

Hardware Accelerated Collision Detection — An Architecture and Simulation Results

ANDREAS
RAABE
Dept. Technical CS

BLAZEJ
BARTYZEL
Dept. Technical CS

GABRIEL
ZACHMANN
Dept. Comp. Graphics

JOACHIM K.
ANLAUF
Dept. Technical CS

University of Bonn, Germany

Abstract: - We present a hardware architecture for a single-chip acceleration of an efficient hierarchical collision detection algorithm as well as simulation results for collision queries using this architecture.

The architecture consists of two main stages, one for traversing simultaneously a hierarchy of k -DOPs, and one for intersecting triangles. Within each stage, the architecture is deeply pipelined and parallelized. For the first stage, we compare and evaluate different traversal schemes for bounding volume hierarchies.

A simulation in VHDL shows that a hardware implementation can offer a speed-up over a software implementation by orders of magnitude. Thus, real-time collision detection of complex objects at rates required by force-feedback and physically-based simulations can be achieved.

Keywords: - Computer Graphics, Bounding Volume Hierarchy.

1 Introduction

Collision detection is a fundamental task in areas like animation systems, virtual reality, games, physically-based simulation, automatic path finding, virtual assembly simulation, and medical training and planning systems.

In many of these systems, collision avoidance is the ultimate goal. Most approaches today are reactive, i.e., they first place objects at a new trial position, then they check for collisions, and then compute new forces or positions, based on physical laws or constraints, so as to remove any collisions.

This approach poses very high demands on collision detection, because it must perform many collision checks per simulation cycle. A particularly demanding application is rendering force-feedback, where updates of about 1000Hz must be done in order to achieve stable force computations. Since collision detection is such a fundamental task, it is highly desirable to have hardware acceleration available just like 3D graphics accelerators. Using specialized hardware, the system's CPU can be freed from computing collisions.

In this paper we present a new efficient architecture for hierarchical collision detection of two rigid objects using high-end ASIC technology.

We also present simulation results concerning its speed and size, which show that an implementation in dedicated hardware can speed up applications by at least an order of magnitude.

2 Related Work

Collision detection in computer graphics has been an active field over the past decade [2, 3, 6, 9, 10]. Most approaches dealing with rigid objects utilize some kind of bounding volume hierarchy, where some of the bounding volumes (BVs) are spheres, axis-aligned bounding boxes (AABB), oriented bounding boxes (OBB), and discretely oriented polytopes (DOP).

However, there is very little literature about the design of hardware dedicated to collision detection. To our knowledge, [11, 12] are the only exception. Their architecture was designed to minimize chip resources, while our architecture is designed to maximize throughput. In addition, they presented only a functional simulation, while we have performed a VHDL simulation.

All other hardware-related research in collision detection so far has tried to utilize existing graphics accelerator boards [1, 4, 5, 7, 8]. While earlier approaches, such as [8], can basically handle only convex objects, later algorithms, such as [1, 7], have extended these to more general cases such as unions of convex objects or closed objects. The general class of "polygon soups" can be handled by [4], but they use a hybrid approach where the graphics hardware only identified potentially colliding sets.

All of the approaches using graphics hardware have the disadvantage that they either compete with the rendering process for the same hardware resource, or an additional graphics board must be spent for collision detection. The former slows down the overall frame rate considerably, while

the latter would be a tremendous waste, since most of the resources of the hardware would not be utilized at all. Furthermore, most of these approaches work in image space, which reduces precision significantly.

3 The Algorithm

3.1 Hierarchical Collision Detection and Bounding Volumes

In this paper we concentrate on hierarchical collision detection. This avoids checking every triangle of an object A for collision with all triangles of object B by hierarchically grouping triangles (or other graphical primitives). This yields a so-called *bounding volume hierarchy* (BVH), where each leaf corresponds to one triangle and inner nodes correspond to groups of triangles. In order to achieve a feasible hardware design, we use a binary tree here, but n-ary trees could be considered as well. Each node of the tree stores a bounding volume (BV) that encloses all triangles in its group. Note that only leaves explicitly store any triangles.

In this work, we use k-DOPs as BVs because they were proven to yield very fast collision queries by extensive benchmarking in software [10], and are likely to perform well in hardware.

If two objects are checked for intersection, both hierarchies are traversed starting at both roots. If their BVs intersect, then the next level of BVs is checked. Since two objects will usually intersect only locally in a very small number of primitives, this yields a significant speed-up in the average case. In practical cases, the complexity is in $O(\log n)$ (n = number of primitives) because only a small diagonal “slice” of constant width down the BVH needs to be visited.

3.2 k-DOPs

As motivated above, we use k-DOPs as BVs. For sake of reference, we give a quick recap of this particular type of BV [10, 11].

A k-DOP consists of k distances d_j along pairwise linearly independent vectors \mathbf{B}_j . Each of these vectors forms the normal of a halfspace. These vectors are chosen such that they form $k/2$ pairs, each of which defines a so-called *slab*. The intersection of these slabs forms the BV:

$$D = \bigcap_{j=1, \dots, k} H_j, \quad H_j : \mathbf{B}_j \mathbf{x} - d_j \leq 0 \quad (1)$$

The orientation matrix B is fixed and equal for all objects. This yields a very space-efficient description for every k-DOP: only the k numbers d_j need to be stored. To avoid rounding errors we use single-precision floating-point numbers.

Each object O has its own reference frame (RF) $F(O)$ that describes its rotation R_O and translation T_O with respect to world coordinates. When an object moves, only R_O and T_O have to be updated. So checking two DOPs for intersection




Figure 1: Our DOP overlap test gains its speed by transforming DOP Q into O 's reference frame $F(O)$. The tightness loss is shown in dark grey. Obviously, each d'_i is determined by exactly three original d 's.

requires transformation of one of them into the RF of the other.

Assume O and Q to be 2 objects. Let $D_Q(Q)$ denote the minimum DOP which bounds Q with respect to the orientation matrix \mathbf{B} in Q 's own reference frame $F(Q)$. Since calculating $D_O(Q)$ is prohibitively expensive, we calculate $D_O(D_Q(Q))$, which is the minimum DOP in $F(O)$ bounding $D_Q(Q)$. Naturally, this incurs a loss of the BVs tightness to the underlying geometric structure.

Assume \mathbf{M} to be the rotation and \mathbf{o} the translation which transforms $F(Q)$ into $F(O)$. Then, we need to find distances d'_i which bound $\mathbf{M} \cdot D_Q(Q) + \mathbf{o}$ minimally.

Applying \mathbf{M} and \mathbf{o} to Equation 1 yields

$$h_j : \mathbf{b}_j \mathbf{x} - d_j + \mathbf{b}_j \mathbf{o} \leq 0, \quad \text{where } \mathbf{b}_j = \mathbf{B}_j \mathbf{M}^{-1} \quad (2)$$

Assume $D_O(D_Q(Q))$ is the intersection of

$$H_i : \mathbf{B}_i \mathbf{x} - d'_i \leq 0, \quad i = 1, \dots, k \quad (3)$$

Then, each d'_i corresponds to exactly one vertex of $D_Q(Q)$ and therefore to three d_j (see Fig. 1). These correspondences are the same for all nodes in an object's DOP hierarchy. So they can be determined at startup.

Let $j_l, 0 \leq l \leq 2$, be the indices corresponding to d_i . Then,

$$\begin{pmatrix} \mathbf{b}_{j_0} \\ \mathbf{b}_{j_1} \\ \mathbf{b}_{j_2} \end{pmatrix} \mathbf{x} - \begin{pmatrix} \mathbf{d}_{j_0} \\ \mathbf{d}_{j_1} \\ \mathbf{d}_{j_2} \end{pmatrix} + \begin{pmatrix} \mathbf{b}_{j_0} \\ \mathbf{b}_{j_1} \\ \mathbf{b}_{j_2} \end{pmatrix} \mathbf{o} = \mathbf{0} \quad (4)$$

$$\mathbf{b}_i \mathbf{x} - d'_i = 0 \quad (5)$$

Equating (4) and (5) yields

$$d'_i = C_{ij_1} d_{j_1} + C_{ij_2} d_{j_2} + C_{ij_3} d_{j_3} + c_i \quad (6)$$

where C and c are chosen to be $C_{ij} := B_i \begin{pmatrix} \mathbf{b}_{j_0} \\ \mathbf{b}_{j_1} \\ \mathbf{b}_{j_2} \end{pmatrix}^{-1}$ and $c_i := \mathbf{b}_i \mathbf{o}$. Both are the same for all nodes in an object's DOP hierarchy; thus, they can be calculated at startup.




Figure 2: Two objects overlap, iff their projections intersect on every axis.

Checking $D_O(D_Q(Q))$ and $D_O(O)$ for intersection amounts to projecting them on the k axes given by B . They overlap if and only if there is no axis on which they are non-intersecting. Since there are always two antiparallel axes, we need to take that into account (see Fig. 2).

If e denotes the k -vector describing $D_O(O)$, then this test can be expressed as

$$\text{overlap} \Leftrightarrow \exists i \in [1, \frac{k}{2}] : e_{i+\frac{k}{2}} > -d'_i \vee d'_{i+\frac{k}{2}} > -e_i \quad (7)$$

Due to space limitations, we refer the interested reader to [11, 12] for the calculations involved in triangle-triangle intersection tests.

4 The Architecture

To achieve maximum possible speed we assume to be using high-end hardware components: our target ASIC technology is a NEC UX5 CB-130 in $0.095\mu\text{m}$ -copper-technology. With a maximum of 61 gates in a row it can establish up to 800 MHz clockrate. Furthermore, we assume DDR2 memory modules.


As BVs we chose to use 24-DOPs because extensive software benchmarking has shown 24 axes to be a good compromise of tightness and effort.

4.1 Design of the DOP Intersection Test

Our DOP intersection test is the combination of criterion (7) with (6), which amounts to the function

$$\tilde{d}'_i = d_k C_{i0} + d_m C_{i1} + d_n C_{i2} + c_i + e_{i+\frac{k}{2}} \quad (8)$$

This can be implemented as the following three-staged macro-pipeline, which we call D_CRITCHK unit:



The macro-pipeline stages were refined furthermore resulting in a pipeline of 15 stages and therefore an initialization delay of 15 clock cycles.

We use 24 D_CRITCHK units in parallel, and their results are NOR-reduced to check the criterion. To fill the

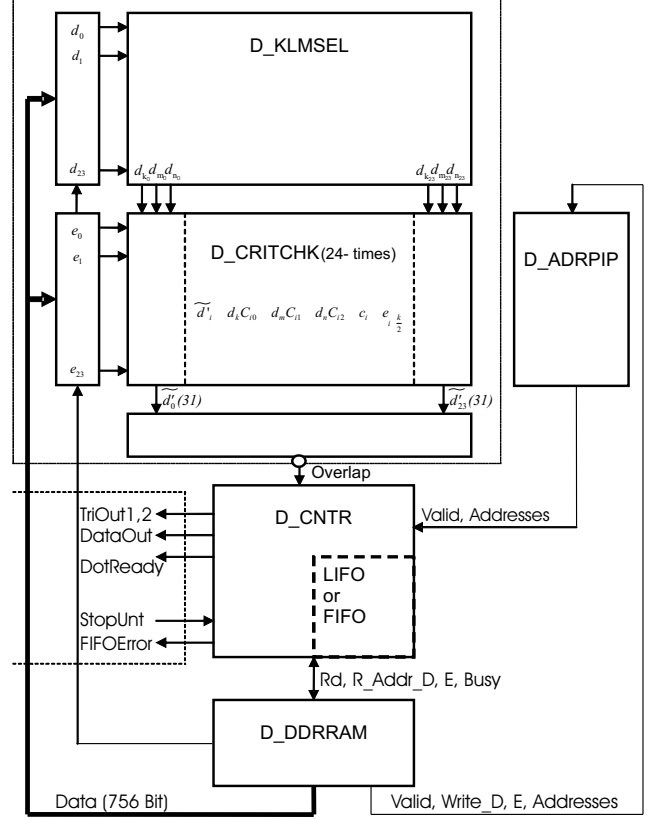


Figure 3: Architecture for DOP intersection testing.

pipeline we use a 756-bit wide bus from the DDR2-RAM. A hypermultiplexer (D_KLMSEL) routes the correct inputs of the DOP to the D_CRITCHK units, which will then be transformed into the reference frame of the other DOP (see Fig. 3). A D_CNTR unit controls which DOP pair will be processed next (see Section 4.3).

4.2 Design of the Triangle Intersection Test

Using homogeneous coordinates, the affine transformations needed for the triangle intersection test (see Section ??) can be represented as 3×4 matrices.

The T_CHECK unit that performs the intersection test gets as input one triangle $V_A^i = [x_i y_i z_i 1]$, $1 \leq i \leq 2$, the pre-computed matrix

$$M_B := [m_{ij}] := \begin{bmatrix} m_{11} & m_{12} & m_{13} & m_{14} \\ m_{21} & m_{22} & m_{23} & m_{24} \\ m_{31} & m_{32} & m_{33} & m_{34} \end{bmatrix}, \quad (9)$$

and $M_{AB} := [b_{ij}]$ that transforms O 's reference frame into Q 's.

Calculating $\tilde{V}_A^i = M_B \times M_{AB} \times V_A^i$ is done in the first two of 5 macro pipeline stages. These will be detailed in the following.

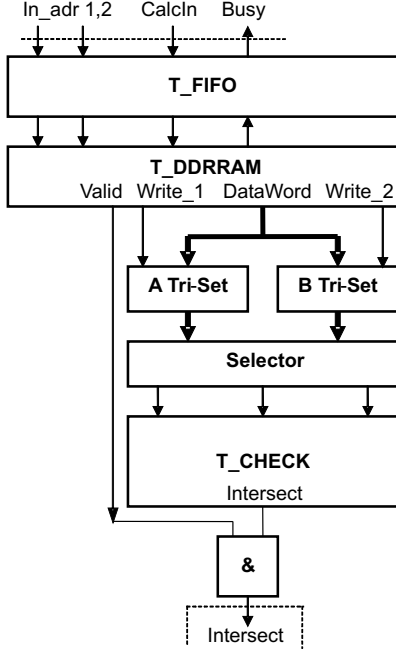


Figure 4: Architecture for Triangle Intersection

1st macro stage: The calculation of $M_b = [b'_{ij}] = M_B \times M_{AB}$ is split into three substages (marked with different brackets):

$$b'_{ij} = \{[(b_{i1}m_{1j}) + (b_{i2}m_{2j})] + (b_{i3}m_{3j})\} \quad (10)$$

$$b'_{i4} = \{[(b_{i1}m_{14}) + (b_{i2}m_{24})] + [(b_{i3}m_{34}) + (b_{i4})]\} \quad (11)$$

$$i, j = 1, \dots, 3$$

The fourth row of the resulting matrix is always [001]. These substages are further divided into microstages to gain maximum clock frequency. With this refinement, the first macro stage consumes 36 multiplication and 24 addition floating point units and takes 15 clock cycles delay to produce the first result.

2nd macro stage: Calculating $\tilde{V}_A^i = M_b \times V_A^i$ works very similarly. The resulting substages are:

$$\tilde{V}^i = \{[(b'_{i1}V_x^i) + (b'_{i2}V_y^i)] + [(b'_{i3}V_z^i) + (b'_{i4})]\} \quad (12)$$

Dividing this into microstages yields 27 multiplications, 27 additions, and another 15 clock cycles delay.

3rd macro stage: Before checking the edges of \tilde{T}_A for intersection with the unit triangle, we need to calculate a and b according to Eq. ?? for all three edges. Therefore, we need to calculate P_xQ_z , Q_xP_z , P_yQ_z , and Q_yP_z first. Additionally, we calculate r_z . This takes 12 multiplications, 3 additions, and 8 clock cycles. Calculating a and b from these terms takes another 6 additions and 4 clock cycles.

4th macro stage: For all three edges we now need to calculate $a + b - r_z$. After division into microstages this consumes 6 additions and 8 cycles.

5th macro stage: Here, we check the signs of a, b , and $a + b - r_z$ for all edges. This needs one clock cycle.

Overall pipeline: Putting all stages together, we get a pipeline with 52 clock cycles delay.

To fill the pipeline with data we need to buffer the triangle addresses, look them up in the DDR2_RAM which contains vertex data and the transformation matrices, and divide the data into two sets (because all edges of T_A have to be checked against T_B and vice versa). The whole TRI_UNIT is presented in Fig. 4.

4.3 Control

When the two object hierarchies are traversed symmetrically, each intersecting DOP pair results in 4 child pairs to be checked for intersection. However, other ways of traversing both BVHs can be more efficient. So far, we have compared those two possibilities that allow a small control unit and reduce the number of necessary memory accesses and transformations.

Traversing the two DOP hierarchies in Fig. 5a symmetrically basically amounts to traversing the “collision tree” shown in Fig. 5b.

Using a FIFO for determining the next DOP pair to be tested for intersection results in a plain breadth-first search on the collision tree. The processed DOP-sequence is **A1 - B2 B3 C2 C3** - D4 D5 E4 E5 - D6 D7 E6 E7 - F4 F5 G4 G5 - F6 F7 G6 G7

This can easily be optimized so that between testing two nodes of the same depth only one DOP needs to be fetched from memory. The resulting sequence is

A1 - B2 B3 C3 C2 - D4 D5 E5 E4 - E6 E7 D7 D6 - F6 F7 G7 G6 - G4 G5 F5 F4

Using a LIFO the sequence depends on the length of the pipeline. If we assume the pipeline to be of length one (for simplicity of presentation) we receive a plain depth first search.

A1 - B2 - D4 D5 E5 E4 - **B3** - E6 E7 D7 D6 - **C2** - G4 G5 F5 F4 - **C3** - F6 F7 G7 G6

If we push and pop nodes pairwise, we can easily reduce the number of necessary memory accesses as we did before.

Note that our DOP transformation pipeline has 15 stages. Thus, the traversal is not plain depth-first in the strictest sense. Instead, it proceeds along several paths in a depth-first manner.

5 Simulation Results

For benchmarking our architecture we used three different objects each of which in several polygon complexities (see Fig. 6). For each of them, two copies are placed at different distances from each other and with different rotations. For each constellation, the collision detection query time is determined.

Comparing the performance of LIFO and FIFO control, we see that finding all intersecting primitives takes equally long (see Fig. 7a).

Since using a LIFO corresponds to depth-first search on the collision tree, finding the first collision is usually much

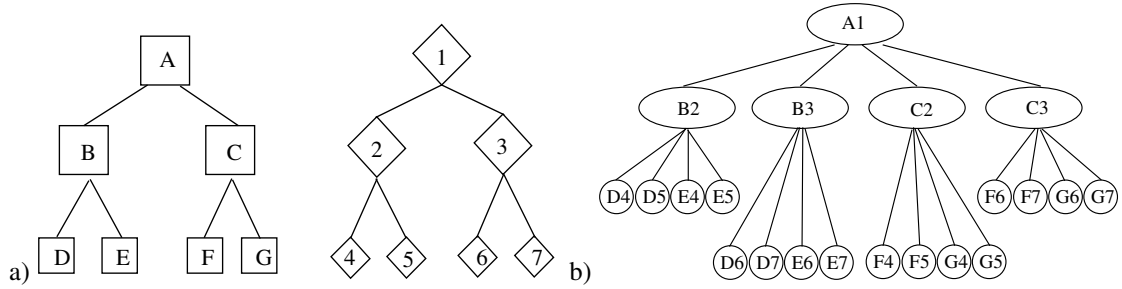


Figure 5: Traversing two BVHs simultaneously amounts to traversing a single “collision tree”. Our new traversal scheme optimizes memory accesses.




Figure 6: For benchmarking and testing, a number of different test objects with several polygon complexities were used.

faster than using a FIFO. Our simulation results verify this (see Fig. 7b).

Another disadvantage of using a breadth-first search is the size needed for the memory structure. In the worst case, when all nodes need to be checked for intersection we must store them all in the FIFO before any leaves are checked and the queue size reduces.

With a strict depth-first traversal, the LIFO would need to be only as large as the depth of the BVH. However, as explained in Section 4.3, our traversal is not strictly depth-first. Fortunately, memory usage of the LIFO in our design seems to behave just as well (see Fig. 8).

Fig. 9 shows that our collision detection architecture is up to 1000 times faster than the software implementation in determining all intersecting triangles of two objects. The software times were obtained on a 1GHz Pentium 3.

6 Conclusion and Future Work

In this paper, we have presented, to our knowledge, the first simulation in VHDL of a hardware architecture for collision detection.

We showed that hardware acceleration can be an effective way to speed up hierarchical collision detection. Using several pipelining and parallelization techniques, we achieved a





Figure 7: (a) Comparing LIFO and FIFO controlled intersection testing shows that finding all intersecting primitives takes equally long. (b) Our comparison shows that LIFO control is far superior to FIFO control when finding the first intersecting primitive. (Object: headlight, 5947 polygons)

speed-up of factor 1000 in VHDL simulations compared to a software implementation.

We also showed that, for simultaneous traversal of BVHs, a LIFO-controlled pipeline is far more space efficient than a FIFO-controlled one, without loss of processing speed.

Since the present paper is one of the first to look at hardware acceleration of collision detection, we believe there are many avenues for further research.

An important concern is the reduction of the bandwidth in order to make the bus from chip to memory smaller. Currently, we are investigating discretization and compression of the BVs. Furthermore, we will evaluate different kinds of primitives and BVs.

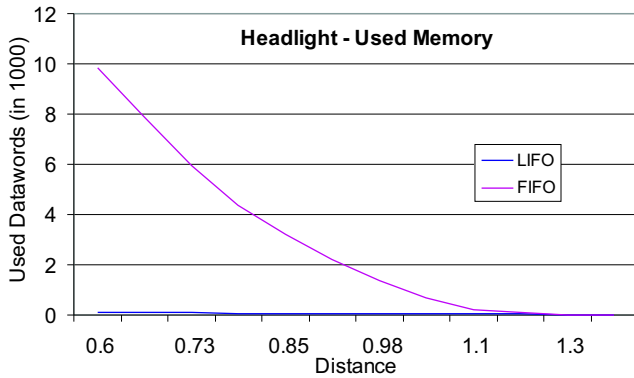


Figure 8: Although our hardware architecture implements a traversal scheme (LIFO) that is not quite a LIFO, it uses still only very few memory, compared to a FIFO. (Object: headlight, 5947 polygons)

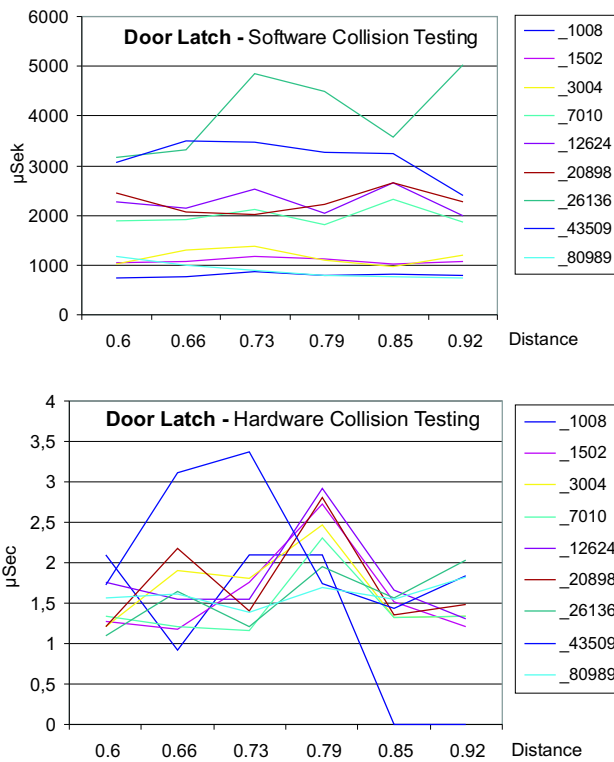


Figure 9: Our collision detection architecture is up to 1000 times faster in determining all intersecting triangles of two objects than the software implementation. (Object: door lock, several polygon counts)

Another important issue is collision detection of deformable objects. It remains an open problem exactly which algorithms and data structures are best suited for hardware implementation.

Since we use a pipelined dataflow architecture, advanced pipelining techniques like speculative execution could be applied to reduce the number of pipeline stalls.

References

- [1] George Baciú, Wingo Sai-Keung Wong, and Hanqiu Sun. RECODE: an image-based collision detection algorithm. *The Journal of Visualization and Computer Animation*, 10(4): 181–192, October - December 1999. ISSN 1049-8907.
- [2] Jens Eckstein and Elmar Schömer. Dynamic Collision Detection in Virtual Reality Applications. In *Proc. The 7-th Int'l Conf. in Central Europe on Comp. Graphics, Vis. and Interactive Digital Media '99 (WSCG'99)*, pages 71–78, Plzen, Czech Republic, February 1999. University of West Bohemia.
- [3] Stefan Gottschalk, Ming Lin, and Dinesh Manocha. OBB-Tree: A Hierarchical Structure for Rapid Interference Detection. In *SIGGRAPH 96 Conference Proceedings*, Holly Rushmeier, Ed., pages 171–180. ACM SIGGRAPH, Addison Wesley, August 1996. held in New Orleans, Louisiana, 04-09 August 1996.
- [4] Naga K. Govindaraju, Stephane Redon, Ming C. Lin, and Dinesh Manocha. CULLIDE: Interactive Collision Detection Between Complex Models in Large Environments using Graphics Hardware. In *Graphics Hardware 2003*, pages 25–32, July 2003.
- [5] Alexander Gress and Gabriel Zachmann. Object-Space Interference Detection on Programmable Graphics Hardware. In *SIAM Conf. on Geometric Design and Computing*, Seattle, Washington, November 13–17 2003.
- [6] P. M. Hubbard. Collision detection for interactive graphics applications. *IEEE Transactions on Visualization and Computer Graphics*, 1(3):218–230, September 1995. ISSN 1077-2626.
- [7] Dave Knott and Dinesh 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.
- [8] Karol Myszkowski, Oleg G. Okunev, and Toshiyasu L. Kunii. Fast collision detection between complex solids using rasterizing graphics hardware. *The Visual Computer*, 11(9):497–512, 1995. ISSN 0178-2789.
- [9] Gino Johannes Apolonia van den Bergen. *Collision Detection in Interactive 3D Computer Animation*. PhD dissertation, Eindhoven University of Technology, 1999.
- [10] Gabriel Zachmann. Rapid Collision Detection by Dynamically Aligned DOP-Trees. In *Proc. of IEEE Virtual Reality Annual International Symposium; VRAIS '98*, pages 90–97, Atlanta, Georgia, March 1998.
- [11] Gabriel Zachmann and Günter Knittel. An Architecture for Hierarchical Collision Detection. In *Journal of WSCG '2003*, pages 149–156, University of West Bohemia, Plzen, Czech Republic, February 3–7 2003. <http://www.gabrielzachmann.org/>.
- [12] Gabriel Zachmann and Günter Knittel. High-Performance Collision Detection Hardware. Tech. Rep. CG-2003-3, University Bonn, Informatik II, Bonn, Germany, August 2003. <http://www.gabrielzachmann.org/>.