# Accelerating Geometric Queries using the GPU Adarsh Krishnamurthy,\* Sara McMains University of California, Berkeley Berkeley, CA, USA Kirk Haller<sup>†</sup> SolidWorks Corporation Concord, MA, USA

Figure 1: Geometric operations such as silhouette curve extraction and minimum distance/closest point computations between NURBS surfaces and complex CAD models accelerated using the GPU.

# Abstract

We present practical algorithms for accelerating geometric queries on models made of NURBS surfaces using programmable Graphics Processing Units (GPUs). We provide a generalized framework for using GPUs as co-processors in accelerating CAD operations. By attaching the data corresponding to surface-normals to a surface bounding-box structure, we can calculate view-dependent geometric features such as silhouette curves in real time. We make use of additional surface data linked to surface bounding-box hierarchies on the GPU to answer queries such as finding the closest point on a curved NURBS surface given any point in space and evaluating the clearance between two solid models constructed using multiple NURBS surfaces. We simultaneously output the parameter values corresponding to the solution of these queries along with the model space values. Though our algorithms make use of the programmable fragment processor, the accuracy is based on the model space precision, unlike earlier graphics algorithms that were based only on image space precision. In addition, we provide theoretical bounds for both the computed minimum distance values as well as the location of the closest point. Our algorithms are at least an order of magnitude faster than the commercial solid modeling kernel ACIS.

**CR Categories:** I.3.3 [Computer Graphics]: Hardware Architecture—Graphics Processors; I.3.3 [Computer Graphics]: Picture/Image Generation—Display Algorithms; I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric Algorithms, Languages, and Systems

Keywords: Silhouette Curves, Minimum Distance, Closest Point, Clearance Analysis, NURBS, GPU, Hybrid CPU/GPU Algorithms

# 1 Introduction

Geometric queries such as evaluating silhouettes and finding the minimum distance to a surface play an important role in many computer aided design and analysis applications that include tolerancing, visibility analysis, and accessibility analysis. Minimum distance queries are especially useful while designing complex assemblies to allow for sufficient clearance between different mechanical components. Such queries are easily answered if the objects or models are made of planar faces and have boxy shapes. However, modern designs make use of curved freeform surfaces; the standard representation of choice being Non-Uniform Rational B-Spline (NURBS) surfaces. Minimum distance queries on such freeform surfaces are currently being solved by commercial solid modeling software by first evaluating and tessellating the surface and then finding the minimum distance to the tessellation vertices [Spatial Corporation 2007]. This approach, in addition to being extremely slow and computationally intensive, is dependent on the tessellation resolution for the accuracy of the solution; the surface has to be very finely tessellated to get the required accuracy.

A technique to accelerate such slow geometric queries is to use programmable GPUs. We have developed a unified framework that uses GPUs as co-processors in accelerating geometric computations; we make use of the fragment processor in a GPU to perform parallel parts of the computations and use the CPU to perform the inherently serial parts. This framework can be extended to solve a wide range of geometric queries; we give a few practical examples of using this framework to answer distance and visibility queries. Previous GPU-based algorithms that render to the screen to per-

<sup>\*</sup>e-mail:{adarsh-mcmains}@me.berkeley.edu

<sup>&</sup>lt;sup>†</sup>e-mail:khaller@solidworks.com

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

<sup>2009</sup> SIAM/ACM Joint Conference on Geometric and Physical Modeling (SPM '09), October 4-9, 2009, San Francisco, CA. Copyright 2009 ACM 978-1-60558-711-0/09/10...\$10.00.

form these computations have restricted accuracy corresponding to the dimensions of the pixel or window. Our framework allows for the GPU algorithms to operate in the model space; therefore, the results of these geometric queries are accurate to any arbitrary userdefined tolerance.

Solid modeling kernels support certain distance queries such as the minimum distance from a point to a surface and the minimum distance between two surfaces. Applications of such distance queries include: finding the closest surface point on a surface to provide haptic feedback; dimensioning and tolerancing of CAD models; and constructing distance fields. In this paper we present an algorithm that uses our hybrid CPU/GPU framework consisting of surface bounding-boxes to accelerate these queries. We provide theoretical bounds on the accuracy of both the computed minimum distance as well as the location of the closest point on the surface, which allow for arbitrary user-defined tolerance values. This is especially important in CAD systems since these distances might be used by the designer to define subsequent features; the model might fail to regenerate if there is an error in the computed distance.

Silhouette curves are defined as a curve of points on the surface whose normal is perpendicular to the view direction. They are used in many visibility and accessibility problems; for example, multiaxis machining, path planning, etc. to separate the visible patches of the surface from the occluded ones. However, silhouette curves depend on the view-direction and have to be recomputed with changes in the orientation of the model. In addition, the topology of the silhouette changes rapidly; hence, we cannot rely on coherence between the silhouette curves at different view directions to accelerate their computation. In our framework, we append the surface normal information to the bounding-box structure to mark points on the silhouette curves using the GPU, resulting in real-time extraction. We can interactively provide visual feedback on the screen while simultaneously computing the silhouette curves in the parametric space that is then used by the modeling system. In addition, we also guarantee theoretical bounds on the extracted curves that comply with user-defined tolerances.

In this paper, we provide a hybrid CPU/GPU framework that can be used to accelerate geometric computations. Our main contributions include:

- A GPU accelerated silhouette curve extraction algorithm that uses surface normals appended to surface bounding-boxes to perform the calculations. The silhouette curves are calculated in the model and parametric spaces while simultaneously being output to the screen in real time.
- A practical GPU algorithm to find the minimum distance to a surface given any point in space. We use our hybrid framework to compute the distances efficiently in parallel using the GPU.
- A fast algorithm that computes the minimum distance between two surfaces or between two solid models represented by B-reps, using bounding-box hierarchies on the GPU. Our algorithm is orders of magnitude faster than the commercial solid modeling kernel ACIS in calculating these distances.
- A unified framework that uses the GPU as a co-processor to improve the performance of algorithms used for solving geometric queries. This framework can be extended to accelerate several related queries that are based on properties of the underlying shapes such as normal or curvature.
- Theoretical guarantees for all of our geometric computations. They allow for user-defined tolerance values that are essential for integrating our algorithms in a CAD system.

### 1.1 Hybrid Framework

We present a hybrid framework that can use both the CPU and GPU to perform geometric computations. The main idea is to split the computations into serial and parallel stages as shown in Fig. 2. To perform the parallel operations on the GPU, we make use of the map-reduce parallelism pattern that consists of assigning the computations to separate non-communicating parallel threads [Mattson et al. 2004]. The inter-communication between the CPU and GPU is shown in Fig. 3. Once the computations are performed, the computed result can be used by the modeling system in the three different ways shown. Read-back is important for integrating the GPU algorithms with traditional modeling systems. In addition, since GPUs are designed for pipelining the data only in one direction from the CPU to the GPU for display, the method of read-back significantly affects the performance of hybrid algorithms. The most efficient method of read-back is reducing the results to a smaller set of values by using operations such as finding the maximum, minimum, sum, or by using non-uniform stream reductions [Sengupta et al. 2007; Blelloch 1990]. The second method is to directly display the output on the screen using the GPU. This is ideal for certain operations that require only visual outputs; for example displaying the evaluated NURBS surface directly. The last and the most expensive method is to read-back all the results from the GPU to the CPU; this might be required for certain computations where the result of a computation is required for further processing on the CPU.



Figure 2: Operation flow for performing geometric computations. The parallel operations are mapped and performed on the GPU while serial operations are performed on the CPU. The intermediate parallel output is reduced and read back to the CPU.

Our operations on the GPU fall into three main types. The first type includes parallel geometric computations that can be performed efficiently on the GPU. The output of such operations are usually numeric values that are then stored in the GPU as textures. If an operation produces more than one output value for each parallel operation, we can store them using separate channels of the same texture or using different textures. The second type of GPU operations is parallel search operations that give a binary output of 0 or 1 based on the type of search; these include operations such as bounding-box intersection tests, finding if a value lies within a given range, etc. The third type of operations is reductions that are performed in multiple passes on the GPU. GPU reductions can in turn be classified into two types. The first type, called standard reductions, include reducing the given input to a single value such as computing the sum, min, max, etc. Standard reduction operations are usually performed in  $O(\log n)$  passes and hence are very efficient. The second type of reductions, called non-uniform stream reductions, reduces the input to a smaller set of values. Non-uniform stream reduction operations are particularly important when the result of a reduction operation is not a single value but multiple values that satisfy a particular criterion. Since the positions of the output elements do not have any fixed correspondence with the positions of the input, the stream-reduction process is considered non-uniform.



**Figure 3:** Schematic showing our hybrid framework that extends traditional geometric computations to use the GPU as a coprocessor to perform some parts of the computations in parallel.

We make use of an O(n) GPU stream-reduction algorithm that we presented in previous work [Krishnamurthy et al. 2008; Krishnamurthy et al. 2009] to perform non-uniform stream reductions.

To perform geometric computations on NURBS surfaces or assemblies, we make use of a surface bounding-box structure to map the computations to the GPU. We make use of Axis-Aligned Bounding-Boxes (AABBs) constructed from an evaluated mesh of points on the NURBS surface to accelerate the computations [Krishnamurthy et al. 2008]. The main advantage of AABBs over Oriented Bounding-Boxes (OBBs) is that several geometric computations such as finding intersections and distances are simpler in the case of AABBs. This is especially important because the efficiency of GPU programs can be reduced dramatically with increases in the complexity of the parallel kernels that are used. The individual computational kernels for OBBs are more complex and contain many branching conditions; the GPU has to wait until the most computationally intensive branch of the kernel in a particular pass is completed before proceeding to the next pass. In addition, since OBB kernels make use of more temporary registers, the number of computations that can be performed simultaneously on the GPU (called fragments in flight) is reduced; it is difficult to hide the memory access latency in this case. Thus, we found that the advantage provided by tight OBBs is offset by the increase in complexity of the algorithms that use them. We achieve better results by using AABBs even if we must decompose the model to a finer resolution with AABBs than OBBs in order to maintain the same tolerance bounds.

# 2 Background and Previous Work

### 2.1 Related Work

Extracting the silhouette curve of an object has been used extensively for rendering [Appel 1968; Saito and Takahashi 1990], nonphotorealistic rendering, and hidden-line removal [Elber and Cohen 1990]. [Krishnan and Manocha 1997] make use of silhouette curves to perform global accessibility and visibility analysis by computing the visible portions of free-form surfaces from a given viewpoint. Silhouette curves have also been used extensively in toolpath generation algorithms [Balasubramaniam et al. 2003] and for NC code verification [Chung et al. 1998].

Extracting the silhouette curves for a complex model made of NURBS surface is a complex process and in addition, the silhouette curves have to be recomputed after a change in the view direction. [Everitt 2002] used the GPU to render the silhouette curves directly while visualizing the model. However, this method does not give the modeling system access to the rendered silhouette curves. [Johnson and Cohen 2001] introduced a method called spatialized

normal cone hierarchies to extract silhouettes from free-form surfaces. They bound the normal directions possible inside a given surface patch and construct a hierarchy of normal directions; they then traverse along this hierarchy to find portions of the surface that are perpendicular to the view direction.

Minimum distance computations are used by many algorithms that generate geometrical constructs such as Voronoi diagram and medial axis transforms. They are also used in path planning and robot motion planning [Gilbert et al. 1988] and for projecting points onto a patch of a CAD model [Henshaw 2002]. Minimum distance computations on curved NURBS surface are very time-consuming and hence, the commercial solid modeling system ACIS makes use of the tessellation of the surface to find the closest vertex or pair of vertices while performing tolerance analysis [Spatial Corporation 2007]. [Johnson and Cohen 1998] gave a unified framework for minimum distance computations, which was later extended to find the closest point for haptics applications [Nelson et al. 2005]. We use a similar method that uses AABBs to find the regions of the model that are likely to contain the closest points. However, the methods they describe were better suited for a serial CPU implementation, since they make use of the convex hull of the freeform surface to iteratively refine the search. In our adaptation, the distance computations and search operations are done in parallel, which is better suited for a GPU implementation. In addition, we also provide theoretical guarantees for the solutions we compute.

[Edelsbrunner 1985] prove that the minimum distance between two convex polygons can be computed in  $O(\log n)$ . However, it is a theoretical algorithm that has large time-constants in practice. [Quinlan 1994] extend minimum distance computations to non-convex objects by first performing a convex decomposition and then using bounding spheres for the convex pieces to create a hierarchy. However this method is not practical for dynamic geometries since the convex decomposition might be expensive. [Chen et al. 2008] compute the minimum distance between a point and a NURBS curve by subdividing the curve into portions that might contain the closest point. Many minimum distance algorithms use Bounding Volume Hierarchies (BVHs) to accelerate the computations. CPU algorithms usually make use of BVHs that are more complex than AABBs. [Gottschalk et al. 1996] make use of OBBs to perform distance computations. [Larsen et al. 2000] perform proximity queries using a construct called a sphere swept volume, which consists of a sphere swept over a point, line or a plane, as primitives of a BVH.

Collision detection and distance field computation are two problems that are closely related to minimum distance computations that have been effectively accelerated using the GPU. Occlusion queries on graphics hardware were used by [Govindaraju et al. 2003] to detect collisions of polygonal meshes in large environments. [Greß et al. 2006] solve the collision detection problem by generating a bounding-box hierarchy for deformable parameterized surfaces and then detect collisions by checking overlap between the boundingboxes using the GPU. [Sud et al. 2004] use the GPU to generate 3D distance fields by first slicing the model into 2D slices and by using culling and spatial coherence to reduce the number of distance computations in each slice. Recently, [Lauterbach et al. 2009] used the GPU to construct BVHs that can then be used to accelerate collision detection.

There has been only limited use of GPUs to perform geometric operations because they are restricted to image-space resolution if the computations are to be performed by rendering on the screen. [Agarwal et al. 2003] make use of the GPU to perform geometric computations on a stream of points by using point-line duality. They compute geometric properties such as diameter and width of a set of points. However, these algorithms are not stable for points that are very close and are limited to image-space resolution. [Hoff et al. 2001] use the GPU to perform fast proximity queries on 2D shapes using a pixel grid to perform distance computations, but their technique does not extend to 3D shapes. Recently, researchers at SINTEF accelerate spline intersections by using the GPU to test for intersections and iteratively subdivide the spline patches until a prescribed accuracy is attained [Briseid et al. 2006; Dokken et al. 2005].

### 2.2 NURBS Evaluation and Modeling

Our minimum distance computation and silhouette extraction algorithms build on our previous papers on GPU NURBS evaluation and modeling. We present a short outline of our GPU algorithms that were explained in detail in [Krishnamurthy et al. 2007; Krishnamurthy et al. 2008; Krishnamurthy et al. 2009]. In our NURBS evaluation paper, we developed a method to directly evaluate a mesh of points on a NURBS surface using the GPU. Our algorithm used a fragment program to evaluate a NURBS surface of arbitrary degree in several passes. After evaluation we have the NURBS surface as 4-component vectors—(x, y, z, w) coordinates—in space stored as a texture on the GPU. While rendering, we interpret these values stored in the texture as vertex coordinates using a Vertex Buffer Object (VBO) and display it directly on the screen.



Figure 4: Surface bounding-boxes constructed from points evaluated on a NURBS surface.

In our NURBS modeling work [Krishnamurthy et al. 2008; Krishnamurthy et al. 2009], we construct surface AABBs that enclose a surface patch having four adjacent surface points as corners (Fig. 4). As a first step in the construction, we find the minimum and maximum coordinates of the four adjacent surface points to fit an AABB. However, the AABBs constructed by this method do not guarantee that the surface patch lies completely inside the constructed bounding-box. In order to guarantee complete coverage of the surface patch, we find the maximum possible deviation K of a curved surface from the linearized approximation, and then expand the bounding-boxes in all the three dimensions by K (Fig. 5). The analytical expression for the factor that can be used to expand the bounding-boxes based on the surface curvature is given by [Filip et al. 1987]. They show that if a parametric  $C^2$  surface is evaluated at  $(n+1) \times (m+1)$  grid of points, the deviation of the surface from the piecewise linear approximation cannot exceed the constant Kdefined by Equations (1) - (4). We use this constant K in computing the bounds for our closest point algorithms.

$$M_{1} = \max_{\forall (u,v)} \left[ \max\left( \left| \frac{\partial^{2} x}{\partial u^{2}} \right|, \left| \frac{\partial^{2} y}{\partial u^{2}} \right|, \left| \frac{\partial^{2} z}{\partial u^{2}} \right| \right) \right]$$
(1)

$$M_{2} = \max_{\forall (u,v)} \left[ \max\left( \left| \frac{\partial^{2} x}{\partial u \partial v} \right|, \left| \frac{\partial^{2} y}{\partial u \partial v} \right|, \left| \frac{\partial^{2} z}{\partial u \partial v} \right| \right) \right]$$
(2)

$$M_{3} = \max_{\forall (u,v)} \left[ \max\left( \left| \frac{\partial^{2} x}{\partial v^{2}} \right|, \left| \frac{\partial^{2} y}{\partial v^{2}} \right|, \left| \frac{\partial^{2} z}{\partial v^{2}} \right| \right) \right]$$
(3)

$$K = \frac{1}{8} \left( \frac{1}{n^2} M_1 + \frac{2}{nm} M_2 + \frac{1}{m^2} M_3 \right)$$
(4)



**Figure 5:** We expand the AABBs by K in all three dimensions to guarantee that the surface patch is completely enclosed.

# 3 Geometric Queries on NURBS Surfaces

We present first geometric queries that are performed on individual NURBS surfaces and later in Sec. 5 extend them to complex objects made up of multiple curved surfaces.

### 3.1 Silhouette Curves Extraction

We first explain how we can extract silhouette curves in real time using our hybrid framework. Using our NURBS evaluator, we evaluate the surface normals at the evaluated points and store them as a separate texture on the GPU. We then send to the GPU the view vector v as a constant. Using the GPU, we compute the sign of the dot-product of v with each of the surface normals at the four vertices that are inside each bounding box in parallel (Fig. 6). We then test if the sign of any of the dot-products is different from the others; we mark the surface patch inside the bounding-box as a patch that contains part of the silhouette curve if they are different. We can check for the variation in sign by only checking the sign of the six possible products of the four dot-products taken two at a time; these operations reduce the number of branching conditions and hence, can be done efficiently in parallel using the GPU.



**Figure 6:** We calculate the dot-product of the view direction with the four surface normals of the points inside the bounding-box. We mark the surface as part of the silhouette curve if the sign of any dot-product differs from the others, as in this example.

Once we have marked all the bounding-boxes that contain the silhouette curve in the texture, we perform non-uniform stream reduction to get a list of bounding-boxes that contain the silhouette curve (Fig. 7). We find the median of the four points that are inside the bounding-box on the CPU and output it as a point on the silhouette curve. Once we have points on the silhouette curve, we use a greedy algorithm that connects closest points [Krishnamurthy et al. 2009] to fit a polyline. This polyline can then be used by the modeling system for further operations. We also simultaneously map the marked texture onto the surface, which renders real-time interactive updates of the silhouette curves. In addition, since each of our points are within the bounding-box, the points we compute on



**Figure 7:** Schematic showing our silhouette extraction algorithm that follows our proposed hybrid framework.

the silhouette curves have known tolerance values both in the model space and in the parametric space. If a higher accuracy is required in any particular region, we re-evaluate only that part of the surface at a higher resolution to meet the required tolerances.

### 3.2 Minimum Distance to a NURBS Surface

The next geometric query we accelerate using the GPU is computing the minimum distance and the closest point on a NURBS surface given any point in space. As a first step, we evaluate the NURBS surface as a grid of points using our NURBS evaluator and construct surface AABBs enclosing four neighboring points. Using these bounding-boxes and the input point, we calculate the range of distances to each bounding box as explained in Sec. 3.3.

Fig. 8 shows how our GPU closest point algorithm fits into our hybrid framework. We first use the GPU to compute the minimum and maximum distance to each AABB efficiently in parallel. These distances are stored using the red and green channels in a min/max texture on the GPU. We then perform a parallel reduction in  $\log n$  passes on the GPU to find the bounding-box with the minimum lower value for the distance range. We read back the range of this particular bounding-box. In the next pass, we use the upper-bound of this particular bounding-boxes. We use the GPU to perform a parallel search on the same min/max texture we computed in the first step to find all the bounding-boxes whose range lie within the upper-bound. This prunes the list of bounding-boxes to search for the closest point; we read back this smaller list by performing non-uniform stream reduction on the results of the search.

Once we read back the potentially close bounding-boxes, we approximate the surface patch inside each of the bounding boxes with two triangles formed from the evaluated surface points. We then find the distance to each of these triangles and finally choose the one with the minimum distance. We also find the point lying on the triangle that has the minimum distance as the closest point on the surface. We prove that the evaluated minimum distance and the calculated closest point lie within theoretical bounds based on the surface approximation and the distance of the closest point from the given point in Section 4.

#### 3.3 Minimum and Maximum Distance to an AABB

The first step of our minimum distance algorithm requires the computation of the minimum and maximum distance between a point and an AABB. Since we want to perform these computations in parallel for each AABB, the computations have to be efficient and



**Figure 8:** Schematic of our closest point algorithm showing the inter-communication between the CPU and GPU. The vertical bars represent the range of minimum and maximum distances from the point to the bounding box.

optimized for the GPU. The maximum distance can be computed in a straight-forward manner by finding the vertex of the boundingbox that is farthest from the given point. However, to compute the minimum distance, we not only need to find the minimum distance to the vertices of the AABB but also to the faces. The number of computations becomes prohibitively many if we have to check all the possibilities.

In order to efficiently compute the minimum and maximum distance, we make use of the fact that the bounding-boxes are axisaligned. This makes the calculations simpler and unified for computing both the minimum and maximum distance simultaneously (Fig. 9). For computing the maximum distance from a point Oto an AABB, we compute the maximum distance along each axis separately and finally take the  $L_2$  norm of the individual maximum distances to find the maximum distance (Equations (5) – (8)). However, if we extend the same method to compute the minimum distance, we have to make sure that the individual distance components are non-zero; if we directly subtract the half bounding-box widths, we will end up with negative distances. To overcome this, we take the minimum distance along a particular direction as zero if it is negative (Equations (9) – (12)).

$$x_{max} = D_{cx} + B_x \tag{5}$$

$$y_{max} = D_{cy} + B_y \tag{6}$$

$$z_{max} = D_{cz} + B_z \tag{7}$$

$$D_{max} = \sqrt{\left(x_{max}^2 + y_{max}^2 + z_{max}^2\right)}$$
(8)

$$x_{min} = \max(D_{cx} - B_x, 0) \tag{9}$$

$$y_{min} = \max(D_{cy} - B_y, 0)$$
 (10)

$$z_{min} = \max(D_{cz} - B_z, 0)$$
 (11)

$$D_{min} = \sqrt{(x_{min}^2 + y_{min}^2 + z_{min}^2)}$$
(12)

This formulation is efficient for GPU implementation, since it has the least number of branches (one for each max while computing the minimum distances). We implement these equation using a single fragment program and output the minimum and maximum distance to a texture using the red and green channels. Thus the minimum and maximum distances are computed simultaneously for all



**Figure 9:** *Efficiently computing the maximum and minimum distance between a point and an AABB. The example shown here is for the 2D case, but the method can be extended to 3D. See Equations* (5) - (12).

AABBs in parallel. We then use these min/max distances as the input texture for finding the minimum distance to a NURBS surface (Fig. 8) as explained in Sec. 3.2.

# 4 Theoretical Bounds for Minimum Distance Computations

In this section we give theoretical bounds for both the computed minimum distance and the location of the closest point on the curved surface given any point in space.

**Theorem 1.** (Minimum Distance Bound) *The computed minimum distance does not deviate from the theoretical minimum distance to the actual surface by more than the surface deviation value K*.

*Proof.* Let *O* be the point from which we want to find the minimum distance to a curved surface patch S showed in green in Fig. 10. Let  $A_1, A_2, A_3$  be three points (of the four points used to construct the bounding-box) evaluated on the surface. The surface can be approximated linearly by triangle  $A_1A_2A_3$ ; the maximum deviation of the linear approximation from the curved surface is K (Eqn. (4)). Let Q be the actual point closest to O on the curved surface and P'be the computed closest point on the triangle. Let P be the closest point to P' on the surface. Since Q is the closest point on the surface from O, OQ < OP. From triangle OPP', by applying triangle inequality to the sides, we get OP < OP' + PP'. Since the maximum deviation of the surface from the triangle is K, distance PP' < K. Combining these inequalities, we get OQ < OP' + Kor OQ - OP' < K. This shows that the distance OQ, the theoretical minimum distance, cannot be larger than the computed distance OP' by more than K.

Now, consider the point on the triangle that is closest to Q, call it Q'. In this case OP' < OQ' since P' is the closest point on the triangle from O. Again from triangle OQQ', we get OQ' < OQ + QQ' and QQ' < K since Q' is the closest point on the triangle from Q. Combining these three inequalities, we get OP' < OQ + K or OP' - OQ < K. This shows that the theoretical minimum distance cannot be smaller than the computed distance by K. Combining the minimum and maximum bound on the distance, we get |OP' - OQ| < K.

Thus, from Theorem 1, we know that the theoretical minimum distance is bounded to lie within the range (d - K, d + K), where dis the computed minimum distance. We now show how we use this bound to prove that the location of the closest point we compute is also bounded.



**Figure 10:** Illustration to prove the bound for the minimum computed distance. The actual surface is shown in green while the linearized approximation is shown in orange.

**Theorem 2.** (Closest Point Location Bound) *The maximum possible distance between the computed closest point and the theoretical one is*  $\sqrt{4Kd + K^2}$  *where d is the computed minimum distance to the surface.* 

Proof. From Theorem 1, the theoretical minimum distance cannot deviate from d by more than K, i.e.  $OQ \in [d - K, d + K]$ . We have two possible cases: the closest point P' computed on the plane lies inside the triangle used to approximate the surface or it lies on one of the edges of the triangle (see Fig. 11(a) and Fig. 11(b), which show a 2D cross-section). In the first case (Fig. 11(a)), the minimum distance bound restricts the theoretical closest point Q to lie in an annular region between spheres with center O and radii d + K and d - K (marked in blue). From our tessellation bound K, we know that the actual surface lies within a region of width 2Kcentered around the approximating triangle (marked in red). Thus the point Q lies in the intersection of these overlapping regions. The maximum possible distance P'Q in this intersecting region is  $\sqrt{4Kd + K^2}$ . In the second case (Fig. 11(b)), the approximating triangle is oriented at an obtuse angle with respect to  $\overline{OP'}$ . In this case, the maximum distance in the overlapping region occurs only when OP' is perpendicular to the triangle; for all other angles of rotation of OP', it is always less than  $\sqrt{4Kd + K^2}$  (please refer to the Appendix for a detailed explanation). Hence, the maximum possible distance between the computed closest point and the theoretical one is always  $\sqrt{4Kd + K^2}$ . 



**Figure 11:** Illustration to evaluate the bound for the computed closest point location when the closest point on the plane lies either (a) inside or (b) on the edge of the triangle approximating the surface.

Thus, both the minimum distance computed and the location of the closest point are bounded. We show in the Results section that these theoretical bounds translate to realistic values that are useful in practice. Next, we extend our minimum distance computations to compute minimum distance between two NURBS surfaces or two complex CAD objects represented as B-reps.

### 5 Clearance Analysis

### 5.1 Minimum Distance Between Two NURBS Surfaces

We use a method similar to finding the minimum distance from a point to a surface to find the minimum distance between two surfaces. However, we cannot use this method directly because the number of distance comparisons increase as  $O(n^2)$ , where *n* is the number of AABBs of each surface. Therefore, we make use of a method that uses bounding-box hierarchies to successively refine the number of potentially-close bounding-box pairs. We show that this approach, which is similar to a breadth-first search, can also be fit into our hybrid framework. We perform the search for potentially-close bounding-box pairs in parallel at each level using the GPU.



Figure 12: We perform minimum distance computation between two NURBS surfaces with the help of AABB hierarchies for both the surfaces. We compute a list of potentially close bounding-boxes at each level using the GPU and then refine on the CPU until we reach a set of potentially close bounding-boxes at the lowest level.



**Figure 13:** Plot showing the actual number of AABB pairs compared during a typical minimum distance computation. The number of pairs being tested in parallel remains almost constant after level 3 of the hierarchy. Note logarithmic scale used for the y-axis.

We first generate a bounding-box hierarchy by recursively combining four AABBs in a level to get a bigger AABB of the next higher level. Thus, we construct an AABB hierarchy starting with the surface bounding-boxes and finally reaching a single, level-0 bounding box. This operation can be effectively performed in  $O(\log n)$ passes using the GPU. We store the bounding-boxes in a manner that optimizes GPU storage space (Fig 12). This process is performed once after evaluating the surface and constructing the surface bounding-boxes. When the model is transformed, we fit new AABBs to the transformed bounding-boxes. However, we still store and use the original AABBs, since if we keep only the new AABBs after every transformation, the bounding-boxes will keep growing in size.

We compute the minimum distance between the surfaces by recursively going down the hierarchy and finding potentially-close bounding-boxes at the base level of the hierarchy. We start at level 1 of the hierarchy where we compute the minimum and maximum distance between four AABBs from surface 1 with each of the four AABBs of level 1 from surface 2; we compute 16 mini-



**Figure 14:** Computing the maximum and minimum distance between two AABBs. The equations are similar to the point-AABB distance case. See Equations (13) - (20)

mum and maximum distance pairs. The method used for finding the minimum and maximum distance between two AABBs is explained in Section 5.2. Once we compute the set of minimum and maximum distances, we prune those AABB pairs that are outside the range similar to our method described in Section 3.2. We get a list of potentially-close AABB pairs for this level of the hierarchy at the end of the search. We then use the GPU to map the next finer level of the hierarchy, in sets of  $4 \times 4$  AABB pairs, and repeat finding the potentially-close AABB pairs in the next finer level on the GPU (Fig. 12). Finally at the end of the recursion, we get a list of potentially-closest AABB pairs in the finest or highest level of the hierarchy of both the surfaces. Using a hierarchy to prune AABBs outside the range keeps the number of potentiallyclose AABB pairs almost constant. Fig 13 shows that the number of pairs to be tested increases at first and after level 3 remains almost constant at a few thousand potentially-close pairs. These computations can be done efficiently by the GPU in parallel at each level, as seen in the Results section.

Finally, once we obtain all the potentially-closest AABB pairs at the highest level, we compute the closest distance between the surface patches enclosed by these AABBs on the CPU since the list of pairs is usually small. We approximate each surface patch with two triangles and then compute the closest distance between the triangles. Similarly we also compute the pair of closest points that have the minimum distance between them.

### 5.2 Minimum and Maximum Distance Between Two AABBs

We extend our method described in Sec. 3.3 to compute the minimum and maximum distance between two AABBs (Fig. 14). Similar to the point case, we compute the minimum and maximum distance along each dimension and then calculate the overall minimum and maximum distances (Equations (13) - (20)). As before, if any component is negative while computing the minimum distance, we take that component as zero.

$$x_{max} = D_{cx} + B_{1x} + B_{2x} \tag{13}$$

$$y_{max} = D_{cy} + B_{1y} + B_{2y} \tag{14}$$

$$z_{max} = D_{cz} + B_{1z} + B_{2z} \tag{15}$$

$$D_{max} = \sqrt{\left(x_{max}^2 + y_{max}^2 + z_{max}^2\right)}$$
(16)

$$x_{min} = \max(D_{cx} - B_{1x} - B_{2x}, 0) \tag{17}$$

$$y_{min} = \max(D_{cy} - B_{1y} - B_{2y}, 0) \tag{18}$$

$$z_{min} = \max(D_{cz} - B_{1z} - B_{2z}, 0) \tag{19}$$

$$D_{min} = \sqrt{(x_{min}^2 + y_{min}^2 + z_{min}^2)}$$
(20)

These equations are implemented using a fragment program on the GPU; we output the values to the red and green channels of a texture. The distances are computed for all potentially-close AABB pairs at a particular level in parallel and are then used for finding the potentially-close AABB pairs in the next level as explained in Sec. 5.1.

#### 5.3 Minimum Distance Between Two Complex Objects

Finally, we extend our minimum distance computations between NURBS surfaces to complex objects made up of many NURBS surfaces. CAD systems have support for this query to give feedback about the clearance between the models in an assembly while the user is manipulating them. However, existing systems are not interactive due to long computation times for performing this query. We perform this query in two stages; in the first-stage we find a list of potentially close surface pairs and in the second-stage we find the minimum distance between the surfaces.

#### Voxel-based First Stage

In the voxel-based approach for the first-stage, we construct a grid of voxels in the region occupied by the object (Fig. 15). We then consider these voxels as individual AABBs to perform the minimum distance computation. We create the voxel representation of the model as a preprocessing step. We first overlay a regular voxel grid that covers the object completely. We then use the coarse tessellation of the object that is used for display to populate the voxel grid. For each triangle in the tessellation, we find the voxels that the triangle intersects and then add a reference in the voxel to the surface to which the triangle belongs. Thus each voxel has information about its minimum and maximum point extents that define the AABB and a list of surfaces that intersect it. Since this is done only once per object when the object is loaded for display, we perform this operation on the CPU. In addition, since this is a linear O(n)operation, where n is the number of triangles in the tessellation, it is fast.



**Figure 15:** A complex model and its voxel representation. We store the surfaces that intersect with a particular voxel to accelerate the minimum distance computation.

As a first step in finding the closest points, we find a set of potentially-close voxel pairs by performing a single pass of minimum distance computation. To perform this operation on the GPU, we map the voxels from the first object to the rows and the voxels from the second object to the columns of a 2D texture (Fig. 16). We compute the minimum and maximum distances for each voxel pair of the two objects and output these distances to the texture. This texture is then used to find the list of potentially-close voxel pairs that lie within the range of the closest voxel pair (as in Fig. 8). We perform non-uniform stream reduction to transfer address information of the potentially-close voxel pairs to the CPU. Since each voxel has information about the surfaces that pass through it, we can create a list of potentially-close surface pairs from these potentially-close voxel pairs. We also make sure that there are no duplicated entries in the surface pairs list, since the same surface can pass through many voxels in the potentially-close voxel pair list.



**Figure 16:** We map the list of voxels of one object to the rows and the other object to the columns of a 2D texture to compute the minimum and maximum distances between the voxels.

#### Surface-based Second Stage

In the second stage, we compute the minimum distance for each surface pair in the potentially-close surface list using our algorithm explained in Sec. 5.1. We can then output the minimum distance or clearance between the two objects as the minimum distance computed from all the surface pairs. We also output the points on each surface as the closest points on the two objects. Even though we use the coarse tessellation for constructing the voxel grid, we do not use it for the minimum distance computations. Our computations are performed using the NURBS surfaces directly and lie within the computed bounds. Hence, they are more accurate than only using the tessellation for the computations. In the Results section we show that our algorithm performs orders of magnitude better than commercial CPU systems tested.

### 6 Results

We timed our GPU-accelerated queries on a 2.66GHz CPU running Windows Vista with 4GB of RAM and an NVIDIA GeForce 9800GX2 GPU with 512MB graphics memory. We compare our timings to perform the geometric queries with those of the commercial solid modeling kernel ACIS (v18).

#### Silhouette Timings

We compared the time taken by ACIS to extract the silhouette curves from a bi-cubic NURBS surface made up of  $100 \times 105$  control points. We evaluated the silhouette curves using the standard ACIS tolerance of  $10^{-3}$  measured in the parametric space  $[0,1] \times [0,1]$ . The GPU accelerated algorithm, in addition to being at least an order of magnitude faster (Table 1), on average output more than 5 times the number of points on the silhouette curves than ACIS. This is because we not only maintain the tolerance on the position of the points but also on the spacing between them, which leads to a better polyline approximation of the silhouette curves.

View	Segments	ACIS Time (s)	GPU Time (s)	Speed-up	
(0, 1, 1)	2	0.802	0.0093	86x	
(1, 1, 0)	3	0.852	0.0091	94x	
(1, 0, 1)	2	0.890	0.0010	89x	
(1, 1, 1)	3	0.918	0.0090	102x	

**Table 1:** Time to extract silhouette curves of a NURBS surface from different views.

### **NURBS Minimum Distance Timings**

We timed our minimum distance computations between two curved NURBS surfaces by interactively translating as well as rotating one surface made of  $199 \times 33$  control points relative to the another surface made of  $100 \times 105$  control points (Fig. 1). Fig. 17(a) shows the interactive computation times recorded during the interaction; the computation times were less than 0.15 seconds for most positions, a near-interactive frame rate of 8 - 10 fps. Fig. 17(b) shows the distance and position tolerances computed corresponding to the runs in Fig. 17(a). Since these tolerance values are dependent on the model size, we report them as a fraction of the model size in order to make them consistent with tolerance definitions used by ACIS [Corney and Lim 2001]; a value of 0.01 corresponds to 1% of the model size. The model size is the length of the diagonal of the smallest AABB that will enclose the model.



Figure 17: Interactive times for evaluating the minimum distance between two NURBS surfaces and the corresponding distance and position tolerances scaled with respect to the maximum model size.

We recorded the time taken by ACIS to compute the minimum distance at some arbitrarily chosen locations of the NURBS surfaces relative to one another. We set the tolerance value for ACIS to be  $4 \times 10^{-2}$ , well looser than our position tolerances reported in Fig 17(b). Table 2 summarizes the results of our NURBS minimum distance computations. The GPU accelerated algorithm is at least two orders of magnitude faster than ACIS. This can be explained by the fact that ACIS first tessellates the object to get a dense mesh of points on the surface and then performs the minimum distance computation on these points. We on the other hand use our fast NURBS evaluator to evaluate the surface and construct surface boundingboxes in real time. In addition, we not only achieve better performance but also a higher accuracy; our results have theoretical bounds that are practical for use in a CAD system.

Position	ACIS Time (s)	GPU Time (s)	Speed-up
1	67.3	0.141	477x
2	68.7	0.097	708x
3	68.8	0.191	681x
4	64.5	0.203	318x

**Table 2:** Time for performing minimum distance computations between two NURBS surfaces.

### **Object Clearance Timings**

We performed object clearance computations using the CAD models listed in Table 3; the models are of approximately the same complexity as standard CAD models used in a mechanical assembly. We used a voxel grid of  $40 \times 40 \times 40$  to perform the first-stage of the minimum distance computations. The objects were also tessellated to a coarse level that is sufficient for display; the number of triangles in this tessellation is given in Table 3.

Minimum distance queries were performed between the object pairs shown in Table 4; the objects were randomly positioned with respect to each other to perform the queries using both ACIS and our GPU accelerated algorithm. It can be seen that the GPU accelerated algorithm is again at least an order of magnitude faster than ACIS. We can even perform these computations interactively with slightly less complex objects or if the user specifies a subset of an object to compute the minimum distance.

Object	Surfaces	Triangles
Scooby	157	72094
Toy Car	127	17170
Car Body	80	7134

**Table 3:** Complexity of the objects used for the minimum distance computations. The number of triangles shown is the default coarse level of tessellation used for display.

# 7 Conclusions

We have developed a hybrid framework that uses GPUs to accelerate geometric computations. We have shown that two important geometric queries, finding the minimum distance between models and silhouette curve extraction, can be effectively performed using our framework. Our algorithms have theoretical bounds and are based on object-space resolution instead of just image-space resolution. They make use of actual surface data and not just the tessellation, which make them independent of tessellation errors. We also show tremendous performance improvements over existing commercial CPU-based systems.

Our framework can be easily extended to solve other CAD operations such as intersection curve evaluation and collision detection. We find that having alternating serial and parallel stages and using the map-reduce motif for parallelism to be ideally suited for developing geometric algorithms that use the GPU. In addition, the parallel stages can be easily modified to be executed on a multi-core CPU in the absence of a powerful GPU. Our framework provides for maximum flexibility and optimized performance in developing fast geometric algorithms.

# Acknowledgments

We would like to thank NVIDIA and AMD for providing us with their hardware. This material is based upon work supported in part by SolidWorks Corporation, UC Discovery under Grant No. DIG07-10224, and the National Science Foundation under Grant No. 0547675.

Object1	Object2	ACIS		GPU		Improvement	
		Time (s)	Tolerance	Time (s)	Tolerance	Time	Tolerance
Toy Car	Car Body	9.73	$10^{-2}$	0.939	$2.5 \times 10^{-5}$	10x	408x
Scooby	Car Body	193.50	$10^{-2}$	1.252	$22.1 \times 10^{-5}$	155x	45x
Scooby	Toy Car	118.63	$10^{-2}$	1.667	$2.8 \times 10^{-5}$	71x	361x
Car Body	Car Body	17.34	$10^{-2}$	0.439	$2.2 \times 10^{-5}$	39x	463x
Scooby	Scooby	71.42*	$10^{-1}$	0.533	$21.7 \times 10^{-5}$	134x	462x

**Table 4:** *Time for performing minimum distance computations between different complex objects.* \**The tolerance in this case was reduced because otherwise the computations did not finish.* 



**Figure 18:** *The different cases that were used for timing the minimum distance computations.* 

### References

- AGARWAL, P., KRISHNAN, S., MUSTAFA, N., AND VENKATA-SUBRAMANIAN, S. 2003. Streaming geometric optimization using graphics hardware. In *11th European Symposium on Algorithms*.
- APPEL, A. 1968. Some techniques for shading machine renderings of solids. In *Proceedings of the AFIFPS Spring Joint Computing Conference*, vol. 32, AFIFPS, 37–45.
- BALASUBRAMANIAM, M., SARMA, S. E., AND MARCINIAK, K. 2003. Collision-free finishing toolpaths from visibility data. *Computer-Aided Design* 35, 4, 359 – 374.
- BLELLOCH, G. E., Ed. 1990. Vector Models for Data-Parallel Computing. MIT Press.
- BRISEID, S., DOKKEN, T., HAGEN, T. R., AND NYGAARD, J. O. 2006. Computational Science - Lecture Notes in Computer Science, vol. 3994/2006. Springer, ch. Spline Surface Intersections Optimized for GPUs, 204–211.
- CHEN, X.-D., YONG, J.-H., WANG, G., PAUL, J.-C., AND XU, G. 2008. Computing the minimum distance between a point and a NURBS curve. *Computer-Aided Design 40*, 10-11, 1051 – 1054.
- CHUNG, Y. C., PARK, J. W., SHIN, H., AND CHOI, B. K. 1998. Modeling the surface swept by a generalized cutter for NC verification. *Computer-Aided Design 30*, 8, 587 – 594.
- CORNEY, J., AND LIM, T. 2001. *3D Modeling with ACIS*. Saxe-Coburg.
- DOKKEN, T., SKYTT, V., HAGEN, T. R., AND NYGAARD, J. O. 2005. US Patent 20080259078 - Apparatus and Method for Determining Intersections.
- EDELSBRUNNER, H. 1985. Computing the extreme distances be-

tween two convex polygons. *Journal of Algorithms* 6, 2, 213–224.

- ELBER, G., AND COHEN, E. 1990. Hidden curve removal for free form surfaces. In *SIGGRAPH 1990*, ACM, 95–104.
- EVERITT, C. 2002. One-pass silhouette rendering with GeForce and GeForce2. White paper, NVIDIA Corporation.
- FILIP, D., MAGEDSON, R., AND MARKOT, R. 1987. Surface algorithms using bounds on derivatives. *Computer Aided Geometric Design 3*, 4, 295–311.
- GILBERT, E. G., JOHNSON, D. W., AND KEERTHI, S. S. 1988. A fast procedure for computing the distance between complex objects in three-dimensional space. *IEEE Journal of Robotics* and Automation 4, 2, 193–203.
- GOTTSCHALK, S., LIN, M. C., AND MANOCHA, D. 1996. OBB-Tree: A hierarchical structure for rapid interference detection. In *ACM SIGGRAPH*, ACM, 171–180.
- GOVINDARAJU, N. K., REDON, S., LIN, M. C., AND MANOCHA, D. 2003. CULLIDE: Interactive collision detection between complex models in large environments using graphics hardware. In ACM SIGGRAPH/EUROGRAPHICS Conference on Graphics Hardware, Eurographics Association, 25–32.
- GRESS, A., GUTHE, M., AND KLEIN, R. 2006. GPU-based collision detection for deformable parameterized surfaces. *Computer Graphics Forum* 25, 3, 497–506.
- HENSHAW, W. D. 2002. An algorithm for projecting points onto a patched CAD model. *Engineering with Computers* 18, 3, 265– 273.
- HOFF, K. E., ZAFERAKIS, A., LIN, M., AND MANOCHA, D. 2001. Fast and simple 2D geometric proximity queries using graphics hardware. In *I3D '01: Proceedings of the 2001 Sympo*sium on Interactive 3D Graphics, ACM, 145–148.



Figure 19: The three different cases that can arise when the closest point computed is on the edge of the triangle.

- JOHNSON, D., AND COHEN, E. 1998. A framework for efficient minimum distance computations. *IEEE International Conference on Robotics and Automation* 4, 3678–3684.
- JOHNSON, D. E., AND COHEN, E. 2001. Spatialized normal cone hierarchies. In *I3D '01: Proceedings of the 2001 Symposium on Interactive 3D Graphics*, ACM, 129–134.
- KRISHNAMURTHY, A., KHARDEKAR, R., AND MCMAINS, S. 2007. Direct evaluation of NURBS curves and surfaces on the GPU. In *ACM Symposium on Solid and Physical Modeling*, ACM, 329–334.
- KRISHNAMURTHY, A., KHARDEKAR, R., MCMAINS, S., HALLER, K., AND ELBER, G. 2008. Performing efficient NURBS modeling operations on the GPU. In ACM Symposium on Solid and Physical Modeling, ACM, 257–268.
- KRISHNAMURTHY, A., KHARDEKAR, R., MCMAINS, S., HALLER, K., AND ELBER, G. 2009. Performing efficient NURBS modeling operations on the GPU. *IEEE Transactions* on Visualization and Computer Graphics 15, 4, 530–543.
- KRISHNAN, S., AND MANOCHA, D. 1997. An efficient surface intersection algorithm based on lower-dimensional formulation. ACM Transactions on Graphics 16, 1, 74–106.
- LARSEN, E., GOTTSCHALK, S., LIN, M., AND MANOCHA, D. 2000. Fast distance queries with rectangular swept sphere volumes. *Proceedings of ICRA '00: IEEE International Conference* on Robotics and Automation 4, 3719–3726.
- LAUTERBACH, C., GARLAND, M., SENGUPTA, S., LUEBKE, D., AND MANOCHA, D. 2009. Fast BVH construction on GPUs. In *Proceedings of Eurographics 2009*, Eurographics Association.
- MATTSON, T. G., SANDERS, B. A., AND MASSINGILL, B. L. 2004. *Patterns for Parallel Programming*. Addison-Wesley.
- NELSON, D. D., JOHNSON, D. E., AND COHEN, E. 2005. Haptic rendering of surface-to-surface sculpted model interaction. In *SIGGRAPH '05: ACM SIGGRAPH Courses*, ACM, 97.
- QUINLAN, S. 1994. Efficient distance computation between nonconvex objects. In *Proceedings of IEEE International Conference on Robotics and Automation*, IEEE, 3324–3329.
- SAITO, T., AND TAKAHASHI, T. 1990. Comprehensible rendering of 3-D shapes. SIGGRAPH Computer Graphics 24, 4, 197–206.
- SENGUPTA, S., HARRIS, M., ZHANG, Y., AND OWENS, J. D. 2007. Scan primitives for GPU computing. In Symposium on Graphics Hardware, Eurographics Association, ACM, 97–106.

- SPATIAL CORPORATION, A. D. S. C. 2007. ACIS Geometric Modeler: User Guide, version 18.0 ed.
- SUD, A., OTADUY, M. A., AND MANOCHA, D. 2004. DiFi: Fast 3D distance field computation using graphics hardware. *Computer Graphics Forum* 23, 10, 557–566.

# Appendix

### **Detailed Explanation for Theorem 2**

We give a detailed explanation for our closest-point bound proved in Theorem 2. From Theorem 2, we know that there are two possible cases. In the first case, the closest point P' lies inside the triangle and the bound can be computed directly to be  $\sqrt{4Kd + K^2}$ . However, in the second case, to find the maximum possible value of P'Q, we have to consider all possible orientations of the triangle with respect to OP'. Let  $\alpha$  denote the angle the triangle makes with OP';  $\alpha$  can vary from 90° to 180° (the two extremes and a general case are shown in Fig. 19). Angle  $\alpha$  cannot be less than 90° because then P' will no longer be the closest point on the triangle. The angle subtended by P'Q at the center of the sphere, denoted by  $\theta$ , monotonically decreases from  $\theta_{max}$  to  $\theta_{min}$ , as  $\alpha$  increases from 90° to 180°. The values of  $\theta_{max}$  and  $\theta_{min}$  can be computed to be  $\sin^{-1}\left(\frac{\sqrt{4dK}}{d+K}\right)$  and  $\sin^{-1}\left(\frac{K}{d+K}\right)$  from Fig 19(a) and Fig 19(b) respectively.

Consider the general case when  $90 < \alpha < 180$ . P'Q can be computed to be  $\sqrt{(d+K)^2 + d^2 - 2d(d+K)\cos\theta}$  from the cosine rule on triangle OP'Q. P'Q will be maximized when the term  $2d(d+K)\cos\theta$  is minimized, since all the other terms in the expression are positive.  $2d(d+K)\cos\theta$  is minimized when  $\theta$  is the maximum possible value in the range  $[\theta_{min}, \theta_{max}]$ . Thus P'Q is maximized when  $\theta = \theta_{max}$ ; the extreme case is shown in Fig 19(a) with maximum value of P'Q again being  $\sqrt{4Kd+K^2}$  as shown in Fig 11(a).