

GRAPHICS IN FLATLAND: A CASE STUDY

Herbert Edelsbrunner, Mark H. Overmars and
Derick Wood

ABSTRACT

Usually in computer graphics, a two-dimensional view of a set of three-dimensional objects is considered. In this article we reduce the dimensionality by one in each case. In other words we study what, for obvious reasons, we call Flatland graphics.

This forms the beginning of a mathematical investigation of computer graphics and, at the same time, provides uniform solutions for a number of computational geometry problems. In particular we study the maintenance of a view during insertion and deletion of objects and the "frame-to-frame" coherence while walking around a set of objects. Both parallel and perspective projections are considered.

Our major concern is convex objects that are simple—in a sense, made precise in this article. However, we will close this article by discussing some possible extensions to nonconvex objects and/or to higher dimensions.

Advances in Computing Research, Volume 1, pages 35–59

Copyright © 1983 by JAI Press Inc.

All rights of reproduction in any form reserved.

ISBN: 0-89232-356-6

The investigation also serves to demonstrate a number of tools that have been developed recently in the context of computational geometry. For example, dynamization and searching.

1. INTRODUCTION

An important and fundamental algorithmic problem in computer graphics is the following: given a set of objects in three-dimensional space, compute the view from some fixed direction or point, that is, compute the picture one sees looking from the direction or point. The main issue is to eliminate all parts of the objects that cannot be seen (i.e., that lie behind some other object). It is a generalization of the hidden line problem in which objects have straight line edges. The problem has numerous applications in areas such as picture processing and the development of movies (see, for example, [15]).

Notwithstanding the problem's importance, there have been few attempts in the literature to investigate it theoretically (see [22]). It is only recently that such investigations have begun to make their appearance (for example, [4, 10, 20, 24, 25]).

We simplify the preceding problem conceptually by only considering graphics in "Flatland." Flatland is a two-dimensional world, described by Edwin A Abbott in his book *Flatland: A Romance of Many Dimensions* [1], with objects like line segments, convex polygons, and circles as inhabitants. We study the generalized hidden line problem in Flatland for three reasons:

1. The results and algorithms are interesting in their own right because of their relation to other planar geometric problems.
2. A number of three-dimensional graphics problems are for the main part two-dimensional, for example, walking in a maze.
3. The results possibly lay a foundation and show directions for further research on the corresponding problems in three- and higher-dimensional space (see Section 6).

Some aspects of the hidden line problem in Flatland have already been studied. Shamos [21] and Lee and Preparata [13] consider the problem of computing the kernel of a simple polygon, that is, the region inside the polygon from which the whole polygon is visible. Shamos [21] and El-Gindy and Avis [8] treat the problem of computing the view of a simple polygon from some interior point in $O(n)$ time. Freeman and Loutrel [9]

consider the more general case in which the point of view may also be outside the polygon, but their algorithm appears to run in $O(n^2)$ time. Avis and Toussaint [3] define the notion of visibility from an edge of a polygon and give an $O(n)$ algorithm for solving the problem. All of these results only deal with the case in which there is a single polygon. In this article we consider the much more general setting of the "hidden line" problem in which we have a set of simple objects. We assume that these objects do not intersect, but we allow them to touch. Hence "larger" objects can be decomposed into simple objects (a simple polygon is in general not a simple object).

We consider the two standard types of views in Flatland. The first type is the view from some direction, a *parallel view*. In this case, the view consists of a line on which the parts of the objects that are visible from the direction are projected (Figure 1). (The reader is advised to look along the paper from the direction of view). For the sake of clarity we number the objects in the set from 1 to n (the total number of objects in the set). The second type of view is the view from a point, a *perspective view*. In this case the view consists of a circle on which the parts of the objects we can see from the given point are projected (Figure 2). Perspective views correspond to our natural way of seeing, although we are only able to see part of the perspective projection.

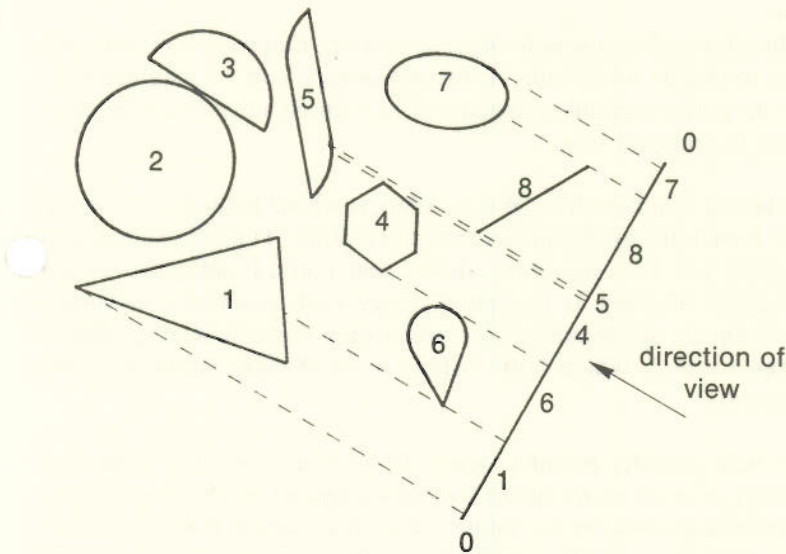


Figure 1. A parallel view of a set of objects.




Figure 2. A perspective view of a set of objects.

Definition. A simple object is a bounded convex object such that

1. Any parallel or perspective view of it can be computed in constant time.
2. If the views of two simple objects overlap, then constant time suffices to decide which one of the two objects can be seen entirely.
3. The up to four common tangents of two simple objects can be computed in constant time.

Two objects that touch each other (i.e., objects whose boundaries overlap but do not cross) are by definition nonintersecting. Typical examples of simple objects are line segments, disks, and convex polygons with a bounded number of edges. In the sequel, objects are assumed to be simple and sets of objects are assumed to contain only nonintersecting objects (see Figures 1 and 2, which both display a set of eight nonintersecting simple objects).

We first will consider parallel views. In Section 2 we will show that parallel views of a set of n objects can be computed in $O(n \log n)$ time, using a divide-and-conquer technique. It is also shown that known dynamization techniques can be applied for maintaining the view from some fixed direction at the cost of $O(n)$ time per insertion or deletion of an object.

An important related problem is the problem of maintaining the view while “walking around” the set of objects. So, starting with the view from some starting direction, we want to maintain it while “walking” in a counterclockwise direction, say. Of course, with each new direction the relative size of the parts of the objects one sees changes, but the order in which one sees these parts only changes when a “critical” direction is crossed. In Section 3 we will give methods for maintaining this order of objects while walking around the set.

In Section 4 we will consider the searching variant of the problem of computing a view: given a set of objects, preprocess them in such a way that views from different directions can be computed efficiently. The problem was recently treated by Fuchs, Kedem, and Naylor [10]. We will follow a completely different approach to obtain trade-offs between query time, preprocessing time, and storage required.

In Section 5 we will examine perspective views. It is shown that the method for computing and maintaining a view, as described in Section 2, carries over. Restricted versions of the walking around problem and the searching variant will be considered as well.

Finally, in Section 6, we will discuss how nonconvex objects might be included, and also we will discuss the problems raised by moving to higher dimensions.

2. COMPUTING AND MAINTAINING A PARALLEL VIEW

A parallel view of a set of objects consists of a partition of a line. Each part of the line corresponds to an object in the set or to a place where one can look through the set (from the direction of view). See Figure 3 for some different parallel views of the same set of objects. To each part of the line we assign the number of the corresponding object. If the part corresponds to a place where one can look through the set we assign 0 to it. It is possible that different parts of the line correspond to the same object, and hence, are assigned the same number (see, for example, view 3 of Figure 3). A *partition point* corresponds to a leftmost point of an object (Figure 4A) or a rightmost point of an object (Figure 4B), with respect to the direction of view. (A part of the line might consist of one point if it corresponds to a line segment in the direction of view. In this case we treat it as a double partition point).

LEMMA 1. *The parallel view of a set of n objects from a fixed direction consists of at most $2n + 1$ parts, that is, has at most $2n$ partition points.*

PROOF. We show that one can assign left and right end points of objects to partition points in a unique way. In Case 1 we assign the corresponding




Figure 3. Three different parallel views of the same set of objects.

left end point to the partition point; in Case 2 we assign the corresponding right end point. As the number of left and right end points is bounded by $2n$, there are at most $2n$ partition points and hence at most $2n + 1$ parts. This completes the argument. \square

For computing the view of a set of objects from a fixed direction we will use a divide-and-conquer technique.




Figure 4. Two partition points: (A) at the leftmost point of an object (Case 1) and (B) at the rightmost point of an object (Case 2).

LEMMA 2. Given a fixed direction and a set of objects $V = \{x_1, \dots, x_n\}$. The view of V can be computed in $O(n)$ time from the views of $A = \{x_1, \dots, x_i\}$ and $B = \{x_{i+1}, \dots, x_n\}$, for any $1 \leq i < n$.

PROOF. Let the parts of the view of A from left to right be numbered a_0, a_1, \dots, a_k and let the partition points in between be p_1, \dots, p_k . Similarly, let the parts of the view of B be numbered $a_0', a_1', \dots, a_{l'}$, and let the partition points in between be $p_1', \dots, p_{l'}$. One easily verifies that partition points of $V = A \cup B$ are partition points of A or B and that if at some location in the view of $A \cup B$ object O is visible, then it is visible at that location in the view of A or in the view of B (see Figure 5 for an example). We will compute the view during one simultaneous walk along the lists for the views of A and B . At each partition point of both views we check whether or not this point is also a partition point of the total view. As a result of the definition of a simple object (condition 2) this checking can be done in constant time, which immediately implies the assertion. This completes the argument. \square

Lemma 2 provides the conquer step we need.

THEOREM 1. Given a set of n objects and a direction, one can compute the parallel view from the given direction in $O(n \log n)$ time, and this is optimal in the worst case.

PROOF. To compute the view of the set of objects V , we split V into two equal-sized subsets, compute the views of the two subsets recursively in the same way, and merge the answers using Lemma 2. Besides the




Figure 5. Three views of a set of objects from a fixed direction.

recursive calls, this method takes $O(n)$ time. Hence, the total time $T(n)$ needed is given by the following recurrence:

$$T(n) = 2T(n/2) + O(n)$$

It is well known that the solution of this recurrence is $T(n) = O(n \log n)$.

Because computing the view can also be used to sort n real numbers, the obtained bound is also optimal in the worst case. This completes the argument. \square

This solves the problem of computing the view of a set of objects from a given direction.

A second problem is the problem of maintaining the view from a fixed direction while objects are inserted into and deleted from the set. We assume that the set will be updated in such a way that objects never intersect. Overmars [17] considered the following class of problems, which are defined for finite sets of objects and hence are called set problems.

Definition. [17] A set problem P is called $O(n)$ -order decomposable if there exists an ordering ORD and a function \square such that for each set $\{a_1, \dots, a_n\}$, ordered according to ORD

$$P(a_1, \dots, a_n) = \square [P(a_1, \dots, a_i), P(a_{i+1}, \dots, a_n)]$$

for all i with $1 \leq i < n$, where \square takes $O(n)$ time to compute.

Using for ORD any ordering we like and using for \square the merging step described in the proof of Lemma 2, it follows from Lemma 2 that the problem of computing the view is $O(n)$ -order decomposable. Using the dynamization method known for $O(n)$ -order decomposable set problems [17], we obtain the following result:

THEOREM 2. *Given a fixed direction, one can maintain the view of a set of objects from that direction at the cost of $O(n)$ time per insertion and deletion of an object, where n is the cardinality of the set.*

The disadvantage of $O(n \log n)$ space required by this method can be remedied using a technique attributable to Gowda and Kirkpatrick [11]. It reduces the space to $O(n \log \log n)$ without changing the update time bounds.

When we are not interested in the view after each update but only want to maintain the set of objects in such a way that the view can be computed efficiently when needed, we can apply another dynamization technique. This technique was originally described for so-called searching problems, that is, a set V of objects is to be stored such that questions dependent

on V and on some query object needed for the formulation of the question can be answered efficiently. For those problems, Overmars [18] defines the following subclass:

Definition. [18] A searching problem P is called $O(n)$ -decomposable if there exists a function \square such that for all sets A and B with $A \cap B = \emptyset$ and for each query object x

$$P(x, A \cap B) = \square [P(x, A), P(x, B)]$$

where \square takes at most $O(n)$ time to compute when the sets contain n objects.

$O(n)$ -decomposability is an extension of the concept of decomposability for searching problems as defined by Bentley [5]. As a consequence of Lemma 2, the problem of computing the view from a fixed direction is $O(n)$ -decomposable. (The problem at hand can be considered to be a searching problem for which the questions do not depend on some query object.) Applying known transformations on structures for $O(n)$ -decomposable searching problems [5, 18] to the structure used in Theorem 2 we obtain the following result:

THEOREM 3. *Given some fixed direction, one can maintain a set of objects at the cost of $O(\log^2 n)$ time per insertion and $O(n)$ time per deletion such that at any moment one can compute the view from the given direction in $O(n)$ time.*

Using the techniques from [11], the amount of storage required can be bounded by $O(n \log \log n)$.

3. COMPUTING ALL PARALLEL VIEWS ("WALKING AROUND THE SET")

Problems in computer graphics often require that the view of a set of objects be maintained while walking around the set. So, we start with the view from some starting direction; and we want to maintain the view while changing the direction counterclockwise, say, until we again reach the starting direction. Of course, with each new direction, the relative size of the parts of the objects we can see is different, and hence it is unlikely that there are truly efficient ways of maintaining the actual view while walking around the set. But the order in which we see parts of the objects only changes at some critical directions. Hubschman and Zucker [12] have considered frame-to-frame coherence for convex objects in three-dimensional space.

Definition. Scanning the view of a set of objects from some fixed direction from left to right, we get a list of at most $2n + 1$ numbers ranging from 0 to n that gives the order in which we see parts of the objects. This list is called the *visibility list* of the set from the direction.

Clearly, given a view one can construct the corresponding visibility list in $O(k)$ time, where k is the length of the list. We will show that one can maintain the visibility list of the set efficiently, while walking around. It is important that we are able to compute the view corresponding to a visibility list and an appropriate direction efficiently.

LEMMA 3. *Given a visibility list for some direction, one can compute the view from this direction in $O(k)$ time, where k is the length of the list.*

PROOF. Let the visibility list be a_1, \dots, a_k . Hence, we have to locate the partition point between each pair a_i, a_{i+1} of parts of the view line. If $a_i = 0$, then the partition point is the projection of the leftmost point of a_{i+1} on the line. Similarly, if $a_{i+1} = 0$, then the partition point is the projection of the rightmost point of a_i on the line. (They cannot both be 0.) Otherwise we compute the view of a_i and a_{i+1} . This takes constant time by definition. In this view there is exactly one point where a_{i+1} starts replacing a_i . This point is the partition point we searched for. Hence, each partition point takes constant time to compute, and, hence, the total time needed is bounded by $O(k)$. This completes the argument. \square

Hence, at each moment while walking around the set, one can obtain the view in $O(k)$ time. The visibility list of a set of objects only changes when we pass some critical direction.

Definition. Given two objects A and B , a *critical direction* of A and B is a direction in which the visibility list of A and B changes.

See Figure 6 for an example of the critical directions two simple objects A and B can have. Those directions are given by the up to four common tangents of the two objects. In between the critical directions, the visibility lists are given. By condition 3 of the definition of a simple object, the critical directions for two such objects can be computed in constant time.

LEMMA 4. *Two simple objects have at most eight critical directions.*

The proof is obvious since the at most four common tangents of two convex objects define, at most, eight directions.




Figure 6. Possible critical directions for two simple objects, A and B.

LEMMA 5. Given a set of n objects, there are at most $4n(n - 1)$ critical directions.

PROOF. There are $n(n - 1)/2$ pairs of objects. Each pair of objects gives rise to at most eight critical directions, which completes the argument. \square

See Figure 7 for an example of all critical directions of a set of three objects.

The following lemma shows that the visibility list of the total set only changes when passing a critical direction of some pair of objects in the set.

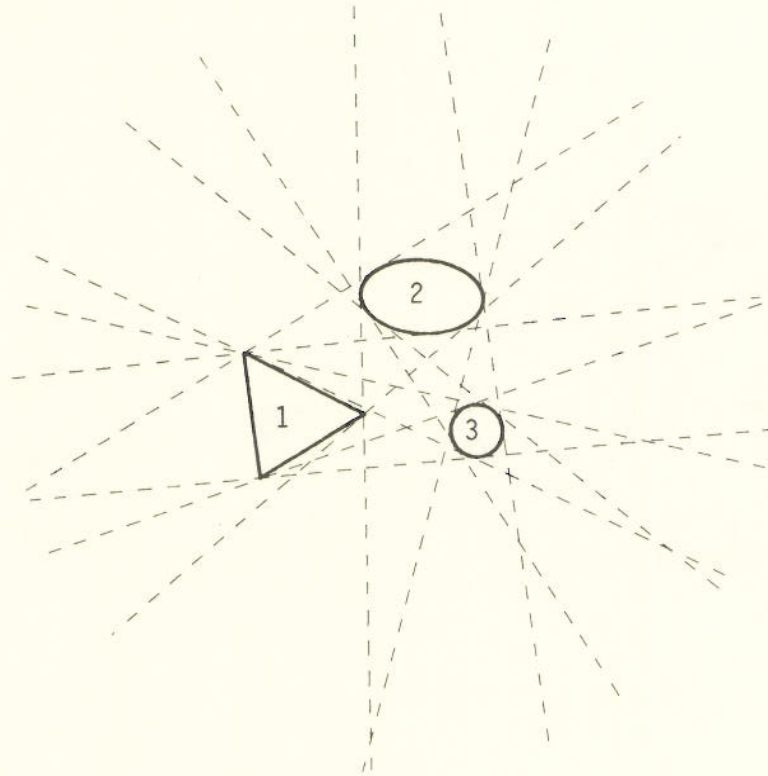


Figure 7. All critical directions of a set of three objects.

LEMMA 6. *Given two directions of view α and β . If there is no critical direction in between α and β , then the visibility lists from both directions are the same.*

PROOF. If the visibility list from α is not equal to the visibility list from β , then there are two objects A and B such that the visibility list from α restricted to A and B is not equal to the visibility list from β restricted to A and B . Hence, there must be a critical direction of A and B in between α and β . This contradicts the assumption and completes the argument. \square

To be able to maintain the visibility list while walking around the set of objects, we have to represent the list in such a way that the "update" necessary when passing critical directions can be made efficiently.

LEMMA 7. *One can represent a visibility list in such a way that the updates needed when passing a critical direction can be made in $O(\log k)$ time, where k is the size of the list.*

PROOF. We store the k elements of the visibility list in a dictionary D , which allows us to locate in $O(\log k)$ time the element C of the visibility list with the following properties: given a direction for which the visibility list is correct and a point p in the view from this direction, then p lies in an interval of the view that corresponds to C . In addition to this, the dictionary D accommodates in $O(\log k)$ time changes that are necessary when a critical direction is passed. (D can be implemented as an AVL-tree storing the elements of the visibility list in its inner nodes. The nodes are also maintained in a doubly linked list that allows us to determine in constant time the predecessor or successor in the symmetric order.)

For convenience, we will interchangeably use D to denote the dictionary that stores a visibility list as well as the visibility list itself. We assume that, with each critical direction, we are given the two objects A and B of which it is a critical direction. Examining the view of A and B before and after the critical direction, we determine how their view changes and where the new or disappearing partition point p lies in the view from the critical direction. We will locate the object C in the visibility list stored in the dictionary D such that p lies in an interval corresponding to C . Examining the view of A , B , and C (C is possibly A or B), we can find out in constant time what change in the visibility list needs to be made (if necessary). This change consists of the insertion or deletion of one object. This clearly takes at most $O(\log k)$ time, which completes the argument. \square

(If two or more critical directions are the same, we treat them separately.) The structure clearly requires $O(k)$ storage.

LEMMA 8. *All critical directions of a set of n objects, in order of occurrence, can be computed in $O(n^2 \log n)$ time.*

PROOF. We start with object 1 and compute all its critical directions with the other objects. This can be done in $O(n)$ time. Next, we compute all critical directions between object 2 and the objects 3, . . . , n , between object 3 and objects 4, . . . , n , and so forth. This takes a total of $O(n^2)$ time. Ordering the $O(n^2)$ critical directions takes $O(n^2 \log n)$ time, which completes the argument. \square

Lemmas 7 and 8 lead to the following solution for the problem of walking around a set of objects:

THEOREM 4. *Given a set of n objects, one can maintain the visibility list while walking around the objects in a total of $O(n^2 \log n)$ time using $O(n^2)$ space.*

PROOF. We first compute the view from the starting direction and represent the corresponding visibility list in a dictionary D as described in the proof of Lemma 7. This takes a total of $O(n \log n)$ time. Next we compute all critical directions ordered counterclockwise from the starting direction. This takes $O(n^2 \log n)$ time according to Lemma 8. Passing each next critical line takes $O(\log k) = O(\log n)$ time for updating the visibility list, according to Lemma 7. As there are $O(n^2)$ critical lines to pass when walking around, this takes a total of $O(n^2 \log n)$ time. The asserted time bound follows. Because we have to store all $O(n^2)$ critical directions, we need $O(n^2)$ storage. This completes the argument. \square

Although the method presented is fast, the required amount of storage is rather large. We will now describe a method that takes more time but uses only linear storage. (The same amount of storage has also been achieved for a special case of our problem by van Leeuwen and Mehlhorn [23]. They considered the walking around problem for a simple polygon and obtained an $O(n \log n)$ time and $O(n)$ space algorithm where n denotes the number of vertices of the polygon.)

Again we start with computing the view from the starting direction and representing the corresponding visibility list in a dictionary D . Next, for each object i we compute the critical direction d_i of i with some other object such that (1) i lies behind the other object where the change in view from d_i occurs (this is needed to prevent a direction from being reported more than once), and (2) d_i lies nearest to the starting direction. This takes $O(n^2)$ time (Figure 8). We know that the first critical direction we have to pass must be among the n directions d_1, \dots, d_n . We store these directions in order of occurrence in a priority queue Q , such that the first critical direction is the first element of Q . Let this direction be d_i . To pass this direction we (1) delete d_i from Q , (2) perform the appropriate update in the visibility list D , (3) compute the new d_i , and (4) insert this new d_i into Q . Steps 1, 2, and 4 take $O(\log n)$ time (see, for example [2]) and step 3 takes $O(n)$ time. Hence, passing a critical direction takes $O(n)$ time. One easily verifies that in this way the next critical direction to be passed is always the first element in Q .

THEOREM 5. *Given a set of n objects, one can maintain the visibility list while walking around the set in a total of $O(n^3)$ time using $O(n)$ storage.*

PROOF. Both Q and D have size $O(n)$, and, hence, the amount of storage required is $O(n)$. The preprocessing takes $O(n^2)$ time. The crossing of a critical direction takes $O(n)$ time, and, as there are $O(n^2)$ critical directions, the total amount of time needed is $O(n^3)$. This completes the argument. \square




Figure 8. Three critical directions of a set of three objects.

The method just described is only one example of a whole class of methods for walking around the set of objects, yielding different trade-offs between the run time and the amount of storage required.

THEOREM 6. *Given a set of n objects, for each function $f(n)$, with $1 \leq f(n) \leq n/\log n$, there exists a method for maintaining the visibility list while walking around the set in a total of $O[n^3/f(n)]$ time using of $O[nf(n)]$ storage.*

PROOF. After computing the visibility list D for the starting direction, we compute for each object i the first $f(n)$ critical directions $d_i^1, \dots, d_i^{f(n)}$ between i and other objects after the starting direction, where i lies behind the other object at the place of the change. This takes $O(n)$ time per object and hence a total of $O(n^2)$ time. Storing the $nf(n)$ critical directions in a priority queue Q takes $O[nf(n) \log n]$ time. We know that the first direction in Q is the first critical direction we have to pass. Passing this direction consists of deleting it from Q and updating the visibility list D . When we pass a $d_i^{f(n)}$, then we also compute the new $d_i^1, \dots, d_i^{f(n)}$ and insert them into Q . In this way, Q clearly always contains the first critical direction. For each object i , we pass at most $O[n/f(n)]$ times a $d_i^{f(n)}$. Hence, the total amount of time needed for passing all critical di-

rections is bounded by $O(n^2 \log n)$ time for updating D and deleting and inserting directions in Q and $O[n^3/f(n)]$ time for computing critical directions. As $1 \leq f(n) \leq n/\log n$, we can estimate these costs by $O[n^3/f(n)]$. The amount of storage required for D and Q is clearly bounded by $O[nf(n)]$. This completes the argument. \square

For example, one can obtain a total run time of $O(n^2 \log n)$ using only $O(n^2/\log n)$ storage (cf. Theorem 4).

There are a number of problems that can be solved using the walking around algorithm described in Theorem 6, for example, the shadow problem attributable to Preparata [19].

COROLLARY 1. *Given a set of n objects, for each function $f(n)$ with $1 \leq f(n) \leq n/\log n$, there exists a method to compute all directions in which the shadows of the objects do not overlap, i.e., in which all objects can be seen separately, in $O[n^3/f(n)]$ time using $O[nf(n)]$ storage.*

PROOF. We use the method of Theorem 6 for walking around the set. With the dictionary D , which contains the visibility list, we maintain the information whether or not the objects can be seen separately. By a proper choice of D , this information can be maintained at a cost of $O(\log n)$ time per crossing a critical direction. The assertion follows. \square

A second example is the stabbing problem [7]. In [6], the general topic of transversals is discussed.

COROLLARY 2. *Given a set of n objects, for each function $f(n)$ with $1 \leq f(n) \leq n/\log n$, there exists a method for computing a line that intersects all objects when it exists) in $O[n^3/f(n)]$ time using $O[nf(n)]$ storage.*

PROOF. One easily verifies that, if there is such a line with direction d and if $d_1 \leq d \leq d_2$ are the two nearest critical directions, then there is such a line with any direction between d_1 and d_2 . There is such a line with direction d if and only if the views of all objects from d overlap at some point. Whether there is such a point can be maintained in D at the cost of $O(\log n)$ time per crossing a critical direction. Details are left to the interested reader. \square

Other problems that can be solved by these methods are, for example, (1) give all directions in which the view contains holes (places where you can look through the set); (2) give all directions in which the view is connected; and (3) give all directions in which you can see only one object.

4. SEARCHING FOR THE PARALLEL VIEW

In this section we treat the searching variant of the problem of computing a view: preprocess the set of objects in such a way that views from different directions can be computed efficiently. Hence, the directions are query objects of the searching problem.

THEOREM 7. *Given a set of n objects, for each function $f(n)$, with $1 \leq f(n) \leq n$, there exists a method of preprocessing the set in time $O[n^2 \log n + n^3/f(n)]$ using $O[n^3/f(n)]$ storage, such that the views from different directions can be computed in $O[f(n) \log n + k]$ time, where k is the number of objects in the view.*

PROOF. We will show how to preprocess the set such that one can obtain the visibility list from any direction efficiently. Computing the view afterward takes $O(k)$ time according to Lemma 3.

In the preprocessing phase, we compute all critical directions between each pair of objects in the set, order them, and store them in a dictionary D . We also assume that constant time suffices to determine the predecessor and the successor of a given critical direction in D . Let d_0 be the first direction in D . We compute the view from d_0 and store the corresponding visibility list in a structure D_0 , which we connect with d_0 . Our goal is to build a visibility list after each $f(n)$ critical directions. To build the second visibility list D_1 connected to $d_{f(n)}$, we first copy the list D_0 and perform the appropriate updates on it while we walk from d_0 to $d_{f(n)}$. In general, to build the visibility list D_i that has to be associated with $d_{if(n)}$, we first copy D_{i-1} and then update the copy while walking from $d_{(i-1)f(n)}$ to $d_{if(n)}$. The building of D_i in this way takes $O(n)$ time for copying and $O[f(n) \log n]$ time for performing the updates. Hence, the building of all $O[n^2/f(n)]$ visibility lists takes $O[n^3/f(n) + n^2 \log n]$ time. As the computation of the critical directions and the building of D take $O(n^2 \log n)$ time, the bound on the preprocessing time follows. As D takes $O(n^2)$ storage and each of the $O[n^2/f(n)]$ visibility lists takes $O(n)$ storage, the total amount of storage needed is $O[n^2 + n^3/f(n)] = O[n^3/f(n)]$ because $f(n) \leq n$.

To perform a query with direction d to compute the visibility list of the set from d , we search with d in D for the direction d_i nearest to and before d . Next we walk along the list of critical directions until we find a direction with an associated visibility list D_j . On D_j we perform the necessary updates while walking back to d_i . D_j now contains the visibility list from d . We copy it and afterward we make all updates to restore it to its proper form at $d_{jf(n)}$. The searching with d in D takes $O(\log n)$ time. D_j can be at most $f(n)$ directions away from d_i , and hence the locating and updating

of D_j takes $O[f(n) \log n]$ time. Copying D_j takes $O(k)$ time. Undoing the actions we performed on D_j also takes $O[f(n) \log n]$ time. The bound for the query time follows, which completes the argument. \square

Since $k = O(n)$, we consider $f(n) = n/\log n$ as a reasonable choice. It yields a data structure that takes $O(n^2 \log n)$ time to construct and uses $O(n^2 \log n)$ storage such that the views from different directions can be computed in $O(n)$ time. If we only want to ask some questions about the view or visibility list, there is in general neither any need to compute the view itself nor to copy the structure D_j . We can answer the question using the updated D_j , after which we undo the actions. In this case the $O(k)$ in the query time need not be paid.

According to Lemma 2 the problem of searching for the view is an $O(n)$ -decomposable searching problem. Using known results for $O(n)$ -decomposable searching problems, we get the following result.

THEOREM 8. *For a set of n objects in the plane, there exists a data structure that requires $O(n^2 \log n)$ storage, $O(n \log n)$ time per insertion, and $O(n)$ time for computing the view from a specified direction.*

5. PERSPECTIVE VIEWS

The perspective view of a set of objects (i.e., the view from some fixed point) consists of a partition of a circle. Each circle segment corresponds to a part of an object that one can see from the fixed point or to a place where one can look through the set. See Figure 9 for some examples of




Figure 9. Perspective views of the same set of objects from different viewpoints.

perspective views of the same set of objects from different viewpoints. One easily verifies that a perspective view contains at most $2n$ partition points and hence at most $2n$ parts. In this section, we will list a few results for computing perspective views. Proofs and details will be omitted because the methods are very similar to those described in Sections 3 and 4 for parallel views.

LEMMA 9. *Given a fixed point and a set of objects $V = \{x_1, \dots, x_n\}$. The view of V from that point can be computed in $O(n)$ time from the views of $A = \{x_1, \dots, x_i\}$ and $B = \{x_{i+1}, \dots, x_n\}$, for all $1 \leq i < n$.*

THEOREM 9. *Given a set of n objects and a fixed point, one can compute the view from that point in $O(n \log n)$ time.*

THEOREM 10. *Given a fixed point, one can maintain the view of a set of objects from that point at a cost of $O(n)$ time per insertion and deletion of an object, where n is the cardinality of the set.*

THEOREM 11. *Given a fixed point, one can maintain a set of objects at a cost of $O(\log^2 n)$ time per insertion and $O(n)$ time per deletion such that at any moment one can compute the view from the given point in $O(n)$ time.*

To adapt the ideas of walking around the set of objects to the case of perspective views, we need an analog of the notion of the visibility list.

Definition. Given a perspective view, the corresponding *visibility cycle* is the cyclic list of the numbers of the objects one can see (0 for a place where one can look through the set) in order of occurrence.

Clearly, the visibility cycle of a circular view can be computed in $O(k)$ time from the view, where k is the length of the list; and given the visibility cycle, the corresponding view from a point for which the cycle is correct can also be computed in $O(k)$ time.

We generalize the notion of walking around in the following way: given some connected curve, called *tour*, through the plane (see Figure 10), we want to maintain the visibility cycle while the point of view walks along the tour. We assume that a tour is simple in the sense that it crosses any line only a constant number of times and that the intersections with a given line can be computed in constant time. A tour is also assumed to intersect no object in the set. For example, a tour might be a circle, a line, or the boundary of some object in the set.




Figure 10. A connected curve, called tour, through a plane.

Definition. Given two objects A and B , a *critical line* of A and B is a line, the crossing of which somewhere causes a change in the visibility cycle.

The critical lines of two simple objects are exactly their common tangents. There are at most four critical lines for any two simple objects and these lines can be computed in constant time. One easily verifies that, while walking along a tour, the visibility cycle does not change as long as we do not cross a critical line of some pair of objects.

THEOREM 12. *Given a set of n objects and a tour. For each function $f(n)$, with $1 \leq f(n) \leq n/\log n$, there exists a method of maintaining the visibility cycle while walking with the viewpoint along the tour in a total of $O[n^2 + k_{cr}n/f(n)]$, time, using $O[nf(n)]$ storage, where k_{cr} is the total number of crossings between the tour and the critical lines of objects in the set.*

PROOF. The method is quite similar to the method described in the proof of Theorem 6. We first compute the visibility cycle from the starting point. One can represent this cycle in a dictionary D such that the updates needed when crossing a critical line can be performed in $O(\log n)$ time. Next, for each object i we compute the first $f(n)$ points of intersection $p_i^1, \dots, p_i^{f(n)}$ between the tour and critical lines of i and other objects

where i lies behind the other object at the place the view changes. Those $nf(n)$ intersection points are stored in a priority queue Q . The topmost point in Q is the first point on the tour where we cross a critical line. Crossing such a line consists of deleting the point from Q and performing the necessary updates on the visibility cycle in D . This takes $O(\log n)$ time. When we cross a $p_i^{f(n)}$ (for some i) we have to compute the next $p_i^1, \dots, p_i^{f(n)}$ and insert them into Q . This takes a total of $O[n + f(n) \log n]$ time. Clearly, Q always contains the first intersection with a critical line. While walking along the tour we cross k_{cr} critical lines, which takes $O(k_{cr} \log n)$ time. Moreover, we have to compute $O[k_{cr}/f(n)]$ times new intersection points (when passing a $p_i^{f(n)}$), which takes a total of $O[k_{cr}n/f(n) + k_{cr} \log n]$ time. Hence, the total amount of time needed is bounded by $O[k_{cr}n/f(n) + n^2]$ because $1 \leq f(n) \leq n/\log n$, which completes the argument. \square

Preprocessing the set of objects in such a way that views from different points can be computed efficiently gives rise to tremendous problems. As stated earlier, the visibility cycle might change when the point of view crosses a critical line. There are $O(n^2)$ critical lines and hence there are $O(n^4)$ regions with possibly different visibility cycles. Preprocessing the set in a similar way to that for parallel views as described in Section 4 would hence take $\Omega(n^4)$ storage and $\Omega(n^4 \log n)$ preprocessing time in the worst case.

We will consider a more restricted version of the searching problem. Given a set of objects and a tour, preprocess the set in such a way that the view from any point on the tour can be computed efficiently. The method for solving this problem is very similar to the method described in the proof of Theorem 7. We first compute all points of intersection of the tour with critical lines. We order those points and store them in a dictionary D . Next, we compute the visibility cycles from the starting point and after each $f(n)$ intersection points in the same way as described in the proof of Theorem 7. To search for the view from some point on the tour we search for the point in D , walk along the list of intersection points to the nearest visibility cycle, walk back while updating the visibility cycle, compute the view from the point using the updated visibility cycle, and restore afterward the visibility cycle in its original form. This leads to:

THEOREM 13. *Given a set of n objects and a tour, for each function $f(n)$, with $1 \leq f(n) \leq n$, there exists a method of preprocessing the set of objects in time $O[n^2 + k_{cr} \log n + k_{cr}n/f(n)]$ using $O[k_{cr}n/f(n) + n]$ space such that the view from any point on the tour can be computed in*

$O[f(n) \log n + k]$ time, where k_{cr} is the number of crossings between critical lines and the tour and k is the length of the view.

Using the fact that the searching problem is $O(n)$ decomposable, we get the following dynamic solution (choosing $f(n) = n/\log n$):

THEOREM 14. *For a set of n objects and a tour, there exists a data structure that requires $O(k_{cr} \log n + n)$ space, $O(n + k_{cr} \log n/n)$ time per insertion, and $O(n)$ time for computing the view from a specified point on the tour.*

6. CONCLUDING REMARKS

In this article we have provided a general setting for solving a number of computational geometry problems in the plane. In particular we have considered the frame-to-frame coherence problem for one-dimensional views of two-dimensional scenes, namely, in Flatland graphics. In this setting the objects are assumed to be convex, bounded, nonintersecting, and stationary. Either the whole scene is rotated in the plane while the viewpoint is fixed or a walk-around or prespecified tour of the scene is taken when the scene is fixed. For these equivalent scenarios we are able to demonstrate a time-space trade-off for computing and updating the visibility status of the objects. These results hold for the parallel and perspective projections.

Our results, while setting a number of questions, raise even more. These can be split, somewhat arbitrarily, into "pure" computational geometry questions and computer graphics questions. We will discuss them in this order.

In Corollary 1 we have provided an $O(n^2 \log n)$ time algorithm for the shadow problem. This problem arises in the article by Lee and Preparata [14], who consider the shortest path between two points in the presence of line segment obstacles. If there is a "shadow direction," then they have an $O(n \log n)$ time algorithm; whereas when no such direction is known, they have an $O(n^2 \log n)$ time algorithm. Thus, solving the shadow problem in substantially less than $O(n^2 \log n)$ time ensures a better algorithm for the shortest path problem when a shadow direction exists. Can the shadow problem be solved in $O(n \log n)$ or even $O(n^2)$ time? Is the shadow problem simpler than the obstacle problem, that is, computing the shortest path?

There are also questions of optimality with respect to Theorem 4. Is $n^2 \log n$ also a lower bound for the time needed to compute a walk-around? It seems that this should be reducible to the sorting of n^2 elements.

Turning to computer graphics-oriented questions, perhaps the first obvious question is whether general two-dimensional objects can be dealt with rather than only convex ones. If we relax the convexity condition, then Lemma 1 no longer holds as stated. For example, the objects given in Figure 11 exhibit the difficulties that now arise. Intuitively, the new difficulties stem from the fact that A may hide part of B while B again may hide part of A . However, these difficulties can be overcome as follows: the view of an object O is no longer a single interval but a sequence of intervals whose end points correspond to extreme points of O with respect to the given direction (see the new view of A in Figure 11).

Obviously, we have to reformulate the definition of a simple object such that the intervals that are the view of a simple object can be computed in constant time. By the same method also the perspective view can now be handled efficiently.

Similarly, the walking around algorithms can be adapted to nonconvex objects. Note that now two simple objects may give rise to more than four critical lines [resp. directions].

Second, we have shown that addition and removal of objects can be dealt with using dynamization methods. However, can moving objects be integrated into our framework? This appears to be difficult to achieve, since even moving points in the plane are nontrivial (see [16]).

Third, because we have claimed that this forms the beginning of a mathematical theory of computer graphics, how can these methods and techniques be extended into three and more dimensions? Recently, Hubschman and Zucker [12] have considered the frame-to-frame coherence




Figure 11. Views of two-dimensional objects.

problem for three-dimensional scenes consisting of stationary, bounded, convex, nonintersecting polyhedra. They make use of critical planes as part of their characterization of how the view changes during rotation. From their five characterizing criteria, they derive an algorithm to generate successive frames of an "animation" sequence. Although they present no complexity analysis at all, they demonstrate that the techniques we have used should be extendable to three and more dimensions, albeit with greatly increasing time and space requirements. One would expect this problem to be simpler for wire frames than for solid polyhedra, but we have no results at this time.

Of course, the hidden line problem of Flatland becomes the hidden surface problem for three-dimensional scenes. In the spirit of this article we are primarily interested in knowing which polyhedra are visible some extend rather than the precise specification of visible surfaces. There are two possible approaches, at least. On the one hand, because polyhedra have polygonal faces, this reduces to knowing which faces are visible to some extent. From this point of view, the three-dimensional scene becomes a collection of stationary, bounded, convex, and nonintersecting two-dimensional objects arbitrarily oriented in three-dimensional space. On the other hand, if the three-dimensional objects are arbitrary, as our two-dimensional ones are, then such a decomposition cannot, in general, be achieved. Following the development in the present article, we need to consider "simple" objects once more.

Which of these two approaches is the more reasonable remains to be seen, but it should be clear that in both cases many interesting questions are raised.

Finally let us state how we see the contents of this article in relation to the research that is done on computer graphics questions. The present research is mainly experimental, and its concerns are:

- how to deal adequately with arbitrary shapes
- how to tint, texture, color, and shade such objects
- how to represent scenes composed of such objects
- how to animate, rotate, and update such scenes

The results in this article can only be considered as the first step toward answering such questions. Whether or not an adequate mathematical theory of present-day computer graphics can be constructed remains to be seen.

ACKNOWLEDGMENTS

H.E. was supported by the Fonds zur Foerderung der wissenschaftlichen Forschung. M.H.O. was supported by the Netherlands Organization for the Ad-

vancement of Pure Research (ZWO). D.W. was supported by the Natural Sciences and Engineering Research Council of Canada, grant A-7700.

REFERENCES

1. Abbott EA: *Flatland. A Romance of Many Dimensions*, ed. 2. Dover Publications, 1884.
2. Aho AV, Hopcroft JE, Ullman JD: *The Design and Analysis of Computer Algorithms*. Addison-Wesley, Reading MA, 1974.
3. Avis D, Toussaint GT: An optimal algorithm for determining the visibility of a polygon from an edge. *IEEE Trans Comp C-30*:910-1014, 1981.
4. Beatty JC, Booth KS, Matthies LH: Complexity of scan line algorithms. Unpublished manuscript, 1981.
5. Bentley JL: Decomposable searching problems. *Info Proc Lett* 8:244-251, 1979.
6. Edelsbrunner H: Finding transversals for sets of simple geometric figures. Submitted for publication, 1982.
7. Edelsbrunner H, Maurer HA, Preparata FP, Rosenberg AL, Welzl E, Wood D: Stabbing line segments. *BIT* 22:276-281, 1982.
8. El-Gindy H, Avis D: A linear algorithm for determining the visibility polygon from a point. *J Algor* 2:186-197, 1981.
9. Freeman H, Loutrel PP: An algorithm for the two-dimensional "hidden line" problem. *IEEE Trans Electr Comp EC-16*:784-790, 1967.
10. Fuchs, H, Kedem ZM, Naylor BF: On visible surface generation by a priori tree structures. *Comp Graph* 14:124-133, 1980.
11. Gowda IG, Kirkpatrick DG: Exploiting linear merging and extra storage in the maintenance of fully dynamic geometric data structures. *Proc 18th Annu Allerton Conf on Communication, Control, and Computers*, pp 1-10, 1980.
12. Hubschman H, Zucker SW: Frame-to-frame coherence and the hidden surface computation: Constraints for a convex world. *Comp Graph* 15:45-54, 1981.
13. Lee DT, Preparata FP: An optimal algorithm for finding the kernel of a polygon. *JACM* 26:415-421, 1979.
14. Lee DT, Preparata FP: Euclidean shortest paths in the presence of rectilinear barriers. *Proc. 7th Conf. on Graphth. Concepts in Comp. Sci.*, Carl Hanser: 303-316, 1981, 1982.
15. Newman WM, Sproull RE: *Principles of Interactive Computer Graphics*, ed 2. McGraw-Hill, New York, 1979.
16. Ottmann T, Wood D: Dynamical sets of points. Unpublished manuscript, 1982.
17. Overmars MH: Dynamization of order decomposable set problems. *J Algor* 2:245-260, 1981.
18. Overmars MH: Searching in the past II: General transforms. Report RUU-CS-81-9, Dept of Comp Sci, University of Utrecht, The Netherlands, 1981.
19. Preparata FP: personal communication, 1981.
20. Schmitt A: Time and space bounds for hidden line and hidden surface algorithms. *Proc Eurographics-81*, North-Holland Publishing Company, Amsterdam, 1981.
21. Shamos MI: Problems in computational geometry. Report, Dept of Comp Sci, Carnegie-Mellon University, Pittsburgh PA, 1977.
22. Sutherland IE, Sproull RF, Shumacker RA: A characterization of ten hidden surface algorithms. *Comput Surv* 6:1-55, 1974.
23. van Leeuwen J, Mehlhorn K: Personal communication, 1981.
24. Wood D: On wire frames and haloes. Unpublished manuscript, 1982.
25. Yao FF: On the priority approach to hidden-surface algorithms. *Proc. 21st Ann. IEEE Symp. on Found. of Comp. Sci.*, pp. 301-307, 1980.

