

COMPUTER GRAPHICS TECHNICAL REPORTS

CG-2004-2

---

# Object-Space Interference Detection on Programmable Graphics Hardware

Alexander Greß

Computer Graphics, Universität Bonn. [gress@cs.uni-bonn.de](mailto:gress@cs.uni-bonn.de)

Gabriel Zachmann

Computer Graphics, Universität Bonn. [zach@cs.uni-bonn.de](mailto:zach@cs.uni-bonn.de)

Institut für Informatik II  
Universität Bonn  
D-53117 Bonn, Germany

© Universität Bonn 2004  
ISSN 1610-8892

# Object-Space Interference Detection on Programmable Graphics Hardware

Alexander Greß

Gabriel Zachmann

June 9, 2004

## Abstract

We present a novel method for checking the intersection of polygonal models on graphics hardware utilizing its SIMD, occlusion query, and floating point texture capabilities. It consists of two stages: traversal of bounding volume hierarchies, thus quickly determining potentially intersecting sets of polygons, and the actual intersection tests, resulting in lists of intersecting polygons. Unlike previous methods, our method does all computations in object space and does not make any requirements on connectivity or topology.

## 1 Introduction

Fast and exact collision detection of polygonal objects undergoing rigid motions is at the core of many simulation algorithms in computer graphics. In particular, virtual reality applications such as virtual prototyping need exact collision detection at interactive speed for very complex, arbitrary “polygon soups”. It is also a fundamental problem of dynamic simulation of rigid bodies, simulation of natural interaction with objects, haptic rendering, path planning, and CAD/CAM.

Currently, the performance of graphics hardware (GPUs) is progressing faster than general-purpose CPUs. The main reason is an architecture that combines *stream processing* [9] and SIMD processing. In addition, the programmability of the GPU has increased drastically over the past few years. Overall, today a programmer can write *kernels* for all stages of the graphics pipeline that are automatically executed in parallel on an indefinite number of processing units. This has led many researchers to investigate exploitation of the GPU for other computations,

such as matrix computations, ray tracing, distance field computation, etc.

Many algorithms have been proposed to utilize graphics hardware for the problem of collision detection. They can be classified into techniques that make use of the depth and stencil buffer tests, and those that compute discrete distance fields. In any case, the problem is approached in *image space*, i.e., it is discretized.

However, to our knowledge, no attempts have been made to utilize the GPU while still performing all computations in *object space*. This is what we address in this paper.

Based on techniques known from traditional CPU based collision detection approaches we develop a new method that utilizes the graphics hardware for hierarchical collision detection. Our algorithm simultaneously traverses a pair of bounding volume hierarchies performing all the necessary computations during this traversal, including the final triangle intersection tests, in vertex and fragment programs on the GPU. The algorithm has no requirements on the shape, topology, or connectivity of the polygonal input models.

## 2 Related Work

GPU based processing has become a trend over the past few years [13]. Generally, the idea is to formulate the given problem such that it can be solved by a number of rendering passes. During each of them, some of the computations are performed by rendering a number of geometric primitives, thereby updating one or more of the available buffers (stencil, color, z-buffer, etc.).

A clever way to utilize graphics hardware was presented by [11]. Based on the observation that an intersection can occur if and only if an edge of one object intersect the other one, they render edges of one object and polygons of the other.

This even works for deformable geometry. Unlike many previous approaches, objects do not need to be convex. However, they must still be closed. Furthermore, it seems to work robustly only for moderate polygon counts.

A hybrid approach was proposed by [7]. Here, the graphics hardware is used only to detect potentially colliding objects, while triangle-triangle intersections are performed in the CPU. While this approach alleviates previous restrictions on object topology, its effectiveness seems to degrade dramatically when the density of the environment increases.

The approach presented by [2] can compute the penetration depth using graphics hardware, but only for convex objects.

Earlier image-based methods include [17, 15, 3, 14, 4].

Virtually all image-based collision detection methods have several drawbacks in common:

1. Their complexity is usually in  $O(n)$ , which is much slower than the complexity of hierarchical, software-based approaches that seems to be in  $O(\log n)$ .
2. Rendering geometric primitives basically amounts to a discretization of the problem and thus introduces geometric errors. These errors depend more or less on the size of the viewport, the internal representation of numbers, and the number of bits per pixel in the z-buffer. The size of the viewport has significant impact on the performance.

Traditionally, rigid collision detection has been solved by simultaneously traversing a precomputed bounding volume hierarchy. A wealth of different BV hierarchies has been explored, such as sphere trees [8, 16], OBB trees [6], DOP trees [10, 19], Boustrees [20, 1], AABB trees [18, 12], and convex hulls [5].

Another rather recent effort is to design hardware specifically for the purpose of collision detection [21, 22]. In the present paper, however, we are concerned with utilizing commodity general-purpose graphics hardware.

To our knowledge, there is no previous work tackling the problem of collision detection in object space on the GPU.

### 3 Triangle Intersection Tests

Assume that two triangle meshes are to be checked for interference. Obviously, the brute-force approach is to check all triangles of the first mesh against all triangles of the second mesh for intersection. Before passing on to the hierarchical approach in the following section, we describe how to realize the straight-forward solution in programmable graphics hardware.

On programmable graphics hardware, rendering a rectangle of size  $m \times n$  corresponds to  $m \cdot n$  invocations of a fragment program. Theoretically, all fragment program invocations could be performed in parallel since the input required for each invocation may not depend upon the output of another fragment program invocation for the same rectangle primitive. This way, all available fragment program execution units can be utilized for such a task.

If we want to check all  $m$  triangles of a mesh against all  $n$  triangles of another mesh, the  $m \cdot n$  triangle intersection tests obviously do not depend on each other and thus can be performed efficiently in graphics hardware by rendering a rectangle of size  $m \times n$  using a fragment program that tests exactly one triangle pair at a time.

Suppose we want to determine the number of intersecting triangle pairs. This can be done using an occlusion query. (Occlusion queries are OpenGL 1.5 core functionality, and are also available in prior OpenGL versions on various graphics hardware via a corresponding extension.) The occlusion query returns the number of pixels actually written to the output buffer. If we use the fragment program to discard all fragments corresponding to non-intersecting triangle pairs, this number exactly corresponds to the number of intersecting triangle pairs. A fragment can be discarded conditionally using the *KIL* instruction in an ARB fragment program or using the *clip* instruction in the high-level shading languages of NVidia and Microsoft.

The input required for the intersection tests has to be stored in graphics memory. This can be done using floating-point textures available on current graphics hardware. An array of vertex positions is represented by a three-component floating-point 1D texture. Using a simple triangle soup representation, three such textures are used for each input model storing the three vertex positions of each triangle.

To be able to check two polygonal models specified in different object-spaces for interference, we need the transformation matrix from one object-space to the other one as further input. Since this matrix is constant for all  $m \cdot n$  intersection tests, it can be passed to the fragment program as program parameter (i.e., as uniform variable in the high-level shading languages). This way, the fragment program first transforms the three vertices of one triangle to the object-space belonging to the other triangle and then does the intersection test, discarding the fragment if there is no intersection. However, this way each triangle of the first mesh is transformed  $n$ -times.

Therefore, it is more efficient to move the transformation into the vertex program, as it is done in the following alternative solution. Instead of rendering a single rectangle primitive of size  $m \times n$  we render  $m$  horizontal line primitives of length  $n$ . This way the vertex program is invoked  $2m$ -times (as a line primitive consists of two vertices), and the fragment program is invoked  $n$ -times per line. The vertex program transforms a triangle of the first mesh into the object-space of the second mesh and passes the transformed triangle as fragment data to the fragment program. The transformation matrix is now a program parameter of the vertex program. Since a vertex program cannot yet access textures on current graphics hardware, the three arrays containing the vertex positions of the first mesh now have to be represented by vertex attribute arrays rather than by textures.

## 4 Hierarchical Interference Detection

Although the simple brute-force approach described in the previous section takes advantage of the parallel architecture of the GPU, it can clearly not outperform a clever hierarchical approach if large objects or scenes are to be tested. Therefore in this section, we propose a method for hierarchical interference detection on the GPU that is based on a bounding volume hierarchy.

### 4.1 Bounding Volume Hierarchy

As in traditional CPU-based approaches one object is to be checked for interference with another one by simultaneously traversing their two bounding volume hierarchies. We use axis-aligned bounding boxes (AABBs) as bounding volumes

since they are suitable for the GPU and still efficient [20].

Therefore, we generate an AABB tree for each object that consists of a bounding box at each inner node and a triangle at each leaf node. This generation is done in a preprocessing step on the CPU.

During the simultaneous traversal of two AABB trees  $S$  and  $T$ , all those pairs of nodes  $(S_i, T_j)$  are to be visited that are on the same hierarchy level in the corresponding trees and for which the parent nodes overlap. Various traditional CPU-based approaches use a depth-first traversal strategy. However, this way the decision whether a pair of nodes has to be visited depends on the result of an overlap test that was performed immediately before. Therefore, this strategy is not suited for execution using an indefinite number of vertex and fragment program units.

Instead, we use a breadth-first traversal scheme, i.e., all node pairs of a certain hierarchy level that have to be visited are processed before any node pair of the succeeding hierarchy level.

To be able to traverse the AABB trees efficiently, the trees have to be balanced. Furthermore, since leaf nodes will be handled differently than inner nodes by our algorithm, we require that there are no leaf nodes in the tree other than at the lowest hierarchy level, even if we construct a tree with a number of leaf nodes that is not a power of two. Therefore, we construct a tree where the lowest hierarchy level has exactly as many nodes as there are triangles in the input model. For any other hierarchy level  $L$  the number of nodes  $n(L)$  equals  $\lceil \frac{n(L+1)}{2} \rceil$ . This way, we yield an AABB tree where each inner node has one or two child nodes. Note, that there is at most one inner node at each hierarchy level that has just one child node.

When traversing two AABB trees simultaneously, in the following we assume that both consist of the same number of hierarchy levels. For two trees of different depths, this requirement can be achieved by adding further levels consisting of a single inner node at the top of one of the two trees.

### 4.2 Outline of the Algorithm

Since we traverse the tree breath-first and since at each hierarchy level only certain node pairs are to be visited, we have to store the indices of these node pairs temporarily during the traversal. For

this purpose we use a 2D buffer which we will refer to in the following as *node pair index map*.

This buffer contains a set of index lists as follows. Let  $L_j = \{i \mid \text{AABB}(S_i) \text{ overlap } \text{AABB}(T_j)\}$ . Putting  $L_j$  in the 2D buffer at row  $j$ , stored successively starting at the first pixel in this row, the complete buffer consists of  $m$  horizontal lines of different lengths. For each row, the length of the corresponding line (possibly 0) is stored in a vertex array (or, more precisely, its start and end points).

In addition, we require a second temporary 2D buffer, that we call *overlap count map*. This buffer consists of multiple levels, exactly as much as there are levels in the AABB trees. Each level consists, analogous to the node pair index map, of  $m$  horizontal lines of different lengths. The contents of the overlap count map are constructed at each hierarchy level  $L$  during the AABB tree traversal as follows.

At first we perform AABB overlap tests for the AABB node pairs that are to be visited at the considered level  $L$ . Each such node pair corresponds to one entry in the overlap count map at level  $L$ . If the AABB overlap test of a certain node pair was positive and thus the corresponding child nodes are to be visited when processing the next hierarchy level, the number of these child nodes is written into the corresponding entry of the overlap count map. Otherwise the entry of the overlap count map is set to 0. How this is done using the GPU is described in Section 4.3.

If this step results in a map containing only 0-entries (which can be determined using occlusion query), all AABB overlap tests have been negative, and therefore the two objects definitively do not collide.

Otherwise the AABB tree traversal is continued as follows. Before the iteration proceeds to the next hierarchy level, the node pair index map has to be updated as well as the vertex array containing the lengths of the horizontal lines contained in this map. First, the new vertex array is updated. In the same step also levels  $0, \dots, L-1$  of the overlap count map are updated. Second, the information contained in levels  $0, \dots, L$  of the overlap count map is used to construct the node pair index map required as input for processing the next hierarchy level. These two steps are explained in detail in Section 4.4.

The whole process is repeated for all hierarchy levels as long as there had been positive AABB

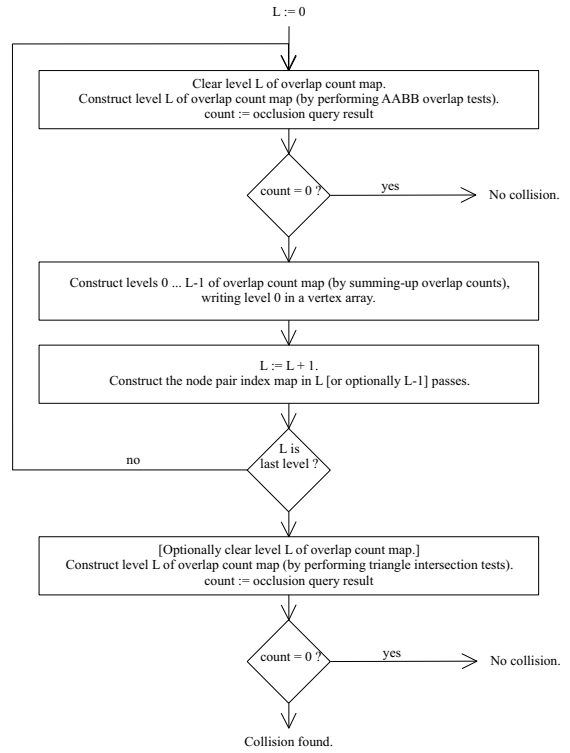


Figure 1: The outline of our approach.

overlap tests. If we reached the last hierarchy level, instead of testing AABB overlaps triangle intersection tests are performed on the GPU for all leaf node pairs that have to be visited according to the node pair index map. Using occlusion query, we obtain the number of intersecting triangle pairs for the considered objects. If required, the actual list of intersecting triangles can be obtained via read-back from graphics memory.

The outline of the overall algorithm is summarized in Fig. 1.

### 4.3 Box Overlap Tests

To perform the overlap tests of all AABB pairs corresponding to the indices contained in the node pair index map, we use a fragment program that is executed by rendering  $m$  horizontal lines, similar to the method described in Section 3. Here however, those lines are of different lengths, which are contained in the vertex array. The fragment program uses the index  $i$  contained in the node pair index map to obtain the AABB of node  $S_i$  from a pair of 1D floating point textures. One of these two textures contains the center point of the AABBs, while the other one contains the cor-

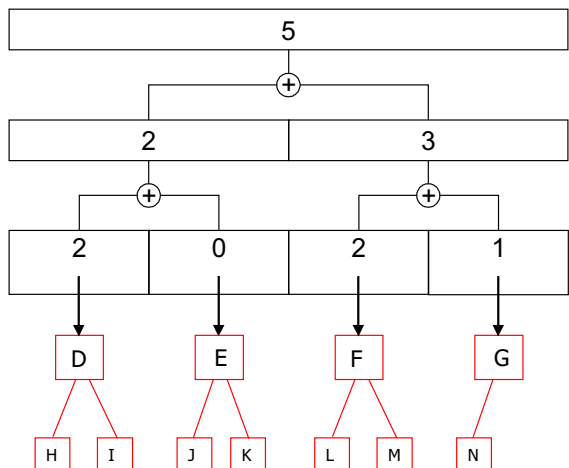


Figure 2: Summing up the overlap counts.

responding box extents. In addition, one of them contains also the number of child nodes of the tree node, which is 1 or 2.

This information is used in the fragment program to check whether the AABB pair overlaps by performing a *SAT lite* test [18]. This is shown in the fragment program subroutine found in Appendix A.

All fragments corresponding to non-overlapping AABB pairs are discarded by this program. Otherwise (not shown in the program), the number of children of  $S_i$  is written to level  $L$  of the overlap count map.

#### 4.4 Generating the Node Pair Index Map

We construct the new node pair index map and the corresponding vertex array in multiple passes using the following technique.

The lengths of the horizontal lines for the new node pair index map corresponds to the number of nodes of level  $L + 1$  for whose parent nodes the AABB overlap test was positive. By construction, this equals to the sum of all values in level  $L$  of the overlap count map at the corresponding row.

Therefore, we can construct the vertex array by summing up those values. In analogy to the construction of 1D MIP maps on the GPU, we do this by constructing the levels  $L - 1, \dots, 0$  of the overlap count map. Each level  $i = L - 1, \dots, 0$  of this map consists of  $2^i \times m$  entries, each of which is calculated by summing up two values from level  $i + 1$ , see Fig. 2.

The contents of level 0 of the overlap count map, which correspond to the totals sums of the values in each row of level  $L$ , are to be stored in a vertex array as required for the AABB overlap tests at hierarchy level  $L + 1$ . Using the upcoming `ARB_super_buffer` extension (which is currently under specification), it will be possible to render them directly into the vertex array. Another, less efficient alternative is to transfer the data from the texture directly to the vertex array using the `EXT_pixel_buffer_object` extension.

Next, we construct the new node pair index map for hierarchy level  $L + 1$  in  $L + 1$  passes. (Note, that the technique presented here is also suited for graphics hardware with dependent texture read limit, like the current ATI GPUs.)

The basic idea is that for every row of the node pair index map the  $n$ th entry corresponds to the  $n$ th AABB tree node of level  $L + 1$  that is to be visited. This node can be found by traversing the AABB tree starting at the root node. Note, that the first entry of level 1 of the overlap count map contains the number of nodes of level  $L + 1$  to be visited that are reached from the root node via its first child node. Therefore, depending on this value, it is clear whether we must proceed to the first or to the second child of the root node to reach the searched node. Then, this step is repeated using levels  $2, \dots, L$  of the overlap count map.

This technique to construct the new node pair index map is realized on the GPU as follows. We need a temporary buffer of the same size as the node pair index map that we are going to construct, consisting of two components per entry. The first component, called *current node index* in the following, is used to store the indices of AABB tree nodes that are visited during the traversal. The second component, called *current child index* in the following, corresponds to  $n$  if we search the  $n$ th node of level  $L + 1$  to be visited that is reached from the node specified by the current node index.

At the beginning, this temporary buffer is initialized as follows: In each row of the texture, the current child index is numbered consecutively starting with 0. The current node index is initialized to 0 and thus addresses the root node. Each row of the temporary buffer is assumed to be of the corresponding length stored in the vertex array.




Figure 3: Construction of the node pair index map in multiple passes. The values on gray background are the overlap counts that are compared to the current child indices shown at the top row.

We use a fragment program that does the following for each pass  $i = 1, \dots, L$ : First, the value from the overlap count map of level  $i$  at the index that equals  $2 \cdot \text{current node index}$  is read. Then, this value (*overlap count*) is compared against the current child index. If it is greater, the current node index is replaced by  $2 \cdot \text{current node index}$  in the temporary buffer, and the current child index remains unchanged. Otherwise, the current node index is replaced by  $2 \cdot \text{current node index} + 1$ , and the overlap count is subtracted from the current child index; see Fig. 3.

After these  $L$  passes one further pass is required to obtain the new node pair index map based on the values in the temporary buffer and the current contents of the node pair index map. For each entry of the map, we identify the actual corresponding AABB tree node by accessing the current contents of the node pair index map at the index specified by the current node index from the temporary buffer. Then we store the index of its first or second child, depending on the current child index from the temporary buffer, into the new node pair index map. This step is shown in the lower right of Fig. 3.

To avoid the additional render pass, the last operation optionally can be incorporated directly

into the fragment shaders used for box overlap and triangle intersection testing.

## 5 A Hybrid Approach Based on Temporal Coherence

It is clear that for a given object pair it might not be most efficient to start the hierarchy traversal at the top levels of the two hierarchies. If more than half of the box pairs of a certain level overlap, then we can save box overlap tests by starting the traversal at that level. Since the traversal itself incurs some further computational overhead, it might be more efficient to start at a certain hierarchy level even for a smaller percentage of overlapping box pairs. However, to decide which level is most efficient to start with, we would have to know the number of overlapping box pairs in advance, which is obviously not possible.

However, in a typical real-time application of collision detection objects are moving on a smooth path. Therefore, we can use a heuristic based on temporal coherence: we simply remember for each object pair the highest hierarchy level at which more than a certain percentage of box pairs overlapped. This percentage is obtained using an occlusion query. For the next frame we start the traversal of the hierarchy at exactly this level. Based on the number of overlaps we determine on this level this time, we adjust the “entrance” level for the next collision check.

## 6 Implementation

We implemented the method described in Section 5 using C++ and OpenGL on a NVidia GeForce FX 5900 GPU.

The AABB tree is stored using 32 bit floating point textures with rectangle texture target (i.e., the entries of a  $m \times n$ -sized texture are addressed using coordinates  $(i + \frac{1}{2}, j + \frac{1}{2})$  for  $i = 0, \dots, m - 1, j = 0, \dots, n - 1$  on texture lookup).

The multi-level overlap count map and the node pair index map are both used as temporary buffers during the described algorithm. Therefore, we require a method that enables these textures to be used as targets for rendering operations. The most efficient technique allowing this would be to use the upcoming super buffer extension. But since it is not yet available for our target system

we had to use the p-buffer and render-to-texture WGL extensions instead.

However, due to driver limitations on this GPU we cannot create integer p-buffers with more than 8 bits per component. As this would not be sufficient in our case, we use 16 bit half-precision floating point p-buffers instead. (Its 11 bit mantissa is sufficient to accurately store values up to 4096, which corresponds to the maximum texture size on the used GPU and thus to our maximum index value.)

The main disadvantage of the fact that we cannot use super buffers is that using p-buffers each render target change causes a GL render context switch which is connected with degraded performance caused by pipeline flushes. Another disadvantage of using p-buffers is that we cannot share occlusion queries between multiple render contexts which would be of great benefit for delayed evaluation of the query result. Fortunately, these disadvantages will disappear as soon as the super buffer extension will be available.

## 7 Results

We tested the performance of our algorithm using a test scenario similar to that of [22]: Two identical objects are positioned at a certain distance from each other. The distance is computed between the centers of the bounding boxes of the two objects; objects are scaled uniformly so they fit into a cube of size  $2^3$ . One of the two objects is rotated around a fixed axis by fixed number of small steps. In each step, the two objects are checked for collision, and the average collision detection time for a complete revolution at that distance is computed. Then the process is repeated with a slightly decreased distance of the two objects.

This test was performed on a set of CAD objects with varying complexities. To compare the performance of GPU and CPU based collision detection methods, we also implemented the AABB tree traversal approach on the CPU using an identical traversal scheme and ran the same tests on that implementation. Fig. 4 shows the results.

Note, that the timings include the determination of all intersecting triangle pairs as well as the read-back of its indices from graphics memory in case of the GPU implementation.

When comparing the performance of the individual steps of the algorithm between its GPU

and CPU implementations, it turns out that in the GPU implementation the box overlap and triangle intersection tests perform up to four times faster especially at the lower hierarchy levels. However, this speed-up is reduced by the overhead of the node pair index map generation which is larger in the GPU implementation. Overall, in general our current GPU implementation is slightly faster than the CPU implementation.

We also made some tests with ATIs current super buffer beta implementation on an ATI Radeon GPU, where we used the super buffer for rendering to our temporary textures as well as for directly rendering into the vertex array. However, due to stability issues, we cannot provide reliable results so far. But as soon as they will be resolved, we expect that by using this technique the overhead introduced by generating the node pair index maps in several passes could be reduced significantly.

## 8 Conclusion and Future Work

We presented a method for interference detection using programmable graphics hardware. Unlike previous GPU-based approaches, there are no requirements on shape, topology, and connectivity on the polygonal input models. Furthermore, all calculations are done in object-space rather than image-space.

Our method is an hierarchical algorithm that borrows ideas from traditional CPU based collision detection approaches. It simultaneously traverses a pair of bounding volume hierarchies consisting of axis-aligned bounding boxes. In addition, the method utilizes temporal coherence when used for collision testing over a period of time.

Although our current implementation could still be improved by using upcoming driver features it is still able to compete with CPU based hierarchical collision detection approaches. This allows using the CPU idle time (for example while waiting for occlusion query results) for other tasks. And with the release of future GPUs our method is potentially able to outperform the CPU based approaches clearly.

As there is no early termination in case of a determined intersecting triangle pair, the main target of this method are applications that need to determine the complete list of intersecting triangles or at least its number.



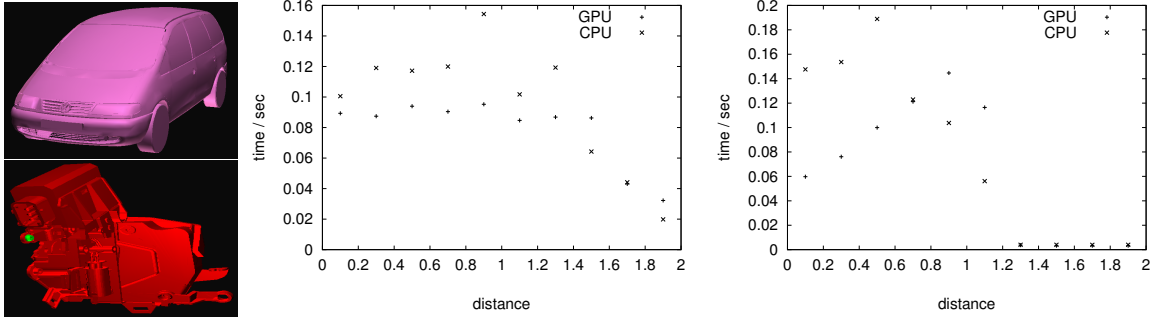


Figure 4: The collision response times for the car are shown in the middle, and those for the door lock at the bottom. The corresponding polygonal models can be seen in the top row. (Data courtesy of VW and BMW.)

There are still possible enhancements and optimization opportunities we want to investigate in the future:

- The algorithm might be enhanced for the application of collision detection in complex environments where multiple polygonal models with its own object-space have to be checked against each other. It might be worthwhile to check if we can speed-up the collision detection in this scenario by starting the bounding volume traversal with processing multiple pairs of AABB trees at once, for example by accessing a matrix palette from the vertex program for handling the individual transformation matrixes.
- Our approach could easily be modified to use hierarchies based on bounding volumes other than AABBs. One future task would be to evaluate the influence of using different bounding volumes on the performance of our GPU based approach.

## A Fragment Program Implementation

*SAT lite* test in high-level shading language:

```
void testBoxOverlap (float4 centerS, float3 extentS,
                   float4 centerT, float3 extentT)
{
    float3 dist;           // distance between centers
    float3 extsum;        // interval radii

    // compute difference of box centers in S space
    float3 distInS = float3(dot(centerT, matC[0]),
                           dot(centerT, matC[1]),
                           dot(centerT, matC[2])
                           - centerS.xyz);
```

```
    // determine three potentially separating axes
    dist      = distInS;
    extsum    = extentS;
    extsum.x += dot(extentT, matAbsC[0]);
    extsum.y += dot(extentT, matAbsC[1]);
    extsum.z += dot(extentT, matAbsC[2]);

    // discard fragment if any component of
    // "extsum - abs(dist)" < 0
    clip(extsum - abs(dist));

    // determine three more potentially separating axes
    dist.x    = dot(distInS, matCInv[0]);
    dist.y    = dot(distInS, matCInv[1]);
    dist.z    = dot(distInS, matCInv[2]);
    extsum.x  = dot(extentS, matAbsCInv[0]);
    extsum.y  = dot(extentS, matAbsCInv[1]);
    extsum.z  = dot(extentS, matAbsCInv[2]);
    extsum   += extentT;

    // discard fragment if any component of
    // "extsum - abs(dist)" < 0
    clip(extsum - abs(dist));
}
```

where the uniform parameter `matC` is the transformation matrix from `T`'s object-space to `S`'s, `matCInv` its inverse, and `matAbsC` and `matAbsCInv` the matrixes obtained by taking the absolute value of each entry of `matC` and `matCInv`, respectively.

## References

- [1] Agarwal, P., M. de Berg, J. Gudmundsson, M. Hammar, and H. Haverkort, *Box-Trees and R-Trees with Near-Optimal Query Time*, Discrete and Computational Geometry, 28 (2002), pp. 291–312.
- [2] Agarwal, P., S. Krishnan, N. Mustafa, and S. Venkatasubramanian, *Streaming Geo-*

- metric Optimization Using Graphics Hardware*, in 11th European Symposium on Algorithms, 2003.
- [3] Baciú, G. and W. S.-K. Wong, *Hardware-assisted self-collision for deformable surfaces*, in Proceedings of the ACM Symposium on Virtual Reality Software and Technology (VRST), 2002, pp. 129–136. ISBN 1-58113-530-0. <http://doi.acm.org/10.1145/585740.585762>.
- [4] Baciú, G. and W. S.-K. Wong, *Image-Based Techniques in a Hybrid Collision Detector*, IEEE Transactions on Visualization and Computer Graphics, 9 (2003), pp. 254–271.
- [5] Ehmann, S. A. and M. C. Lin, *Accurate and Fast Proximity Queries Between Polyhedra Using Convex Surface Decomposition*, in Computer Graphics Forum, vol. 20, 2001, pp. 500–510. ISSN 1067-7055.
- [6] Gottschalk, S., M. Lin, and D. Manocha, *OBB-Tree: A Hierarchical Structure for Rapid Interference Detection*, in SIGGRAPH 96 Conference Proceedings, Rushmeier, H., ed., ACM SIGGRAPH, Aug. 1996, pp. 171–180. held in New Orleans, Louisiana, 04-09 August 1996.
- [7] Govindaraju, N., S. Redon, M. C. Lin, and D. Manocha, *CULLIDE: Interactive Collision Detection Between Complex Models in Large Environments Using Graphics Hardware*, in Proc. of Graphics Hardware, San Diego, California, July 2003. <http://graphics.stanford.edu/papers/photongfx/>.
- [8] Hubbard, P. M., *Approximating Polyhedra with Spheres for Time-Critical Collision Detection*, ACM Transactions on Graphics, 15 (1996), pp. 179–210. ISSN 0730-0301.
- [9] Kapasi, U. J., S. Rixner, W. J. Dally, B. Khailany, J. H. Ahn, P. Mattson, and J. D. Owens, *Programmable Stream Processors*, IEEE Computer, (2003), pp. 54–61.
- [10] Klosowski, J. T., M. Held, J. S. B. Mitchell, H. Sowizal, and K. Zikan, *Efficient Collision Detection Using Bounding Volume Hierarchies of k-DOPs*, IEEE Transactions on Visualization and Computer Graphics, 4 (1998), pp. 21–36.
- [11] Knott, D. and D. K. Pai, *CInDeR: Collision and Interference Detection in Real-Time Using Graphics Hardware*, in Proc. of Graphics Interface, Halifax, Nova Scotia, Canada, June 11–13 2003.
- [12] Larsson, T. and T. Akenine-Möller, *Collision Detection for Continuously Deforming Bodies*, in Eurographics, 2001, pp. 325–333. short presentation.
- [13] Lin, M. C. and D. Manocha, *Interactive Geometric Computations Using Graphics Hardware*, in SIGGRAPH 2002 Course Notes Number 31, July 2002.
- [14] Lombardo, J.-C., M.-P. Cani, and F. Neyret, *Real-time collision detection for virtual surgery*, in Proc. of Computer Animation, Geneva, Switzerland, May 26–28 1999.
- [15] Myszkowski, K., O. G. Okunev, and T. L. Kunii, *Fast collision detection between complex solids using rasterizing graphics hardware*, The Visual Computer, 11 (1995), pp. 497–512. ISSN 0178-2789.
- [16] Palmer, I. J. and R. L. Grimsdale, *Collision Detection for Animation using Sphere-Trees*, Computer Graphics Forum, 14 (1995), pp. 105–116. ISSN 0167-7055.
- [17] Shinya, M. and M.-C. Forgue, *Interference detection through rasterization*, The Journal of Visualization and Computer Animation, 2 (1991), pp. 132–134. ISSN 1049-8907.
- [18] van den Bergen, G., *Efficient Collision Detection of Complex Deformable Models using AABB Trees*, Journal of Graphics Tools, 2 (1997), pp. 1–14.
- [19] Zachmann, G., *Rapid Collision Detection by Dynamically Aligned DOP-Trees*, in Proc. of IEEE Virtual Reality Annual International Symposium; VRAIS '98, Atlanta, Georgia, Mar. 1998, pp. 90–97.
- [20] Zachmann, G., *Minimal Hierarchical Collision Detection*, in Proc. ACM Symposium on Virtual Reality Software and Technology (VRST), Hong Kong, China, Nov. 11–13 2002, pp. 121–128. <http://www.gabrielzachmann.org/>.

- [21] Zachmann, G. and G. Knittel, *An Architecture for Hierarchical Collision Detection*, in Journal of WSCG '2003, University of West Bohemia, Plzen, Czech Republic, Feb.3–7 2003, pp. 149–156. <http://www.gabrielzachmann.org/>.
- [22] Zachmann, G. and G. Knittel, *High-Performance Collision Detection Hardware*, Tech. Rep. CG-2003-3, University Bonn, Informatikk II, Bonn, Germany, Aug. 2003. <http://www.gabrielzachmann.org/>.