

----- Original Message -----

Subject: Fwd: GD/SPM 2011 notification - #146  
Date: Tue, 7 Jun 2011 12:17:51 -0700  
From: Adarsh Krishnamurthy <adkrishnamurthy@ucsd.edu>  
To: Sara McMains <mcmains@me.berkeley.edu>, Iddo Hanniel  
<iddohanniel@gmail.com>

Begin forwarded message:

> From: mskim.at.snu.ac.kr@gmail.com  
> Date: June 7, 2011 9:52:12 AM PDT  
> To: adarshk@berkeley.edu  
> Cc: mskim.at.snu.ac.kr@gmail.com  
> Subject: GD/SPM 2011 notification - #146  
> Reply-To: mskim.at.snu.ac.kr@gmail.com  
>  
> Dear Adarsh Krishnamurthy -  
>  
> Congratulations! We are pleased to inform you that your paper  
>  
> 146 - GPU-Accelerated Hausdorff Distance Computation Between Dynamic Deformable NURBS Surfaces  
>  
> has been accepted to the 2011 SIAM Conference on Geometric and Physical Modeling (GD/SPM11) as a regular paper.  
>  
> We are attaching the final reviews of your paper below. Please take note of the following:  
> \* A revised version of your paper is due no later than July 1. You should respond to the comments offered by reviewers as much as possible. When si  
> \* The final version must not be longer than 12 pages, and should follow the Elsevier format for Computer-Aided Design.  
> \* For your paper to be published, at least one of the authors must attend the conference in order to present. Registration for the conference will c  
>  
> We will send specific information about the process for submitting your final version, copyright forms, abstract, etc. at a later date.  
>  
> Congratulations again, and we look forward to receiving the final version of your paper and seeing it presented at the conference.  
>  
> Sincerely,  
>  
> Chandrajit Bajaj, Stefanie Hahmann, and Myung-Soo Kim, Program co-Chairs  
> bajaj@cs.utexas.edu, Stefanie.Hahmann@imag.fr, mskim@snu.ac.kr  
>

# GPU-Accelerated Hausdorff Distance Computation Between Dynamic Deformable NURBS Surfaces

Adarsh Krishnamurthy<sup>a,\*</sup>, Sara McMains<sup>a</sup>, Iddo Hanniel<sup>b</sup>

<sup>a</sup>University of California, Berkeley  
Berkeley, USA

<sup>b</sup>Technion, Israel Institute of Technology  
Haifa, Israel

---

## Abstract

We present a parallel GPU-accelerated algorithm for computing the directed Hausdorff distance from one NURBS surface to another, within a bound. We make use of axis-aligned bounding-box hierarchies that bound the NURBS surfaces to accelerate the computations. We dynamically construct as well as traverse the bounding-box hierarchies for the NURBS surfaces using operations that are optimized for the GPU. To compute the Hausdorff distance, we traverse this hierarchy after culling bounding-box pairs that do not contribute to the Hausdorff distance. Our contribution includes two-sided culling tests that can be performed in parallel using the GPU. The culling, based on the minimum and maximum distance ranges between the bounding boxes, eliminates bounding-box pairs from both surfaces that do not contribute to the Hausdorff distance simultaneously. We calculate accuracy bounds for our computed Hausdorff distance based on the curvature of the surfaces. Our algorithm runs in real-time with very small guaranteed error bounds for complex NURBS surfaces. Since we dynamically construct our bounding-box hierarchy, our algorithm can be used to interactively compute the Hausdorff distance for models made of dynamic deformable surfaces.

*Keywords:* Hausdorff Distance, NURBS, GPU, Geometric Algorithms

---

## 1. Introduction

The Hausdorff distance is a useful measure of the similarity between geometric objects. There are many applications that benefit from an efficient computation of the Hausdorff distance, including shape matching [Alt et al., 1995], mesh simplification [Guthe et al., 2005b], geometric approximation [Varadhan and Manocha, 2006], and penetration depth calculation for physically based animation [Tang et al., 2009]. As a result, the Hausdorff distance computation has attracted considerable research attention in computer graphics, computational geometry, and geometric modeling. However, most previous work focused on computing the Hausdorff distance for polygons and polygonal meshes.

This paper introduces a GPU-accelerated algorithm for the computation of the Hausdorff distance for NURBS surfaces, which remain the de facto standard for CAD models. Our algorithm runs in real-time with very small guaranteed error bounds for complex NURBS surfaces. Since we do not rely on pre-processing the input, our algorithm can be used to interactively compute the Hausdorff distance for models made of dynamic deformable surfaces.

### 1.1. Contributions

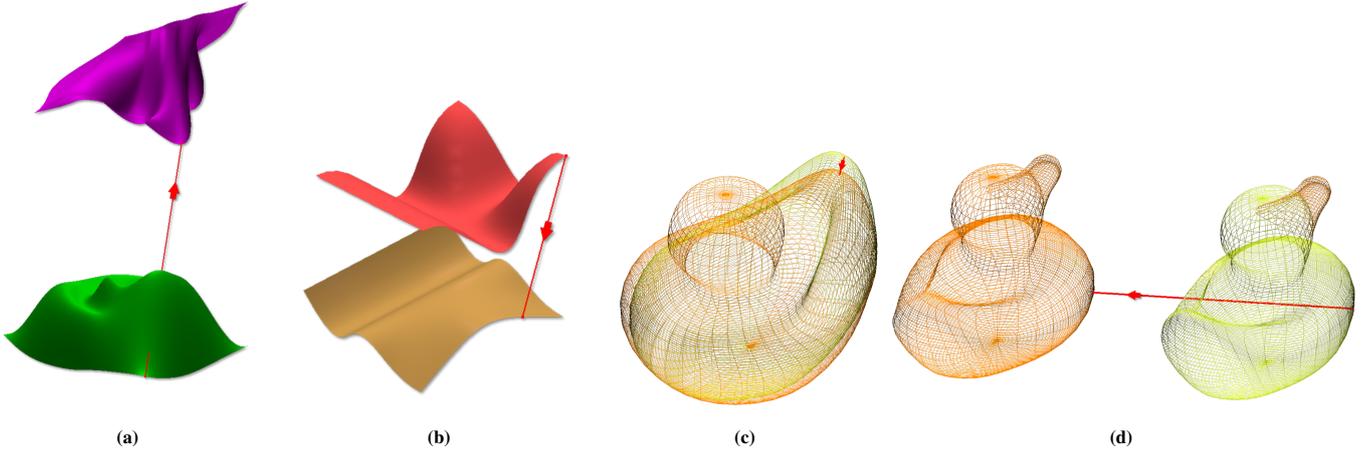
In this paper, we introduce a GPU-accelerated algorithm for computing the directed Hausdorff distance from one NURBS surfaces to another. Our main contributions include:

- A GPU-accelerated Hausdorff distance computation for NURBS surfaces. We perform GPU traversal of a bounding-box hierarchy to compute the Hausdorff distance after selectively culling parts of the surface that do not contribute to the Hausdorff distance. To the best of our knowledge, this is the first completely implemented algorithm to compute the Hausdorff distance for pairs of complex curved surfaces without requiring pre-tessellation.
- Novel culling tests that cull bounding-box pairs that do not contribute to the Hausdorff distance based on the range of minimum and maximum distance between the pairs. We perform two culling tests that bound both the minimum and maximum value of the Hausdorff distance for two-sided culling.
- Ability to efficiently compute a tight range for the Hausdorff distance and also the locations on the surfaces where the Hausdorff distance is within this bound.
- Theoretical bounds for the Hausdorff distance computations. We can guarantee user-defined tolerance values based on curvature-based bounds on the NURBS surfaces.
- Interactive computation of the Hausdorff distance between models made of dynamic deformable surfaces.

---

\*Corresponding Author

*Email addresses:* [adarshk@berkeley.edu](mailto:adarshk@berkeley.edu) (Adarsh Krishnamurthy),  
[mcmains@me.berkeley.com](mailto:mcmains@me.berkeley.com) (Sara McMains),  
[ihanniel@technion.ac.il](mailto:ihanniel@technion.ac.il) (Iddo Hanniel)



**Figure 1:** Computing the Hausdorff distance for NURBS surfaces; the arrows connect the locations on the two models where the Hausdorff distance is achieved. The tail of the arrow indicates the surface over which the minimum distances are maximized.

## 2. Related Work

The Hausdorff distance computation has attracted considerable research attention in computer graphics, computational geometry, and geometric modeling. Atallah [1983] developed a linear time algorithm for computing the Hausdorff distance between planar convex polygons. For simple non-convex polygons with  $m$  and  $n$  vertices, Alt et al. [1995] presented an algorithm, which is based on Voronoi diagrams, with  $O((n + m) \log(n + m))$  running time.

In  $\mathcal{R}^3$ , for polygonal meshes with  $O(n)$  triangles, Alt et al. [2003] presented theoretical and randomized algorithms. Their deterministic algorithm (a special case of a general theorem they prove for  $\mathcal{R}^d$ ) runs in  $O(n^5)$  asymptotic time. Using a randomized lower envelope algorithm, they arrive at an  $O(n^{3+\epsilon})$  randomized algorithm, and with minor additional requirements on the input they achieved a sub-cubic randomized algorithm. These algorithms are based on sophisticated data structures and algorithms and therefore are of a theoretical nature. To the best of our knowledge, they have never been implemented. Recently, Bartoň et al. [2010] presented an  $O(n^4 \log n)$  deterministic algorithm for computing the precise (up to floating point) Hausdorff distance between polygonal meshes. However, the high runtime complexity makes their implementation too slow (seconds) for large meshes.

Due to the complexity of computing the Hausdorff distance exactly, approximate practical algorithms have been proposed. Llanas [2005] proposed two approximate algorithms for non-convex polytopes, based on random covering. Guthe et al. [2005b] proposed an approximate algorithm for meshes that finds high-distance regions by performing an octree search. These regions are then further subdivided and regions that cannot attain the Hausdorff distance are purged away.

In a recent paper, Tang et al. [2009] implemented an approximate algorithm, which is similarly based on a Bounding Volume Hierarchy (BVH) that is computed during preprocessing. Using the BVH, many of the triangles are culled away based

on upper and lower bounds that are updated as the algorithm progresses. The triangles that pass the culling are further subdivided until the difference between the upper and lower bound on the Hausdorff distance is smaller than a user-defined tolerance. This results in an approximate distance, which is within the user-defined tolerance of the actual Hausdorff distance between the two meshes. Their implementation is very fast in practice, running at interactive speed for meshes of practical size.

On computing the Hausdorff distance between freeform surfaces, there are fewer previous results. For freeform curves, Alt and Scharf [2004, 2008] presented an algorithm for the Hausdorff distance computation, based on a characterization of the possible points where the distance can be attained. Chen et al. [2010] presented an algorithm for computing the Hausdorff distance between two B-Spline curves. The algorithm improves the one from Alt and Scharf [2008] by using a pruning technique to save computation time. Kim et al. [2010] presented a real-time algorithm for planar curves. The algorithm assumes that the planar curves are approximated by biarcs in a preprocessing stage and makes use of the GPU depth buffer to detect high distance regions efficiently. Recently, Bai et al. [2011] have computed an approximate Hausdorff distance between planar free-form curves by approximating the input curves with polylines and then computing the Hausdorff distance between the line segments.

Elber and Grandine [2008] extended the results from Alt and Scharf [2004] to Hausdorff distance computations, not only between planar curves, but also between curves and surfaces in  $\mathcal{R}^3$ . The performance of their algorithm is quite efficient for the case of planar/space curves. However, for the Hausdorff distance between a freeform surface and another curve or surface, the algorithm takes several seconds. To the best of our knowledge, this is the only published previous work that discusses the problem of this paper.

### 3. Overview

In this section, we first mathematically define the Hausdorff distance between two sets in  $\mathcal{R}^3$ . Subsequently, we present an overview of our algorithm highlighting the different steps. We then explain the steps in detail in the following sections.

#### 3.1. Mathematical Preliminaries

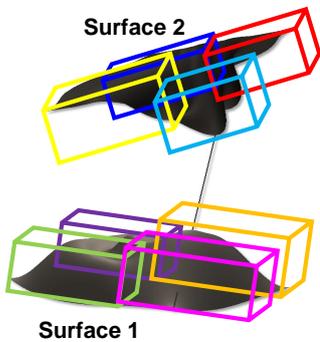
**Definition 1.** (Hausdorff Distance) Given two compact sets  $\mathcal{A}$  and  $\mathcal{B}$  in  $\mathcal{R}^3$ , the Hausdorff distance from  $\mathcal{A}$  to  $\mathcal{B}$  is given by Equation (1), where  $\mathbf{d}(\cdot, \cdot)$  is the Euclidean distance in  $\mathcal{R}^3$ .

$$\mathbf{h}(\mathcal{A}, \mathcal{B}) = \max_{\mathbf{a} \in \mathcal{A}} \left( \min_{\mathbf{b} \in \mathcal{B}} \mathbf{d}(\mathbf{a}, \mathbf{b}) \right) \quad (1)$$

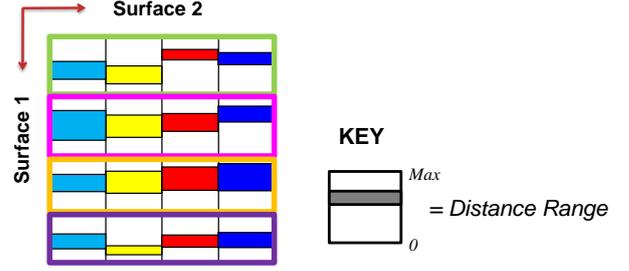
Sometimes this anti-symmetric Hausdorff distance is referred to as the “directed” or “one-sided” Hausdorff distance, but we will simply refer to it as the Hausdorff distance for the remainder of this paper.

#### 3.2. Algorithm Overview

Our algorithm to find the Hausdorff distance consists of three steps. We first construct bounding-box hierarchies (Section (4)) that enclose sub-patches of the NURBS surfaces using the GPU (Figure (2)). If the surfaces are not changing, this step can be performed during initialization. In the second step, we traverse this hierarchy on the GPU, building a virtual 2D array of min-max distance ranges for pairs of bounding-boxes from the two input surfaces (Figure (3)), starting at the coarsest level of the hierarchy. At each level, we perform two tests to cull pairs that could not contain sub-patches potentially contributing to the Hausdorff distance (Section (5.1)). For the first test, we calculate per-row Hausdorff distance maximums and cull ranges with a minimum distance value greater than the corresponding maximum; for the second test, we calculate a global Hausdorff distance minimum and cull rows containing ranges with a maximum distance value less than this minimum. An efficient implementation of this hierarchical culling with virtual array packing is detailed in Section (5.2). Finally, in the last step, we compute the Hausdorff distance between the surface patches that lie inside the potential bounding-box pairs in the finest level of the hierarchy.



**Figure 2:** Surface bounding-boxes constructed at the coarsest level of the hierarchy constructed for two sample input surfaces.

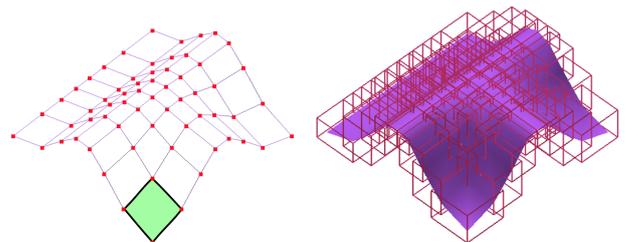


**Figure 3:** Virtual 2D representation of the min-max distances between the bounding-box pairs. The colors correspond to the bounding-boxes in Figure (2).

### 4. Hierarchical Bounding Boxes for NURBS

We employ an axis-aligned bounding-box (AABB) hierarchy for the NURBS surfaces as an accelerating data structure to compute the Hausdorff distance using the GPU. We chose AABBs over a more complex BVH such as spherical shells or oriented bounding boxes (OBBs) [Krishnan et al., 1998] 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 typically must wait until the most computationally intensive branch of the kernel in a particular block is completed before proceeding to the next block. In addition, since OBB kernels make use of more temporary registers, the number of computations that can be active (“in flight”) simultaneously on the GPU is reduced; it may be difficult to hide the memory access latency in this case. Thus, we found that the advantage provided by tight OBBs might be offset by the increase in complexity of the kernels that use them.

We evaluate the NURBS surface on the GPU in a uniformly sampled regular grid in parametric space and then construct AABBs that enclose the surface sub-patches that lie between each group of four adjacent surface points, similar to the method explained in Krishnamurthy et al. [2009] (Figure (4)). We used an optimized CUDA implementation of the surface point evaluation to evaluate the NURBS surface. (An alternate GPU NURBS evaluation method such as the one described by Kanai [2007] could also be used, but that approach requires writing different GPU kernels for different-order surfaces.)



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

Since the surface patches enclosed by the bounding boxes are curved, some part of the surface might penetrate out of the bounding box if we were to merely use the coordinates of the four evaluation points to construct the bounding boxes (see Figure (5)). To guarantee that the bounding boxes enclose their surface patches, we expand the bounding boxes based on the curvature of the surface patch. Using this method, the bounding boxes automatically become tighter when we evaluate the surface at a finer resolution.

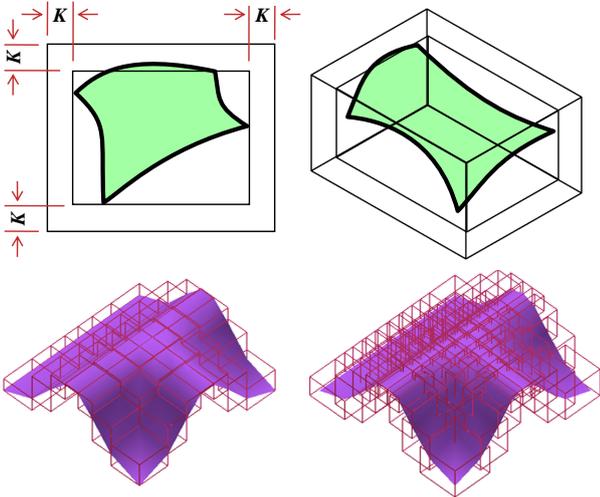
The analytical expression for the factor that can be used to expand the bounding boxes based on the surface curvature was originally given by Filip et al. [1987]. They show that if a parametric  $C^2$  surface  $\mathbf{S}$  is evaluated at  $(n + 1) \times (m + 1)$  grid points, the deviation of the surface from the piecewise linear approximation cannot exceed a constant  $K$ , which is a function of the magnitude of the maximum curvature vector on the surface, as defined by Equations (2)–(5). We use this  $K$  to derive bounds for our computed Hausdorff distance.

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

$$M_1 = \max_{\forall(u,v)} \left\| \frac{\partial^2 \mathbf{S}}{\partial u^2} \right\| \quad (3)$$

$$M_2 = \max_{\forall(u,v)} \left\| \frac{\partial^2 \mathbf{S}}{\partial u \partial v} \right\| \quad (4)$$

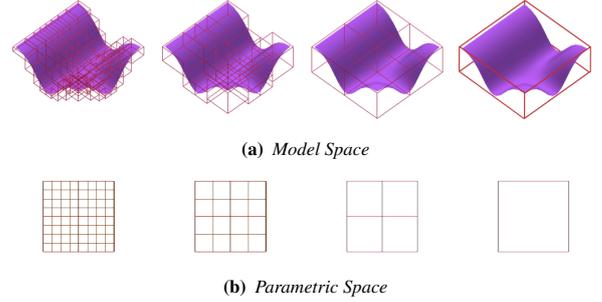
$$M_3 = \max_{\forall(u,v)} \left\| \frac{\partial^2 \mathbf{S}}{\partial v^2} \right\| \quad (5)$$



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

Once we construct the bounding boxes at the finest level of the hierarchy, we use a bottom-up approach to construct the BVH. We construct the bounding-box hierarchy by recursively combining four adjacent bounding boxes at a finer level to get the next coarser level. We repeat this operation until we get a single zeroth-level bounding box that encloses the whole surface. This bottom-up approach to hierarchy construction allows

the use of the tighter curvature bounds from the densest sampling throughout the hierarchy. Lauterbach et al. [2009, 2010] have recently developed a GPU algorithm where the BVH is constructed top-down, but a top-down approach is not as attractive for freeforms. Our approach allows using tight bounds since in a top-down method, one can only calculate very loose bounds because estimated local-curvature accuracy can only be refined after more evaluation points are available; thus the initial higher-level bounding boxes cannot be as tight as with a bottom-up approach.



**Figure 6:** We construct the bounding-box hierarchy by recursively combining four bounding boxes from the finer level to get the next coarser level.

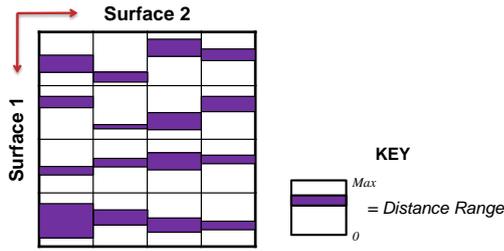
## 5. Hierarchy Traversal on the GPU

Once we have a bounding box hierarchy, we traverse this hierarchy on the GPU by selectively culling bounding-box pairs that cannot contain the surface sub-patches corresponding to the Hausdorff distance. To aid in the culling process, we allocate another array on the GPU that we will refer to as the min-max “distance matrix” that contains the distance ranges between pairs of bounding boxes, calculating the distance ranges in parallel using the GPU-friendly approach detailed in [Krishnamurthy et al., 2009]. If the bounding boxes overlap, we take the minimum distance between them as zero. We store the minimum and maximum distance in the global memory of the GPU, since the sizes of these arrays can be larger than what can fit into shared memory. We then use these min-max distances for computing the Hausdorff distance.

### 5.1. Culling Tests

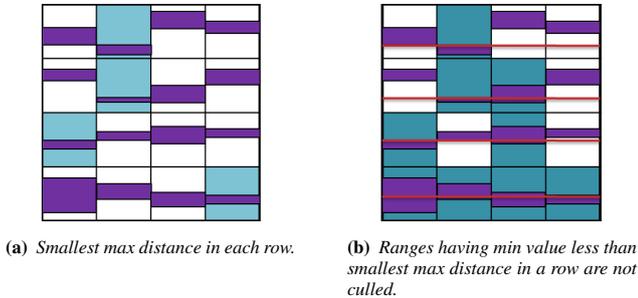
Bounding box hierarchy culling for Hausdorff distances is more complex than for traditional distance metrics because of the presence of both a minimum and a maximum in the formula (Equation (1)), which must be coordinated with the distance ranges for the bounding boxes. We will explain the culling portion of our algorithm using as a concrete example the level 1 (top) of the hierarchy, containing 4 bounding boxes from each surface (as illustrated in Figure (2)). We compute the min-max distances for all 16 pairs of bounding boxes and store them as a  $4 \times 4$  distance matrix (Figure (7) for this example). In this matrix, as in Figure (3), each row corresponds to distance ranges between a *surface1* sub-patch’s bounding box and the

bounding boxes of the different sub-patches of *surface2*; similarly, each column contains distance ranges between a *surface2* sub-patch's bounding box and the bounding boxes of the different sub-patches of *surface1*. Thus, each entry is the distance range for a pair of bounding boxes. If the entries of the matrix were exact distance values instead of ranges, to find the Hausdorff distance, we would only need to find the minimum value for each row and then find the maximum of these minimum values. However, since the entries are ranges, we cull those ranges that cannot contribute to the Hausdorff distance. We apply two culling tests to cull the bounding-box pairs that cannot contribute to the Hausdorff distance.



**Figure 7:** The minimum and maximum distance ranges for the bounding-box pairs. The range of distance is represented by the colored regions for each pair.

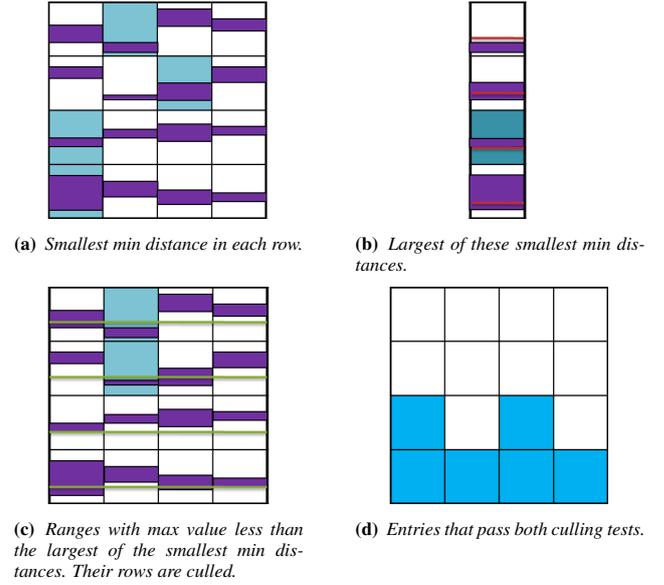
Keeping in mind that the Hausdorff distance is the maximum of the minimum distances from points on *surface1*, we first cull based on the bound on the maximum possible value of the minimum distance from each *surface1* bounding box, which corresponds to the maximum possible minimum distance per row. Since the actual minimum distance between a pair of bounding boxes can be anywhere in the range pictured, we can only exclude those bounding-box pairs whose range lies completely above the smallest maximum value in a row. To perform this culling, we first find the smallest maximum value in each row as shown by the shaded boxes in Figure (8(a)). We then use this value as a cut-off and cull all the bounding-box pairs whose minimum value is greater than the smallest maximum value in each row. In Figure (8(b)) the candidate bounding-box pairs that are not culled for the example case are shaded.



**Figure 8:** The first culling test.

The second culling test is based on a global bound on the minimum value of the Hausdorff distance. We calculate this bound by first finding the smallest minimum distance in each

row (Figure (9(a))). We then find the largest of these smallest minimum distance values as shown in Figure (9(b)). This value gives a lower bound on the Hausdorff distance. We use this bound to find all bounding-box pairs whose range lies completely below this lower bound value (Figure (9(c))). Such ranges indicate that the *surface1* bounding-box corresponding to that row has a smaller minimum distance to *surface2* than the Hausdorff distance lower bound calculated in Figure (9(b)), so it cannot contribute to the Hausdorff distance. Thus we can cull all pairs in such rows.

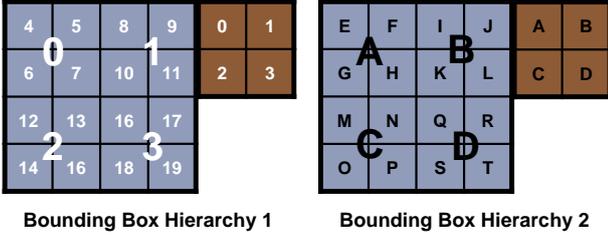


**Figure 9:** The second culling test. The Hausdorff distance is greater than the green line in Figure (9(c)), so its first two rows are culled. The shaded pairs in Figure (9(d)) pass both this and the previous culling tests.

We apply both these culling tests simultaneously and only those bounding-box pairs that pass both the tests are finally passed on for subdividing and further testing at the next level of refinement in the hierarchy. Figure (9(d)) shows the result of applying both the culling tests to the example case.

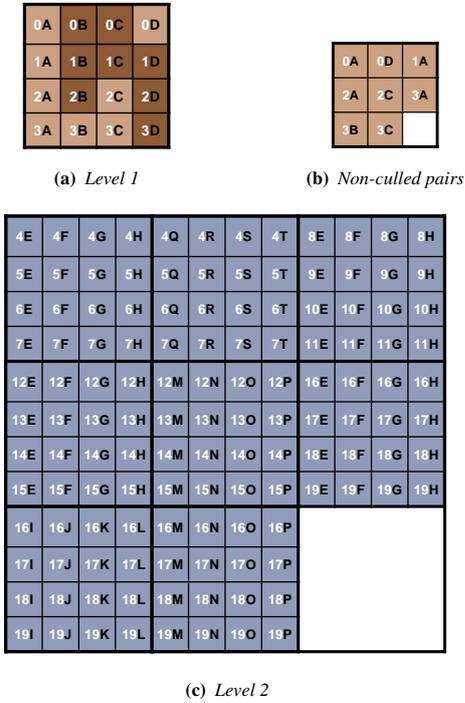
## 5.2. GPU Hierarchy Traversal Implementation

Our GPU implementation of the hierarchy traversal consists of two steps for each hierarchy level. In the first step, at any given level, we construct the min-max distance matrix consisting of the bounding-box pairs that remain potential candidates for the Hausdorff distance. At finer levels of the hierarchy, there will be too many bounding-box pairs to devote a separate row to each unique bounding-box from *surface1*. In addition, since more and more pairs will also be culled away, we must pack the remaining pairs into the matrix. In the second step, we apply the culling tests on these pairs. In order to perform the culling tests correctly on the packed min-max distance matrix, we also need to maintain bounding-box labels that indicate which original bounding-box of *surface1* defines each distance range, so that distance ranges from the same original *surface1* bounding



**Figure 10:** Two levels of the bounding-box hierarchy stored on the GPU. We refer to this hierarchy while explaining our hierarchy traversal algorithm.

box can continue to be grouped together (the row groupings in Figures (8) and (9)). We allocate an “address” array to aid in both the construction of the packed min-max distance matrix and the culling operation. Finally, we collect all the bounding-box pairs that are not culled and pass them to the next hierarchy level.

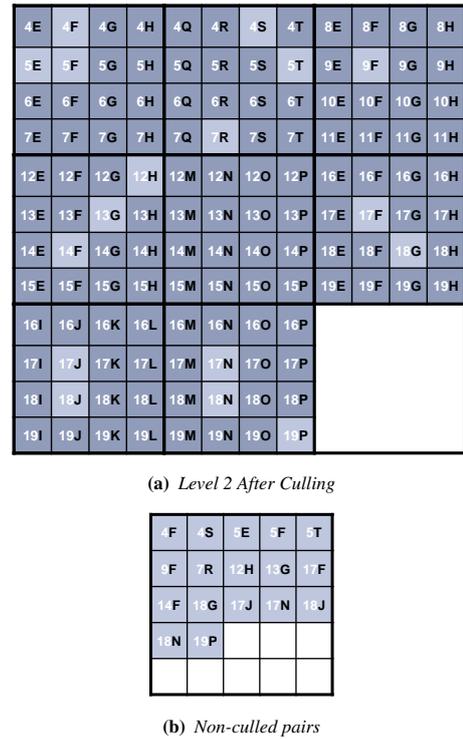


**Figure 11:** Address arrays used to aid in the hierarchy traversal. Based on bounding-box pairs not culled in level 1, the level 2 bounding-box pairs are constructed.

We will explain the two steps of the process using an example starting with the min-max distance matrix generation. Figure (10) is the original bounding box hierarchy from which the min-max distance matrix will be calculated. We start with the bounding-box pairs that are culled in level 1. In Figure (11(a)), the lighter-shaded boxes correspond to the non-culled bounding-box pairs at this level. We collect these pairs and pack them into a square array that serves as an address template to the bounding boxes in the hierarchies of both surfaces to perform the refinement (Figure (11(b))). Using this infor-

mation, and knowing the hierarchical layout pattern from Figure (10), we create the address array for the next level of the hierarchy (Figure (11(c))). Once we have the address array, we can construct the corresponding min-max distance matrix with the same structure. We compute the min-max range of distances between the bounding-box pairs specified by each entry of the address array.

We now proceed to perform the culling operations using this min-max distance matrix. In order to perform the culling, we need to compute the bounds for the Hausdorff distance (explained in Section (5.1)) based on this min-max distance matrix. For applying the first culling test, this is a “segmented reduction” of the min-max distance matrix to compute the smallest maximum distance for each group of pairs corresponding to the same *surface1* bounding box. In order to perform this operation on the GPU, we would first need to sort the bounding-box pairs with respect to their *surface1* bounding-box labels and perform the segmented reduction using segmented scan [Sengupta et al., 2007]. Both these operations are supported by *Thrust*<sup>1</sup>, a CUDA library of parallel algorithms with an interface resembling the C++ Standard Template Library (STL).



**Figure 12:** Creating the address array for the next finer level based on the non-culled bounding-box pairs in a level.

In our current implementation, we perform the “segmented reduction” operation on the CPU. We then create a new array with the smallest maximum distance value for each unique *surface1* bounding-box label and send it back to the GPU to perform the culling. The disadvantage of performing this step on

<sup>1</sup><http://code.google.com/p/thrust>

the CPU is the extra data-transfer step that could potentially be costly. On the other hand, by performing this operation on the CPU, we can avoid the sort operation by maintaining a hash table that is based on the *surface1* bounding-box label. We scan through all the min-max distance ranges and maintain the current bounds on the smallest maximum distance for each *surface1* bounding-box in this hash table. (Note that each min-max distance range is tagged with a *surface1* bounding-box label).

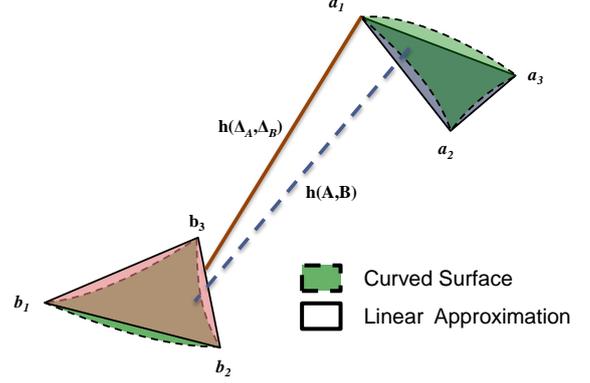
We also compute the global largest of the smallest minimum distance for the second culling operation (Figure 9(a)-9(c)) in a similar manner. We then perform the combined culling operation shown in Figure 9(d). After the culling operation, we again collect the bounding-box pairs that are not culled and refine them to the next level of the hierarchy (Figure 12(b)). (Note that by now, the references to *surface1* bounding boxes are not only out of order, but they are no longer contiguous; more coherence is lost at each step of refinement.)

Once we reach the finest level of the hierarchy, we get the list of bounding-box pairs that could potentially contribute to the Hausdorff distance. We then find the (approximated) Hausdorff distance between the surface sub-patches contained in these pairs of bounding-boxes by first decomposing the surface sub-patches into two triangles each, using the four evaluated points on each of *surface1* and *surface2*, and then finding the Hausdorff distance between the triangles. We find the minimum distance between the triangles (2 for each sub-patch resulting in 4 different triangles pairs) in each sub-patch pair in the list in parallel using the GPU. However, typically the list will have multiple entries for the same *surface1* sub-patch corresponding to pairing it with different *surface2* sub-patches. For each *surface1* sub-patch that is in the list, we group all its entries to find the minimum of all the calculated triangle minimum distances that correspond to that *surface1* sub-patch. To perform this operation, we again make use of a hash table based on *surface1* labels (note that these labels now correspond to individual sub-patches since we are in the finest level of the hierarchy). Finally, to find the Hausdorff distance, we find the maximum of these minimum distances using a standard reduction operation.

## 6. Theoretical Bounds for Hausdorff Distance

In this section, we make use of the curvature-based bounds for the NURBS surfaces to derive the theoretical bounds for our Hausdorff distance computations. Given a point in space, the error in computing the minimum distance to a linear approximation of a curved surface, instead of to the surface itself, is bounded by the maximum deviation  $K$  (Equation (2)) of the approximation to the surface. We can extend this bound to the Hausdorff distance from one surface to another.

Consider two surface patches  $\mathbf{A}$  and  $\mathbf{B}$ , having bounds  $K_A$  and  $K_B$  respectively, with respect to their linearized approximations. The surface patches cannot deviate by more than their respective bounds from any triangles,  $\Delta_A$  and  $\Delta_B$ , from their respective linearized approximations (Figure 13)). We shall prove that  $|\mathbf{h}(\Delta_A, \Delta_B) - \mathbf{h}(\mathbf{A}, \mathbf{B})| < K_A + K_B$  and therefore the Hausdorff distance found by our algorithm is bounded.



**Figure 13:** The Hausdorff distance we compute is within the sum of the distance bounds for each surface.

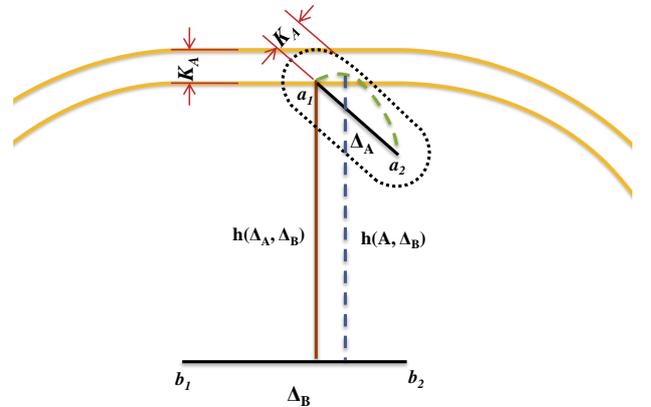
First, given any point  $a \in \mathbf{A}$ , the difference between  $\min_{b \in \mathbf{B}} \mathbf{d}(a, b)$  and  $\min_{b \in \Delta_B} \mathbf{d}(a, b)$  is bounded by the value  $K_B$ . Hence,  $|\mathbf{h}(\mathbf{A}, \Delta_B) - \mathbf{h}(\mathbf{A}, \mathbf{B})| < K_B$ .

We now bound the difference in the Hausdorff distances  $|\mathbf{h}(\Delta_A, \Delta_B) - \mathbf{h}(\mathbf{A}, \Delta_B)|$ . Following the notation in Alt et al. [2003], we first define a  $\delta$ -neighborhood (Equation (6)).

$$nh_\delta(Q) = \{p \in \mathbb{R}^3 \mid d(p, Q) \leq \delta\} \quad (6)$$

For a triangle  $\Delta$ , the  $nh_\delta(\Delta)$  is an offset of the triangle, consisting of the union of three  $\delta$ -radius balls centered at the vertices, three  $\delta$ -radius cylinders around the edges of the triangle, and a triangular prism of height  $2\delta$  containing  $\Delta$  in its center ( $nh_\delta(\Delta)$  can also be viewed as the swept volume of a  $\delta$ -ball over  $\Delta$ ). By definition, any point within  $nh_\delta(\Delta_B)$  has a minimum distance to  $\Delta_B$  that is smaller than  $\delta$  and any point on the boundary of  $nh_\delta(\Delta_B)$  has a minimum distance exactly  $\delta$  from  $\Delta_B$ .

Figure 14) shows a schematic 2D diagram of the triangles  $\Delta_B$  and  $\Delta_A$  and surface patch  $\mathbf{A}$  (shown as a dotted green line). We denote by  $hd_\Delta = \mathbf{h}(\Delta_A, \Delta_B)$ , the Hausdorff distance between the triangles. The  $hd_\Delta$  is always achieved at a vertex of



**Figure 14:** Figure to prove the bound on Hausdorff distance on surface  $\mathbf{A}$ ,  $|\mathbf{h}(\Delta_A, \Delta_B) - \mathbf{h}(\mathbf{A}, \Delta_B)| < K_A$ .

$\Delta_A$ , call it  $a_1$  (see Figure (14)). The orange curves in the figure (consisting of arcs and line segments) represent portions of the boundary of  $nh_{(hd_\Delta)}(\Delta_B)$  (the inner curve) and  $nh_{(hd_\Delta+K_A)}(\Delta_B)$  (the outer curve), i.e., the triangle  $\Delta_B$  offset by  $hd_\Delta$  and by  $(hd_\Delta + K_A)$  respectively.

By definition, triangle  $\Delta_A$  is completely contained in  $nh_{(hd_\Delta)}(\Delta_B)$  (with  $a_1$  on its boundary). Therefore,  $nh_{(hd_\Delta+K_A)}(\Delta_B)$  (which is an offset of  $nh_{(hd_\Delta)}(\Delta_B)$  by  $K_A$ ) contains all of  $nh_{(K_A)}(\Delta_A)$  (denoted by dotted lines in Figure (14)). This means that any point inside the neighborhood  $nh_{(K_A)}(\Delta_A)$  cannot have distance from  $\Delta_B$  that is greater than  $(hd_\Delta + K_A)$  and in particular  $\mathbf{h}(\mathbf{A}, \Delta_B) < \mathbf{h}(\Delta_A, \Delta_B) + K_A$ .

Thus we have:

$$\begin{aligned} |\mathbf{h}(\Delta_A, \Delta_B) - \mathbf{h}(\mathbf{A}, \mathbf{B})| &= \\ |\mathbf{h}(\Delta_A, \Delta_B) - \mathbf{h}(\mathbf{A}, \Delta_B) + \mathbf{h}(\mathbf{A}, \Delta_B) - \mathbf{h}(\mathbf{A}, \mathbf{B})| &\leq \\ |\mathbf{h}(\Delta_A, \Delta_B) - \mathbf{h}(\mathbf{A}, \Delta_B)| + |\mathbf{h}(\mathbf{A}, \Delta_B) - \mathbf{h}(\mathbf{A}, \mathbf{B})| &< \\ K_A + K_B, \end{aligned}$$

and our bound is established.

## 7. Results

We tested our Hausdorff distance computation algorithm on a 3.33GHz Intel CPU equipped with an NVIDIA GeForce GTX480 GPU running CUDA 4.0. Along with the Hausdorff distance, we also compute the theoretical guarantees based on the curvature of the surface. We tested our algorithms on several different NURBS surfaces that varied in complexity, at a large number of positions, since the algorithm is very sensitive to the relative positions of the input. Table (1) lists the specifications for all the NURBS surfaces tested and the short names used to refer to them later. The surface size, given as the maximum dimension of the smallest bounding-box that encloses it, is used to measure the relative accuracy of our computations.

Surface	Name	Control Points	Size
Green Surface	Green	298 × 313	501.14
Purple Surface	Purple	595 × 97	541.77
Red Surface	Red	6 × 6	722.84
Brown Surface	Brown	6 × 6	734.85
Yellow-DuckBody	Yellow	14 × 13	535.45
Orange-DuckBody	Orange	14 × 13	552.98

**Table 1:** Surfaces used for timing our Hausdorff distance algorithm. The surfaces are shown in Figure (1).

### 7.1. Culling Performance

One of the first measures of the performance of our algorithm is the number of pairs of bounding boxes that are culled using our culling tests. The culling performance that is summarized

in Table (2) is for the surface pair (Green-Purple) at positions shown in Figure (1(a)). The first column in Table (2) gives the number of AABB-pairs that would need to be tested in each level to compute the Hausdorff distance without culling. We can see that we cull more than 99.5% from the fifth level of the hierarchy on. This is typical of our culling performance.

Hierarchy Level	Hierarchy Pairs	Non-culled	Culling %
1	16	16	0.000%
2	256	256	0.000%
3	4096	2377	41.968%
5	1×10 <sup>6</sup>	3791	99.638%
7	268×10 <sup>6</sup>	10803	99.996%
9	68×10 <sup>9</sup>	46920	99.999%

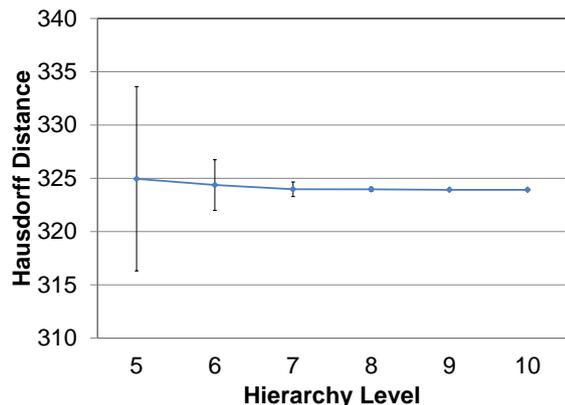
**Table 2:** The performance of our culling tests based on the level of the hierarchy of the surface. It can be seen that the culling rate improves with more levels of the hierarchy.

### 7.2. Accuracy

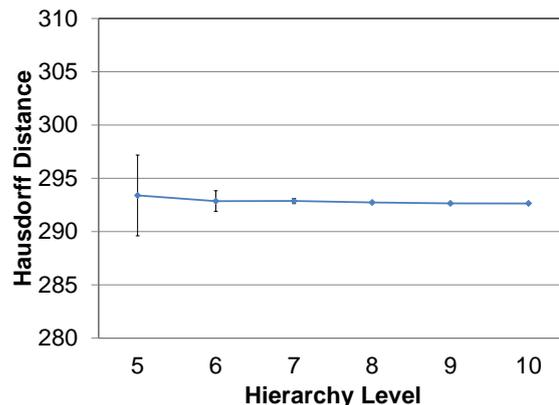
We make use of our guarantees to compute the accuracy of our Hausdorff distance computations for different surface pairs shown in Figure (1). The error bars on each plot in Figure (15) show actual values of our guarantees. We get extremely accurate results (error < 1% of input model size) even with just six levels of the hierarchy. Table (3) gives the relative Hausdorff distance bound as a percentage of the average size of the two input models at various levels of the hierarchy. The fourth column in Table (3) gives the equivalent number of triangles that are required in a tessellation of the same surfaces to achieve the same accuracy guarantees.

Hierarchy Level	Green-Purple Accuracy	Yellow-Orange Accuracy	Equivalent Triangles
5	1.643%	3.683%	2,048
6	0.451%	1.260%	8,192
7	0.131%	0.368%	32,768
8	0.038%	0.098%	131,072
9	0.009%	0.025%	524,288
10	0.002%	0.006%	2,097,152

**Table 3:** The guaranteed bounds of our Hausdorff distance algorithm based on the actual Hausdorff distance as a percentage of the average size of the two input models. The fourth column gives the equivalent number of triangles that would be required in a tessellated model of each input to achieve the same guarantees.



(a) Surface Pair Green-Purple



(b) Surface Pair Red-Brown

**Figure 15:** Accuracy of our Hausdorff distance computation based on the number of levels of hierarchy used for the computation. The error bars represent the guaranteed bounds calculated.

### 7.3. Timing Results

Figure (16) shows the computation time and the Hausdorff distance for the three different pairs of rigidly transforming surfaces (Green-Purple, Red-Brown, and Yellow-Orange), interactively moving the surfaces while computing to level 6 of the hierarchy. The timings shown include the time taken to construct the bounding-boxes and the hierarchy at each position, which takes from 5-10ms at this resolution. This time also includes the time taken to evaluate the second derivative and the bounds for the computation. It can be seen that the average time when the surfaces are far apart is less than 20ms. However, when the surfaces overlap, the computation time increases due to more unculled AABB-pairs to be tested at each level.

We also timed our algorithm while computing the Hausdorff distances between both rigidly transforming and deforming surfaces while varying the number of levels of hierarchy used in the computations. Table (4) summarizes the results for the surface pairs tested. Since the computation is very sensitive to the relative positions of the two surfaces, we report the average, minimum, and the maximum timings of about 100 different runs. To compute these timings for the first three non-deforming pairs, we initially positioned the surfaces with separation distance approximately equal to their respective sizes, then moved them closer to each other until they overlapped, and finally continued the motion until they were again separated in the opposite direction. For the final deforming pair, the ducks were placed as shown in the movie, with separation distance about equal to their size, and then the control points were manipulated with the mouse. It can be seen that on average, our algorithm takes less than a hundred milli-seconds for computing the Hausdorff distance to a very high accuracy.

### 7.4. Timing Breakdown

The time differences between the different steps of our algorithm are affected mainly by culling differences. The computation of the ranges for performing the culling and the actual

culling operations dominate the computation time in our current implementation at high resolutions (Table (5)). For overlapping surfaces, distance computation takes a larger proportion of the time at low resolutions (5 level BVH) because culling is not very effective. However, as resolution increases (> 7 level BVH), since culling often becomes more effective even as its running time increases, this tends to reduce the final distance computation time (occasionally more accurate results at higher resolutions are actually faster to compute!).

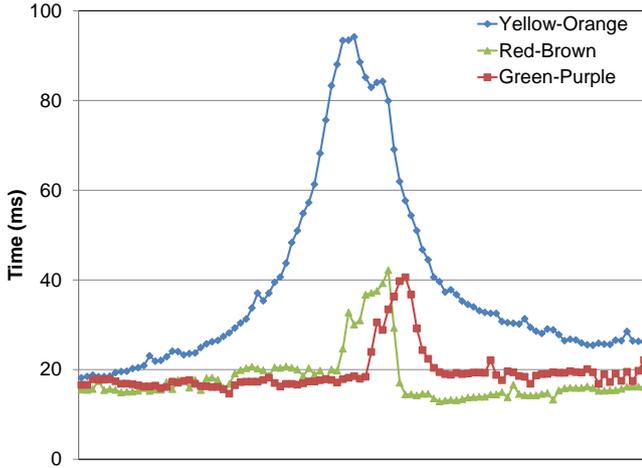
### 7.5. Comparison With Previous Algorithms

The algorithm by Elber and Grandine [2008] was implemented on the IRIT modeling environment <sup>2</sup> using the IRIT multivariate polynomial solver. In their paper, they report running times of several seconds for curve/curve Hausdorff distance in  $\mathbb{R}^3$ . For the problem of this paper, the surface/surface Hausdorff distance, there are no implementation examples in that paper, and there is only a partial implementation in IRIT, which computes only the Hausdorff distance attained at antipodal points (the most common case). The partial implementation takes several seconds to compute the Hausdorff distance between two surfaces.

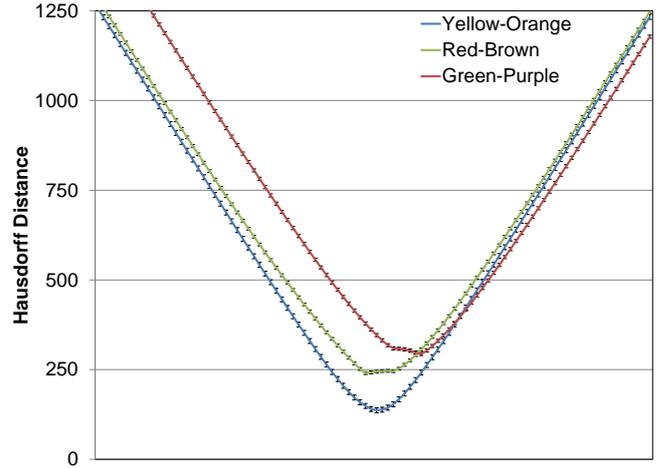
We cannot directly compare our performance with what is to our knowledge the best-performing approximate polyhedral Hausdorff distance software [Tang et al., 2009], since our algorithm requires no separate preprocessing to tessellate the surface into polygons or for one-time BVH construction. Hence, it can handle deformable models, unlike algorithms such as [Guthe et al., 2005a] and [Tang et al., 2009]. However, we can perform a rough comparison of performance on non-deforming models computed to similar accuracy as reported in the latter paper as follows.

Tang et al. [2009] computed results on a 2.6GHz Intel CPU for tessellated models with absolute error bounds of  $10^{-4}$  on model sizes that appear to be about 0.4 or 0.5, for an error of

<sup>2</sup>[www.cs.technion.ac.il/irit](http://www.cs.technion.ac.il/irit)



(a) Computation Time



(b) Hausdorff Distance

**Figure 16:** Computation time for computing the Hausdorff distance between different surface pairs while interactively translating them with keyboard inputs. The surfaces were initially far apart and were moved closer so that they overlapped; they then continued to be moved until they are far apart again on the opposite side. The error bars in Figure (16(b)) correspond to the calculated bounds.

Hierarchy	Green-Purple			Red-Brown			Yellow-Orange (Ducks)			Yellow-Orange Deform		
	Avg	Min	Max	Avg	Min	Max	Avg	Min	Max	Avg	Min	Max
5	13.43	7.82	22.85	12.40	9.13	22.81	26.29	15.40	57.89	40.95	36.01	55.13
6	17.86	10.66	38.21	20.03	13.27	44.39	39.39	19.82	93.77	41.07	32.78	73.91
7	27.34	16.41	60.41	37.21	23.10	75.82	72.51	27.17	227.02	79.01	49.71	159.73
8	55.33	29.34	114.21	72.47	37.19	154.22	176.14	45.73	671.69	106.48	55.18	352.41

**Table 4:** Total time taken in milli-seconds to compute the Hausdorff distance for the different pairs of surfaces.

Operation	Overlapping	Far
Evaluation, BVH	5.28	6.17
Min Max Distance Matrix	12.07	7.34
Culling Ranges Computation	73.23	27.44
Culling	51.82	23.79
Final Distance Computation	11.02	14.79
Total Time	175.85	96.02

**Table 5:** Breakdown of the timing in milli-seconds for different stages of our algorithm using a 7-level hierarchy. The overlapping surfaces are timed as shown in Figure (1(c)).

0.02–0.03% of model size. This is on top of tessellation error, which probably contributed more to the overall error given how many triangles were there in their input. The most challenging case they reported was for the Hausdorff distance between 16.7K and 69.7K triangle bunny models; this corresponds to the number of triangles in tessellated surfaces computed between hierarchy levels 6&7 and 7&8 respectively in our system. Positioning these surfaces with high overlap, their two-

sided Hausdorff distance computation took 948ms (not including pre-processing time to construct the BVH). Our worst-case timing computed at hierarchy level 7 was 227ms, including BVH construction, for the one-sided Hausdorff distance for the more severely overlapped Yellow-Orange DuckBody models (Figure (1(c))).

## 8. Limitations and Future Work

Our implementation of the hierarchy traversal is not optimized for performance. We have not compared the performance of our hybrid CPU-GPU implementation with a pure GPU implementation, which could avoid reading back the min-max distance matrix by using a segmented scan on the GPU.

Another limitation of our algorithm is that our error bounds will be large when the NURBS surface has a very high curvature. While it is true that higher curvature increases error, our bottom-up BVH construction allows the use of the local error derived at the finest level of detail throughout the hierarchy, and thus this effect is not very large in practice. For a top down approach, this would have been a more significant limitation. However, if the input parameterization is also very non-uniform, combined with high local curvature, the error bounds

will be very large. In such cases, the program could use the error bounds to flag the result as potentially ambiguous.

One potential application of our interactive algorithm is to geometric surface optimization problems. For example, if the desired final surface shape is known and the effect of forces on the input surface shape has been extensively modeled but not in an invertible manner, the question is how to find the initial surface shape that yields the final desired surface shape. For this application, measuring shape similarity with our algorithm rather than a Hausdorff distance algorithm that operates on tessellated meshes is preferred since the surface to be optimized is best represented using a high-level representation (such as control points in the case of NURBS) instead of a tessellated mesh. This is because the latter has far more degrees of freedom than optimal for the optimization. Our algorithm could efficiently evaluate the Hausdorff distance between the dynamic surface being optimized and the desired surface, with the Hausdorff distance functioning as the objective function to minimize. Finally, for directing the optimization, our algorithm could find all the surface locations within the error bounds where the “bounded” Hausdorff distance may be attained (unlike Hausdorff distance between meshes, which will typically only find a single location), which should significantly accelerate convergence.

## 9. Conclusions

We have presented the first practical, fully implemented algorithm to compute the Hausdorff distance from one NURBS surface to another, within computed bounds. Because our algorithm has no separate preprocessing, it can compute Hausdorff distance between dynamically deforming surfaces. Our method provides guaranteed bounds for the computations that can be used to match user-defined tolerance values. The user can also loosen the tolerance to increase interactivity or tighten it when accuracy rather than performance is critical. Our implementation is fast enough to calculate the two-sided Hausdorff distance interactively for most practical cases. Our implementation is both more accurate and has better performance than previous algorithms for computing Hausdorff distances between non-deforming, pre-tessellated models.

## Acknowledgments

The models used in the paper were downloaded from 3D Content Central [2009]. We would like to thank Wei Li for his comments. We would also like to thank the reviewers for their valuable suggestions. This material is based upon work supported in part by UC Discovery under Grant No. DIG07-10224, and the National Science Foundation under CAREER Award No. 0547675.

## References

- 3D Content Central, 2009. <http://www.3dcontentcentral.com>.
- Alt, H., Behrends, B., Blömer, J., 1995. Approximate matching of polygonal shapes. *Annals of Mathematics and Artificial Intelligence* 13 (3-4), 251–265.
- Alt, H., Braß, P., Godau, M., Knauer, C., Wenk, C., 2003. Computing the Hausdorff distance of geometric patterns and shapes. In: Aronov, B., Basu, S., Pach, J., Sharir, M. (Eds.), *Discrete and Computational Geometry. The Goodman–Pollack Festschrift*. Vol. 25 of *Algorithms and Combinatorics*. Springer, pp. 65–76.
- Alt, H., Scharf, L., 2004. Computing the Hausdorff distance between curved objects. In: *Proceedings of the 20th European Workshop on Computational Geometry*. pp. 233–236.
- Alt, H., Scharf, L., 2008. Computing the Hausdorff distance between curved objects. *International Journal of Computational Geometry and Applications* 18 (4), 307–320.
- Atallah, M. J., 1983. A linear time algorithm for the Hausdorff distance between convex polygons. *Information Processing Letters* 17 (4), 207–209.
- Bai, Y.-B., Yong, J.-H., Liu, C.-Y., Liu, X.-M., Meng, Y., June 2011. Polyline approach for approximating Hausdorff distance between planar free-form curves. *Computer Aided Design* 43, 687–698.
- Bartoň, M., Hanniel, I., Elber, G., Kim, M.-S., 2010. Precise Hausdorff distance computation between polygonal meshes. *Computer Aided Geometric Design* 27, 580–591.
- Chen, X.-D., Ma, W., Xu, G., Paul, J.-C., December 2010. Computing the Hausdorff distance between two B-spline curves. *Computer Aided Design* 42, 1197–1206.
- Elber, G., Grandine, T., 2008. Hausdorff and minimal distances between parametric freeforms in R2 and R3. In: *Proceedings of the 5th international conference on advances in geometric modeling and processing*. Springer-Verlag, pp. 191–204.
- Filip, D., Magedson, R., Markot, R., 1987. Surface algorithms using bounds on derivatives. *Computer Aided Geometric Design* 3 (4), 295–311.
- Guthe, M., Balázs, A., Klein, R., 2005a. GPU-based trimming and tessellation of NURBS and T-spline surfaces. *ACM Transactions on Graphics* 24 (3), 1016–1023.
- Guthe, M., Borodin, P., Klein, R., 2005b. Fast and accurate Hausdorff distance calculation between meshes. *Journal of WSCG* 13, 41–48.
- Kanai, T., 2007. Fragment-based evaluation of Non-Uniform B-spline surfaces on GPUs. *Computer-Aided Design and Applications* 4 (3), 287–294.
- Kim, Y.-J., Oh, Y.-T., Yoon, S.-H., Kim, M.-S., Elber, G., June 2010. Precise Hausdorff distance computation for planar freeform curves using biarcs and depth buffer. *The Visual Computer* 26, 1007–1016.
- Krishnamurthy, A., McMains, S., Haller, K., 2009. Accelerating geometric queries using the GPU. In: *SIAM/ACM Joint Conference on Geometric and Physical Modeling*. ACM, pp. 199–210.
- Krishnan, S., Gopi, M., Lin, M., Manocha, D., Pattekar, A., 1998. Rapid and accurate contact determination between spline models using shelltrees. *Computer Graphics Forum* 17 (3), 315–326.
- Lauterbach, C., Garland, M., Sengupta, S., Luebke, D., Manocha, D., 2009. Fast BVH construction on GPUs. *Computer Graphics Forum* 28 (2), 375–384.
- Lauterbach, C., Mo, Q., Manocha, D., 2010. gProximity: Hierarchical GPU-based operations for collision and distance queries. *Computer Graphics Forum* 29 (2), 419–428.
- Llanas, B., 2005. Efficient computation of the Hausdorff distance between polytopes by exterior random covering. *Computational Optimization and Applications* 30, 161–194.
- Sengupta, S., Harris, M., Zhang, Y., Owens, J. D., 2007. Scan primitives for GPU computing. In: *Symposium on Graphics Hardware*. ACM, Eurographics Association, pp. 97–106.
- Tang, M., Lee, M., Kim, Y. J., 2009. Interactive Hausdorff distance computation for general polygonal models. *ACM Transactions on Graphics* 28, 74:1–74:9.
- Varadhan, G., Manocha, D., 2006. Accurate Minkowski sum approximation of polyhedral models. *Graphical Models* 68 (4), 343–355.