

VISIBILITY WITH A MOVING POINT OF VIEW

Marshall Bern
David Dobkin
David Eppstein
Robert Grossman

CS-TR-235-89

November 1989

Visibility with a Moving Point of View

Marshall Bern * David Dobkin † David Eppstein * Robert Grossman ‡

Abstract

We investigate 3-d visibility problems in which the viewing position moves along a straight flightpath. Specifically we focus on two problems: determining the points along the flightpath at which the topology of the viewed scene changes, and answering ray-shooting queries for rays with origin on the flightpath. Three progressively more specialized problems are considered: general scenes, terrains, and terrains with vertical flightpaths.

1. Introduction

In recent years computer-generated images have grown commonplace, but computer-generated animations—sequences of images—are still prohibitively expensive for all but a few uses. For the most part, this disparity is inherent: high-quality animation uses at least 12 distinct images per second. On the other hand, this disparity is partially due to a lack of algorithms. Successive images are typically treated independently, even though they may differ only slightly.

In this paper we investigate a very simple type of animation: a fixed three-dimensional scene is viewed from a sequence of different points of view. More specifically, successive images correspond to (perspective or orthographic) views of a polygonal scene from sample points along a straight trajectory, or *flightpath*. Though this problem is quite basic, it is also widely applicable in flight simulation and data visualization.

We assume that scenes are to be computed in *object-space*, that is, output is given as device-independent 2-d coordinates, rather than pixel-by-pixel [SSS]. While there already exist satisfying image-space solutions both in theory (binary space partitions [FKN, PY]) and practice (hardware z-buffers), these solutions suffer

from aliasing. Object-space solutions as proposed here would produce higher quality images.

In a sequence of views of a static scene, transitions between viewpoints will typically be smooth—and rapidly computable—transformations. However, at certain points along the flightpath *topology changes* occur—for example, when an object first peeks around the edge of a closer object—and the visible scene is not so easily computed. We develop algorithms for discovering topology changes (critical flightpath points as well as the corresponding changes to the visible scene). We also describe data structures that answer *ray-shooting queries*, that is, given a ray r with origin on the flightpath and arbitrary direction, return the first polygon struck by r . This type of query is fundamental to the technique of ray-tracing.

The running times of our algorithms depend on three parameters: n , the total number of edges in all objects; p , the number of *transparent* topology changes (that is, the number of different scene topologies visible along the flightpath, assuming that all objects are transparent); and k , the number of opaque topology changes. A major open problem in this area is to replace dependence on p by dependence on k , which is typically much smaller. In general, $0 \leq k \leq p < n^3/3$. We obtain the following results for finding topology changes. In the first case we find all transparent—including opaque—topology changes; in the other two we discover only opaque topology changes.

- For general polygonal scenes, a simple algorithm with running time $O((n^2 + p) \log n)$ and more complicated algorithms with times $O(n^2 \log n + p)$ (using randomization) and $O(n^2 + p \log n)$.
- For terrains, an algorithm with time $O((n + k)\lambda_3(n) \log n)$. (A *terrain* is a polyhedral surface intersected at most once by any line parallel to the z -axis. The functions $\lambda_i(n)$ are slightly superlinear for each i [HS].)
- For terrains with vertical flightpaths, an algorithm with time $O(n\lambda_4(n) \log n)$, matching an earlier result that uses a different technique [CS].

*Xerox PARC, 3333 Coyote Hill Rd., Palo Alto, CA 94304.

†Dept. of Computer Science, Princeton University, Princeton, NJ 08544, Supported in part by NSF Grant CCR87-00917 and a Guggenheim Fellowship, work done while visiting Xerox PARC

‡Dept. of Mathematics, U. of Illinois - Chicago, Chicago, IL 60680, work done while visiting Xerox PARC

Techniques used in our algorithms include geometric sweeps and transforms similar to skewed projection [JK]. There are relationships between finding topology changes and two planar problems: the well-studied problem of line segment intersection and the problem of finding the external contour of a union of polygons.

The ray-shooting problem is, in a sense, a special case of point location in a 3-d subdivision (the visible scene cross time). For this problem we obtain the following results. In the first case, space improvement is possible if queries are ordered by time.

- For the general problem, a data structure of size $O(n^2 + k)$ with query time $O(\log^2 n)$.
- For terrains with vertical flightpaths, a data structure of size $O(n\lambda_4(n))$ with query time $O(\log n)$, improving upon a known $O(\log^2 n)$ [CS].

There has been surprisingly little work on these two problems directly, though there has been a fair amount of related work. Cole and Sharir solve a number of visibility problems on terrains, including finding topology changes and ray-shooting for the special case of vertical flightpaths [CS]. Hubschman and Zucker treat convex objects [HZ]. Swart considers the problem of viewing independently and linearly moving objects with trajectories that can be dynamically changed [S]. His running times, however, depend on events such as changes in the x -coordinate order of vertices in a projection of the scene. Plantinga and others give algorithms that compute “aspect graphs” and “aspect representations” for orthographic views of an object [PD, PDS]. These data structures have vertices or regions for each of the $O(n^4)$ topologically distinct views of an object. Translating our results into their terminology, we show that to determine all views along a given flightpath, only a small portion of the aspect representation need be computed.

2. Preliminaries

Assume we have a set S of polygons, nonintersecting except along boundaries, and a line segment f , the flightpath, in 3-space. Let f be parametrized by “time” t , running from 0 to 1. The point on f with parameter value t will be denoted $p(t)$.

For perspective views of S , we imagine projecting all polygons in S from a given point $p(t)$ on f onto a sphere centered at $p(t)$ that is large enough to contain S . (Of course, all our results also apply to the case in which we view only polygons within a *viewing pyramid*. In this case, we would project S onto a rectangular backplane.) One can view this projection as an embedding of a planar graph G_t , which has vertex set containing




Figure 1. Skewed projection of a polygonal scene.

all intersection points of edges and the obvious edge set. Similarly, the projection of S from $p(t)$ with hidden lines removed defines a subgraph of G_t called H_t . A *transparent (opaque) topology change* occurs at t if the graph G_t (respectively, H_t) changes, that is for each $\epsilon > 0$, $G_{t-\epsilon}$ and $G_{t+\epsilon}$ are nonisomorphic graphs. The following lemma is immediate.

Lemma 1. *A transparent topology change occurs at time t if and only if there are three edges e_1 , e_2 , and e_3 of (not necessarily distinct) polygons in S such that there is a line that intersects $p(t)$, e_1 , e_2 , and e_3 . An opaque topology change occurs at time t if in addition there is a line segment with one endpoint at $p(t)$ that intersects e_1 , e_2 , and e_3 and passes through no polygon interiors. \square*

If more than three edges are colinear as viewed from a point $p(t)$, then we consider $p(t)$ to be the site of a number of topology changes. An example is the case of a vertex of a polyhedron coming into view from behind a closer object.

Now let e be a line segment (not lying on the same line as fixed flightpath f), parametrized by u running between 0 and 1. Let T be the interior of the tetrahedron defined by all line segments with one endpoint on e and one on f . We define a mapping $sp_e : T \rightarrow [0, 1] \times [0, 1]$ as follows: a point $p \in T$ maps to (u, t) , where $p(u)$ and $p(t)$ are the points on e and f with parameter values u and t , and are the endpoints of the (unique) line segment l passing through p with endpoints on e and f . If e_1 is a line segment in T then $sp_e(e_1)$ is either a line segment or a connected piece of a hyperbola in $[0, 1] \times [0, 1]$.

If e and f were complete lines rather than segments, sp_e could be extended to a map from \mathbb{R}^3 to $\mathbb{R}^2 \cup \{\infty\}$. This extension is essentially the same as the *skewed projection* introduced by Jaromczyk and Kowaluk [JK].

It is not hard to see that a transparent topology change involving edge e corresponds exactly to the intersection of two curves $sp_e(e_1)$ and $sp_e(e_2)$ in the




Figure 2. Orthographic skewed projection or_e .

skewed projection of $S \cap T$. The next lemma relates opaque topology changes to the external contour of the union of skewed projections of polygons in S . See Figure 1.

Lemma 2. *An opaque topology change occurs at time t iff there exists an edge e of S and $u \in [0, 1]$ such that (u, t) is a vertex of the boundary of $sp_e(S)$.*

Proof: If (u, t) is a vertex of the boundary of $sp_e(S)$, then the line segment with endpoints $p(t)$ on f and $p(u)$ on e intersects 3 edges and the flightpath, but no polygon interior. Conversely, if there is a line segment that intersects the points $p(t)$ and $p(u)$ and two edges in T , then (u, t) must be the intersection of two curves in the skewed projection. If in addition this segment intersects no polygon interiors, then (u, t) must be a boundary vertex. \square

For orthographic views of S , flightpath f is a piece of a “circle at infinity”; f can be given by a plane P and a range of angles. Lines of sight are parallel to each other and to P . As time t varies from 0 to 1, the lines of sight vary linearly through the given range of angles. Angle is measured by translating the projection of a line of sight to an arbitrary origin on P .

We define a mapping $or_e : \mathbb{R}^3 \rightarrow [0, 1] \times [0, 1]$ as follows: a point p maps to (u, t) , where u is the parameter value of the endpoint on e and t corresponds to the angle on P of the (unique) ray r with endpoint on e that passes through p and is parallel to P . See Figure 2.

A version of Lemma 2 holds for or_e as well. In the following sections we concentrate on perspective views, but all results hold for orthographic views as well. The case of a terrain with a vertical flightpath becomes the case of a terrain with a flightpath that is part of a “vertical circle at infinity”, that is, one contained in a vertical plane. Further details for orthographic projections will be given in the full paper.

3. General Scenes

We first give a simple, practical algorithm and then two asymptotically faster algorithms, neither of which dominates the other.

Theorem 1. *All topology changes for a general scene with a line segment flightpath can be computed in time $O((n^2 + p) \log n)$ and space $O(n)$, where n is the total number of edges in all polygons and p is the number of transparent topology changes.*

Proof: For each edge e of a polygon in S , we perform a rotational sweep around e , similar to Bentley and Ottmann’s line segment intersection algorithm [BO]. Let T_t be the triangle with base equal to flightpath e and apex at point $p(t)$ on f . A *pierce point* of T_t is the intersection of T_t and an edge of a polygon in S .

The sweep proceeds from $t = 0$ to $t = 1$ as shown in Figure 3. During the sweep, a balanced binary tree maintains the pierce points of T_t sorted by angle around $p(t)$. A priority queue maintains future events by increasing t . The events to be handled are: (1) An endpoint of an edge is reached, (2) a polygon edge intersects an edge of T_t , thereby entering or leaving the sweep tetrahedron, and (3) two adjacent pierce points exchange position in the angular order. There are at most $2n$ events of types (1) and (2); scheduling these events is straightforward. The lines passing through e , f , and any other segment define a quadratic surface S (see [JK]). A fourth segment can intersect S in at most two points; thus the number of events of type (3) for a fixed edge e is at most $n(n - 1)$. Scheduling an event of type (3) amounts to finding the minimum future t at which e , f , and two other given segments are colinear. This computation— a bit of analytic geometry that we omit— takes time $O(1)$. After an event of any of the three types, at most two future events of type (3)— the upcoming colinearities of the newly adjacent pairs— must be scheduled and inserted into the priority queue. After events of types (1) or (2), at most two future events that have already been scheduled must be deleted.

A priority queue with $O(\log n)$ update times results in $O((n + p_e) \log n)$ time for a sweep around edge e , where p_e is the number of transparent topology changes discovered. The sum of all p_e values is p . \square

Theorem 2. *All topology changes for a general scene with a line segment flightpath can be computed in time $O(n^2 + p \log n)$ and space $O(n^2)$ or, alternatively, in expected time $O(n^2 \log n + p)$ and space $O(n^2)$.*

Proof: The first algorithm performs a rotational sweep around f . The configuration of pierce points of segments of S can be represented by its dual arrangement of lines, a data structure of size $O(n^2)$. Events are (1)




Figure 3. Sweep-plane algorithm for general scenes.

the appearance or disappearance of a line (corresponding to reaching a vertex of S), or (2) three lines becoming coincident (corresponding to transparent topology changes).

One possible representation of the arrangement is a graph with a node for each border segment of a face and edges between borders that share an endpoint. Each intersection of lines in the arrangement is the meeting of 8 border segments; the edges between their corresponding nodes are augmented with directional information so that faces may be traced either clockwise or counter-clockwise. We also provide pointers so that the border segments incident to an intersection can be found in $O(1)$ time given the identifiers of the two intersecting lines. Events of type (1) necessitate $O(n)$ work in updating this data structure, corresponding to the total complexity of all faces bordering the line that is inserted or deleted [EOS]. Events of type (2) necessitate $O(1)$ work as only $O(1)$ border incidences are changed.

A priority queue (implemented as a heap) holds a schedule of possible future events, including the times at which each triangular cell in the arrangement degenerates to a point. Notice that the initial $O(n^2)$ possible events can be formed into a heap in $O(n^2)$ time. As triangles “invert”, future events must be inserted or deleted, resulting in the $O(p \log n)$ part of the running time.

In the second algorithm, we perform the following steps for each polygon edge e . Compute the projection $sp_e(Q)$ of each polygon Q intersecting tetrahedron T . Each $sp_e(Q)$ will be a “curved polygon”, one with sides that are portions of hyperbolas. Next use an adaptation of Mulmuley’s randomized line segment intersection algorithm to compute all intersections in $sp_e(S)$ in expected time $O(n \log n + p_e)$, where p_e is the number of intersections [Mul]. The expectation is over the randomization used in the algorithm, not over a distribution of inputs. The sum of all p_e is p . \square

The second algorithm above explicitly computes the

points of intersection of a set of curved polygons. By Lemma 2, the computation of all opaque topology changes can be reduced to n computations of the external contour of a union of curved polygons. We expect that an improved algorithm to compute the external contour of a union of ordinary polygons should have implications for the case of curved polygons.

4. Terrains

A *terrain* is a polyhedral surface that is intersected at most once by any line parallel to the z -axis [CS, RS]. Thus the projection of a terrain onto the xy -plane is a planar subdivision. In this section S denotes a terrain with n edges. The advantage of a terrain is given by the following lemma, in which a *forward* ray with origin on flightpath f is one that has positive dotproduct with f oriented in the direction of increasing t .

Lemma 4. *In time $O(n \log n)$, the edges of S can be ordered e_1, e_2, \dots, e_n , such that if there is a forward ray from a point on f that intersects first e_i and next e_j , then $i < j$.*

Proof: Let S^* (respectively f^*) be the projection of S (f) onto the xy -plane. As in Lee and Preparata’s point location algorithm [LP, PS], the edges of S^* can be assigned to polygonal chains, monotone with respect to lines perpendicular to f^* . (A *polygonal chain* is a path of line segments connected only at successive endpoints; it is *monotone* with respect to a line l if its intersection with any line perpendicular to l is at most one point. See [PS] for more on these fundamental notions.) Chains can be ordered front to back with respect to f^* , where front is the direction of decreasing t . Within chains edges may be ordered arbitrarily. It is easy to confirm that this ordering has the desired property. \square

We define the i -th *silhouette* $S_t(i)$ to be the “horizon line” at flightpath point t , considering only the first i




Figure 4. Events in the view of a terrain

edges. That is, $S_t(i)$ is an ordered set of segments, each of which is a piece of an edge of index at most i , such that no line of sight through a segment of $S_t(i)$ passes below an edge e_j , $j \leq i$. $S_t(i)$ is monotone with respect to a horizontal line in the viewed scene. For each t and i the silhouette $S_t(i)$ has at most $\lambda_3(i)$ vertices [CS]. The function $\lambda_3(n)$ is known to be $\Theta(n\alpha(n))$ [HS].

Theorem 3. *All k opaque topology changes for a terrain with an arbitrary flightpath can be computed in time $O((n+k)\lambda_3(n)\log n)$ and space $O(n\lambda_3(n))$.*

Proof: We show how to discover topology changes that are visible along forward rays in order along the flightpath as t increases. Running the algorithm twice, once with time reversed, computes all topology changes.

The first step is to compute all silhouettes for $p(0)$ using an $O(n^2 \log n)$ -time hidden surface algorithm. The algorithm will maintain an unordered set $E_t(i)$ of all polygon edges that contribute at least once to silhouette $S_t(i)$ and an ordered list $V_t(i)$ of all vertices of the silhouette, implemented as a binary search tree. The edges can be specified simply by index, while the vertices are specified by ordered pairs of indices with the order implying the segments of $S_t(i)$. As a practical matter, these edges and lists can be maintained by a similar lists or persistent data structure, though this is not necessary for the bounds of the theorem.

As in the algorithm of Theorem 1, a priority queue maintains future events. The priority queue contains future events of two kinds sorted by increasing t : (1) some future point $p(t)$ on f , an endpoint of edge e_i , and a point on some edge in $E_t(i-1)$ are colinear, and (2) some future $p(t)$ on f , some point on edge e_i , and a vertex in $V_t(i-1)$ are colinear. See Figure 4. Given a vertex (either an endpoint of an edge e_i or a vertex of $V_t(i-1)$) and an edge, it is possible to determine their next colinearity in $O(1)$ time, since in the viewed scene vertices have either linear or quadratic apparent motion.

Notice that for each endpoint of edge e_i , each colinearity (not just the one that occurs first) with an edge

of $E_t(i-1)$ is queued. Similarly, for each vertex of $V_t(i-1)$ each colinearity is queued. Thus throughout the algorithm, the priority queue contains $O(n\lambda_3(n))$ events.

An event of type (1) may not actually be an opaque topology change, as the edge of $E_t(i-1)$ involved in the colinearity may not be part of $S_t(i-1)$ at that intersection point. An event of type (2) always will be an opaque topology change, and all opaque topology changes will be of one type or the other. (Events of type (1) are each reported at least twice, once for each of the edges of S sharing the endpoint. A minor modification avoids this redundancy.) When an event of type (1) occurs, it is tested to see whether it is an opaque topology change. An event involving an endpoint of edge e_i and edge $e_j \in E_t(i-1)$ can be tested in time $O(\log n)$ by searching within the list $V_t(i-1)$ and checking whether the endpoint of e_i lies on $S_t(i-1)$ at the current time t . There are at most $O(n^2)$ events of type (1) that are not opaque topology changes, since each vertex and edge combine to produce at most one.

In the case of an event involving e_i that is also an opaque topology change, we update each $E_t(j)$ and $V_t(j)$, $j \geq i$, along with the priority queue. For each j , the list $E_t(j)$ (respectively $V_t(j)$) is updated by inserting or deleting $O(1)$ affected edges (vertices). The priority queue is updated by deleting all events involving a vertex of $V_t(j)$ (respectively, edge of $E_t(j)$) that no longer exists and inserting all events involving a new vertex of $V_t(j)$ (new edge of $E_t(j)$). In order to find the events that must be deleted, a dictionary into the priority queue to look up events by vertex (edge) must be provided.

The number of events of type (1) scheduled at $t=0$ is bounded by $2n^2$ since (in the absence of degeneracies) each endpoint and edge uniquely specifies a future event time. The number of initial events of type (2) is bounded by $2n\lambda_3(n)$ since each vertex in $V_t(i-1)$ may combine with e_i to produce at most two events. Events that are also opaque topology changes incur extra work of time $O(\lambda_3(n)\log n)$ in inserting and deleting $O(\lambda_3(n))$ events from the priority queue. \square

It may be possible to improve Theorem 3 with a sweep algorithm that, for each edge e_i , queues only its next event, rather than all future events with the current silhouette $S_t(i-1)$. A difficult data structure problem arises in attempting such an improvement: a query asks for the earliest intersection of a line segment, each endpoint of which has linear motion, and a polygonal chain, each vertex of which has quadratic motion. The solution should be dynamic, allowing fairly rapid updates of the polygonal chain.

Cole and Sharir give an example showing k may be $\Theta(n^3)$ for terrains with arbitrary flightpaths [CS]. Thus

the algorithms of Section 3 are preferable in the case of large k .

5. Terrains with Vertical Flightpaths

In this section \mathcal{S} is a terrain and f is a flightpath parallel to the z -axis. Let e be a line segment, not lying on the same line as f .

Lemma 5. *Each vertical (constant u) line in $[0, 1] \times [0, 1]$ intersects the boundary of $sp_e(\mathcal{S})$ at most once.*

Proof: Assume two points (u, t_1) and (u, t_3) both lie outside $sp_e(\mathcal{S})$, but some point (u, t_2) with $t_1 < t_2 < t_3$ lies inside. Then the interior of the vertical triangle in 3-space with vertices at point u on e and points t_1 and t_3 on f intersects \mathcal{S} , but the lower edge of this triangle does not intersect \mathcal{S} . This contradicts the fact that \mathcal{S} is a terrain. \square

Theorem 4. *All $O(n\lambda_4(n))$ opaque topology changes for a terrain with a vertical line segment flightpath can be computed in time $O(n\lambda_4(n)\log n)$ and space $O(\lambda_4(n))$.*

Proof: The algorithm goes as follows. For each edge e of \mathcal{S} , we repeat the following steps. We compute the image $sp_e(g)$ of each edge g of \mathcal{S} . By Lemmas 2 and 5 topology changes occur at exactly the vertices of the pointwise maximum of the curved segments $sp_e(g)$. To compute the pointwise maximum we use the divide-and-conquer method of Atallah: recursively compute the pointwise maximum of two halves of the set of curved segments and then merge these maxima [A]. The pointwise maximum has complexity $O(\lambda_4(n))$, and the divide-and-conquer algorithm takes time $O(\lambda_4(n)\log n)$. \square

Cole and Sharir give an example to show that the number of opaque topology changes for terrains with vertical flight paths may be $\Omega(n\lambda_3(n))$. It is unknown whether the number of topology changes may be as high as $\Theta(n\lambda_4(n))$.

6. Ray-shooting for General Scenes

In this section we sketch a data structure to answer ray-shooting queries for a general polygonal scene with an arbitrary flightpath. In the next section, we specialize to the case of terrains with vertical flightpaths. In the first case, we use a direct approach, that is, we maintain the visible scene as a subdivision of a 2-sphere and treat ray-shooting queries as point location queries. In the second case we use the dual approach of Cole and Sharir [CS].

An interesting feature of this problem is that the subdivision is dynamic in two senses. At topology changes edges must be inserted or deleted; between topology changes the subdivision transforms continuously. Preparata and Tamassia have recently considered the problem of monotone planar subdivisions dynamic in the first sense [PT]; we make use of their results.

Assume without loss of generality that line segment f lies along the z -axis. Sphere S_t will be centered at point $p(t)$ on f ; each S_t is the same size and large enough that it contains all of \mathcal{S} . Assume that S_t is parametrized by spherical coordinates ϕ (latitude) and θ (longitude) with f lying along its polar axis. Thus lines parallel to the z -axis project to constant- θ lines (meridians).

The first step is to compute the projection of \mathcal{S} onto the initial sphere S_0 . Next hidden lines are removed, giving an initial view of the scene that may be considered as a planar graph H_0 or as a polygonal subdivision of S_0 . The polygonal subdivision can be made monotone with respect to latitude lines (that is, the intersection of any cell with a meridian is a single segment) by adding some *artificial edges* that extend latitudinally (along constant- ϕ lines) from *interior cusps*, as shown in Figure 5. We then compute Lee and Preparata's *chain tree* in order to answer point location queries in this subdivision (see [PS] or [LP]). A chain tree stores a monotone polygonal chain at each node. Each edge of the subdivision is explicitly listed in only one chain, though we may think of each chain as completely dividing the subdivision into higher-latitude and lower-latitude parts. Because we have fixed the orientation of the scene by choosing f to lie along the polar axis, some "monotone" chains may include meridial segments; this degeneracy does not cause any real difficulties. (We call a line segment *meridial* if it lies along a meridian.)

Notice that there is a one-to-one correspondence between point location queries in the subdivision and ray-shooting queries with origin at $p(0)$. The following lemma assures us that a chain that is monotone with respect to latitude remains monotone as we vary t , so long as its topology remains unchanged. Notice that under a smooth transformation an edge must become meridial before it "bends backwards".

Lemma 6. *If edge e projects to a meridial segment from some point along f , then e projects to a meridial segment from every point along f .*

Proof: If edge e projects to a meridial segment from some point $p(t)$ along f , then e is contained in a plane containing f . \square

Notice that the chain tree, unlike other planar point location data structures, does not need to change as the subdivision transforms smoothly while remaining

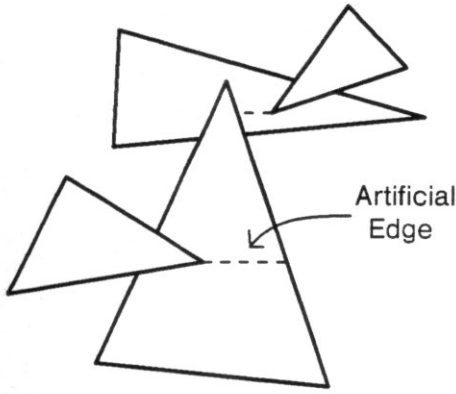


Figure 5. Making a subdivision monotone.

monotone. That is, comparing a query ray (given by time t and spherical coordinates ϕ and θ) against a chain \mathcal{C} still takes only $O(\log n)$ time, since the spherical coordinates of a given vertex or edge of \mathcal{C} at time t can be computed in $O(1)$ time.

Each topology change necessitates the addition or deletion of $O(1)$ edges and vertices from the polygonal subdivision. When an interior cusp first comes into view an artificial edge must also be added. Each addition or deletion is an update that can be handled by the methods of Preparata and Tamassia [PT]; in fact, our updates are local, special cases. Thus we can update the chain tree in time $O(\log^2 n)$. By using the persistence methods of [DSST] to maintain “old versions” of the chain tree, we can answer ray-shooting queries with arbitrary origins on f . If ray-shooting queries are ordered by time, then we may update the chain tree nonpersistently instead.

In addition to handling topology changes, however, we must also handle *artificial topology changes*, that is, points along f at which graph H_t changes because an artificial edge a of H_t intersects a vertex v not previously on a . At artificial topology changes we must add a new vertex v' to the subdivision (at first coincident with v) and redefine the artificial edge to lie between v' and the interior cusp. The next lemma shows that the number of artificial topology changes is not excessive.

Lemma 7. *There are $O(n^2)$ artificial topology changes along f .*

Proof: Assume artificial edge a lies within a polygonal face F in the embedding of H_t and that a intersects a vertex v of H_t at time t but not at any prior time after the last topology change. Then v must be a vertex of the boundary of F at which the interior angle is reflex; hence v must be the projection of a vertex of a polygon of S . Thus at time t , two vertices of S —the one that induces artificial edge a and the one corresponding to v —project to the same ϕ -coordinate, and these vertices

do not project to the same ϕ -coordinate at all times. There are $O(n^2)$ such t . \square

Theorem 5. *For general scenes with arbitrary flightpaths, a data structure of space $O(n^2 + k)$ that answers ray-shooting queries in time $O(\log^2 n)$ can be built in preprocessing time $O((n^2 + k)\log^2 n + p \log n)$. If queries are ordered by time, then the space can be reduced to the maximum complexity of a visible scene along f .*

Proof: We first run the algorithm of Theorem 1 and remember all opaque topology changes. We also compute all artificial topology changes in time $O(n^2)$ by testing each pair of vertices of S . We then follow the method given above: compute the initial scene with hidden lines removed, build a chain tree, and persistently update the chain tree through topology changes. The preprocessing time follows from Theorem 1, the query time from [LP] and [PT], and the space bound for unordered queries from [DSST]. \square

7. Ray-shooting for Terrains with Vertical Flightpaths

Assume S is a terrain and f is a segment along the z -axis. For simplicity, assume f is the entire z -axis. Below we describe a data structure that answers ray-shooting queries for rays with origin on f in time $O(\log n)$. As above, a ray is given by a triple (t, θ, ϕ) , where $t = z$ is a parameter running along the flightpath, θ is longitude around sphere S_t , and ϕ is latitude.

We briefly describe the method of Cole and Sharir [CS]. Consider the intersection of S with the vertical half-plane with boundary f and a fixed longitude θ_0 . The intersection is a polygonal chain \mathcal{C} as shown in Figure 6(a). If points in the vertical half-plane are given by cylindrical coordinates (r, z) , then a ray with origin $p(t)$ on f and longitude θ_0 can be specified by an equation $z = ar + t$, $r \geq 0$. A duality mapping takes such a ray to a point $(-a, t)$. Each polygon P_i in Figure 6(b) consists of exactly those points that are dual to rays that first strike a given segment of the chain in 6(a). Polygons in 6(b) are unbounded, since one can see the entire terrain from a sufficiently high viewpoint. (Think of the horizontal axis as ϕ , though ϕ varies nonlinearly with horizontal distance.) Furthermore, each edge of the polygonal subdivision $D(\theta_0)$ in 6(b) lies on a ray r_i formed by the union of edges of $D(\theta_0)$. (Rays r_i are the duals of viewing rays through a vertex of \mathcal{C} .)

Point location on $D(\theta_0)$ answers ray-shooting queries with longitude θ_0 . What happens to this polygonal subdivision as θ varies? Between two successive critical longitudes, the topology of subdivision $D(\theta)$ remains constant. There are two types of critical longitudes:




Figure 6. (a) Cross-section of S at θ_0 . (b) Dual subdivision $D(\theta_0)$. (c) A topology change in $D(\theta)$.

(C1) the longitudes of vertices of S , and (C2) longitudes at which 3 vertices of C and flightpath f can be connected by a straight segment that passes through no interiors of edges of C . There are at most n critical longitudes of type (C1) and $O(n\lambda_4(n))$ of type (C2) [CS]. At a critical longitude of type (C2), two vertices v_i of $D(\theta)$ pass through each other as shown in Figure 6(c). Below we view such a topology change as a rotation in a binary tree.

The crux of the ray-shooting problem is to give a planar point location method that works for varying θ . Cole and Sharir use chain trees. In the proof below we describe a faster method that exploits the fact that for each θ the edges of $D(\theta)$ form a tree.

Theorem 6. *For terrains with vertical flightpaths, a data structure with space complexity $O(n\lambda_4(n))$ that answers ray-shooting queries in time $O(\log n)$ can be built in preprocessing time $O(n\lambda_4(n) \log n)$.*

Proof: We first divide S into wedge-shaped strips by cutting outwards from f along a plane of constant θ through each vertex of S . We shall build a separate search structure for each strip. Building an initial search structure for a strip can be accomplished in time $O(n \log n)$ and finding the strip for a given ray-shooting query takes time $O(\log n)$, so we may treat strips separately. (A unified structure, however, would be an improvement in practice.)

Now consider the polygonal subdivision $D(\theta_0)$ in the dual space of rays for the minimum longitude θ_0 in a strip as in Figure 6(b). $D(\theta_0)$ gives an (unbalanced) binary search tree T_{θ_0} as shown in Figure 7(a). (Think of hanging $D(\theta_0)$ upside-down.) At each node of T_{θ_0} , an $O(1)$ -time test determines whether a query point (t, ϕ) lies above or below the line through the named ray of $D(\theta_0)$. Notice that such a search tree remains invariant as $D(\theta)$ transforms smoothly.

We now show how to create a balanced search tree

using a technique familiar in the design of parallel algorithms. Following Miller and Reif, we define an operation *Rake* on rooted trees that merges each leaf with its parent [MR]. Call a connected set of degree-2 nodes in a tree a *path*; a node is called a path node if it lies on a path. Define an operation *Compress* that merges adjacent pairs of path nodes simultaneously all over the tree. Any set of adjacent pairs may be chosen, so long as any set of 4 successive vertices along a path contains a pair that merge. This is a nondeterministic, generalized version of *Compress*; the ordinary version merges successive pairs. The proof given in [MR] generalizes to show that any n -node tree is reduced to a single node after at most $c \cdot \log n$ alternating applications of *Rake* and *Compress*, where c is a constant.

We alternately apply *Rake* and *Compress*, starting with *Rake*, to T_{θ_0} until we obtain a single supernode, as shown in Figure 7(b). Here a dashed oval represents a merging due to *Rake* and a solid oval a merging due to *Compress*; numbers indicate the order in which supernodes merge. For simplicity, the *Rake* operation numbered 1 is not shown; *Compress* operations 6 and 8 do nothing.

We can define a new search tree level-by-level by considering each combined supernode as the parent of the combining supernodes. Each internal node in the new search tree R_{θ_0} results from the merger of two supernodes along an edge of T_{θ_0} or from the merger of two leaves and their parent. Thus at least one of the child supernodes corresponds to a proper subtree of T_{θ_0} . A proper subtree of T_{θ_0} corresponds, in turn, with a roughly wedge-shaped unbounded polygon in D_{θ_0} . This polygon has a lower boundary that is a ray and an upper boundary that is a convex chain. For example, the root of R_{θ_0} in Figure 7(c) corresponds to merger 9 in Figure 7(b), which is along the edge between the nodes labeled r_1 and r_8 in Figure 7(a), which in turn corresponds to the edge between v_1 and v_8 in




Figure 7. (a) Polygonal subdivision tree T_{θ_0} . (b) As merged by *Rake* and *Compress*. (c) Balanced search tree R_{θ_0} .

Figure 6(b). The associated wedge has vertex v_8 and an upper boundary formed by r_8 and r_{11} .

We augment each internal node of R_{θ_0} with the following information: (I1) the coordinates (t, ϕ) of the leftmost vertex v_i of the corresponding wedge-shaped polygon in $D(\theta_0)$ (as named in Figure 7(c)), (I2) the slope of the polygon's lower boundary, and (I3) the largest slope of a boundary segment of the wedge-shaped polygon. (Notice that (I1), (I2), and (I3) will vary predictably with θ once longitude is unfixed.) This information allows an $O(1)$ -time "within-wedge" test to determine whether a given point query (t, ϕ) lies in the left or right subtree of a node in R_{θ_0} . Points in the polygon P_i immediately above the wedge-shaped polygon may go either way in this test. For example, a point just above the line segment between v_6 and v_7 in Figure 6(b) may go either way when tested at the node marked v_3 , depending on whether it falls to the right or left of a line through v_3 with the same slope as r_7 . Say this point tests inside v_3 's and v_5 's wedges, outside v_7 's wedge, and finally inside v_6 's wedge; then a single, final test determines whether the point lies in P_6 or P_8 . These extra tests are indicated at the leaves in Figure 7(c); thus the number of tests needed for point location may be one more than the height of R_{θ_0} . In Figure 7(c), i marks the direction to take if a point tests *in* the wedge. Altogether point location for queries at longitude θ_0 can be accomplished using tree R_{θ_0} in $O(\log n)$ time.

Search tree R_{θ_0} is actually valid for all θ until the next critical longitude. At a critical longitude, either the strip ends or a rotation occurs in tree T_{θ} . We now show that by changing only $O(\log n)$ nodes and edges of R_{θ} at a rotation of T_{θ} , we can maintain the invariant that R_{θ} is a tree that could have resulted from T_{θ} by an alternating sequence of *Rake* and *Compress* operations.

A generic rotation is shown in Figure 8, with the T_{θ}




Figure 8. Before and after a rotation in T_{θ} .

trees shown before and after a critical longitude. (Of course, before and after could be reversed.) After an alternating sequence of *Rake* and *Compress* operations, call a supernode in the left tree *clean* if it contains neither y nor z and is not the parent of a supernode containing y .

Assume inductively that each clean supernode on the left (except at most one) has a *counterpart* on the right, that is, a supernode containing exactly the same set of original nodes of T_{θ} . This condition certainly holds before any *Rake* and *Compress* operations have been performed. Now consider applying *Rake* to both the left and right trees. The counterparts of each pair of clean supernodes that merge on the left will merge on the right, since the adjacencies of clean supernodes and their counterparts are identical.

Now consider applying *Compress* to the left. We assert that there exists a valid *Compress* for the right tree that maintains counterparts for each clean supernode. We join the counterparts of each merging pair of clean supernodes in this *Compress*. The pairing of other supernodes on the right is then dictated by this earlier pairing. For example, if A and B are both single nodes in Figure 8, then the first *Compress* on the left may

combine x and y but x' may have to remain unchanged on the right. The next merger above x and y , however, can be mimicked on the right. As in this example, the pairing on the right may leave "gaps", that is, the merging pairs may be unsuccessful, but gaps of one are legal in our nondeterministic version of *Compress*.

There is also the case that the *Compress* on the right must merge the counterparts of a pair that did not merge on the left. Thus a single clean supernode on the left can lose its counterpart on the right. This loss cannot be repeated, however, until it has been reversed (i.e., until every clean supernode on the left has regained a counterpart). Thus after any number of *Rake* and *Compress* operations, there is a one-to-one mapping that takes all but one clean supernode on the left to a counterpart on the right; except for $O(1)$ nodes on the right this mapping is onto.

Altogether we conclude that only $O(1)$ supernodes in each level of search tree R_θ must change at a critical longitude. Furthermore, information (I1), (I2), and (I3) can be updated in time $O(1)$ per changed supernode by consulting the children of the changing supernode.

Finally we claim that the changes to R_θ can be performed persistently using the methods of [DSST]. Thus we obtain an $O(\log n)$ -time search structure for each strip of the scene. \square

8. Conclusions

We have given algorithms for some natural computer graphics problems that have not received sufficient attention. There are numerous possibilities for improvements to our algorithms. We list some specific open questions that we find intriguing.

- Can all transparent topology changes for general scenes be found in deterministic time $O(n^2 \log n + p)$? In (randomized or deterministic) time $O(n^2 + p)$?
- Can the external contour of a union of triangles (or "curved triangles") be found in time faster than the total number of intersections of sides? (It appears that Mulmuley's randomized methods give a positive answer to these questions, with running time proportional to a sum in which each intersection contributes the reciprocal of one more than the number of polygons strictly containing it [Mu2]. This would improve the running time of the $O(n^2 \log n + p)$ algorithm of Theorem 2.)
- Can all opaque topology changes for general scenes be found in time sensitive to k ? (The analogous question for static viewpoints is the longstanding open question of finding an output-sensitive hidden line removal algorithm.)
- Can the "sensitivity"—i.e., the term involving k —of our algorithm for terrains with arbitrary flight-paths be improved?
- Can ray-shooting queries for general scenes be answered in time $O(\log n)$? Even in the special case of no opaque topology changes along f ?
- Can our results be generalized to linearly moving objects?

References

- [A] M. J. Atallah, *Dynamic Computational Geometry*, Proc. 24th IEEE Foundations of Comp. Science, 1983, 92-99.
- [BO] J. L. Bentley and T. A. Ottmann, *Algorithms for Reporting and Counting Geometric Intersections*, IEEE Trans. on Computers 28 (1979), 643-647.
- [CS] R. Cole and M. Sharir, *Visibility Problems for Polyhedral Terrains*, J. Symbolic Computation 7 (1989), 11-30.
- [DSST] J. R. Driscoll, N. Sarnak, D. Sleator, and R. E. Tarjan, *Making Data Structures Persistent*, Proc. 18th ACM Symp. on Theory of Computing, 1986, 109-121.
- [EOS] H. Edelsbrunner, J. O'Rourke, and R. Seidel, *Constructing Arrangements of Lines and Hyperplanes with Applications*, SIAM J. Computing 15 (1986), 341-363.
- [FKN] H. Fuchs, Z. M. Kedem, and B. F. Naylor, *On Visible Surface Generation by A Priori Tree Structures*, Computer Graphics 14 (1980), 124-133.
- [HS] S. Hart and M. Sharir, *Nonlinearity of Davenport-Schinzel Sequences and of Generalized Path Compression Schemes*, Combinatorica 6 (1986), 151-177.
- [HZ] H. Hubschman and S. Zucker, *Frame-to-Frame Coherence and the Hidden Surface Computation: Constraints for a Convex World*, Computer Graphics 15 (August 1981), 45-54.
- [JK] J. W. Jaromczyk and M. Kowaluk, *Skewed Projections with an Application to Line Stabbing in R^3* , Proc. 4th ACM Symp. on Comp. Geometry, 1988, 362-370.

- [LP] D. T. Lee and F. P. Preparata, *Location of a Point in a Planar Subdivision and its Applications*, *SIAM J. on Computing* 6 (1977), 594-606.
- [MR] G. L. Miller and J. H. Reif, *Parallel Tree Contraction and its Applications*, *Proc. 26th IEEE Foundations of Comp. Science*, 1985, 478-489.
- [Mu1] K. Mulmuley, *A Fast Planar Partition Algorithm, I*, *Proc. 29th IEEE Foundations of Comp. Science*, 1988, 580-589.
- [Mu2] K. Mulmuley, *On Obstructions in Relation to a Fixed Viewpoint*, *Proc. 30th IEEE Foundations of Comp. Science*, 1989.
- [PY] M. Paterson and F. F. Yao, *Binary Partitions with Applications to Hidden Surface Removal and Solid Modelling*, *Proc. 5th ACM Symp. on Comp. Geometry*, 1989, 23-32.
- [PD] W. H. Plantinga and C. R. Dyer, *An Algorithm for Constructing the Aspect Graph*, *Proc. 27th IEEE Foundations of Comp. Science*, 1986, 123-131.
- [PDS] W. H. Plantinga, C. R. Dyer, and B. Seales, *Real-time Hidden-Line Elimination for a Rotating Polyhedral Scene Using the Aspect Representation*, manuscript, 1988.
- [PS] F. P. Preparata and M. I. Shamos, *Computational Geometry: An Introduction*, Springer-Verlag, 1985.
- [PT] F. P. Preparata and R. Tamassia, *Fully Dynamic Techniques for Point Location and Transitive Closure in Planar Structures*, *Proc. 29th IEEE Foundations of Comp. Science*, 1988, 558-567.
- [RS] J. H. Reif and S. Sen, *An Efficient Output-Sensitive Hidden-Surface Removal Algorithm and its Parallelization*, *Proc. 4th ACM Symp. on Comp. Geometry*, 1988, 194-200.
- [SSS] I. E. Sutherland, R. F. Sproull, and R. A. Schumacker, *A Characterization of Ten Hidden-Surface Algorithms*, *Computing Surveys* 6 (1974), 1-25.
- [S] G. R. Swart, *A Schema for Real Time Hidden Line Removal*, Tech. Report, Dept. of Computer Science, U. of Washington, 1984.