem [2]. A proof sketch is as follows [14].

It is first established that the number of vertices of each piece of a cover can be bounded by a function of n. This ensures that a searching algorithm need only try pieces up to a certain size. Secondly it is shown that the property of being a valid cover is equivalent to a set of algebraic equations in the vertex coordinates. Finally Tarski's decision procedure for first-order sentences in the field of real numbers is invoked to determine whether the set of equations has a solution.

Thus to determine whether a polygon has a cover of K or fewer pieces, a generic cover of K pieces is constructed, and Tarski's procedure is used to see if there exist vertex coordinates that satisfy the covering equations..

5.2 Bounded Size of Covers

First it is shown that the number of pieces needed to cover a polygon is never more than $O(n)$.

Lemma 1. A polygon (perhaps multiply-connected) with n vertices has a cover with $\leqslant 5n/3-4$ convex pieces.

Proof Sketch. It is easily established by induction that a polygon with n vertices and h holes can be triangulated into $n-2+2h$ triangles. Since each hole must have at least 3 vertices, $n \geqslant 3+3h$. Substitution then establishes that the number of triangles satisfies the bound in the Lemma. □

Next it is shown that any piece of a convex cover never needs to have more than $O(n)$ vertices. This is established by expanding a given piece as much as possible.

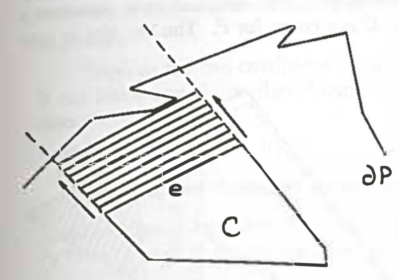

Figure 6

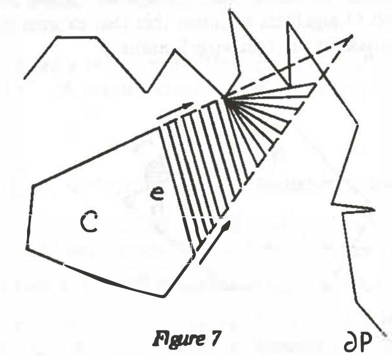

Figure 7

Lemma 2. Any convex polygonal subset C of an n vertex polygon P is itself enclosed within a subset C' that has no more than n vertices.

Proof. Perform the following expansion procedure for each edge e of C. Let ∂P be the boundary of P. If e does not touch ∂P, move it parallel to itself away from the centroid of C, adjusting its length to mate with extensions of its adjacent edges (see Figure 6). Continue this motion until either its length is reduced to zero or it touches ∂P. If it touches ∂P at a single point that is not an endpoint of e or at two or more points, then halt. If it touches ∂P at just one endpoint, then swing e about this endpoint, again adjusting its length to match the extension of its neighboring edge. Stop when another point of ∂P is hit or e becomes collinear with its adjacent edge (see Figure 7).

This expansion procedure halts either when e disappears or it abuts against ∂P. The possible "abutments" are the vertices and edges of P. No two non-collinear edges can be stopped in the expansion process by the same edge or vertex, nor even by an adjacent edge and vertex. Since there are n vertices and edges in ∂P, there can be no more than n non-collinear edges in the expanded polygon C'.

Clearly the method of construction guarantees that $C \subseteq C' \subset P$. □.

This bound is tight: an n edge convex polygon P has a convex subset, namely itself, that is not enclosed in a polygon of fewer edges.

5.3 The Covering Property

The proceeding two lemmas establish that only covers of at most $O(n)$ pieces, each of at most $O(n)$ vertices, need to be considered. Constructing an expression to represent the property of covering will be simpler if the cover is first *refined* or fragmented. Given any cover, introduce vertices wherever two edges of P or any C_i in the cover intersect (and there is not already a vertex). Make the introduced vertices part of both intersecting parts. This fragmentation ensures that every edge is entirely enclosed within at lease one convex piece. After refinement, each of the $O(n)$ pieces has at most $O(n^2)$ vertices, and the polygon P has at most $O(n^2)$ vertices.

Let a refined polygon P have vertices $(a_1,b_1),(a_2,b_2), \cdots ,(a_m,b_m)$, with $m = O(n^2)$, and let a refined cover C be represented by $C_1,C_2, \cdots ,C_k$, where each C_i is a convex polygon with vertices $(x_1^i,y_1^i),(x_2^i,y_2^i), \cdots ,(x_{m_i}^i,y_{m_i}^i)$, with $m_i = O(n^2)$. Boundary lists are assumed to be oriented so that traversal always keeps the interior to the left. Thus outer boundary chains are oriented counterclockwise, and hole chains clockwise. All sets will be considered closed regions, including their boundaries, unless otherwise mentioned. The next task is to construct a set of algebraic relations that that express the fact that **C** is a cover for P. The key idea is contained in the following Lemma.

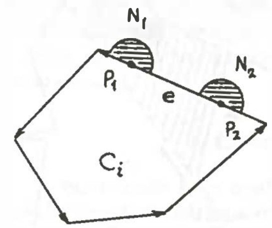

Figure 8

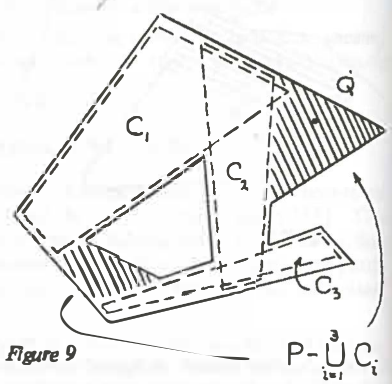

Figure 9

Lemma 3. The following two conditions are necessary and sufficient for **C** to be a refined cover for a refined polygon P:

(1) For every edge e of some $C_i \in \mathbf{C}$, either

  (a) e is equal to some edge of P, oriented in the same direction, or

  (b) e is inside of some C_j (different from C_i), but if it lies on ∂C_j, it is oriented in the opposite direction.

(2) None of the vertices of P are inside the *interiors* of any of the C_i.

Proof.

Necessity. Suppose **C** is a refined cover for P. We will first show that condition (1) holds for an arbitrary edge e of $C_i \in \mathbf{C}$. Let p_1 and p_2 be two arbitrary points on the interior of e, and let N_1 and N_2 be neighborhoods of p_1 and p_2 outside of C_i (see Figure 8). It cannot be the case that N_1 is inside P and N_2 is outside or vice versa, because then ∂P would cross e and the refinement procedure would have split e. Consider each remaining case separately.

Case 1: N_1 *not* $\subseteq P$ and N_2 *not* $\subseteq P$.

Because $p_1 \in C_i$ and $p_2 \in C_i$, and **C** is a cover, $p_1 \in P$ and $p_2 \in P$. Since the neighborhoods are outside of P, p_1 and p_2 lie on ∂P. They must lie on the same edge of P, and e must match this edge exactly, otherwise e would have been refined. The orientations of the edges must match to keep the inside of C_i consistent with the inside of P. Thus clause (a) of condition (1) holds.

Case 2: $N_1 \subseteq P$ and $N_2 \subseteq P$.

Since the points in the neighborhood are in P, and **C** is a valid cover, they must be inside some C_j's. Each neighborhood must be entirely enclosed within one C_j, and in fact the N_1 and N_2 must be enclosed within the same C_j, otherwise e would have been refined. This implies that all of e is enclosed within C_j. Since this C_j encloses neighborhoods outside of C_i, if e is on ∂C_j, then the edge orientations are in opposite directions. Thus clause (b) of condition (1) holds.

Concerning condition (2), if any vertex of P is in the interior of some C_i, then since a neighborhood of this vertex includes points outside of P, C_i encloses points outside of P, contradicting the assumption that **C** is a valid cover.

Sufficiency.

Suppose the two conditions hold true, but **C** is not a valid cover. Then either $\exists p \in P$ that is not inside any C_i, or $\exists p \notin P$ that is inside some C_i. A contradiction will be derived in both cases.

Case 1: $\exists p \in P$ *s.t.* $p \notin C_i$ $\forall i$.

Let Q be the connected component of $P - \bigcup_{i=1}^{k} C_i$ in which p lies. Q's boundary is composed of edges from either P or the C_i's (see Figure 9). Suppose all the edges of Q come from P. Then Q must in fact be the same as P, and all the C_i are outside P. But then the edges that form the boundary of $\bigcup_{i=1}^{k} C_i$ do not satisfy either clause (a) or (b) of condition (1). So suppose at least one edge e of Q comes from C_i. Clause (a) of condition (1) cannot be true for e, because the subtraction $P - C_i$ would eliminate e from Q's boundary. Similarly clause (b) cannot be true because the subtraction $P - C_i - C_j$ would eliminate e. Thus all possibilities contradict condition (1).

Case 2: $\exists p \notin P$ *s.t.* $p \in C_i$ for some i.

Let Q be the connected component of $\bigcup_{i=1}^{k} C_i - P$ that includes p. If all the edges of Q come from P, then Q coincides with a hole of P. Any vertex on the boundary of this hole must be inside some C_i, contradicting condition (2). So suppose at least one edge e of Q comes from C_i. This edge cannot satisfy either clause (a) or (b) of condition (1), because in the first case it would be removed by the subtraction of P, and in the second case it would be absorbed by the union with C_j. Thus all cases contradict the lemma. □

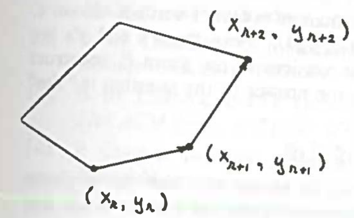

Figure 10

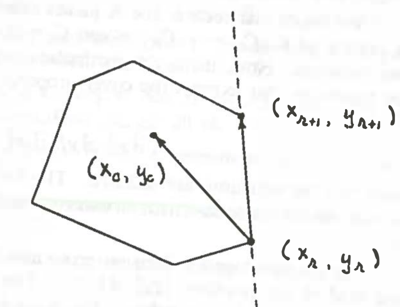

Figure 11

A set of relations for expressing the proposition that **C** is a refined cover for P must then check three conditions:

(1) Each piece C_i must be convex.

(2) Each edge of **C** must satisfy condition (1) of Lemma 3.

(3) Each vertex of P must satisfy condition (2) of Lemma 3.

The first check can be accomplished by computing the cross products of successive edge vectors (see Figure 10). Thus C_i is convex and oriented counterclockwise iff (superscript i's will be dropped when clear from the context)

$$(x_{r+1}-x_r)(y_{r+2}-y_{r+1})-(x_{r+2}-x_{r+1})(y_{r+1}-y_r) \geqslant 0, \ \forall r, \ 1 \leqslant r \leqslant m_i \qquad (1)$$

Clause (a) of condition (1) merely requires a check for equal edges. The edge $e=((x_1,y_1),(x_2,y_2))$ equals an edge of P iff

$$\exists r \ s.t. \ x_1=a_r, \ y_1=b_r, \ x_2=a_{r+1}, \ y_2=b_{r+1}, \qquad (2)$$

where $((a_r,b_r),(a_{r+1},b_{r+1}))$ is an edge of P.

Clause (b) requires determining whether an edge is inside a convex polygon C_i. It is iff its two endpoints are. A point is inside C_i iff it is inside every half-plane determined by C_i's edges. This half-plane check can be performed by a cross-product. Thus the point (x_0,y_0) is inside C_i (see Figure 11) iff

$$(x_{r+1}-x_r)(y_0-y_r)-(x_0-x_r)(y_{r+1}-y_r) \geqslant 0, \ \forall r, \ 1 \leqslant r \leqslant m_i \qquad (3)$$

The check that no vertex of P should be inside any C_i can be made by slight modification of equation (3).

It should now be clear that the property that C is a cover for P can be expressed by a series of algebraic equations of the form (1), (2), or (3), coupled by the logical connectives *and* and *or*.

5.4 Tarski's Decidability Theorem

The bounds derived in Section 3 and the equations constructed in Section 4 can be combined to write a single algebraic relation that precisely captures the question "Is there a cover of P with K or fewer convex pieces?"

We know that each of the K pieces cannot have more than $m=O(n^2)$ vertices. So set C to consist of $C_1,C_2,\cdots,C_K$, where $C_i=((x_1^i,y_1^i),\cdots,(x_m^i,y_m^i))$. Here the x's and y's are real variables. Now, using the particular integer coordinate vertices of the given P, construct the equations that express the cover property for C. Then the answer to the question is "Yes" iff

$$\exists x_1^1 \ \exists y_1^1 \ \exists x_2^1 \ \exists y_2^1 \cdots \exists x_K^m \ \exists y_K^m$$

such that the equations are satisfied. The fact that there may be covers with less than K pieces or with pieces using less than m vertices clearly will not affect the validity of a cover of the size of C.

Tarski developed a decision procedure for all first-order sentences of no free variables in the field of real numbers [22] [4] [21]. The expression above is such a sentence, and is therefore decidable by his procedure. The procedure eliminates quantifiers and reduces the sentence to a single polynomial relation of no variables, which is easily checked for validity. Modern formulations of Tarski's procedure [10] appear to require $O(2^{2^n})$ time. Despite this slowness, the procedure establishes the decidability of the minimum convex cover problem:

Theorem. The minimum convex cover problem is decidable.

6. Discussion

The proof of the preceding section can easily be modified to hold under any of the various restrictions we have discussed, such as to simply-connected polygons or to rectangular pieces. Thus all these minimum covering (and partition) problems are decidable.

The current state of knowledge on the complexity of the various problems is displayed in the Table below.

Polygon Class	Components	Cover	Partition
multiply-connected	convex	NP-hard	NP-complete
simply-connected	convex	?	$O(n^3)$
rectlinear, multiply-connected	rectangular	NP-complete	polynomial
rectilinear, simply-connected	rectangular	?	$O(n^3)$

It is an outstanding open problem to determine whether a simply-connected polygon can be minimally covered in polynomial time, even when restricted to the rectilinear case. It would also be of interest to derive some restrictions on the class of convex pieces that are required in general minimum covers. For example, if the vertices of the original polygon have rational coordinates, are irrational coordinates ever required in a minimum cover?

Acknowledgements

I am indebted to Kenneth J. Supowit and the participants of a McGill University Computational Geometry seminar for clarifying discussions.

REFERENCES

[1] S. Chaiken, D. J. Kleitman, M. Saks, and J. Shearer, "Covering Regions by Rectangles," *SIAM Journal on Algebraic and Discrete Methods*, Vol. 2, No. 4, Dec. 1981, pp. 394-410.

[2] B. M. Chazelle, "Computational geometry and convexity," PhD Dissertation, Carnegie-Mellon University, Technical Report CMU-CS-80-150, 1980, pp. 140-145.

[3] B. M. Chazelle and D. Dobkin, "Decomposing a polygon into its convex parts," *Proc. of 11th ACM Symp. on Theory of Computing*, Atlanta, Georgia, pp. 38-48, 1979.

[4] P. Cohen, "Decision procedures for real and p-adic fields," *Communications on Pure and Applied Mathematics*, Vol. 22, 1969, pp. 131-151.

[5] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, San Francisco: W. H. Freeman, 1979.

[6] D. S. Johnson, "The NP-Completeness Column: An Ongoing Guide," *Journal of Algorithms*, June 1982.

[7] H. R. Lewis, *Unsolvable Classes of Quantification Formulas*, Reading, Mass.: Addison-Wesley, 1979.

[8] A. Lingas, "The Power of Non-Rectilinear Holes," to appear in the *Proc. of the 9th Colloquium on Automata, Languages, and Programming*, Aarhus, 1982.

[9] W. J. Masek, "Some NP-complete set covering problems," unpublished manuscript, Aug. 1979; quoted in [5, p. 232].

[10] L. Monk, "Elementary-recursive decision procedures," PhD Thesis, University of California, Berkeley, 1975.

[11] J. O'Rourke, "Polygon decomposition and switching function minimization," *Computer Graphics and Image Processing*, Vol. 19, pp. 384-391, Aug. 1982.

[12] J. O'Rourke, "Minimum Convex Covers: Some Counterexamples," Johns Hopkins University, Department of Electrical Engineering and Computer Science, Technical Report JHU-EECS 82-1, Jan. 1982.

[13] J. O'Rourke, "A Note on Minimum Convex Covers for Polygons," Johns Hopkins University, Department of Electrical Engineering and Computer Science, Technical Report JHU-EECS 82-3, Mar. 1982.

[14] J. O'Rourke, "The Decidability of Covering by Convex Polygons," Johns Hopkins University, Department of Electrical Engineering and Computer Science, Technical Report JHU-EECS 82-4, May 1982.

[15] J. O'Rourke and K. J. Supowit, "Some NP-hard polygon decomposition problems," to appear in *IEEE Trans. on Information Theory*, Mar. 1983.

[16] L. Pagli, E. Lodi, F. Luccio, C. Mugnai, and W. Lipski, "On two dimensional data organization 2," *Fundamenta Informaticae*, vol. 2, pp. 211-226, 1979.

[17] T. Pavlidis, "Analysis of set patterns," *Pattern Recognition*, vol. 1, pp. 165-178, 1968.

[18] T. Pavlidis, "Representation of figures by labelled graphs," *Pattern Recognition*, vol. 4, pp. 5-17, 1972.

[19] T. Pavlidis, "Structural pattern recognition: Primitives and juxtaposition relations," in *Frontiers of Pattern Recognition* (Ed. S. Watanabe), New York: Academic Press, pp. 421-451, 1972.

[20] T. Pavlidis, *Structural Pattern Recognition*, pp. 236-241, Berlin-Heidleberg-New York: Springer Verlag, 1977.

[21] M. O. Rabin, "Decidable theories," in *The Handbook of Mathematical Logic*, ed. J. Barwise, North-Holland, 1978, pp. 595-629.

[22] A. Tarski, "A decision method for elementary algebra and geometry," second edition, revised, University of California press, 1951.

[23] G. T. Toussaint, "Pattern recognition and geometrical complexity," *Proc. of 5th International Conference on Pattern Recognition*, Miami Beach, Florida, pp. 1324-1347, Dec. 1980.

AN OPTIMAL ALGORITHM TO CONSTRUCT ALL VORONOI DIAGRAMS FOR K NEAREST NEIGHBOR SEARCHING IN THE EUCLIDEAN PLANE

Frank Dehne
Lehrstuhl für Informatik I
Bayerische Julius-Maximilians-Universität Würzburg
D-8700 Würzburg

ABSTRACT

This paper presents an algorithm, that constructs all Voronoi diagrams for k nearest neighbor searching simultaneously. Its space and time complexity of $O(N^4)$ is shown to be optimal.

1. INTRODUCTION

In [6] Shamos and Hoey introduce the idea of generalized Voronoi diagrams to get an optimal solution of the k nearest neighbor problem and give an O(N logN) algorithm to construct the order one diagram.
Lee ([2]) extends this to an algorithm, that computes an order k diagram in $O(k^2N \log N)$.
To answer k nearest neighbor queries with arbitrary k we now want to construct all Voronoi diagrams.
This paper presents a simple solution of this problem.
The given algorithm has time and space complexity $O(N^4)$ and is shown to be optimal. Its implementation is not very difficult and the constant factors for the complexity are expected to be quite good.