Aperture-Angle Optimization Problems in 3 Dimensions*

Elsa Omaña-Pulido
Departamento de Matemáticas
Universidad Autónoma de México
Iztapalapa, México

Godfried T. Toussaint
School of Computer Science
McGill University
3480 University Street
Montreal, Quebec, Canada H3A 2A7

Abstract

Let $[a, b]$ be a line segment with end points a, b and ν a point at which a viewer is located, all in $\mathbf{R}^3$. The aperture angle of $[a, b]$ from point ν, denoted by $\theta(\nu)$, is the interior angle at ν of the triangle $\Delta(a, b, \nu)$. Given a convex polyhedron P not intersecting a given segment $[a, b]$ we consider the problem of computing $\theta_{max}(\nu)$ and $\theta_{min}(\nu)$, the maximum and minimum values of $\theta(\nu)$ as ν varies over all points in P. We obtain two characterizations of $\theta_{max}(\nu)$. Along the way we solve several interesting special cases of the above problems and establish linear upper and lower bounds on their complexity under several models of computation.

1 Introduction

Visibility plays an important role in the manufacturing industry in such problems as accessibility analysis in machining [28], [23], [6] and visual inspection [22] as well as computer graphics, robotics, computer vision, operations research and several other disciplines of computing science and computer engineering [17], [21]. The traditional model of visibility investigated in computational geometry allows for a guard or camera to "see in all directions,"; i.e., the aperture angle is idealized to be 360 degrees. More recently, computational geometry research has begun investigating more realistic models of visibility where the aperture angle (or field-of-view angle as it is called in robotics [7], [8]) is restricted to be some angle θ less

*Research supported by NSERC Grant no. OGP0009293 and FCAR Grant no. 93-ER-0291.

than 360 degrees. For example, given a convex polygon in the plane and a camera with aperture angle θ situated outside the polygon, Teichman [24] computes a description of all the points in the plane where a camera may be placed in such a way that the polygon lies completely in the field of vision of a camera with aperture angle θ. A member x of a set of points S is said to be *θ-visible* if a camera with aperture angle θ can be placed on x in such a way that no other member of S lies in the camera's field of vision. Avis, et al. [1] obtained optimal algorithms for finding all the θ-visible points in a set S. Devroye and Toussaint [10] investigate the cardinality of the θ-visible points among a set of special points which are the intersections of a set of random lines. Finally, in another variant of the problem Bose, et al., [2] have shown that n cameras, each with its own fixed specified aperture angle not exceeding 180 degrees, can be placed at n fixed locations in the plane to see the entire plane if and only if the aperture angles sum to at least 360 degrees. All these results are restricted to the plane.

The simplest problem of this kind is often found as a exercise in calculus texts and is called the "picture-on-the-wall" problem (see for example [20], p. 427, problem # 20). In this problem a picture hangs on the wall in a museum above the level of an observer's eye. How far from the wall should the observer stand to maximize the angle subtended at the observer's eye by the top and bottom of the picture? While this problem is easily solved with calculus, an elegant solution that does not use calculus has been known for some time [16]. This same solution holds for the more general problem where the picture may not be orthogonal to the floor [27].

Bose, Hurtado-Diaz, Omaña-Pulido and Toussaint [3], [4] considered a generalization of the "picture-on-the-wall" problem, namely, the problem of computing the aperture angle of a camera that is allowed to travel in a convex region in the plane and is required to maintain some other convex region within its field of view at all times. More specifically, let P and Q be two disjoint convex polygons in the plane with n and m vertices, respectively. Given a point x in P, the aperture angle of x with respect to Q is defined as the angle subtended by the cone that: (1) contains Q, (2) has apex at x, and (3) has its two rays emanating from x tangent to Q. They gave an $O(n+m)$ time algorithm for computing the minimum aperture angle with respect to Q when x is allowed to vary in P. They also presented algorithms with complexities $O(n+m)$ and $O(n \log m)$ for computing the maximum aperture angle with respect to Q. Thus, when $m = o(n \log n)$ the first algorithm is faster than the second one. However, if $m = \Omega(n \log^{1+\epsilon} n)$, for any $\epsilon > 0$ the second one is faster. Finally, they proved an $\Omega(n)$ time lower bound for the maximization problem and an $\Omega(n+m)$ time bound for the minimization problem establishing the optimality of the algorithm for the latter problem.

In this paper we consider aperture angle optimization problems in 3 dimensions. Such problems have a long history dating at least as far back as the mathematician Regiomontanus (Johannes Muller) [11]. In 1471 Regiomontanus posed the following problem (now known as Regiomontanus' Maximum Problem). At what point of the earth's surface does a perpendicularly suspended rod appear longest? In other words, at what point is the aperture angle at a maximum? A solution of this problem was published by Ad. Lorsch in vol. XXIII of *Zeitschrift fur Mathematik und Physik*, assuming the earth is flat.

Here we consider generalizations of Regiomontanus' maximum problem and the problems considered in [3] and [4] to three dimensions. In particular, let P be a convex polyhedron in $\mathbf{R}^3$. Let $[a,b]$ be a line segment with end points a, b and ν a point at which a viewer is located, also in $\mathbf{R}^3$. The aperture angle of $[a,b]$ from point ν, denoted by $\theta(\nu)$, is the interior angle at ν of the triangle $\Delta(a,b,\nu)$. Given a polyhedron P not intersecting a given segment $[a,b]$ we consider the problem of computing $\theta_{max}(\nu)$ and $\theta_{min}(\nu)$, the maximum and minimum values of $\theta(\nu)$ as ν varies over all points in P. We obtain two characterizations of $\theta_{max}(\nu)$. Along the way we solve several interesting special cases of the above problems and establish linear upper and lower bounds on their complexity under several models of computation.

2 Viewing a segment from a plane

Let $[a,b]$ be a line segment and ν a point at which a viewer is located in $\mathbf{R}^3$. Let the length of $[a,b]$ be l and denote the distance between ν and the end-points of $[a,b]$ by l_a and l_b, respectively. Then, by the law of cosines the aperture angle is given by:

$$\theta = \arccos\left(\frac{{l_a}^2 + {l_b}^2 - l^2}{2 l_a l_b}\right)$$

In this planar-constraint version of the problem we are interested in determining the maximum aperture angle $\theta_{max}(\nu)$ where ν varies over the entire xy-plane. This version of the problem is not interesting for $\theta_{min}(\nu)$. Trivially, $\theta_{min}(\nu)$ is equal to zero and occurs at the intersection of the xy-plane with the line through $[a,b]$ as well as points at infinity. Therefore we restrict our attention to the computation of $\theta_{max}(\nu)$. Furthermore, if $[a,b]$ intersects the xy-plane then $\theta_{max}(\nu)$ is equal to π and occurs at such an intersection point. Therefore, without loss of generality, we assume the segment $[a,b]$ lies above the xy-plane. Let the end points of segment $[a,b]$ have coordinates $a = (a_1, a_2, a_3)$ and $b = (b_1, b_2, b_3)$ and let the point ν on the xy-plane have coordinates $\nu = (\nu_1, \nu_2)$. Then $\theta(a,b,\nu)$ is given by the above equation where:

$$l_a = \sqrt{(a_1 - \nu_1)^2 + (a_2 - \nu_2)^2 + (a_3)^2}$$

$$l_b = \sqrt{(b_1 - \nu_1)^2 + (b_2 - \nu_2)^2 + (b_3)^2}$$

and

$$l = \sqrt{(a_1 - b_1)^2 + (a_2 - b_2)^2 + (a_3 - b_3)^2}$$

For a fixed segment $[a,b]$ in general position (neither vertical nor horizontal) and ν restricted to the xy-plane Fig. 1 illustrates the aperture angle function $\theta(\nu(x,y))$ as a function of the x and y coordinates of the viewer. In general the function looks like a "volcano" with a "crater" that reduces down to a point where the line through $[a,b]$ intersects the xy-plane. Furthermore, the top of the crater is not level but tilted in the direction of the line segment. Since this function contains a singularity it is difficult to obtain approximate

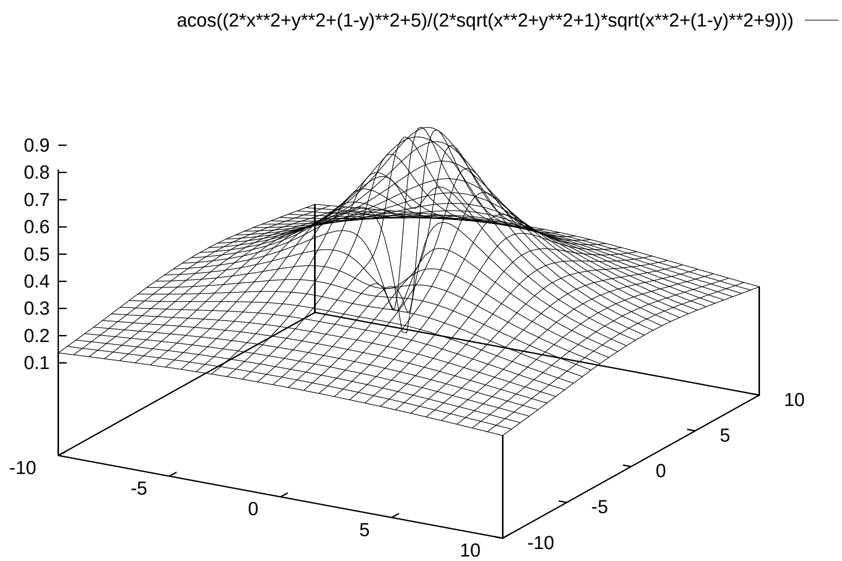

Figure 1: Illustrating the aperture angle function as a function of x and y for a non-vertical, non-horizaontal line segment.

solutions with constrained numerical optimization methods since these techniques require the function to be twice differentiable [13], [18]. Fortunately as in the two dimensional version of this problem [3], [4] we can obtain (without calculus) a simple characterization of the solution that allows us to compute it efficiently without having to resort to iterative numerical methods.

Consider first the special case when $[a, b]$ is vertical, i.e., parallel to the z-axis. Construct any plane through $[a, b]$, say H. This plane is orthogonal to the xy-plane and intersects it at a line L_H. On this line L_H we have that $\theta_{max}(\nu)$ is realized at two points ν_1* and ν_2* determined by the points of tangency where the two circles passing through a and b are tangent to L_H [3], [4]. Let c denote the vertical projection point of the line through $[a, b]$ on the xy-plane. Let $r*$ denote the distance from c to ν_1*, i.e., $r*= d(c, \nu_1*) = d(c, \nu_2*)$, where d denotes Euclidean distance. By symmetry this observation holds for all planes H. Therefore we have the following proposition.

Proposition 2.1 *Assume $[a, b]$ is vertical and that its projection on the xy-plane is the point c. Then $\theta_{max}(\nu)$ has an infinite number of solutions which are characterized by the circle on the xy-plane centered at c with radius equal to $r*$.*

In order to simplify visualization and discussion as well as proofs, it is instructive to examine the locus of points in space that have the same fixed aperture angle θ. With this in mind let us turn first to the simpler situation in $\mathbf{R}^2$. Subsequently we will generalize the concepts to $\mathbf{R}^3$. Consider then a line segment $[a, b]$ in $\mathbf{R}^2$ and refer to Fig. 2. Construct

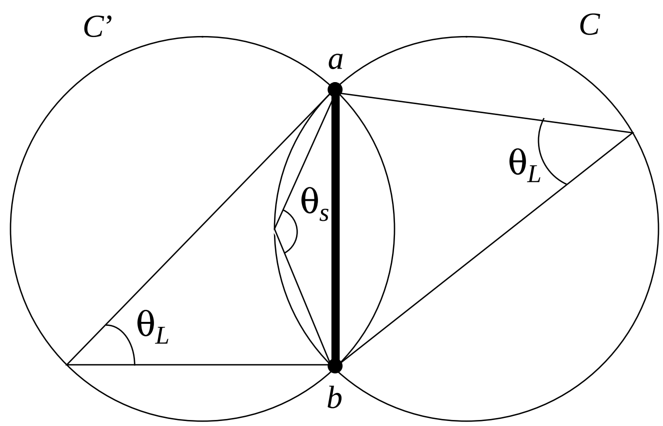

Figure 2: Illustrating the locus of constant aperture angle for a segment in the plane.

any circle C through a and b with some fixed diameter, $diam(C)$, greater than $d(a,b)$. Then $[a,b]$ is a chord of C and it partitions C into two arcs. Denote the large and small arcs by C_L and C_s, respectively. All points ν on C_s have the same aperture angle θ_s. Similarly, all points ν on C_L have the same aperture angle θ_L. Now construct the reflection of C about a line through $[a,b]$ and denote the resulting circle by C'. It follows that $[a,b]$ splits C' into a large arc C'_L and a small arc C'_s. Furthermore, by symmetry, all points ν on $C_L \cup C'_L$ have the same aperture angle θ_L and all points ν on $C_s \cup C'_s$ have the same aperture angle θ_s. Let D and D' denote the discs determined by C and C', respectively. Also we define $D \cap D'$ when $\theta > \pi/2$, $D \cup D'$ when $\theta < \pi/2$ and D when $\theta = \pi/2$, as the lune of segment $[a,b]$. Furthermore, if we need to distinguish between lunes that have $\theta > \pi/2$ and $\theta < \pi/2$ then we shall use the term acute lune when $\theta > \pi/2$ and obtuse lune when $\theta < \pi/2$. We then have the following propositions.

Proposition 2.2 *When $diam(C) \neq d(a,b)$ the locus of points in $\mathbf{R}^2$ of constant aperture angle greater than $\pi/2$ are determined by the boundary of $D \cap D'$. The locus of points of constant aperture angle less than $\pi/2$ are determined by the boundary of $D \cup D'$. When $diam(C) = d(a,b)$ this locus is C itself.*

Proposition 2.3 *The locus of points in $\mathbf{R}^2$ of constant aperture angle θ consists of an open line segment (a,b) when $\theta = \pi$, the boundary of a continuously growing acute lune as θ is decreased from π to $\pi/2$, the boundary of a continuously growing obtuse lune as θ is decreased from $\pi/2$ towards zero, and the union of points at infinity and the line through $[a,b]$ less the closed segment $[a,b]$ when θ equals zero. Furthermore, whenever $\theta_1 < \theta_2$ then the lune for θ_1 contains the lune for θ_2.*

We will call such loci for different fixed values of θ the *iso-aperture-angle contours* (IAA contours). Fig. 3 illustrates a line segment and some of its iso-aperture-angle contours. Such

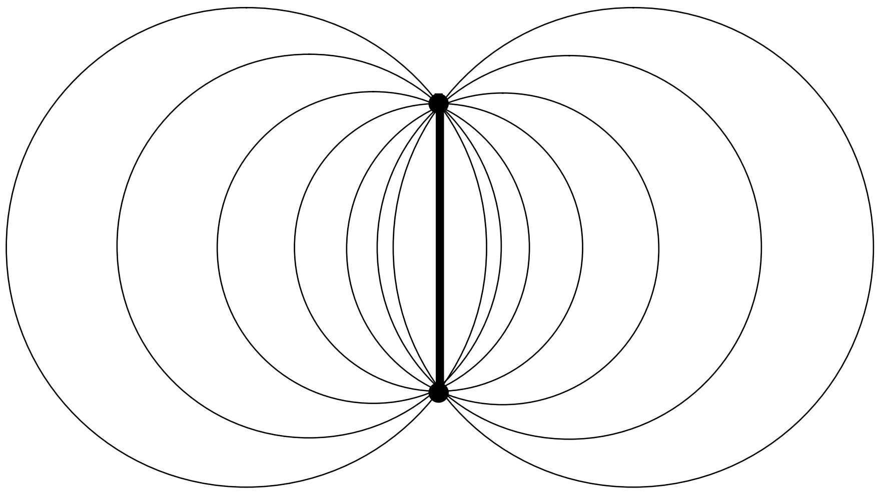

Figure 3: Illustrating the iso-aperture angle contours for a line segment in the plane.

a figure will be referred to as an IAA-diagram. The IAA-diagram visually characterizes the aperture angle function constrained to lie along a specified line. For example, consider a line parallel to $[a, b]$ such as L_1 in Fig. 4. The maximum aperture angle $\theta_{max}(L_1)$ determined by $[a, b]$ for a viewer on L_1 is uniquely determined by point x_1, i.e., the point on L_1 corresponding to the orthogonal projection of the midpoint of $[a, b]$ onto L_1. Alternately, we can view it as the first point of contact of the growing lune as it grows from a line segment to an infinite lune. Furthermore, the IAA-diagram and Proposition 2.3 imply that $\theta(\nu)$ is a *unimodal* function as ν travels along L_1. Consider a line not intersecting $[a, b]$ and orthogonal to it, such as L_2 in Fig. 4. The maximum aperture angle $\theta_{max}(L_2)$ is determined by two points x_2 and x_3, those points where (simultaneously) the obtuse lune makes contact with L_2 as it grows. Furthermore, the IAA-diagram and Proposition 2.3 imply that $\theta(\nu)$ is a bimodal function as ν travels along L_2. Finally, consider a line not intersecting $[a, b]$ and neither orthogonal nor parallel to it such as L_3 in Fig. 4. Then the maximum aperture angle $\theta_{max}(L_3)$ is determined by a unique point x_4, again, the first point of contact of the growing lune with L_3. A second (smaller) local maximum will always occur at x_5, the point where another portion of the boundary of the lune at x_4 later (in the growing process) becomes tangent to L_3 as it grows into a sufficiently large obtuse lune. The IAA-diagram not only effectively expresses the bimodality of $\theta(\nu)$ as ν travels along L_3 but it also shows that the larger local maximum occurs on the portion of L_3 that makes an angle, at the intersection point of the extension of $[a, b]$ with L_3, which is less than $\pi/2$.

We now generalize these ideas to three dimensions. Consider then a line segment $[a, b]$ in $\mathbf{R}^3$. For any plane H containing $[a, b]$ the IAA-contours on H must look like those illustrated in Fig. 3. Therefore the loci of constant aperture angle are IAA-surfaces in $\mathbf{R}^3$

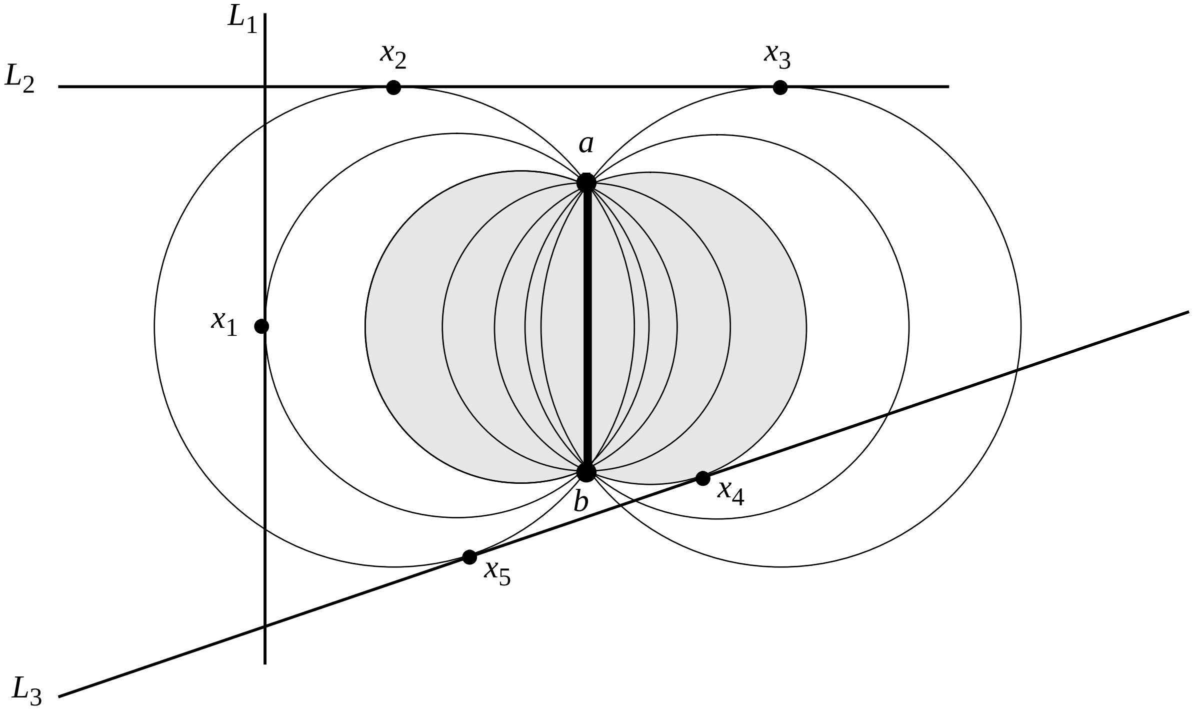

Figure 4: Illustrating how the iso-aperture-angle diagram of a line segment characterizes the aperture-angle function of a view-point travelling along a line.

and are obtained by revolving the IAA-contours, on a plane H containing $[a,b]$, about an axis which is the line containing $[a,b]$. We will refer to the solid of revolution obtained by revolving a lune about the line containing $[a,b]$ as a *lunoid*. Furthermore, if the lune is acute the resulting solid will be called a *pointed lunoid* and if the lune is obtuse the resulting solid will be called a *toroidal lunoid*. We therefore have the following proposition.

Proposition 2.4 *The locus of points in $\mathbf{R}^3$ of constant aperture angle θ consists of an open line segment (a,b) when $\theta=\pi$, the surface of a continuously growing pointed lunoid as θ is decreased from π to $\pi/2$, the surface of a continuously growing toroidal lunoid as θ is decreased from $\pi/2$ towards zero, and the union of points at infinity and the line through $[a,b]$ less the closed segment $[a,b]$ when θ equals zero. Furthermore, whenever $\theta_1<\theta_2$ then the lunoid for θ_1 contains the lunoid for θ_2.*

Proposition 2.4 and the IAA-diagram afford simple proofs of several properties of the aperture angle function in $\mathbf{R}^3$. For example, if $[a,b]$ is orthogonal to and not intersecting a plane H in $\mathbf{R}^3$, as the lunoid grows the first contact with H occurs simultaneously at a circle whose center is the orthogonal projection of $[a,b]$ onto H. Therefore we have a crisp and more direct visual verification of Proposition 2.1. The above observations imply the following lemmas the proofs of which are straightforward and therefore omitted.

Lemma 2.5 *Let $[a, b]$ be parallel to the xy-plane (and above it) and ν a viewing point on the xy-plane. Then $\theta_{max}(\nu)$ occurs at a unique point ν^* which is the orthogonal projection of the mid-point of $[a, b]$ onto the xy-plane.*

Lemma 2.6 *Let $[a, b]$ be non-vertical (not parallel to the z-axis) and not horizontal (not parallel to the xy-plane). Let c denote the point of intersection of the xy-plane with the line containing $[a, b]$. Then $\theta_{max}(\nu)$ occurs at the unique point ν^* which is the point of tangency of the circle (determined by a, b, and ν^*) with the xy-plane such that (1) the circle also lies in the vertical plane containing $[a, b]$ and (2) the angle at c of triangle $\Delta(\nu^*, c, a)$ is less than $\pi/2$.*

In the following lemma as in the rest of the paper we assume in our computational complexity discussion (unless otherwise specified) that our model is the extended standard real RAM (Random Access Machine) [19]. In this model we allow the computation of k'th roots and trigonometric functions in constant time. From lemma 2.6 it follows that we may convert this three-dimensional problem to a two-dimensional problem by finding the plane through $[a, b]$ orthogonal to the xy-plane and solving the problem on this plane with the technique of [3]. Since both steps are $O(1)$ the lemma follows.

Lemma 2.7 *Let $[a, b]$ be a line segment above the xy-plane. Then the point ν^* on the xy-plane that yields $\theta_{max}(\nu)$ may be computed in $O(1)$ time.*

3 Viewing a segment from a line

Let $[a, b]$ be a line segment and L be a line not intersecting $[a, b]$ in $\mathbf{R}^3$ on which the viewer is constrained to patrol. In this constrained version of the problem we are interested in determining the maximum aperture angle $\theta_{max}(\nu)$ where ν varies over the entire line L. This version of the problem is, like the problem where the viewer is constrained to roam a plane, also not interesting for $\theta_{min}(\nu)$. Trivially, $\theta_{min}(\nu)$ is also equal to zero and occurs at the points on L at infinity. Therefore when the viewer is allowed to move over the entire range of the line we restrict our attention to the computation of $\theta_{max}(\nu)$. To make the problem of computing $\theta_{min}(\nu)$ interesting and in order to eventually solve the general problem when the viewer is constrained to move in a given convex polyhedron, we constrain the viewer to move on a given interval on the line L. First we investigate the $\theta_{max}(\nu)$ problem.

3.1 Computing $\theta_{max}(\nu)$

When $[a, b]$ and L are coplanar this problem reduces to the planar problem considered in [3] and [4]. There, a very simple characterization of $\theta_{max}(\nu)$ was employed to find the solution. Given three points $\{a, b, \nu\}$ (the two end points of the segment $[a, b]$ and the view point ν on L), let us denote the unique circle that they determine by $C_{ab\nu}$ and let us call it the *aperture circle*. As can be seen from Fig. 4, the point on L that yields $\theta_{max}(\nu)$ is determined by the

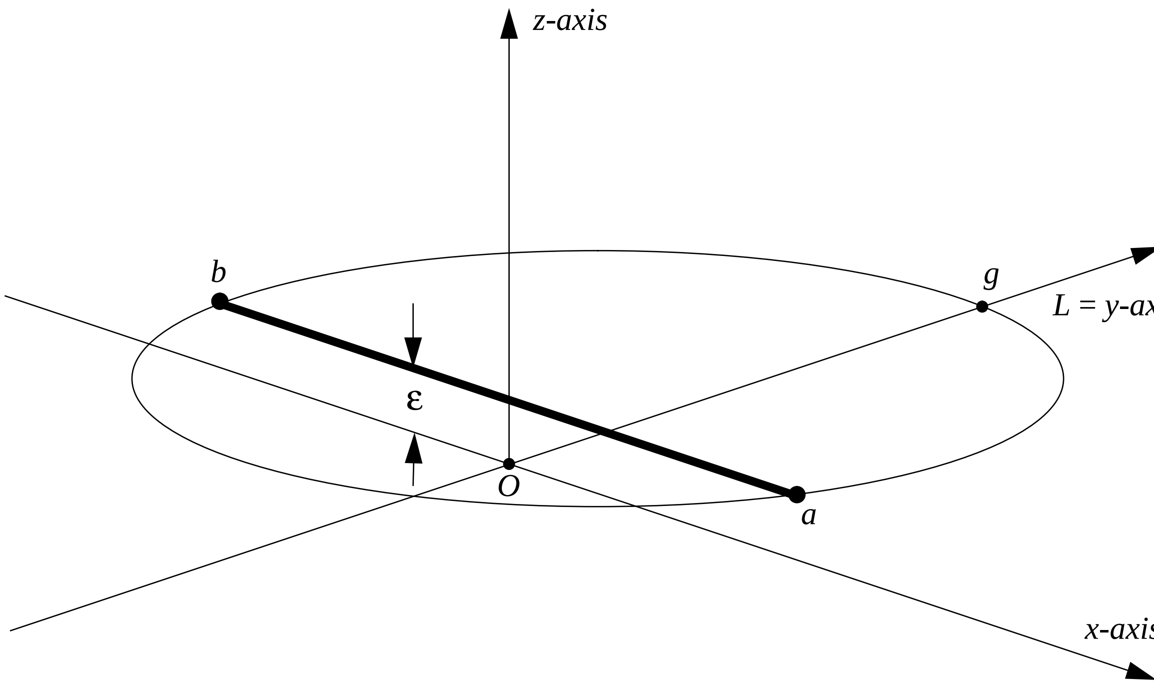

Figure 5: The largest aperture angle of a segment $[a, b]$ from a line L is not necessarily determined by the smallest aperture circle tangent to L.

smallest (there are two of them) aperture circle $C_{ab\nu}$ that intersects L. Furthermore, the solution is realized by the point of intersection. This point may be computed by solving the roots of a quadratic equation and can thus be done in $O(1)$ time with an extension of the basic real RAM model of computation that also allows square root finding as a primitive. One may wonder if this same elegant characterization also holds for the three dimensional problem considered here. It turns out that this is not the case as we now demonstrate.

Let the line L be the y-axis and let the segment $[a, b]$ lie parallel to the x-axis with coordinates $a = (1, 0, \epsilon)$ and $b = (-1, 0, \epsilon)$, where ϵ is a very small positive constant, i.e., $\epsilon \ll 1$ (refer to Fig. 5). Then there are two smallest aperture circles intersecting L, one is C_{abg} where g lies on the positive y-axis and the other is $C_{abg'}$ where g' lies on the negative y-axis symmetrically opposite to g. Both of these aperture circles are close to the circle with $[a, b]$ as diameter and have aperture angles at g and g' close to 90 degrees. The maximum aperture angle $\theta_{max}(\nu)$ however is close to 180 degrees, occurs at the origin O and is determined by the aperture circle C_{abO} which is much larger than either C_{abg} or $C_{abg'}$.

Nevertheless, we show in this section that by introducing additional constraints on the aperture circles, two different more complex characterizations are possible that allow us to efficiently compute $\theta_{max}(\nu)$ in $O(1)$ time with more powerful extensions of the basic real RAM. To set the stage we begin by examining a couple of special cases where the characterization is straight forward and a solution can be obtained with only square roots added to the basic real RAM. It is assumed without loss of generality that L lies in the xy-plane and $[a, b]$ lies above this plane.

The first special case occurs when L is parallel to $[a, b]$. Indeed, this is a special case of the planar problem considered in [3] and [4]. A different and more interesting special case is when $[a, b]$ is orthogonal to the xy-plane as in the following lemma which is an easy consequence of the way in which the toroidal lunoid grows. We therefore state the lemma without proof.

Lemma 3.1 *Let $[a, b]$ be orthogonal to, and at some fixed distance above, the xy-plane such that the line through $[a, b]$ does not intersect L. Then if $[a, b]$ is sufficiently far from L, $\theta(\nu)$ is unimodal and the point on L that realizes $\theta_{max}(\nu)$ occurs at the unique point ν^* which is the point of intersection of L with the vertical plane through $[a, b]$ that is orthogonal to L. If $[a, b]$ is sufficiently near L, $\theta(\nu)$ is* bimodal *and the two points on L that realize $\theta_{max}(\nu)$ occur at the points of intersection of L with the contact circle that the toroidal lunoid (as it grows) determines when it first makes contact with the xy-plane.*

Note that $C(d(a))$ can be obtained by solving the planar problem on any plane containing $[a, b]$, and therefore the solution point ν^* may be computed with square roots as the only primitive added to the basic real RAM.

We now describe our first characterization of the solution for $\theta_{max}(\nu)$ when L and $[a, b]$ are in arbitrary disjoint position in $\mathbf{R}^3$. For a viewpoint ν we define the aperture plane as the plane determined by points $\{a, b, \nu\}$ and denote by c the center of the aperture circle $C_{ab\nu}$.

Lemma 3.2 *$\theta_{max}(\nu)$ is determined by the point ν^* in L that yields the smallest aperture circle with the property that segment $[c, \nu^*]$ is orthogonal to L.*

Proof: We first show that the aperture circle corresponding to $\theta_{max}(\nu)$ belongs to the family of aperture circles that exhibits the orthogonality property. Then we show that the smallest of these yields $\theta_{max}(\nu)$. Recall that the solution point ν^* is determined by the first point of contact of the growing lunoid of $[a, b]$ with L. Denote this lunoid by LUNOID(a, b, ν^*). Let $H_{ST}(\nu^*)$ denote the supporting tangent plane to LUNOID(a, b, ν^*) at ν^*. By lemma 2.6 the aperture plane at ν^* is orthogonal to the supporting tangent plane $H_{ST}(\nu^*)$ and contains the aperture circle determined by ν^*. Furthermore this aperture circle is tangent to $H_{ST}(\nu^*)$. Therefore segment $[c, \nu^*]$ is orthogonal to $H_{ST}(\nu^*)$. But $H_{ST}(\nu^*)$ consists of all lines in $\mathbf{R}^3$ that are tangent to LUNOID(a, b, ν^*) at ν^*. Therefore $[c, \nu^*]$ is orthogonal to L.

Now we show that for any configuration of $[a, b]$ and L no more than three aperture circles satisfy the orthogonality property and the smallest of these yields $\theta_{max}(\nu)$. First it is useful to distinguish between two classes of *aperture circles.* Consider a toroidal lunoid T and its convex hull $CH(T)$. We call a point p on $bd(T)$ a *convex boundary point* if p lies on $bd(CH(T))$. Otherwise we call p a *saddle boundary point.* We also distinguish between two types of contact points between a line and a toroidal lunoid. We say that a line L is globally tangent to a toroidal lunoid T if L intersects T and no point of L lies in the interior of T. On the other hand, we say that a line L is locally tangent to T at a point x, if there exists a sufficiently small, but positive, real number ϵ such that the intersection of L with a sphere of radius ϵ centered at x yields a line segment s such that either no point of s lies in the interior of T or no point of s lies in the exterior of T. In other words, a line L that is locally tangent to T at x does not properly cross the boundary of T in a local neighborhood of x but it may do so at other points. A line that is globally tangent to T does not properly cross the boundary of T at any point.

Consider a line segment $[a, b]$ and some view point ν in space as well as the resulting $LUNOID(a, b, \nu)$ and assume the lunoid is toroidal. A property of all lunoids is that every point p on $bd(\text{LUNOID}(a, b, \nu))$ other than a and b admits a line $L(p)$ that is *locally* tangent to the lunoid at p. This implies that every such point p on $bd(\text{LUNOID}(a, b, \nu))$ admits an aperture circle C_{abp} such that the segment $[c, p]$ is orthogonal to $L(p)$, where c is the center of aperture circle C_{abp}.

Let p be a point on $bd(\text{LUNOID}(a, b))$ as it grows. If L is locally tangent to LUNOID(a, b) at p and p is a convex point we call L a *convex tangent line* and if p is a saddle point we call L a *saddle tangent line*. Now consider the LUNOID(a, b) as it grows from $[a, b]$ without bound. Depending on the configuration of $[a, b]$ and L we have either one orthogonal aperture circle or three of them. Case (1) occurs when the growing LUNOID(a, b) first touches L. Therefore it is determined by a convex point and L is a global tangent line of LUNOID(a, b), i.e., L does not intersect the interior of the lunoid. In case (2) we have two sub-cases: in sub-case 2.1 we have two convex and one saddle tangent lines and in sub-case 2.2 we have one convex and two saddle tangent lines. Case 2.1 can happen when L is orthogonal to $[a, b]$. Case 2.2 can happen when the first tangent line is convex but as the LUNOID(a, b) grows, L enters the region outside LUNOID(a, b) but interior to $CH(\text{LUNOID}(a, b))$. These cases are illustrated in Fig. 6. Since the convex aperture circles are smaller because they are detected before the saddle circles and since $\theta_{max}(\nu)$ is determined by a convex circle, the solution is determined by the smallest of these three circles. ■

The characterization of the solution provided by lemma 3.2 allows us to compute $\theta_{max}(\nu)$ in constant time if the basic real RAM is extended to allow the computation of cube roots and trigonometric functions as primitives, as we now demonstrate. We will show that the three aperture circles discussed above can be obtained from the roots of a cubic equation which can be derived from the constraints of the problem imposed by the characterization. Expressions for computing these roots may be obtained from Cardano's formulas and, depending on the configuration of $[a, b]$ and L, can sometimes be represented as algebraic functions of nested radicals and in some cases trigonometric functions [26], [14].

Lemma 3.3 *Let $[a, b]$ be a line segment and L be a line not intersecting $[a, b]$ in $\mathbf{R}^3$. Then the point ν^* on L that realizes $\theta_{max}(\nu)$, as well as the corresponding value of $\theta_{max}(\nu)$, may be computed in $O(1)$ time under the extended real RAM model of computation.*

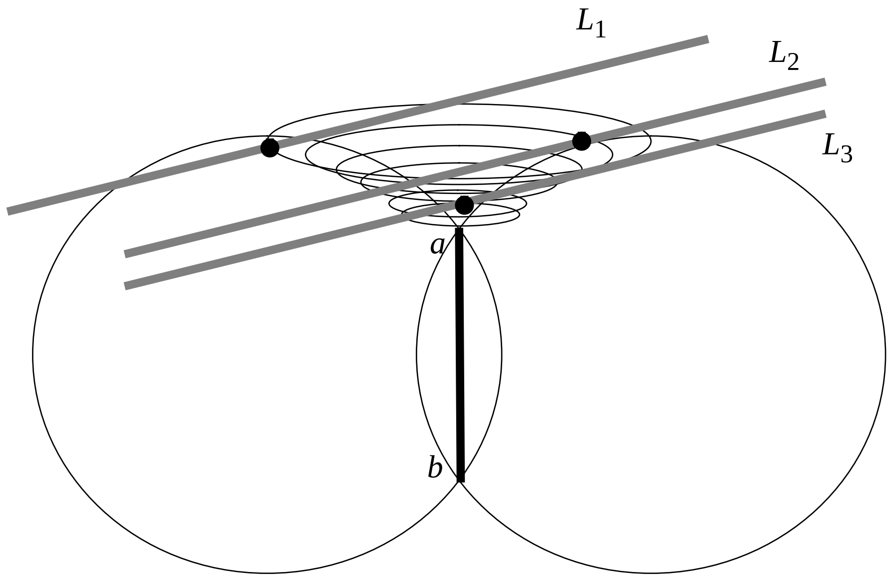

Figure 6: Illustrating the three types of aperture circles that exhibit the orthogonality property.

Proof: Without loss of generality assume that the line L is the x-axis and refer to Fig 7. From lemma 3.2 it follows that we wish to find the value of the x-coordinate of the viewpoint ν^* on L such that the aperture circle C with center o, determined by the three points $a = (a_1, a_2, a_3)$, $b = (b_1, b_2, b_3)$ and $\nu = (x, 0, 0)$, where ν is a variable, is as small as possible with the constraint that the segment $[o, x]$ be orthogonal to the x-axis. Let $o = (h, k, l)$ denote the coordinates of the center of aperture circle C. Since a and b are points on C it follows that the Euclidean distance $d(o, a) = d(o, b)$. This constraint leads to the following equation:

$$(h - b_1)^2 + (k - b_2)^2 + (l - b_3)^2 = (h - a_1)^2 + (k - a_2)^2 + (l - a_3)^2$$

Similarly, since b and ν are points on C it follows that the Euclidean distance $d(o, b) = d(o, \nu)$ and, by substituting $\nu = (x, 0, 0)$ for $\nu = (\nu_1, v_2, v_3)$, this constraint leads to:

$$(h - b_1)^2 + (k - b_2)^2 + (l - b_3)^2 = (h - x)^2 + (k)^2 + (l)^2$$

The constraint that the segment $[o, x]$ be orthogonal to the x-axis implies $h = x$. Substituting $h = x$ in the two above equations, expanding terms and rearranging leads to two linear equations in the three unknowns, k, l and x (linear because all the quadratic terms are present in both sides of both equations and they cancel each other out). Solving this system of equations for k and l as a function of x yields the following two equations:

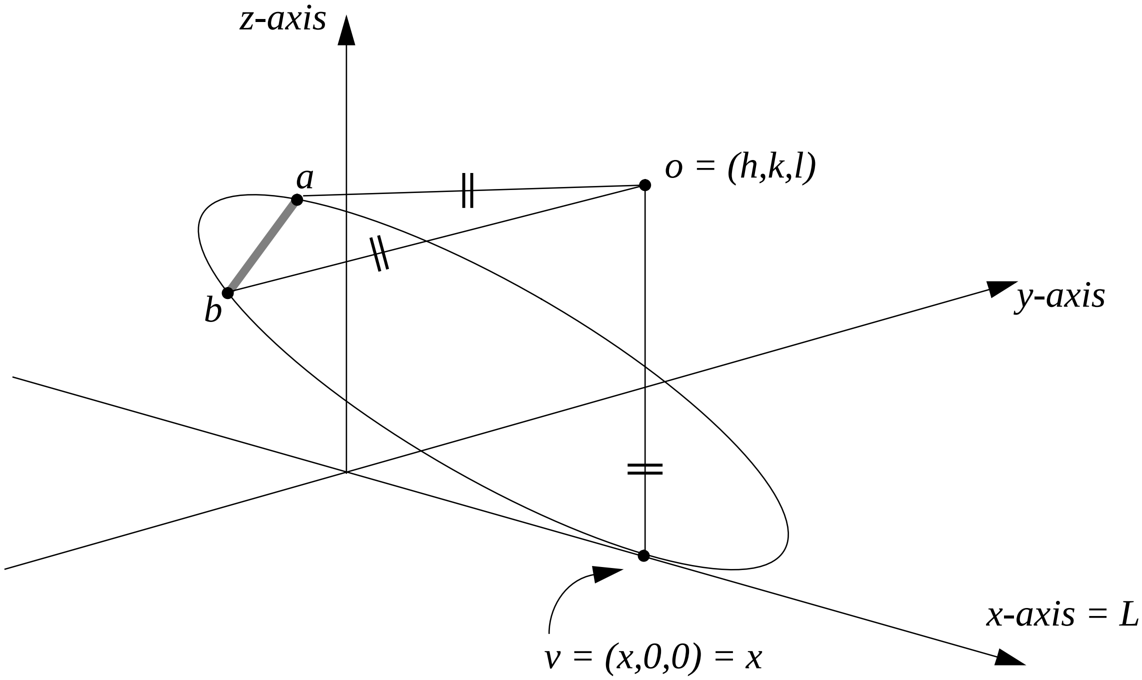

Figure 7: Computing the location of $\theta_{max}(\nu)$ on line L.

$$k(x) = \left(-\frac{1}{\Delta}\right)\left\{\frac{(a_3-b_3)}{2}x^2 + (a_1b_3 - a_3b_1)x + \frac{(\|b_2\|^2a_3 - \|a_2\|^2b_3)}{2}\right\}$$

and

$$l(x) = \frac{1}{\Delta}\left\{\frac{(a_2-b_2)}{2}x^2 + (a_2b_1 - a_1b_2)x + \frac{(\|b_2\|^2a_2 - \|a_2\|^2b_2)}{2}\right\}$$

where $\Delta = a_3(a_2 - b_2) - a_2(a_3 - b_3) \neq 0$, if and only if a, b and L are non-coplanar. If a, b and L are coplanar, a condition that we test at the beginning, then the solution is found by the simpler methods of [3], [4].

There are still an infinite number of locations in space for o that satisfy the above three constraints as Fig 7 illustrates. These locations are determined by the line orthogonal to and through the center of the circle determined by a, b and x. To make sure o is the center of the aperture circle determined by a, b and ν we make sure that the four points a, b, ν and o are co-planar. We do this by invoking the co-planarity determinant equation:

$$\det \begin{bmatrix} a_1 & a_2 & a_3 & 1 \\ b_1 & b_2 & b_3 & 1 \\ h & k & l & 1 \\ x & 0 & 0 & 1 \end{bmatrix} = -x \det \begin{bmatrix} a_2 & a_3 & 1 \\ b_2 & b_3 & 1 \\ k & l & 1 \end{bmatrix} + \det \begin{bmatrix} a_1 & a_2 & a_3 \\ b_1 & b_2 & b_3 \\ h & k & l \end{bmatrix} = 0$$

Substituting the equations for k and l as a function of x in the above determinant equation, setting $h = x$, expanding and rearranging terms yields the following cubic equation in x:

$$Ax^3 + \frac{3}{2}Bx^2 + (C_1 + C_2)x + D = 0$$

where:

$$A = \frac{(a_3 - b_3)^2}{2} + \frac{(a_2 - b_2)^2}{2}$$

$$B = (a_3 - b_3)(a_1 b_3 - b_1 a_3) + (a_2 - b_2)(a_1 b_2 - b_1 a_2)$$

$$C_1 = \left((a_1 b_3 - b_1 a_3)^2 + \left(\|b_2\|^2 a_3 - \|a_2\|^2 b_3 \right) \frac{(a_3 - b_3)}{2} \right)$$

$$C_2 = \left((a_1 b_2 - b_1 a_2)^2 + \left(\|b_2\|^2 a_2 - \|a_2\|^2 b_2 \right) \frac{(a_2 - b_2)}{2} \right)$$

and,

$$D = (a_1 b_3 - b_1 a_3)^2 \frac{(\|b_2\|^2 a_3 - \|a_2\|^2 b_3)}{2} + (a_1 b_2 - b_1 a_2)^2 \frac{(\|b_2\|^2 a_2 - \|a_2\|^2 b_2)}{2}$$

There are three aperture circles that satisfy the above four constraints and the tangent points on the x-axis of these circles are given by the three roots of the above cubic equation. The smallest circle will in turn lead to the maximum aperture angle. Computational formulas for these roots were derived by Cardano (also known as Cardan's formulas) and involve either nested radicals with square and cube roots or trigonometric functions, depending on whether D is positive or negative. The expressions are lengthy and involved functions of A, B, C_1, C_2, and D and will not be reproduced here as they can be found in Uspensky's book [26]. By substituting the above expressions for A, B, C_1, C_2, and D into Cardan's formulas the three roots may be computed. Therefore the appropriately extended RAM will find the location of $\theta_{max}(\nu)$ in $O(1)$ time. Furthermore if the RAM is also afforded inverse trigonometric functions then the actual value of $\theta_{max}(\nu)$ may also be computed in constant time. ∎

We now describe our second characterization of the solution for $\theta_{max}(\nu)$ when L and $[a, b]$ are in arbitrary disjoint position in $\mathbf{R}^3$. The first characterization presented above involved *four* constraints that must be satisfied by the aperture circle. The second characterization involves minimizing the radius of a sphere subject to *three* constraints. This characterization has no computational advantage over the previous one because, as we shall see, it leads to computing the roots of exactly the same cubic equation. However, it is in a sense a more natural generalization of the two-dimensional characterization to three dimensions and, together, the two characterizations provide a deeper understanding of the problem.

Lemma 3.4 *$\theta_{max}(\nu)$ is determined by the point ν^* in L that is the point of tangency of L with the smallest sphere S containing a and b on its boundary.*

Proof: First observe the fact that given a circle C_r of radius r, the smallest sphere S that can contain C_r on its boundary must have radius at least as large as r. Furthermore, if the radius of S equals r then C_r is a great circle of S.

From lemma 3.2 we have that the point ν^* on L that maximizes $\theta(\nu)$ is determined by a circle C with center o such that C contains a and b, C intersects L at ν^*, and the segment $[o, \nu^*]$ is orthogonal to L. Let the sphere S_c be the sphere determined by the aperture circle C that determines ν^* and such that C is a great circle of S_c. Therefore the sphere S_c contains a, b and ν^*. Furthermore, since o is the center of the great circle C it must also be the center of S_c. Therefore S_c is tangent to L. Finally, S_c must be the smallest such sphere because a smaller sphere would imply that there exists an aperture circle smaller than C, a contradiction. ∎

Let us examine the method of computing $\theta_{max}(\nu)$ that is suggested by lemma 3.4. Let S be the sphere in question with center $o = (h, k, l)$. The equation of a sphere is then given by:

$$(x-h)^2 + (y-k)^2 + (z-l)^2 = r^2$$

We know that the sphere S must pass through the three points a, b and ν. These three constraints yield the following three constraint equations:

$$(a_1-h)^2 + (a_2-k)^2 + (a_3-l)^2 = r^2$$

$$(b_1-h)^2 + (b_2-k)^2 + (b_3-l)^2 = r^2$$

$$k^2 + l^2 = r^2$$

The third equation requires $[o, v]$ to be orthogonal to L. The reader may verify that there are still an infinite number of spheres that satisfy the above three constraints. Furthermore the third constraint that the segment $[o, \nu]$ be orthogonal to L does not imply the resulting sphere is tangent to L, the last requirement of the lemma. Observe however that the only way to make a sphere that intersects L be tangent to L is to make it the smallest such sphere. Lemma 3.4 then implies we can find $\theta_{max}(\nu)$ by minimizing r^2 subject to the three constraints described above. Now minimizing r^2 is the same as minimizing $k^2 + l^2$. Since $h = x$ the first two constraints yield a pair of linear equations in x, k and l. These can be solved as before for k and l as functions of x and substituted in $r^2 = k^2 + l^2$ yielding a function $r^2(x)$. To minimize $r^2(x)$ we may take the derivative and set it equal to zero. When this is done we obtain the same cubic equation that was obtained from the first characterization using the co-planarity condition of points a, b, ν and o. The three roots correspond to the three possibly distinct and locally minimal spheres that a line segment admits such that the sphere is tangent to a given line. In the first characterization of the solution of $\theta_{max}(\nu)$ the three corresponding aperture circles are the three great circles of the three respective spheres that are obtained in the second characterization.

3.2 Computing $\theta_{min}(\nu)$

In order to solve the problem of computing $\theta_{min}(\nu)$ when the viewer is allowed to move in a given convex polyhedron, we will need a solution to a special case of the above problem in which we constrain the viewer to move on a given interval $[c, d]$ on the line L. From the previous lemmas concerning the computation of $\theta_{max}(\nu)$ we know that depending on the configuration of the segment $[a, b]$ and the line L in $\mathbf{R}^3$, the function $\theta(\nu)$, as ν traverses L, is either unimodal or bimodal. Furthermore, these two situations correspond, respectively, to the cases when the cubic equation that yields $\theta_{max}(\nu)$ as one of its roots, has either one or three real roots. When the equation has one real root it corresponds to the maximum of the unimodal function. When the equation has three real roots, two of them correspond to the two local maxima of the bimodal function and the third corresponds to the local minimum lying in between the two local maxima on L. These observations point the way to using the same cubic equation for solving the problem of computing $\theta_{min}(\nu)$. If we have one real root the function is unimodal on L and in this case $\theta_{min}(\nu)$ is realized by one of the end-points of the constraint interval $[c, d]$. Hence all we need to do is compute the aperture angle at points c and d and select the smaller of the two values. If on the other hand we have three real roots then let r_2 denote the point on L (root) that lies in between the two other points (roots) r_1 and r_3. The points r_1 and r_3 correspond to the local maxima and r_2 to the local minimum. We test to determine if r_2 lies in the interior of segment $[c, d]$. If it does not then we compute the aperture angle at points c and d and select the smaller of the two values as before. If r_2 lies in the interior of segment $[c, d]$ then we compute the aperture angle at points c, d and r_2 and select the smaller of the three values. We have therefore proved the following lemma under the extended RAM model of computation that allows square and cube roots as well as trigonometric and inverse trigonometric functions.

Lemma 3.5 *The point ν^* on L that realizes $\theta_{min}(\nu)$, as well as the corresponding value of $\theta_{min}(\nu)$, may be computed in $O(1)$ time.*

4 Viewing a segment from a convex polygon

Let P be a convex polygon in $\mathbf{R}^3$ not intersecting $[a, b]$. Furthermore, assume that P lies on the xy-plane and that $[a, b]$ lies above the xy-plane. It is also assumed that the polygon is stored in an array. Here we are interested in determining $\theta_{max}(\nu)$ and $\theta_{min}(\nu)$ where ν varies over the polygon P. We consider two versions of the problem: (1) ν varies over the entire polygon P including its interior and (2) ν varies only over the boundary of polygon P. Problems (1) and (2) correspond to optimization problems of $\theta(\nu)$ with *linear inequality* and *linear equality* constraints, respectively.

First we prove a basic lemma which concerns the determination of whether a given point lies in the interior of a given n-vertex convex polygon in $O(\log n)$ time. It is well known that given a convex n-vertex polygon stored in an array, it is possible to construct from it a data structure, in $O(n)$ time and space, such that subsequently point inclusion queries can be determined in $O(\log n)$ time. Two such data structures are the star-decomposition [19] and the balanced hierarchical decomposition [15]. We show here that $O(\log n)$ time suffices even if the polygon is stored simply as an array, thus strengthening the results found in [19] and [15]. In all our results we assume the extended real RAM model of computation.

Lemma 4.1 *Given a convex polygon P of n vertices in $\mathbf{R}^2$ stored in an array and a point c, whether c lies in P can be determined in $O(\log n)$ time.*

Proof: Let L and R denote the vertical lines through the leftmost and right-most points of P. These points partition P into the upper and lower chains P_{up} and P_{down}, respectively. If c lies to the left of L or to the right of R then c lies outside P and we are done. Therefore assume c lies in between L and R. Perform binary search among the x coordinates of the vertices of P_{up} to locate c within a slab determined by an edge of P_{up}. Let p_i and p_{i+1} be the vertices of the slab thus found. Find the intersection points of the vertical projections of p_i and p_{i+1} onto P down and call these points z_1 and z_2. Now perform binary search on the section of P_{down} that lies between z_1 and z_2 to locate c inside a slab determined by an edge of P_{down}. If c lies between the two edges thus found it lies in P and otherwise it lies outside P. Lines L and R as well as the intersection points z_1 and z_2 can be found in $O(\log n)$ time with the algorithms of Chazelle and Dobkin [5]. ∎

4.1 Computing $\theta_{max}(\nu)$

We turn now to the problem of computing $\theta_{max}(\nu)$ in $\mathbf{R}^3$ where ν is allowed to move in a convex polygon P (the boundary of P as well as its interior). We first consider the special cases where the line segment $[a, b]$ is vertical (orthogonal to the xy-plane). This special case admits a very simple algorithm for solving the problem.

Lemma 4.2 *Let $[a, b]$ be orthogonal to the xy-plane (and above it) and ν a viewing point on the xy-plane constrained to lie in polygon P. Then computing $\theta_{max}(\nu)$ has time complexity $\Theta(n)$.*

Proof: By proposition 2.1 let C denote the circle that describes the infinite set of unconstrained maxima values of $\theta(\nu)$ on the xy-plane. Let c denote the center of C. First we check every edge of P to see if the boundary of P intersects C. If it does then the intersection points of P on C describe all the arcs of solution values in P and we are done. Therefore assume that C does not intersect the boundary of P. Two cases arise: case 1: C lies completely in the interior of P or, case 2: C lies completely in the exterior of P. In case 1 the entire circle C forms the solution. Therefore consider case 2 when C lies in the exterior of P. There are two sub-cases: case 2.1: P lies in the interior of C, and case 2.2: P lies exterior to the circle C.

Case 2.1: As the lunoid grows the circle C divides into two circles. These two circles define the intersection points of the boundary of the lunoid with the xy-plane. One grows larger to infinity playing no role in the solution and the other shrinks until it touches P. The first points of contact of the shrinking circle with P form the solution. Such a point (or points) may be found by computing the furthest points of the boundary of P from c, the center of C.

Case 2.2: As the lunoid grows the circle C again divides into two circles. One grows smaller to zero playing no role in the solution and the other expands until it touches P. The first point of contact that the expanding circle makes with P is the unique solution. This point may be found by computing the nearest point of the boundary of P from c.

Consider now the complexity of the above construction. Checking the intersection of every edge of P with C costs $O(n)$ time. If the boundary of P does not intersect C then to distinguish between case 1 and case 2, i.e., to determine whether C lies completely inside P or not, it suffices to take any point on C and test it for inclusion in P which by lemma 4.1 can be done in $O(\log n)$ time. To distinguish between case 2.1 and case 2.2, i.e., to determine whether P lies completely inside C or not, it suffices to take any vertex of P and test it for inclusion in C, which can be done in $O(1)$ time. In case 2.1 the furthest points of the boundary of P from c can be computed in $O(n)$ time by examining every vertex. In case 2.2 the nearest point of the boundary of P from c may be computed in $O(\log n)$ time with the algorithm of Edelsbrunner [12]. Therefore $O(n)$ suffices for the entire algorithm. Since the problem of computing $\theta_{max}(\nu)$ is equivalent to computing the maximum distance from a point to a convex polygon in the xy-plane (case 2.1) and the latter problem has complexity $\Omega(n)$ when the polygon is stored in an array [12], the lemma follows. ∎

Lemma 4.3 *Let $[a, b]$ be orthogonal to the xy-plane (and above it) and ν a viewing point on the xy-plane constrained to lie on the boundary of polygon P. Then computing $\theta_{max}(\nu)$ has time complexity $\Theta(n)$.*

Proof: If P lies completely in C or the boundary of P intersects C we proceed as in lemma 4.2. Therefore assume C lies completely in P. As the lunoid grows the circle C splits into two circles. One shrinks to zero and the other grows until it touches P. The first points of contact form the solution. Such a point may be found by computing the nearest point of the boundary of P from c. Therefore we can test every edge of P to find a point that minimizes the distance. Therefore $O(n)$ suffices to find $\theta_{max}(\nu)$. To show that the lower bound on the time complexity of computing $\theta_{max}(\nu)$ is $\Omega(n)$ it suffices to see that there exist convex polygons easily constructed by perturbing a regular polygon such that every edge must be visited to ensure $\theta_{max}(\nu)$ is not missed in the search. ■

Consider now the general problem of computing $\theta_{max}(\nu)$ in $\mathbf{R}^3$ where ν is allowed to move in a convex polygon P (the boundary of P as well as its interior) on the xy-plane and the line segment lies in general position above the xy-plane. The intersection curves of the lunoid with the xy-plane are no longer circles and we cannot reduce the problem to simple distance computations on the plane as we did in the special cases discussed above. Nevertheless, we may compute $\theta_{max}(\nu)$ in optimal linear time, as we now show, by using the cubic equations derived in the previous section.

Theorem 4.4 *$\theta_{max}(\nu)$ may be computed in time $\Theta(n)$.*

Proof: First we solve the problem for the xy-plane, disregarding the polygon P altogether. Let ν^* be the point on the xy-plane that realizes $\theta_{max}(\nu)$. Next we test to see if ν^* lies in P. If it does we are done because no other point in P can have a larger angle. If ν^* lies in the exterior of P then the corresponding lunoid must grow further before it touches the boundary of P. The point at which it touches P will be the unique solution. Therefore in this case we know the solution lies on some edge of P. Hence for each edge of P we consider the line that contains it and solve the problem for the line by computing the roots of the cubic equation of section 3. Consider one such line L containing some edge e of P. Let x be the solution obtained on L when e is ignored. If x lies in e we are done. If x does not lie in e we choose the maximum aperture angle obtained by comparing the angles that the other roots (that lie in e) yield with the angles obtained at the end-points of e. Thus each edge of P yields a candidate for the global solution. Finally, we select the maximum among all the candidates. $O(n)$ time clearly suffices and the $\Omega(n)$ bound follows from the above lemmas.

■

4.2 Computing $\theta_{min}(\nu)$

We turn now to the problem of computing $\theta_{min}(\nu)$ in $\mathbf{R}^3$ where ν is allowed to move in a convex polygon P (the boundary of P as well as its interior) and the line segment is in general position. In fact we assume that P lies in the xy-plane, the segment $[a, b]$ lies above the xy-plane and the line through $[a, b]$ does not intersect P. The case when the line through $[a, b]$ intersects P is trivial since in this case $\theta_{min}(\nu) = 0$. Furthermore, this case can be tested for and dispensed with in $O(\log n)$ time by lemma 4.1 where the point c is the intersection point of the line through $[a, b]$ with the xy-plane.

Lemma 4.5 *$\theta_{min}(\nu)$ is determined by a point on the boundary of P.*

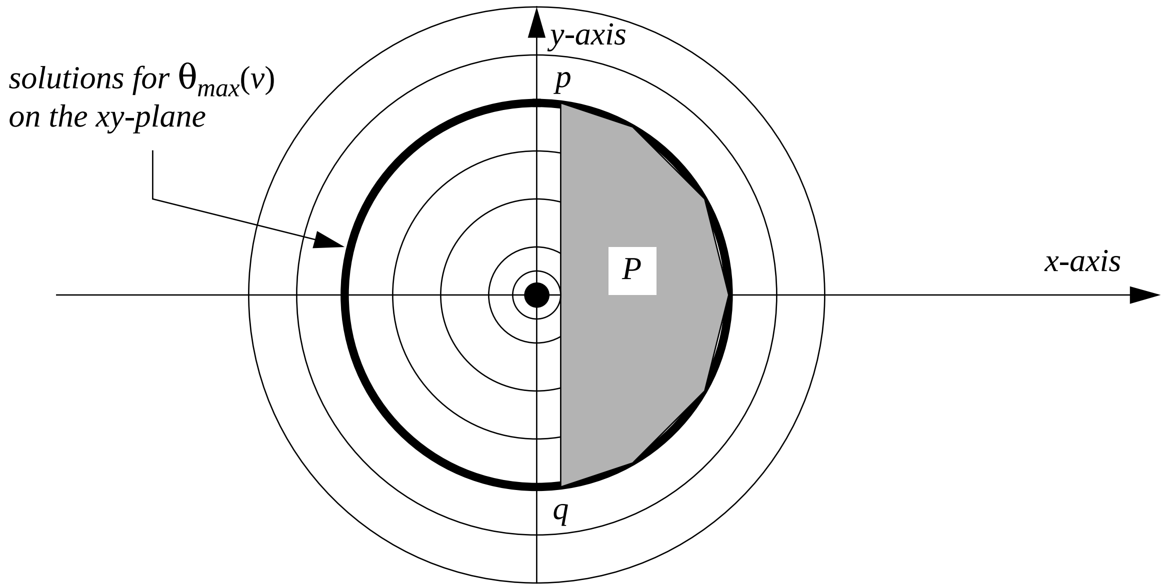

Figure 8: $\theta_{min}(\nu)$ need not be determined by a vertex of P.

Proof: (by contradiction) Assume $\theta_{min}(\nu)$ is determined by a point ν in the interior of P. Then the lunoid that determines $\theta_{min}(\nu)$ contains ν on its surface and there exist points of P outside this lunoid. Therefore there exists a larger lunoid determined by some point of P that yields a value of θ smaller than $\theta_{min}(\nu)$, a contradiction. ■

Note that this lemma establishes that in the case of $\theta_{min}(\nu)$ there is no distinction between the equality constraints and the inequality constraints. Allowing ν the additional freedom to move in the interior of P does not change the solution obtained when ν is constrained to lie on the boundary of P.

In the 2-dimensional version of this problem considered in [3] and [4] it was shown that $\theta_{min}(\nu)$ is determined by a vertex of P. However, this property fails to carry over to the 3-dimensional problem considered here as the following example demonstrates.

Let the line segment $[a, b]$ be positioned vertically over the origin of the xy-plane at some fixed height above it and refer to Fig. 8. The solution region for $\theta_{max}(\nu)$ when ν is constrained to lie on the xy-plane then consists of a circle C max centered at the origin (heavy circle in Fig. 8). We construct our convex polygon P on the xy-plane such that all its vertices lie on C max and have positive x-coordinate. Furthermore we choose the highest and lowest vertices, p and q respectively in Fig. 8, to be very close to the y-axis. Since all points on the xy-plane and not on C max have smaller aperture angles it follows that $\theta_{min}(\nu)$ is not determined by a vertex of P. In the example of Fig. 8 $\theta_{min}(\nu)$ is realized by an interior point of the edge $[p, q]$ of P, and, in general, $\theta_{min}(\nu)$ may be realized by a point anywhere on the boundary of polygon P. Nevertheless, we can still prove a linear lower bound on

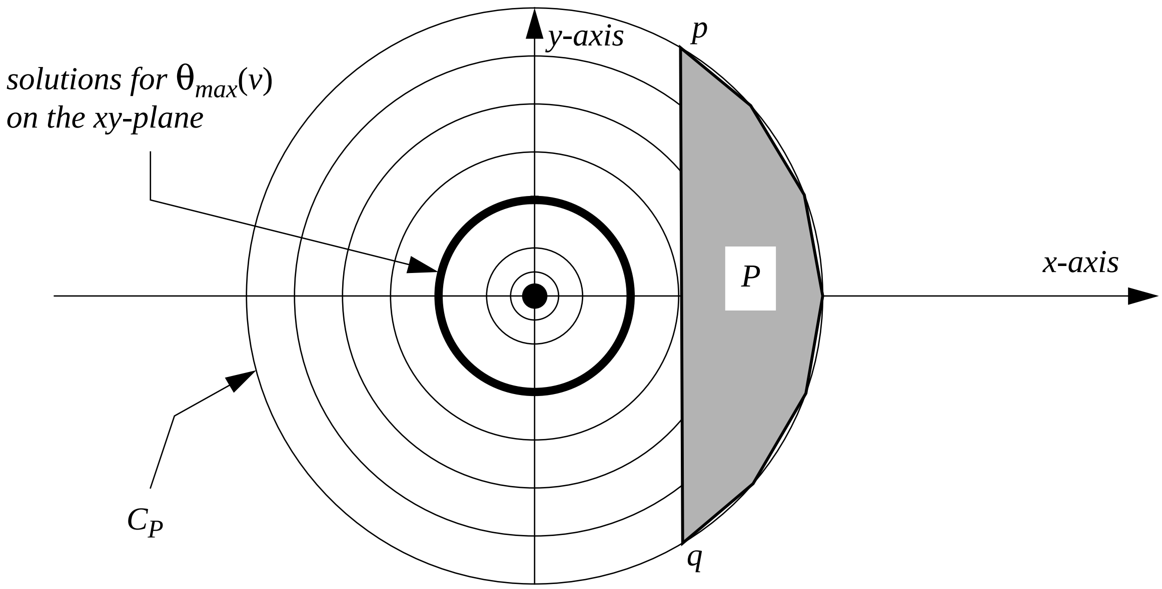

Figure 9: $\theta(\nu)$ may have $\Omega(n)$ local minima.

the complexity of computing $\theta_{min}(\nu)$ for this problem because there exist instances where all vertices of P must be visited in order to find the solution, as we now demonstrate. The lower bound relies again on an adversary argument that exploits on the natural representation of a polygon by a linear array.

Lemma 4.6 *$\theta(\nu)$ may have $\Omega(n)$ local minima.*

Proof: Let the segment $[a, b]$ be located vertically (parallel to the z-axis) some fixed height above the xy-plane and let C_{max} denote the circular set of solutions that realize $\theta_{max}(\nu)$ on the xy-plane and refer to Fig. 9 where C_{max} is shown in bold. We now construct the convex polygon P such that all its vertices lie on a circle C_P centered at the origin and greater than C_{max} and such that the vertices of P with largest and smallest y-coordinates (vertices p and q respectively) lie on a vertical line with x-coordinate greater than the radius of C_{max}. It follows that $\theta_{min}(\nu)$ is non-unique and realized by each and every vertex of P. ∎

Theorem 4.7 *Computing $\theta_{min}(\nu)$ has complexity $\Omega(n)$.*

Proof: Lemma 4.6 shows a construction such that $\theta_{min}(\nu)$ is realized by each and every vertex of P. To finish the construction we take an arbitrary vertex z of P and pull it slightly outside the circle C_P while preserving the convexity of P. This arbitrary vertex now contains the global minimum. Therefore if an algorithm does not check every vertex of P it may miss the global minimum. ■

We turn now to algorithms for computing $\theta_{min}(\nu)$. We first consider the special case where the line segment $[a, b]$ is vertical (orthogonal to the xy-plane). This special case admits a simple algorithm, similar to its $\theta_{max}(\nu)$ counterpart, that does not require the computation of cube roots or trigonometric functions.

Lemma 4.8 *Let $[a, b]$ be orthogonal to the xy-plane (and above it) and ν a viewing point on the xy-plane constrained to lie in polygon P. Then computing $\theta_{min}(\nu)$ has time complexity $\Theta(n)$.*

Proof: By proposition 2.1 let C denote the circle that describes the infinite set of unconstrained maxima values of $\theta(\nu)$ on the xy-plane. Let c denote the center of C. Let p_f be the point of P furthest (according to the Euclidean distance measure) from c. Similarly let p_n be the point of P nearest to c. Then it suffices to observe that $\theta_{min}(\nu) = \min\{\theta(p_f), \theta(p_n)\}$. The correctness of this solution follows from arguments similar to those in lemma 4.3. Since p_f and p_n can be computed in $O(n)$ time [12] $\theta_{min}(\nu)$ can be computed in $O(n)$ time. The $\Omega(n)$ lower bound for this special case of the problem was established in theorem 4.7. ■

Consider now the general case where the line segment $[a, b]$ is in general position.

Theorem 4.9 *$\theta_{min}(\nu)$ may be computed in time $\Theta(n)$.*

Proof: By lemma 4.5 we know that $\theta_{min}(\nu)$ is determined by an edge of P. Therefore, for each edge of P we solve the problem using the method of section 3.2 to obtain a candidate for the global minimum. Finally, we select the smallest of the candidate values for the global solution. Since each candidate is computed in $O(1)$ time, the entire solution is $O(n)$. The $\Omega(n)$ bound follows from theorem 4.7. ■

5 Viewing a segment from a convex polyhedron

Let P be a convex polyhedron in $\mathbf{R}^3$ not intersecting the line segment $[a, b]$. In this section we are concerned with the problems of determining $\theta_{max}(\nu)$ and $\theta_{min}(\nu)$ where ν varies over the entire polyhedron P. We assume that the polyhedron is represented in the form of a *doubly-connected-edge-list* (DCEL). This is a simple and natural data structure employed widely for representing planar graphs [19] and since a convex polyhedron is isomorphic to a planar graph it is suited for our problem. Furthermore this data structure is easily implemented with six linear arrays and thus corresponds very much with our data structure used in the problems concerning convex polygons in the previous sections of this paper. We first examine the $\theta_{max}(\nu)$ problem for which we assume that the line $L(a, b)$ through the segment $[a, b]$

does not intersect the interior of polyhedron P. If it does (which can be determined in $O(n)$ time) we decompose P into two convex polyhedra by cutting P (also in $O(n)$ time) with any plane containing $[a, b]$. We then solve for $\theta_{max}(\nu)$ in each of the two resulting polyhedra and select the greater of the two results as the global maximum.

5.1 Computing $\theta_{max}(\nu)$

First we show that only the boundary of P need be searched.

Lemma 5.1 *$\theta_{max}(\nu)$ is determined by a point ν^* on $bd(P)$.*

Proof: Let ν^* be the point in P that determines $\theta_{max}(\nu)$ and assume $\nu^* \in int(P)$. The segment $[a, \nu^*]$ must intersect $bd(P)$ at some point, say a^*, for which $\angle(a, a^*, b) > (a, \nu^*, b)$, a contradiction. ■

In fact we may restrict further our search for $\theta_{max}(\nu)$ to a smaller region of the boundary of P "facing" $[a, b]$, a region we refer to as the inner facets of P with respect to segment $[a, b]$. To make this notion more precise it is useful to introduce the notion of visibility and to regard P as an opaque object that is a barrier to straight lines of sight. Accordingly we say that two points x and y in $\mathbf{R}^3$ are *visible* (or *see* each other) if the interior of $[x, y]$ does not intersect the interior of P. We say that a point x sees a set S if x sees every element of S. We now relate the solution region to the subset of $bd(P)$ that has the property that every point of this subset sees $[a, b]$.

Lemma 5.2 *$\theta_{max}(\nu)$ is determined by a point ν^* that sees $[a, b]$.*

Proof: (by contradiction) Let ν^* determine $\theta_{max}(\nu)$ and assume that ν^* does not see every point of $[a, b]$. Then there exists a point $u \in [a, b]$ such that $int[u, \nu^*]$ intersects $int(P)$. Let w be such an intersection point. Since (1) $w \in int(P)$, (2) $w \in \Delta(a, b, \nu^*)$ and (3) $w \neq \nu^*$ it follows that $\angle(a, w, b) > \angle(a, \nu^*, b)$, a contradiction. ■

Let us now consider the problem of computing this restricted region of P. Let x be a point exterior to P and let $P_V(x)$ denote the set of points of $bd(P)$ that see x. Let $P_V[a, b]$ denote the set of points of $bd(P)$ that see segment $[a, b]$, i.e., each point in $P_V[a, b]$ sees every point in $[a, b]$ (see Fig. 10). Note that in general $P_V[a, b]$ is not equal to the intersection of $P_V(a)$ and $P_V(b)$. It is possible that $P_V(a) \cap P_V(b) = bd(P)$ and yet $bd(P)$ may contain many points that do not see $[a, b]$.

A *common tangent* plane of x and P is a plane that contains x, contains at least one point of P (touches P) and has $int(P)$ to one side of it. Let $H(P)$ be a closed half-space in $\mathbf{R}^3$ that contains P. Define the *support cone* of P from x as the intersection of all half-spaces $H(P)$ such that x is on each plane that determines each $H(P)$. Denote such a polyhedral support cone by $SC_P(x)$. Note that the set $\{bd[SC_P(x)]\} \cap \{bd(P)\}$ partitions $bd(P)$ into two connected components: the visible set $P_V(x)$ and the invisible set $P_I(x)$. The set $P_I(x)$

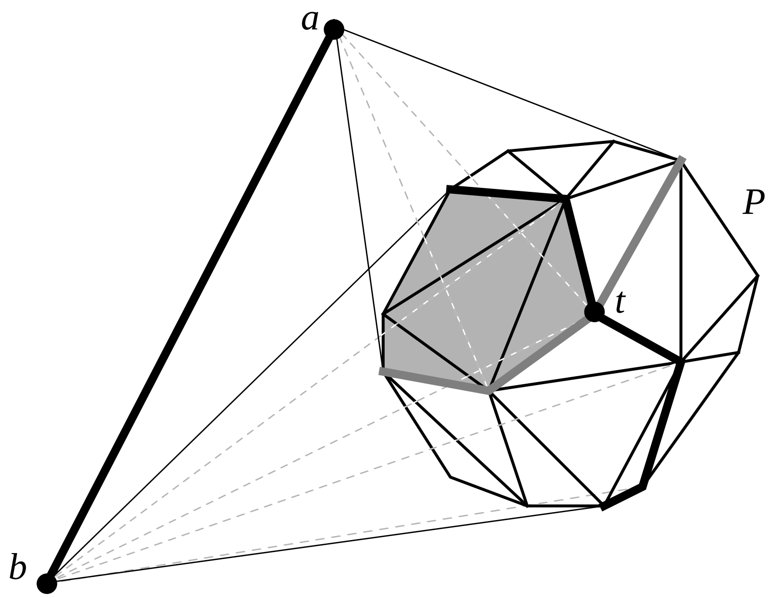

Figure 10: Illustrating the region $P_V[a,b]$ (shaded) from which $[a,b]$ is visible.

consists of those points of P that do not see x. For algorithmic purposes it is further useful to define the notion of a *separating support cone* of P with respect to the segment $[a,b]$. A separating support plane of segment $[a,b]$ and polyhedron P is a plane that contains at least one point of $[a,b]$, at least one point of P, and is such that $int[a,b]$ is not on the same side of the plane as $int(P)$, i.e., the plane is tangent to both objects and separates their interiors. A separating support plane partitions $\mathbf{R}^3$ into two separating support half-spaces. The *separating support cone* of P with respect to the segment $[a,b]$, denoted by $SSC_P[a,b]$, is the unbounded convex polyhedral set obtained by intersecting all separating support half-spaces that contain P. Note that removing P from $SSC_P[a,b]$ partitions the remaining portion of $SSC_P[a,b]$ into two components: the bounded portion and the unbounded one. Similarly, the set $\{bd[SSC_P[a,b]\} \cap \{bd(P)\}$ partitions $bd(P)$ into two connected components: the set facing the bounded component of $SSC_P[a,b]$ and the set facing the unbounded component of $SSC_P[a,b]$. We denote the portion of $bd(P)$ facing the bounded component by $P_R[a,b]$ because we may compute it (as we shall see later) by "rotating" a separating plane between $[a,b]$ and P. First we establish an equivalence relation between $P_R[a,b]$ and $P_V[a,b]$.

Lemma 5.3 *For a given line segment $[a,b]$ and disjoint convex polyhedron P, $P_V[a,b] \equiv P_R[a,b]$.*

Proof: We prove the lemma by showing that each set is contained in the other, i.e., (i) $P_V[a,b] \subseteq P_R[a,b]$ and (ii) $P_R[a,b] \subseteq P_V[a,b]$.

(i) Let ν be a point in $P_V[a,b]$, i.e., on the boundary of P such that it sees $[a,b]$. Then $int(P)$ cannot intersect $int[\Delta(a,\nu,b)]$, for otherwise there would exist a point in $int[a,b]$ not visible from ν, a contradiction. Since both $\Delta(a, \nu, b)$ and P are convex sets with only point ν in common, it follows that there exists a plane through ν supporting both $\Delta(a, \nu, b)$ and P which separates $\Delta(a, \nu, b)$ from $int(P)$. Call this tangent plane H_ν. If H_ν also passes through a or b then H_ν belongs to $P_R[a, b]$. If H_ν does not intersect any point of $\Delta(a, \nu, b)$ other than ν then it can be rotated in any direction, while remaining tangent to P throughout the rotation, until it supports both P and $[a, b]$, thus ensuring that ν is in $P_R[a, b]$. Since ν is arbitrary it follows that $P_V[a, b] \subseteq P_R[a, b]$.

(ii) Let r be a point in $P_R[a, b]$. If $r \in bd(P_R[a, b])$ then by construction r admits a plane corresponding to a facet of the separating support cone $SSC_P[a, b]$ that supports P, contains r and separates $int(P)$ from $int[\Delta(a, r, b)]$. Therefore r sees $[a, b]$ and $r \in P_V[a, b]$. If $r \in int(P_R[a, b])$ then consider any plane H_r tangent to P at r. By the construction of $SSC_P[a, b]$ H_r must cut $SSC_P[a, b]$ into two components one of which is bounded and therefore H_r separates $int(P)$ from $[a, b]$ and hence from $\Delta(a, r, b)$. Therefore r sees $[a, b]$ and $r \in P_V[a, b]$. Therefore $P_R[a, b] \subseteq P_V[a, b]$. ■

Consider now the complexity of computing for a given line segment $[a, b]$ and disjoint convex polyhedron P the region $P_V[a, b]$ from which at each location we can see $[a, b]$ completely and where the solution for $\theta_{max}(\nu)$ can be found. By lemma 5.3 this problem is equivalent to computing $P_R[a, b]$. If we consider the more general problem where the segment $[a, b]$ is also a convex polyhedron then $P_R[a, b]$ may be obtained by computing all extreme separating planes of the two polyhedra. For this an $O(n)$ time algorithm due to Davis [9] is available which uses a projective transformation to convert this problem to a convex hull problem of two new disjoint convex polyhedra. However, since in our problem $[a, b]$ is only a line segment we can dispense with the projective transformation altogether and use a simpler $O(n)$ time algorithm which we describe next.

Assume for this description that no four vertices of P and $[a, b]$ are co-planar. This assumption places no limitations on our methods or complexity results and simplifies the presentation. Furthermore we assume without loss of generality that $[a, b]$ is vertical, i.e., parallel to the z-axis and that a lies above b. $P_R[a, b]$ is determined by all the separating tangent planes of $[a, b]$ and P. The extremal bounds of this set consist of two planes that contain $[a, b]$ and are tangent to vertices of P. Let these two vertices of P be denoted by t and t'. Let HU_b denote the intersection of the set of half spaces (determined by non-vertical planes) that contain P, such that their bounding planes lie above P, contain b and are tangent to P. Similarly, let HL_a denote the intersection of the set of half spaces (determined by non-vertical planes) that contain P, such that their bounding planes lie below P, contain a and are tangent to P. Then the intersection of HU_b and HL_a is equal to $SSC_P[a, b]$. Furthermore, $bd(HU_b) \cap P$ together with t and t' define a polygonal chain of vertices and edges of P, $T_b = (t', \ldots, t)$. Similarly, $bd(HL_a) \cap P$ together with t' and t define a polygonal chain of vertices and edges of P, $T_a = (t, \ldots, t')$. Therefore the concatenation of T_b and T_a encloses the region $P_R[a, b]$ and to compute $P_R[a, b]$ it suffices to compute T_b and T_a . Since P is represented by a DCEL data structure we may use the "rotating calipers" [25] to construct these chains. Consider computing T_b (the computation of T_a is analogous). First

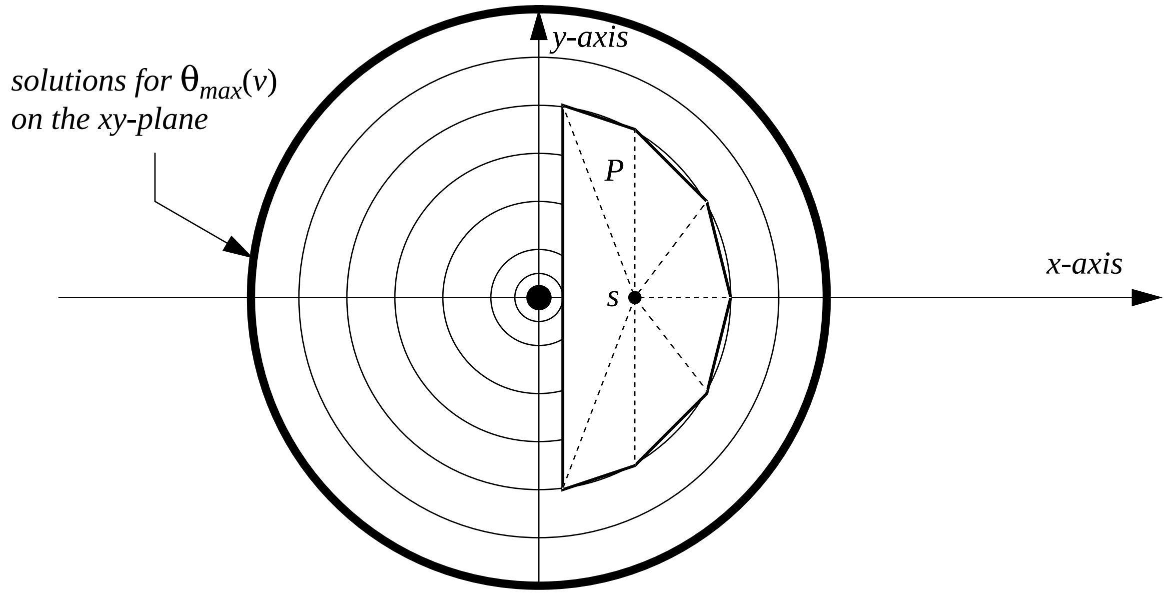

Figure 11: $\theta(\nu)$ may have $\Omega(n)$ local maxima realized at vertices of polyhedron P.

the points t and t' can be found in $O(n)$ time by finding the common tangents from the projection point of $[a, b]$ to the projection of polyhedron P, on the xy-plane. We initialize the rotating plane in the vertical position through a, b, and t. Then we rotate the plane in a counter-clockwise direction as viewed from b while maintaining the rotating plane tangent to $[a, b]$ and P. To find the next contact vertex of P we need only compute the dihedral angles (or functions monotonic to these) of the vertices adjacent to the vertex being considered. Since the DCEL allows us to find these in constant time per vertex, each rotation can be done in $O(d)$ time where d is the degree of the vertex being considered. We have therefore established the following result.

Lemma 5.4 *Given a line segment $[a, b]$ and a disjoint convex n-vertex polyhedron P, the region $P_R[a, b]$ may be computed in $O(n)$ time.*

Lemma 5.5 *Given a line segment $[a, b]$ and a disjoint convex n-vertex polyhedron P, $\theta(\nu)$ may have $\Omega(n)$ local maxima.*

Proof: We modify the example of Fig. 8 for the segment-polygon case in $\mathbf{R}^3$. Let the line segment $[a, b]$ be positioned vertically over the origin of the xy-plane, as before, but now at some fixed height above it so that the solution region for $\theta_{max}(\nu)$ when ν is constrained to lie on the xy-plane consists of the circle centered at the origin (heavy circle in Fig. 11) such that it contains the polygon in its interior. To convert the polygon to a polyhedron P we add a new vertex s that has xy-coordinates in the interior of the polygon and that lies very close to, but below, the xy-plane. Finally we connect s by edges to all the vertices of the original polygon that now is the top face of the convex polyhedron P. As the lunoid grows further its boundary surface will intersect the xy-plane in two circles. One will grow without bound. The other will eventually intersect simultaneously all the vertices of P but s. Thus each vertex of the top face of P will yield a local maximum. ∎

Theorem 5.6 *Given a line segment $[a, b]$ and a disjoint convex n-vertex polyhedron P, $\theta_{max}(\nu)$ may be computed in optimal $\Theta(n)$ time.*

Proof: By lemmas 5.1, 5.2, and 5.3, $\theta_{max}(\nu)$ must lie on a face of P that belongs to the set $P_V[a, b]$. Furthermore, each such face f_i has the property that segment $[a, b]$ lies on one side of the plane containing f_i. Therefore each face f_i of $P_V[a, b]$ is an instance of the segment-polygon problem considered in section 4. Therefore for each such face f_i we can use theorem 4.4 to obtain a candidate solution in $O(m_i)$ time, where m_i is the number of edges of f_i. The maximum of these candidates is the global solution. Since the summation of m_i over all i is $O(n)$, and by lemma 5.4 the computation time of $P_V[a, b]$ is $O(n)$, it follows that the entire algorithm runs in $O(n)$ time. Finally, lemma 5.5 allows us (by perturbing a vertex of P) to apply an adversary argument to obtain an $\Omega(n)$ bound on the complexity. ∎

5.2 Computing $\theta_{min}(\nu)$

We assume that the line $L(a, b)$ through the segment $[a, b]$ does not intersect polyhedron P. In this case the solution of $\theta_{min}(\nu)$ is zero and occurs at any intersection point. Furthermore this case can be dispensed with in $O(n)$ time by testing the intersection of $L(a, b)$ with every facet of P. In the previous section we saw that $\theta_{max}(\nu)$ may be realized by a point in the interior of some facet. On the other hand the $\theta_{min}(\nu)$ problem is fundamentally different in the sense that the solution must lie on the boundary of a facet of P. Furthermore, the solution need not lie on a vertex of P, as it necessarily must in the 2-dimensional version of this problem. In 3-dimensions the solution may lie in the interior of an edge of P.

Lemma 5.7 *$\theta_{min}(\nu)$ is determined by a point on an edge of P.*

Proof: (by contradiction) First we show that $\theta_{min}(\nu)$ is determined by a point on the boundary of P. Assume that $\theta_{min}(\nu)$ is determined by a point ν^* on the interior of a facet f of P. Construct a plane H determined by the three points $\{a, b, \nu^*\}$ and let P_H be the convex polygon which is the intersection of P with H. Two cases arise: either P_H is a facet of P or it is not. In the first case which arises if ν^* lies on a facet F that lies in H, we may consider the two-dimensional problem with ν^* in the interior of P_H. But then we can find a point on the boundary of P_H with smaller aperture angle, a contradiction. Therefore assume P_H is not a facet of P. In this case we may consider the two-dimensional problem with ν^* in the interior of an edge e of P_H. But then since the aperture angle function $\theta(\nu)$

is an upwards unimodal function on e, one of the endpoints of e must have smaller aperture angle than $\nu*$, a contradiction. ■

A natural question is whether $\theta_{min}(\nu)$ can be restricted further to a subset of the edges of P. We now answer this question in the affirmative by showing that $\theta_{min}(\nu)$ is determined by an edge of the "outer facets" of P with respect to segment $[a, b]$. First we make more precise what we mean by the outer facets of P. Let $B(P)$ denote the portion of $bd(P)$ on $CH(P \cup [a, b])$, where CH denotes the convex hull. We call the facets of $B(P)$ the outer facets of P with respect to $[a, b]$.

Lemma 5.8 *$\theta_{min}(\nu)$ is determined by a point on $B(P)$.*

Proof: (by contradiction) By lemma 5.7 $\theta_{min}(\nu)$ is determined by a point on an edge of P. Let $B'(P)$ denote the complement of $B(P)$, i.e., the part of $B(P)$ that lies in $int[CH(P \cup [a, b])]$. Now assume $\theta_{min}(\nu)$ is determined by a point ν on some edge of P in $B'(P)$. Then we may extend segment $a\nu$ in the direction of ν until it intersects $bd(P)$ at some point, say ν'. Now $\angle a\nu' b < \angle a\nu b$, a contradiction. Therefore $\theta_{min}(\nu)$ is determined by a point in the complement of $B'(P)$, i.e, $B(P)$. ■

Lemma 5.9 *$\theta(\nu)$ may have $\Omega(n)$ local minima.*

Proof: The proof of lemma 4.6 (Fig. 9) can be applied directly by converting the polygon to a polyhedron as was done in lemma 5.4 (see Fig. 11).

■

Theorem 5.10 *Given a line segment $[a, b]$ and a disjoint convex n-vertex polyhedron P, $\theta_{min}(\nu)$ may be computed in optimal $\Theta(n)$ time.*

Proof: By theorem 5.6 and lemma 5.7, $\theta_{min}(\nu)$ must lie on an edge of P that belongs to the set $B(P)$. Each edge of $B(P)$ is an instance of the problem solved in section 3.2 and can thus be solved in $O(1)$ time to give a candidate solution. Since there are $O(n)$ edges in $B(P)$ and $B(P)$ can be computed in $O(n)$ time using the "rotating calipers" the entire algorithm runs in $O(n)$ time. Finally, lemma 5.9 implies an $\Omega(n)$ lower bound via an adversary argument. ■

6 Conclusion

We have obtained theoretically optimal and exact algorithms for computing the maximum and minimum aperture angles that a line segment determines for a camera as it travels in a region bounded by a convex polyhedron in space. A natural question is whether the methods developed here are to be preferred over the approximate numerical global optimization methods traditionally used to solve optimization problems [13]. In practice the algorithms proposed here are of course not exact. Our methods are exact combinatorial decompositions of the original problems into $O(n)$ fundamental primitive computations that include: square roots, cube roots, trigonometric functions and inverse trigonometric functions. However, in practice a computer uses approximate numerical optimization methods to compute these primitives. Nevertheless, these simpler primitive problems (particularly square roots) are much better understood than the general global numerical optimization methods. Therefore from the practical point of view we have proposed new approximate numerical optimization methods for solving these problems. Which of the two approaches is more accurate, stable and faster in practice remains to be empirically determined.

Much theoretical work also remains unexplored. A line segment is the simplest object that determines an aperture angle. One may be interested in viewing more general objects such as convex polyhedra. A variety of generalizations of aperture angle are also possible. For example, one may be interested in the solid angle of the smallest circular cone that contains the object. Such problems remain open.

7 Acknowledgments

The authors would like to thank Chee Yap and Sue Whitesides for discussions and references on the theory of cubic equations and the complexity of computing their roots as well as Chris Paige and Henry Wolkowicz for discussions and references on numerical optimization techniques.

References

[1] D. Avis, B. Beresford-Smith, L. Devroye, H. ElGindy, E. Guevremont, F. Hurtado, and B. Zhu. Unoriented q-maxima in the plane. In *Proc. 16th Australian Conference in Computer Science*, Brisbane, 1993.

[2] P. Bose, L. Guibas, A. Lubiw, M. Overmars, D. Souvaine, and J. Urrutia. The floodlight problem. In *Fifth Canadian Conference on Computational Geometry*, pages 399–405, Waterloo, Canada, August 1993.

[3] P. Bose, F. Hurtado-Diaz, E. Omaña-Pulido, and G.T. Toussaint. Some aperture angle optimization problems. Technical Report No. SOCS 93.14, School of Computer Science, McGill University, December 1993.

[4] P. Bose, F. Hurtado-Diaz, E. Omaña-Pulido, and G.T. Toussaint. Some aperture angle optimization problems. In *892nd Meeting of the American Mathematical Society*, Polytechnic University, Brooklyn, N.Y., April 8-10 1994.

[5] B. Chazelle and D.P. Dobkin. Intersection of convex objects in two and three dimensions. *Journal of the ACM*, 34(1):1–27, January 1987.

[6] L.-L. Chen and T.C. Woo. Computational geometry on the sphere with application to automated machining. *Transactions of the ASME*, 114, June 1992.

[7] J. Colly, M. Devy, M. Ghallab, and L.-H. Pampagnin. Sensory control for 3d environment modeling. In *IARP Workshop on Multisensor Fusion and Environment Modeling*, Toulouse, 16-17 October 1989.

[8] C.K. Cowan. Model-based synthesis of sensor location. In *Proc. IEEE International Conference on Robotics and Automation*, pages 900–905, Philadelphia, April 1988.

[9] G. Davis. Computing separating planes for a pair of disjoint polytopes. In *Proc. ACM Symposium on Computational Geometry*, pages 8–14, Baltimore, Maryland, 5-7, June 1985.

[10] L. Devroye and G.T. Toussaint. Convex hulls for random lines. *Journal of Algorithms*, 14:381–394, 1993.

[11] H. Dörrie. *Triumf der Mathematik*. Physica-Verlag, Würzburg, 1958.

[12] H. Edelsbrunner. Computing the extreme distances between two convex polygons. *Journal of Algorithms*, 6:213–224, 1985.

[13] P.E. Gill, W. Murray, M.A. Saunders, and M.H. Wright. Linearly constrained optimization. In M.J.D. Powell, editor, *Nonlinear Optimization 1981*, pages 123–139. Academic Press, 1982.

[14] N. Jacobson. *Basic Algebra I*. W.H. Freeman & Co., San Francisco, 1974.

[15] K. Mehlhorn. *Multi-dimensional Searching and Computational Geometry*. Springer-Verlag, 1984.

[16] I. Niven. *Maxima and Minima Without Calculus*. The Mathematical Association of America, 1981.

[17] J. O'Rourke. *Art Gallery Theorems and Algorithms*. Oxford University Press, 1987.

[18] M.J.D. Powell. Introduction to constrained optimization. In P.E. Gill & W. Murray, editor, *Numerical Methods for Constrained Optimization*, pages 1–28. Academic Press, London, 1974.

[19] F.P. Preparata and M.I. Shamos. *Computational Geometry*. Springer-Verlag, 1985.

[20] A. Schwartz. *Analytic Geometry and Calculus*. Holt, Rinehart & Winston, 1960.

[21] T. Shermer. Recent results in art galleries. In *Proceedings of the IEEE*, volume 80, pages 1384–1399, September 1992.

[22] A.J. Spyridy and A.G. Requicha. Accessibility analysis for the automatic inspection of parts. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 1284–1289, Cincinnati, OH, USA, May 13-18 1990.

[23] K. Tang, T. Woo, and J. Gan. Maximum intersection of spherical polygons and workpiece orientation for 4- and 5-axis machining. *Journal of Mechanical Design*, 114:477–485, September 1992.

[24] M. Teichman. Wedge placement optimization problems. Master's thesis, School of Computer Science, McGill University, 1989.

[25] G.T. Toussaint. Solving geometric problems with the "rotating calipers". In *Proc. of IEEE MELECON*, Athens, Greece, May 1983.

[26] J.V. Uspensky. *Theory of Equations.* McGraw-Hill, 1948.

[27] N.B. Vasilev and V.L. Gutenmajer. *Rectas y Curvas.* MIR Publishers, Moscow, 1980. Spanish translation.

[28] T.C. Woo. Visibility maps and spherical algorithms. *Computer Aided Design*, 26(1):1–16, January 1994.