Stabbing Orthogonal Objects in 3-Space*

David M. Mount
Fan-Tao Pu

Institute for Advanced Computer Studies
and Department of Computer Science
University of Maryland
College Park, MD 20742
{mount,ftpu}@cs.umd.edu

October 29, 1996

Abstract

We consider a problem that arises in the design of data structures for answering *visibility range queries*, that is, given a 3-dimensional scene defined by a set of polygonal patches, we wish to preprocess the scene to answer queries involving the set of patches of the scene that are visible from a given range of points over a given range of viewing directions. These data structures recursively subdivide space into cells until some criterion is satisfied. One of the important problems that arise in the construction of such data structures is that of determining whether a cell represents a nonempty region of space, and more generally computing the size of a cell.

In this paper we introduce a measure of the *size* of the subset of lines in 3-space that stab a given set of n polygonal patches, based on the maximum angle and distance between any two lines in the set. Although the best known algorithm for computing this size measure runs in $O(n^2)$ time, we show that if the polygonal patches are orthogonal rectangles, then this measure can be approximated to within a constant factor in $O(n)$ time.

Keywords: lines in 3-space, Plücker coordinates, approximation algorithms, orthogonal polygons, computer graphics

*The support of the National Science Foundation under grant CCR–9310705 is gratefully acknowledged.

1 Motivation

Computer graphics is the source of many interesting and challenging applications for the design of geometric algorithms and data structures. Applications in global illumination simulation and radiosity [8] have motivated our study of problems involving lines in 3-space. In particular we are interested in geometric data structures for answering *visibility range queries*. Given a 3-dimensional scene defined by a set of polygonal patches, we wish to preprocess the scene to answer queries involving the set of patches of the scene that are visible from a given range of points over a given range of viewing directions. Queries of this form are central to global illumination simulation. This has led us to the study of data structures for storing and accessing subdivisions over *directed-line space*, where each point in this space is associated with a directed line in 3-space.

In the search for these data structures, we have come across an important problem. Our solution to this problem shows that choices for the data structures used to store the 3-dimensional scene may have a considerable impact on the design of data structures and algorithms for answering visibility range queries.

Data structures based on recursive subdivisions of space are popular methods for answering range queries. These data structures recursively subdivide space into *cells* until some criterion is satisfied. Such criterion typically involve the number of data objects that intersect the cell, but in the case of approximations may also involve the size of the cell, measured by its diameter or volume [7]. The problems of determining the size of a cell and whether a cell represents a nonempty region of space are implicit (and usually trivial) in the design of algorithms for building these data structures.

In the case of directed lines in 3-space, there is a natural way of subdividing space. Given two directed lines, we may assign them an *orientation*, which is either positive, zero, or negative according to the direction of one line relative to a viewer looking along the direction of the other. Thus

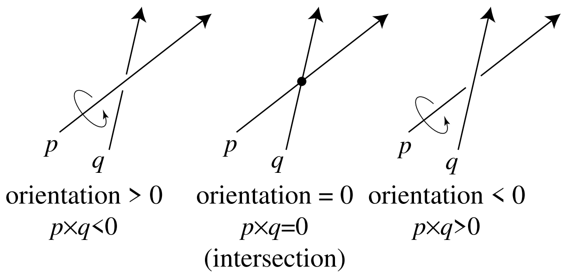

Figure 1: Orientations.

we may use directed lines in 3-space as a basis for subdividing directed-line space. An equivalent formulation, which may be more natural in some circumstances, is to consider subdivisions of line space according to whether lines intersect an oriented polygonal patch. In particular, given a convex polygonal patch in 3-space, whose boundary is oriented in some direction, we say that a directed line *stabs* the patch if it intersects the interior of patch and it crosses the oriented patch according to a right-handed transversal. Equivalently the line has a positive orientation with respect to every directed edge of the patch. Such a line is also called a *transversal*.

Consider a set of convex, oriented polygonal patches in 3-space. Let n denote the total number of edges in all the patches. Teller showed that in $O(n^2)$ time it is possible to determine the existence of a transversal of these patches, and more generally to compute a complete description of the set of stabbing lines [10]. Amenta [2] showed that if the patches consist of a set of orthogonal (that is, axis-aligned) rectangles, then the existence of a transversal can be determined in more efficiently in $O(n)$ time by a variation of linear programming. Megiddo [5] showed that the orthogonal case can be solved in $O(n)$ time in all fixed dimensions through linear programming. This suggests one advantage of using orthogonal rectangles and lines in 3-space (as arise in quad-trees, k-d trees, and R-trees [7], for example) as the basis of subdivisions of directed line space.

An extension to this problem is that of determining not just the existence of a stabbing line, but various properties and functions of this set of lines. Examples of this include computing random samples over this set, balanced spatial decompositions, and measurements of the size or diameter of this set. It is the latter problem that we consider here. We introduce a way of measuring the *size* of a set of directed lines that stab a set of n polygonal patches in 3-space. We show that if the patches are orthogonal rectangles, then this measure can be approximated in $O(n)$ time.

2 The Size of a Set of Lines

Let P denote a set of oriented convex polygonal patches in 3-space. Let $S(P)$ denote the set of *stabbers*, that is, the set of directed lines in 3-space that stab every patch in P. Given two lines $\ell_p, \ell_q \in S(P)$, let $\mathrm{dist}(\ell_p, \ell_q)$ denote the minimum distance between these lines, and let $\mathrm{ang}(\ell_p, \ell_q)$ denote the minimum angle of rotation to make ℓ_p parallel to ℓ_q. Intuitively, these two quantities will be small if patches of P define a narrow "tunnel" through which the stabbers must pass.

Because the distance and angle taken from different domains, and since our size will depend on both, we normalize them over the interval $[0, 1]$ as follows. We assume that we are given an orthogonal cube C in 3-space that encloses the scene and all of the patches of P. Let D denote the diameter of C. Every line in $S(P)$ intersects C, and hence the closest distance between any pair of stabbers is at most D. Define the *normalized distance* $\mathrm{dist}'(\ell_p, \ell_q)$ to be $\mathrm{dist}(\ell_p, \ell_q)/D$. We also assume that the lines of $S(P)$ lie within a region of angular diameter $\pi/2$, and thus the angle between ℓ_p and ℓ_q can be at most $\pi/2$. We define the *normalized angle* $\mathrm{ang}'(\ell_p, \ell_q)$ to be $\sin \mathrm{ang}(\ell_p, \ell_q)$. Clearly both normalized quantities lie in the interval $[0, 1]$. Finally we define the *size* of $S(P)$ to be

$$\mathrm{size}(P) = \sup_{\ell_p, \ell_q \in S(P)} \max(\mathrm{ang}'(\ell_p, \ell_q), \mathrm{dist}'(\ell_p, \ell_q)).$$

Intuitively, if the set P admits a set stabbers that have a very narrow range of angles and distances, then $\mathrm{size}(P)$ will be small.

Although we do not know of an efficient way to compute $\mathrm{size}(P)$ without of computing a complete representation of $S(P)$ (which can be done in roughly $O(n^2)$ time [10]), we claim that we can approximate the size in $O(n)$ time. Our approach combines the use of Plücker coordinates and variants of linear programming spaces of fixed dimension.

3 Plücker Coordinates

Perhaps the most elegant method for representing directed lines in 3-space is through the use of Plücker coordinates. Let ℓ be a directed line in 3-space, and let a and b be any two points on ℓ, such that ℓ is directed from a to b. Let $[a_0, a_1, a_2, a_3]$ and $[b_0, b_1, b_2, b_3]$ be the homogeneous coordinates of a and b with $a_0, b_0 > 0$ be the homogenizing coordinates (so that the Cartesian coordinates of a, for example, are $(a_1/a_0, a_2/a_0, a_3/a_0)$). The *Plücker coordinates* of ℓ are the six real numbers

$$[\pi_{01}, \pi_{02}, \pi_{03}, \pi_{23}, \pi_{31}, \pi_{12}],$$

where $\pi_{ij} = a_i b_j - a_j b_i$. It is easy to see that any positive scalar multiple of the coordinate vector represents the same line. Thus we can regard the Plücker coordinates of a directed line as a point in projective 5-space $\mathcal{P}^5$ [9]. Not all Plücker coordinates represent lines in 3-space. The coordinates of a line in 3-space must satisfy the quadratic equation

$$\pi_{01}\pi_{23} + \pi_{02}\pi_{31} + \pi_{12}\pi_{03} = 0.$$

The locus of points satisfying this equation is called the *Grassmann manifold* (also called the Plücker hypersurface).

Consider two directed lines ℓ_p and ℓ_q in 3-space, and let p and q denote their respective Plücker coordinates. Define $p \times q$ to be

$$p_{01}q_{23} + p_{02}q_{31} + p_{03}q_{12} + p_{23}q_{01} + p_{31}q_{02} + p_{12}q_{03}.$$

Define the *orientation* of ℓ_p relative to ℓ_q to be the negation of the sign of $p \times q$. Observe that the set of lines ℓ_p that have a particular orientation relative to a fixed line ℓ_q is the intersection of a halfspace in 5-dimensional Plücker space with the Grassmann manifold.

Given a set P of oriented polygonal patches in 3-space, two lines ℓ_p and ℓ_q are in the same *orientation class* (*o-class*) relative to P if both have the same orientations with respect to all the lines of P. The set $S(P)$ is the particular o-class that we are interested in. These two lines are in the same *isotopy class* (*i-class*) relative to P if there is way of continuously mapping ℓ_p to ℓ_q (a connected path on the Grassmann manifold between their respective Plücker coordinates) such that all the points on this path are in the same o-class. A striking feature of lines in 3-space is that two lines may be in the same o-class relative to P, but are not in the same i-class [4]. We say that two directed lines ℓ_p and ℓ_q are in the same *orthogonal sign class* (*s-class*) if the signs of the components of the 3-dimensional directional vectors for the two lines are equal. There are 8 orthogonal sign classes, one for each of the coordinate octants.

Given a set of n oriented patches P, it follows that $S(P)$ is the intersection of the Grassmann manifold together with a set of n halfspaces (one for each line of P) in Plücker space. An i-class of $S(P)$ is any connected component of this intersection. Consider a set of patches P consisting of a set of oriented orthogonal rectangles in 3-space, and suppose that we consider only lines in a single orthogonal sign class. (This will be the case if for each coordinate axis there is at least one rectangle in P orthogonal to this axis.) Amenta has shown in this case $S(P)$ is connected [1], that is, if it is nonempty, then it has a single i-class. This is one of the important properties possessed by orthogonal constraints.

4 The Approximation

For our approximation we assume that we are given a set P of n oriented orthogonal rectangles in 3-space, and a cube C enclosing the rectangles of P. To restrict ourselves to a single orthogonal sign class, we assume that P contains at least one rectangle orthogonal to each of the three coordinate axes. If this is not the case, then we may add an appropriate face of C to enforce this constraint. We may assume without loss of generality that the lines of $S(P)$ are directed into the positive octant.

Our approximation is based on two observations. The first is that the first three Plücker coordinates indicate the direction of a line. In particular, it is easy to verify that if $(\pi_{01}, \pi_{02}, \pi_{03})$ are the first three coordinates of a directed line ℓ (not a line at infinity), then this vector is the directional vector for ℓ. Because all scalar multiples of Plücker coordinates represent the same line, we may apply a suitable normalization by requiring that $\pi_{01}+\pi_{02}+\pi_{03}=1$. From our assumption that lines are directed into the positive octant, these coordinates will all lie within the interval $[0,1]$. By computing bounds on these three normalized components for all lines in $S(P)$, we can bound the range of possible line angles.

The second observation is that, although it is difficult to bound the normalized distances between lines, it is a consequence of the fact that P consists of orthogonal rectangles that if there are two lines of $S(P)$ with a large normalized distance, then there are two lines of $S(P)$ whose normalized angle is proportionately as large. This does not hold for arbitrary convex polygonal patches. Thus, it suffices to compute an approximation to the maximum normalized angle.

Here is the approximation procedure. Given the set of patches P, convert each directed side of each rectangle of P into a halfspace in 6-dimensional Plücker space. Add to this set of halfspaces the constraint $\pi_{01}+\pi_{02}+\pi_{03}=1$. Observe that any line directed into the positive octant intersects this plane in a point whose coordinates are in the interval $[0,1]$. It follows from Amenta's results [1] that the resulting system of $O(n)$ linear inequalities is of LP-type [3]. For each axis, $i \in \{1,2,3\}$, invoke any linear time procedure for generalized linear programming (in dimension 6) to determine lower and upper bounds on the i-th coordinate, denoted $\pi_{0,i}^-$ and $\pi_{0,i}^+$, respectively. If the system is infeasible, then return an indication of this. Otherwise return the maximum width,

$$\Delta = \max_i \left(\pi_{0,i}^+ - \pi_{0,i}^-\right)$$

as the approximation. Observe that $0 \le \Delta \le 1$. The running time is dominated by the $O(n)$ time needed for generalized linear programming in dimension 6 [3].

5 Analysis

In this section we show that the above algorithm returns a value that is within a constant factor of size(P). In particular, we prove the following result.

Theorem 5.1 *Given a set of oriented orthogonal patches P, let Δ denote the value returned by the size approximation algorithm of the previous section. Then*

$$\frac{1}{3\sqrt{2}} \le \frac{\text{size}(P)}{\Delta} \le 4\sqrt{5}.$$

The proof begins with a series of technical lemmas. The first lemma provides bounds on the angle between two vectors, in terms of bounds on the differences in their coordinates.

Lemma 5.1 *Let p_0 and p_1 be two points in 3-space that lie in the positive octant on the plane $T: x+y+z=1$, and let*

$$\Delta = \max(|p_0.x - p_1.x|, |p_0.y - p_1.y|, |p_0.z - p_1.z|).$$

Letting o denote the origin, consider the two vectors $\vec{op}_0$ and $\vec{op}_1$, and let θ denote the angle between these vectors. Then

$$\frac{\Delta}{3\sqrt{2}} \le \sin\theta \le \Delta\sqrt{8}.$$

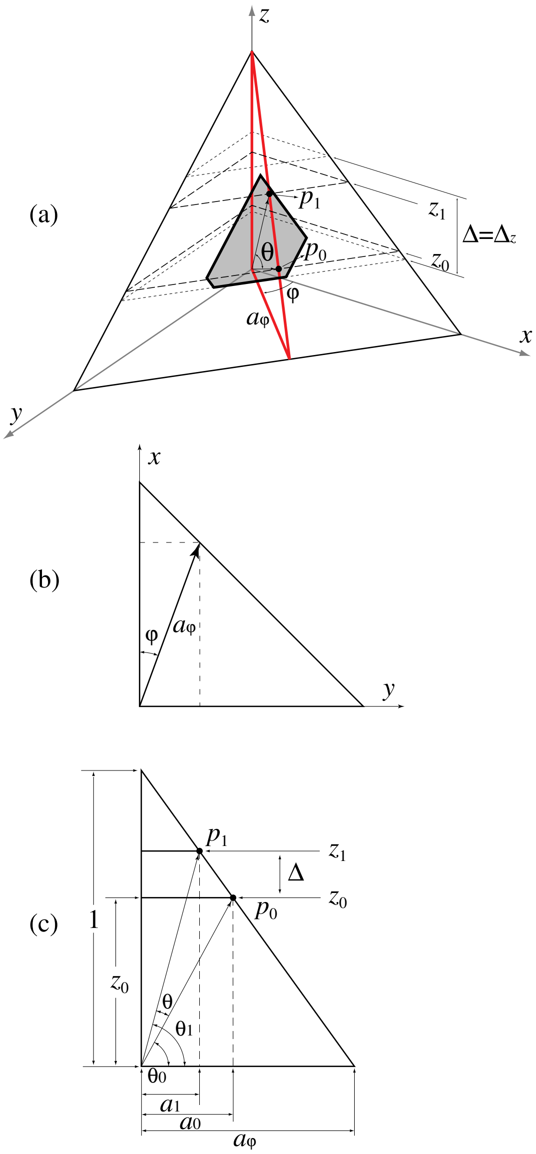

Figure 2: The lower bound.

Proof: We present the lower and upper bounds separately. Define $\Delta_x = |p_0.x - p_1.x|$, and define Δ_y and Δ_z analogously.

Lower Bound. Assume without loss of generality that $\Delta = \Delta_z$. Consider, the triangle formed by the intersection of T with the positive octant and the trapezoidal region on this triangle shown in Figure 2(a). For a fixed Δ, the smallest angle between any two vectors, occurs when points p_0, p_1 lie on a plane containing the z-axis. Let φ denote the angle between this plane and the x-axis, and let a_φ denote the length of the segment on this plane between the origin and T. (See Figure 2(b).) We see that $a_\varphi \cos\varphi + a_\varphi \sin\varphi = 1$, which gives $a_\varphi = 1/(\cos\varphi + \sin\varphi)$ and $1/\sqrt{2} \le a_\varphi \le 1$. Consider a_0, a_1, and a_φ from Figure 2(c). By similar triangles, we have the relations

$$\frac{1}{a_\varphi} = \frac{1-z_0}{a_0} = \frac{1-(\Delta+z_0)}{a_1}.$$

Thus $a_0 = a_\varphi(1-z_0)$ and $a_1 = a_\varphi(1-(\Delta+z_0))$. The angle θ between the two directions can be expressed as $\theta = \theta_1 - \theta_0$. The value of $\tan\theta = \tan(\theta_1 - \theta_0)$ is then

$$\tan\theta = \frac{a_\varphi \Delta}{a_\varphi^2(1-z_0)(1-(\Delta+z_0)) + z_0(\Delta+z_0)},$$

where $\tan\theta_0 = z_0/a_0 = z_0/a_\varphi(1-z_0)$ and $\tan\theta_1 = z_1/a_1 = (\Delta+z_0)/a_\varphi(1-(\Delta+z_0))$. The value of z_0 ranges from 0 to $1-\Delta$. To get the minimum value of $\tan\theta$, we find the maximum value of its denominator. The denominator of $\tan\theta$ is a quadratic in z_0 whose leading coefficient is $1 + a_\varphi^2 > 0$. Thus the maximum value of the denominator occurs at one of the extreme values $z_0 = 0$ or $z_0 = 1-\Delta$. We have $\tan\theta$ equal to the value $\Delta/a_\varphi(1-\Delta)$ for $z_0 = 0$ and $a_\varphi\Delta/(1-\Delta)$ for $z_0 = 1-\Delta$, where

$$\frac{a_\varphi \Delta}{1-\Delta} \le \frac{\Delta}{a_\varphi(1-\Delta)}$$

because Δ is fixed and $1/\sqrt{2} \le a_\varphi \le 1$. Thus, the lower bound of $\tan\theta$ is

$$\frac{\Delta}{\sqrt{2}(1-\Delta)} \le \frac{a_\varphi\Delta}{1-\Delta} \le \tan\theta.$$

Upper Bound. Let $l(p_0p_1)$ be the length of segment p_0p_1. In Figure 3(a), each one of ranges Δ_x, Δ_y, or Δ_z will result in a segment of length $\sqrt{2}\Delta_x$, $\sqrt{2}\Delta_y$ or $\sqrt{2}\Delta_z$ along the plane T. The maximum length of segment connecting any two points of the intersection of T and the positive octant happens when two or three of $\Delta_x, \Delta_y, \Delta_z$ are equal to Δ, and l is the longest diagonal of the hexagon illustrated in Figure 3(b). Each pair of parallel edges of the hexagon are separated by distance at most $\sqrt{2}\Delta$. This implies that the length of its edge is at most $\sqrt{2/3}\Delta$, and the length of the diagonal is at most $\sqrt{8/3}\Delta$. Thus, the bound of $l(p_0p_1)$ is $l(p_0p_1) \le \sqrt{8/3}\Delta$.

Let d be the distance from the origin to any point on the triangle defined by plane T. Then d is bounded as $1/\sqrt{3} \le d \le 1$, where $d = 1$ happens at the corners of this triangle, $d = 1/\sqrt{3}$ happens at the triangle's centroid $(1/3, 1/3, 1/3)$. The relation between angle θ of vectors $\vec{op}_0$, $\vec{op}_1$ and the length of segment p_0p_1, as depicted in Figure 3(c), is

$$\|\vec{op}_0\| \tan\theta \le \|\vec{op}_1 - \vec{op}_0\| \le \|\vec{op}_1\| \sin\theta.$$

Combining with inequalities on $l(p_0p_1)$ and d, we have the upper bound on of $\tan\theta$, $\|\vec{op}_0 - \vec{op}_1\| = l(p_0p_1)$,

$$\frac{\tan\theta}{\sqrt{3}} \le \|\vec{op}_0\| \tan\theta \le l(p_0p_1) \le \sqrt{\frac{8}{3}}\Delta.$$

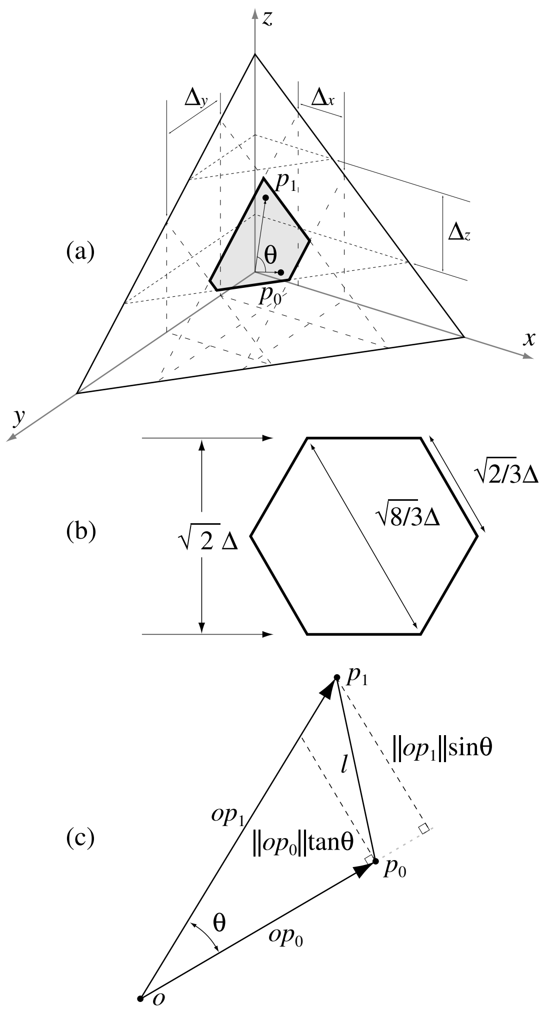

Figure 3: The upper bound.

That is, $\tan\theta \le 2\sqrt{2}\Delta$. The quantity $\sin\theta$ has the properties

$$\sin\theta \le \tan\theta \qquad \text{and} \qquad \sin\theta = \tan\theta/\sqrt{1+\tan^2\theta}.$$

It follows that the bounds on $\sin\theta$ are

$$\frac{\Delta}{3\sqrt{2}} < \frac{\Delta}{(1-\Delta)\sqrt{2(8\Delta^2+1)}} \le \sin\theta < 2\sqrt{2}\Delta,$$

since $(1-\Delta) < 1$ and $(8\Delta^2+1) < 9$. This completes the proof. □

For the next lemma, define $\mathrm{dist}_\infty(p,q)$ to be the L_∞ distance between two points, that is, the maximum absolute difference between corresponding coordinates of p and q. It is easy to see that

$$\frac{\mathrm{dist}(p,q)}{\sqrt{3}} \le \mathrm{dist}_\infty(p,q) \le \mathrm{dist}(p,q).$$

We can define the L_∞ distance between any two sets as the minimum L_∞ distance between any pair of points from each of the sets. The next lemma states that if two lines are distance d apart then there is a projection to a coordinate plane that attests to this separation.

Lemma 5.2 *Let ℓ_1 and ℓ_2 be two lines in 3-space, let $d = \mathrm{dist}(\ell_1,\ell_2)$ and let p be any point on ℓ_1. There exists a coordinate plane such if ℓ_1', ℓ_2' and p' denote the orthogonal projections of these elements onto this plane, then*

$$\mathrm{dist}_\infty(p',\ell_2') \ge \frac{d}{\sqrt{3}}.$$

Proof: Let $d' = d/\sqrt{3}$. Suppose to the contrary that in all three orthogonal coordinate projections, the distance from p' to ℓ_2' is less than d'. Then it follows that $\mathrm{dist}_\infty(p,\ell_2)$ (in 3-space) would be less than d'. This implies that $\mathrm{dist}(p,\ell_2)/\sqrt{3} < d'$, and hence

$$\mathrm{dist}(\ell_1,\ell_2) \le \mathrm{dist}(p,\ell_2) < d,$$

a contradiction. □

The third lemma provides a bound on the angle between two lines in the plane, given conditions on the vertical distance between the lines at some point.

Lemma 5.3 *Consider two directed lines ℓ_1 and ℓ_2 in the plane both with positive slopes, and one whose slope is at most 1. Suppose that they meet at some point q, and within horizontal distance at most h from q the lines are vertically separated by distance at least v. Let θ denote the angle between these lines. Then*

$$\sin\theta \ge \min\left(\frac{1}{\sqrt{3}}, \frac{v}{h\sqrt{10}}\right).$$

Proof: It is easy to see that the minimum value of θ is achieved when the horizontal distance is h, the vertical distance is v, and the line with the smaller slope has a slope of 1. The line with the larger slope has a slope of at least $(v+h)/h$. By basic trigonometry (the law for the tangent of the difference of two angles, in particular) it follows that

$$\tan\theta \ge \frac{(v+h)/h - 1}{1+(v+h)/h} = \frac{v}{2h+v}.$$

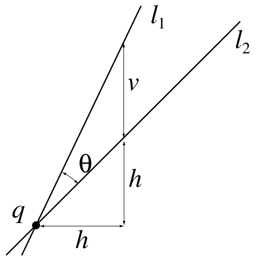

Figure 4: Lemma 5.3.

If $v \leq h$ then we have $\tan\theta \geq v/(3h)$ implying that $\sin\theta \geq v/(h\sqrt{10})$. On the other hand, if $v > h$, then we have $\tan\theta \geq h/2h = 1/2$, from which we have $\sin\theta \geq 1/\sqrt{3}$. □

The analysis of the size approximation is based on proving the following main lemma. Intuitively it states that if there are two lines in $S(P)$ that are separated by a large normalized distance, then there are two lines in $S(P)$ that are separated by a proportionately large normalized angle. Thus, in order to approximate the size of P, it suffices to approximate the range of angles. Recall that we assume that P has been so constrained that the lines of $S(P)$ are directed into the positive octant, and that all the patches of P lie within a cube C.

Lemma 5.4 (Main Lemma) *Let $\ell_1, \ell_2 \in S(P)$. Then there exist lines $\ell_3, \ell_4 \in S(P)$ such that*

$$\mathrm{ang}'(\ell_3, \ell_4) \geq \frac{\mathrm{dist}'(\ell_1, \ell_2)}{\sqrt{10}}.$$

Assuming this result for now, we can prove Theorem 5.1 as follows. Consider the value Δ that is returned by the approximation. Thinking of the directional component of the Plücker coordinates, $(\pi_{01}, \pi_{02}, \pi_{03})$, as (x, y, z) coordinates, recalling the normalization $\pi_{01} + \pi_{02} + \pi_{03} = 1$, and the fact that we consider lines directed into the positive octant (implying that $\pi_{0i} \geq 0$), we may apply Lemma 5.1 to infer that there exist two lines $\ell_3, \ell_4 \in S(P)$ such that

$$\frac{\Delta}{3\sqrt{2}} \leq \mathrm{ang}'(\ell_3, \ell_4) \leq \Delta 2\sqrt{2}.$$

By its definition, size(P) is at least as large as $\mathrm{ang}'(\ell_3, \ell_4)$, implying that

$$\mathrm{size}(P) \geq \frac{\Delta}{3\sqrt{2}}.$$

If size(P) is determined by the normalized angle between two lines $\ell_3, \ell_4 \in S(P)$, then we have

$$\mathrm{size}(P) = \mathrm{ang}'(\ell_3, \ell_4) \leq \Delta 2\sqrt{2} \leq \Delta 4\sqrt{5}.$$

If, on the other hand, size(P) is determined by the normalized distance between two lines $\ell_1, \ell_2 \in S(P)$ then from the Main Lemma there exist two lines $\ell_3, \ell_4 \in S(P)$ such that $\mathrm{ang}'(\ell_3, \ell_4) \geq \mathrm{dist}(\ell_1, \ell_2)/\sqrt{10}$. From this and the inequality above it follows that

$$\mathrm{size}(P) = \mathrm{dist}'(\ell_1, \ell_2) \leq \mathrm{ang}'(\ell_3, \ell_4)\sqrt{10} \leq \Delta 4\sqrt{5},$$

which establishes Theorem 5.1.

6 Proof of the Main Lemma

Consider lines $\ell_1, \ell_2 \in S(P)$. We assume that these lines are in general position. Let $d = \mathrm{dist}(\ell_1, \ell_2)$. Let C_{lo} and C_{hi} denote the low and high endpoints of the enclosing cube C. Thus, for example, all points in C have x-coordinates satisfying $C_{lo}.x \leq x \leq C_{hi}.x$. Given any object p in 3-space, let p^x denote its orthogonal projection onto the zy-coordinate plane. Define p^y and p^z analogously.

Consider any point on ℓ_1 within the cube C. Apply Lemma 5.2. Without loss of generality, we may assume that the axes have been labeled so that the projection described in the lemma is on the xy-coordinate plane, and that the projected line with the smaller slope has slope at most 1 (by swapping the x and y axes if necessary). Let $d' = d/\sqrt{3}$. From the lemma it follows that at some point within the cube the L_∞ distance between this point to ℓ_2 is at least d', implying that the vertical distance between this point and ℓ_2 is at least d'.

Consider the vertical strip on the xy-coordinate plane defined by the projections of the sides of the cube of C,

$$C_{lo}.x \leq x \leq C_{hi}.x.$$

Consider ℓ_1^z and ℓ_2^z, the respective projections of ℓ_1 and ℓ_2 onto the xy-plane. We define two lines ℓ_3^z and ℓ_4^z on the xy-plane. The lines ℓ_3 and ℓ_4 will be constructed so that these are their projections onto the xy-plane. We consider three cases, as illustrated in Fig. 5.

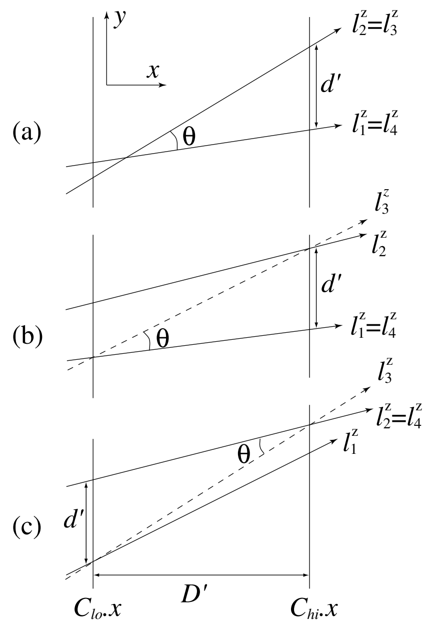

Figure 5: Projection onto the xy-plane.

Case (a) If ℓ_1^z and ℓ_2^z intersect within the strip, then we let $\ell_3^z = \ell_2^z$ and $\ell_4^z = \ell_1^z$.

Otherwise one of the lines lies above the other (has a larger y-coordinate as shown in Fig. 6) throughout the strip. We may assume without loss of generality that ℓ_2 lies above ℓ_1. Let ℓ_3^z be the line connecting the intersection of ℓ_1^z with the left side of the strip to the intersection of ℓ_2^z with the right side of the strip.

Case (b) If ℓ_1 has the smaller slope, then let ℓ_4^z be ℓ_1^z.

Case (c) If ℓ_1 has the larger slope, then let ℓ_4^z be ℓ_2^z.

Let θ denote the angle between ℓ_3^z and ℓ_4^z. The diameter of C is D, implying that the width of the strip is $D' = D/\sqrt{3}$. Observe that within horizontal distance at most D' of the intersection of ℓ_3^z and ℓ_4^z, the vertical distance between the lines is at least d', and that ℓ_4^z has slope at most 1. By applying Lemma 5.3 (where ℓ_4^z and ℓ_3^z are the ℓ_1 and ℓ_2 in Lemma 5.3, respectively) it follows that

$$\sin\theta \;\geq\; \min\left(\frac{1}{\sqrt{3}}, \frac{d'}{D'\sqrt{10}}\right) \;=\; \min\left(\frac{1}{\sqrt{3}}, \frac{d}{D\sqrt{10}}\right).$$

Since $d \leq D$, the second term is always smaller, implying that $\sin\theta \geq d/(D\sqrt{10})$.

Observe that if we can find two lines $\ell_3, \ell_4 \in S(P)$ whose projections are ℓ_3^z and ℓ_4^z, then the angle between these two lines in 3-space will be at least as large as θ, and so establish the desired bound of Lemma 5.4 and completing the proof. The remainder of the proof is concerned with finding these lines.

First we observe that in Case (a) we may select $\ell_3 = \ell_2$ and $\ell_4 = \ell_1$ to complete the proof. Thus it suffices to consider cases (b) and (c). Since ℓ_4^z is equal to either ℓ_1^z or ℓ_2^z, if we choose ℓ_4 to be equal to ℓ_1 or ℓ_2, respectively, then $\ell_4 \in S(P)$. Thus it suffices to find ℓ_3.

Let H denote the plane formed by extruding the line ℓ_3^z parallel to the z-direction. Also consider the extrusion of the strip as well, to the region bounded between two planes that are orthogonal to the x-axis. By choosing any line (in general position) on H, it will follow that its xy-projection is ℓ_3^z. We may think of each patch in P as consisting of four directed, axis-parallel *constraint lines*. Let ℓ_r denote such a line that is parallel to the z-axis.

To establish whether any given line ℓ in 3-space satisfies this constraint, it suffices to consider whether the projection ℓ^z of the line ℓ onto the xy-plane lies above or below (depending on the direction of ℓ_r) the point r onto which ℓ_r projects. Since all patches lie within C, whose projection lies within the strip, r lies within in the strip. By our choice of ℓ_3^z in either of cases (b) or (c), if both ℓ_1^z and ℓ_2^z lie below (above) point r, then ℓ_3^z lies below (above) r as well. Thus by choosing ℓ_3 to lie on H, it follows that it will satisfy any z-axis parallel constraints that both ℓ_1 and ℓ_2 satisfy. Also observe that we have chosen ℓ_3^z so that it is directed into the positive quadrant, and so it satisfies the sign-class constraint (at least with respect to x and y).

We will apply this same analysis to the other two orthogonal projections. First consider y-axis parallel constraint lines. For $i = 1, 2$, see Fig. 6, let a_i and b_i denote the intersections of line ℓ_i with the lower (low x) and upper (high x) sides of the strip, respectively. Let h_a and h_b denote the lines along which H intersects the lower and upper sides of the strip.

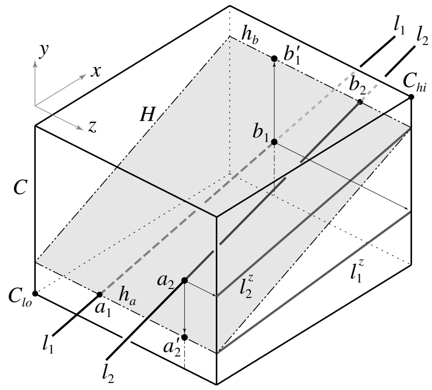

Figure 6: The enclosing cube.

By our assumption that ℓ_1^z lies below ℓ_2^z within the strip, and our assumption on the sign class of the lines it follows that a_1 lies on h_a and b_2 lies on h_b and

$$a_1^x.y \leq \{a_2^x.y, b_1^x.y\} \leq b_2^x.y.$$

(where $a \leq \{b, c\}$ means that $a \leq b$ and $a \leq c$). We also have, for $i = 1, 2$, $a_i^x.z \leq b_i^x.z$, by the sign assumption. Let a_2' and b_1' denote the y-parallel projections of a_2 and b_2 onto h_a and h_b respectively.

Consider the projection of the strip onto the zx-coordinate plane. (See Fig. 7.) The segments a_1a_2 and b_1b_2 project onto segments $a_1^ya_2^y$ and $b_1^yb_2^y$. Letting ℓ_3^y denote the projection of ℓ_3 onto the zx-plane, to guarantee that ℓ_3 satisfies all y-parallel constraints that ℓ_1 and ℓ_2 do, we should select ℓ_3 so that ℓ_3^y intersects both of these segments. Furthermore, if ℓ_1^y and ℓ_2^y intersect at some point p^y within the projected strip, then ℓ_3^y should also intersect this point (for example, as the dashed line in the figure does).

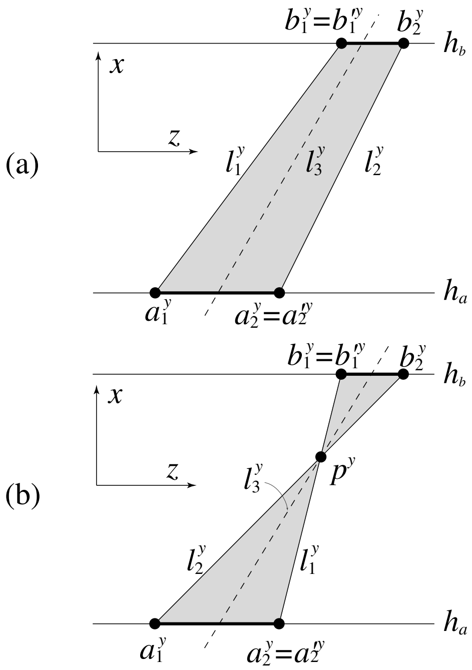

Figure 7: Projection onto the zx-plane.

We consider two cases illustrated in Fig. 7.

Case Fig. 7(a) In the first case, ℓ_1^y and ℓ_2^y do not intersect within the strip. Consider the projections onto the zy-plane (see Fig. 8). Our requirement that ℓ_3^y intersects the segments $a_1^ya_2^y$ and $b_1^yb_2^y$, implies that in this projection, the line ℓ_3^x intersects the segments $a_1^xa_2'^x$ and $b_1'^xb_2^x$ (shown in heavy lines in the figure).

If $a_1^x.z > a_2^x.z$ (see Figs. 8(a)) then because the y-projections of the lines do not intersect within the strip, we have $b_1^x.z > b_2^x.z$. Let ℓ_3 be the line passing through a_2' and b_1'. It is easy to see that ℓ_3 satisfies the sign-class constraints as well as all x-parallel constraints because it lies entirely to the right of the line ℓ_2^x and to the left of the line ℓ_1^x while in the strip.

On the other hand, if $a_1^x.z \leq a_2^x.z$ (see Figs. 8(b)-(c)) then we have $b_1^x.z \leq b_2^x.z$. There are two subcases to consider.

(I) If line ℓ_1^x intersects the segment $b_1'^xb_2$ or line ℓ_2^x intersects segment $a_1^xa_2'^x$ (the former occurs in Fig. 8(b)) then let ℓ_3 be the line, either ℓ_1 or ℓ_2, whose projection satisfies this condition. (We would take $\ell_3 = \ell_1$ in the figure.) Since ℓ_3 is equal to either ℓ_1 or ℓ_2 it is in $S(P)$.

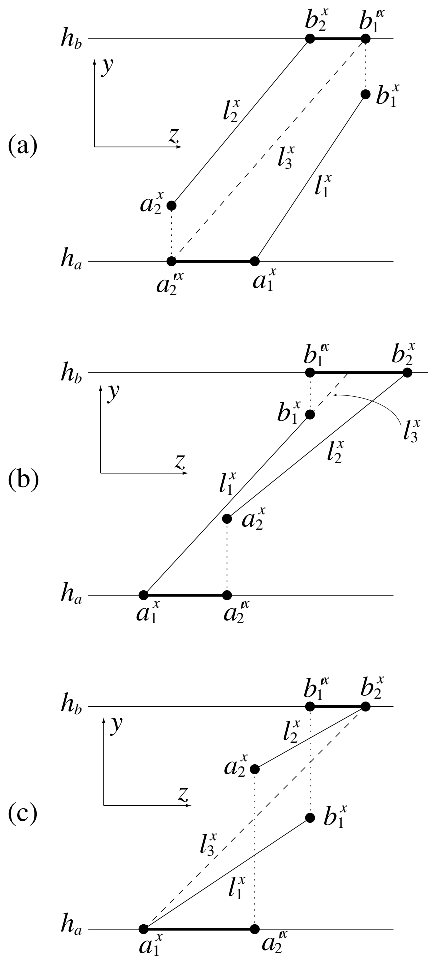

Figure 8: Projection onto the zy-plane.

(II) If both lines ℓ_1^x and ℓ_2^x fail to intersect segments $a_1^x a_2'^x$ and $b_1'^x b_2$ (see Fig. 8(c)), then it follows that ℓ_2^x intersects h_a to the left of $a_1^x a_2'^x$ and ℓ_1^x intersects h_b to the right of $b_1'^x b_2^x$. Let ℓ^3 be the line extending through $a_1 b_2$. This line satisfies all x-parallel constraints because it lies in the portion of the strip the right of ℓ_2^x and to the left of ℓ_1^x.

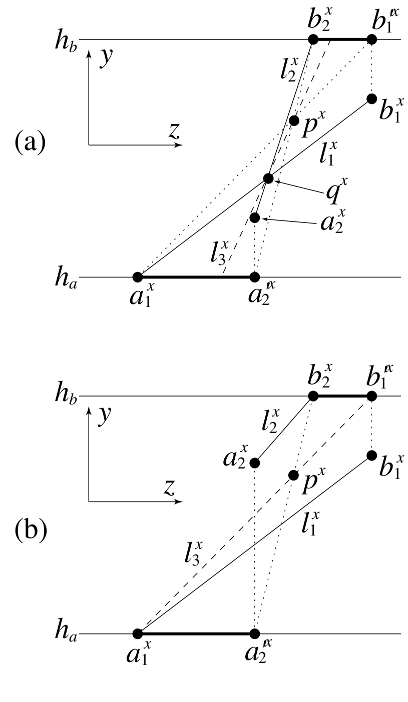

Figure 9: Projection onto the zy-plane.

Case Fig. 7(b) In the second case, ℓ_1^y and ℓ_2^y intersect at some point p^y within the strip. As before, consider the projections onto the zy-plane (see Fig. 9). As before, there are two subcases. Either $a_1^x.z \leq a_2^x.z$, implying that $b_1^x.z > b_2^x.z$, or $a_1^x.z > a_2^x.z$, implying that $b_1^x.z \leq b_2^x.z$. These two cases are symmetric with respect to a 180 degree rotation and a reversal of a's and b's, 1's and 2's, so it suffices to consider just the former case.

The projected intersection point p^y defines a y-parallel line (passing through ℓ_1 and ℓ_2) and intersecting H at some point p. The zy-projection of p, denoted p^x, is the intersection of the line segments $a_1^x b_1'^x$ and $a_2'^x b_2^x$. (This is because these are zy-projections of unique lines on H whose zx-projections are $a_1^y b_1^y$ and $a_2^y b_2^y$, respectively.) Let q^x denote the intersection of lines ℓ_1^x and ℓ_2^x.

(I) If line ℓ_2^x intersects the segment $a_1^x a_2'^x$ (see Fig. 9(a)), then we claim that the line extending the segment $q^x p^x$ intersects both the segments $a_1^x a_2'^x$ and $b_1'^x b_2^x$. This is because q^x lies within the double wedge whose apex is p and whose extreme lines pass through the endpoints of these segments. In this case let ℓ_3 be the unique line on H projecting onto the segment $q^x p^x$. Because it lies on H and intersects segments $a_1^x a_2'^x$ and $b_1'^x b_2^x$, it satisfies the z- and y-parallel constraints of P. Because it passes through q^x it satisfies the x-parallel constraints of P. It satisfies the sign-class constraints because the slope of the line lies between the (positive) slopes of $a_1^x b_1'^x$ and $a_2'^x b_2^x$.

(II) On the other hand, if line ℓ_2^x does not intersect the segment $a_1^x a_2'^x$ (see Fig. 9(b)), then it follows from slope considerations that it intersects h_a to the left of this segment. Let ℓ_3 be the unique line on H that projects onto the segment $a_1^x b_1'^x$. (Observe that q^x lies outside of the strip in this case.) This line passes through p^x, and intersects segments $a_1^x a_2'^x$ and $b_1'^x b_2^x$, and hence it satisfies the z- and y-parallel constraints of P. It satisfies the x-parallel constraints because it lies in the portion of the strip to the left of ℓ_1^x and the right of ℓ_2^x. Finally, it satisfies the sign-class constraints because $a_1^x.z \leq b_1'^x.z$.

This completes the case analysis for the construction of ℓ_3. Because we showed in all cases that ℓ_3 satisfies all three orthogonal constraints as well as the sign-class constraints, it is in $S(P)$, and this completes the proof of the Main Lemma and, hence, the analysis of the approximation algorithm.

References

[1] N. Amenta. Finding a line transversal of axial objects in three dimensions. In *Proc. 3rd ACM-SIAM Sympos. Discrete Algorithms*, pages 66–71, 1992.

[2] N. Amenta. Helly-type theorems and generalized linear programming. *Discrete Comput. Geom.*, 12:241–261, 1994.

[3] J. Matoušek, M. Sharir, and E. Welzl. A subexponential bound for linear programming. In *Proc. 8th Annu. ACM Sympos. Comput. Geom.*, pages 1–8, 1992.

[4] M. McKenna and J. O'Rourke. Arrangements of lines in 3-space: a data structure with applications. In *Proc. 4th Annu. ACM Sympos. Comput. Geom.*, pages 371–380, 1988.

[5] N. Megiddo. Finding a ling of sight thru boxes in d-space in linear time. Unpublished manuscript, 1996.

[6] M. Pellegrini and P. Shor. Finding stabbing lines in 3-space. *Discrete Comput. Geom.*, 8:191–208, 1992.

[7] H. Samet. *The Design and Analysis of Spatial Data Structures.* Addison-Wesley, Reading, MA, 1990.

[8] F. Sillion and C. Puech. *Radiosity Global Illumination.* Morgan Kauffmann, 1994.

[9] J. Stolfi. *Oriented Projective Geometry: A Framework for Geometric Computations.* Academic Press, New York, NY, 1991.

[10] S. J. Teller. Computing the antipenumbra of an area light source. *Comput. Graph.*, 26(4):139–148, July 1992. Proc. SIGGRAPH '92.