




Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich



CS Technical Report #395, April 8, 2003

Volumetric Collision Detection for Deformable Objects

Bruno Heidelberger, Matthias Teschner, Markus Gross

Computer Science Department
ETH Zurich, Switzerland
e-mail: {heidelberger, teschner, grossm}@inf.ethz.ch
<http://graphics.ethz.ch>

Volumetric Collision Detection for Deformable Objects

Bruno Heidelberger

Matthias Teschner

Markus Gross

Computer Graphics Laboratory
ETH Zurich




Figure 1: This sequence shows two deforming objects with 150k and 5k faces. An explicit representation of the red intersection volume can be computed at interactive frame rate.

Abstract

We present a new algorithm for the efficient detection of collisions of geometrically complex objects. Our approach requires neither expensive setup nor sophisticated spatial data structures and is hence specifically suitable for handling deformable objects with arbitrarily shaped, closed surfaces. The algorithm is based on a Layered Depth Image (LDI) decomposition of the intersection volume. Currently, we have implemented two types of collision queries. The first one comprises an explicit representation of the intersection volume. The second one computes vertex-in-volume tests. All queries are processed on the LDI-based volume representation. Our algorithm is very easy to implement and can be accelerated in graphics hardware by using OpenGL.

Keywords: Collision Detection, Computational Geometry, Layered Depth Image, Shape Approximation, Deformable Modeling, Physically-based Modeling, Graphics Hardware

1 Introduction

The detection of collisions of geometric objects constitutes a fundamental problem in computer graphics, computational geometry, modeling, robotics, and many other areas. Most research is devoted to the design of algorithms for the efficient computation of queries, such as intersecting surfaces or penetration depth. For complex objects, such queries are computationally expensive, because they require the testing of many individual graphics primitives. Hence, efficient collision detection algorithms are accelerated by spatial data structures including bounding-box hierarchies, distance fields, or other ways of spatial partitioning. Such object representations are built in a pre-processing stage, and once they are created, perform very well for rigid objects.

Modern physics-based simulation, as for instance used in game engines or computational surgery, demand the detection of collisions of objects deforming over time. In the case of deformable objects, the aforementioned acceleration structures have to be updated to adapt to changes of the geometry. While some of these

algorithms have been successfully employed in such cases, the acceleration structures constitute an additional overhead in memory consumption and computational efficiency.

We propose a simple and efficient algorithm based on Layered Depth Images (LDI) [Shade et al. 1998] as the fundamental representation to accelerate collision detection of rigid and deformable objects. While most existing approaches consider object surfaces, our method is inherently volumetric. Rather than explicitly detecting surface intersections, the LDI provides a discrete representation of the intersection volume and allows for volume-based collision queries as shown in Fig. 1. The generation of LDIs is extremely efficient and can be accelerated by graphics hardware. Various optimizations minimize buffer read-backs.

Our algorithm proceeds in three stages. First, we calculate the intersection of axis-aligned bounding boxes of pairs of objects. If an intersection is non-empty, a second stage computes an LDI representation of each object within the bounding-box intersection. Finally, we obtain the overall intersection volume by simple Boolean operations on the LDI representation.

The simplicity of the algorithm, the absence of expensive setup and the ease of implementation in OpenGL makes it especially attractive for real-time simulation and computer games.

1.1 Contributions

The contributions of this work can be summarized as follows:

- We present a new algorithm for the computation of volumetric penetrations of geometrically complex, deformable objects in real time. Our method handles arbitrarily shaped, closed surfaces of manifold geometry.
- We currently process two different types of collision queries: The first one includes the explicit computation of the intersection volume of two objects. The second one performs vertex-in-volume tests. Our approach is also robust in cases, where one object lies completely inside another object.
- Unlike most existing approaches, our method does not require an involved setup or the implementation of a sophisticated

spatial data structure. It is memory efficient, computationally economical and performs on deformable models of up to hundred thousands of primitives.

1.2 Background and Related Work

Due to the importance of the topic, the interference detection of geometrically complex objects has been investigated extensively in various areas.

Collision algorithms. In graphics, efficient collision detection is an essential component in physically-based simulation or animation [Baraff 2001], [DeBunne et al. 2001], including cloth modeling [Bridson et al. 2002]. Further applications can be found in robotics, computer animation, medical simulations, computational biology, and games.

Collision detection algorithms based on bounding-volume (BV) hierarchies have proven to be very efficient and many types of BVs have been investigated. Among the acceleration structures we find spheres [Hubbard 1993], [Quinlan 1994], axis-aligned bounding boxes [Hughes et al. 1996], [van den Bergen 1997], object-oriented bounding boxes [Gottschalk et al. 1996], and discrete-oriented polytopes [Klosowski et al. 1996].

Initially, BV approaches were designed for rigid objects. In this context, the hierarchy is computed in a pre-processing step. In the case of deforming objects, however, this hierarchy must be updated at run time. While effort has been spent to optimize BV hierarchies for deformable objects [Larsson and Akenine-Moeller 2001], they still pose a substantial computational burden and storage overhead for complex objects. As an additional limitation for some applications, BV approaches typically detect intersecting surfaces. The computation of the intersection volume requires an additional step.

As an alternative to object partitioning, other approaches employ discretized distance fields as volumetric object representation for collision detection. The presented results, however, suggest that this data structure is less suitable for real-time processing of geometrically complex objects [Hirota et al. 2000].

Recently, various approaches have been introduced that employ graphics hardware for collision detection. In [Baciu et al. 1999], a multi-pass rendering method is proposed for collision detection. However, this algorithm is restricted to convex objects. In [Lombardo et al. 1999], the interaction of a cylindrical tool with deformable tissue is accelerated by graphics hardware. The method is restricted to collisions of simplified surgical tools with object surfaces. [Hoff III et al. 2001] proposes a multi-pass rendering approach for collision detection of 2-D objects, while [Kim et al. 2002] and [Kim et al. 2003] perform closest-point queries using bounding-volume hierarchies along with a multipass-rendering approach. The aforementioned approaches decompose the objects into convex polytopes.

Layered Depth Images (LDI). LDIs have been introduced as an efficient image-based rendering technique. The LDI data structure essentially stores multiple depth values per pixel. Thus, an LDI can be used to approximate the volume of an object.

A method for generating LDIs is presented, for instance, in the depth-peeling approach to order-independent transparency [Everitt 2001]. This implementation, however, demands a complex setup including two active depth buffers and texture shaders. In contrast, our approach to LDI generation is much simpler by employing a more efficient rendering setup based on plain OpenGL 1.4.

Hardware-based approaches to interactive CSG rendering, e. g. [Goldfeather et al. 1986] and [Rappoport and Spitz 1997], could also be employed for LDI generation. However, since such CSG rendering implementations have to handle more complex modeling operations they feature more involved rendering setups. Further-

more, these CSG rendering approaches do not explicitly compute the resulting solid.

2 Method

This section presents an overview of our algorithm followed by a detailed description of the three stages.

Input. Our approach takes 3-D closed objects of manifold geometry as an input. Although the approach is not confined to triangular meshes, we require a watertight object surface to perform volumetric collision queries. In addition, a rendering method for the object is needed for LDI generation.

Output. The algorithm computes an explicit representation of the intersection volume as an output. Alternatively, the LDI structure facilitates the fast detection of vertices which penetrate the object.

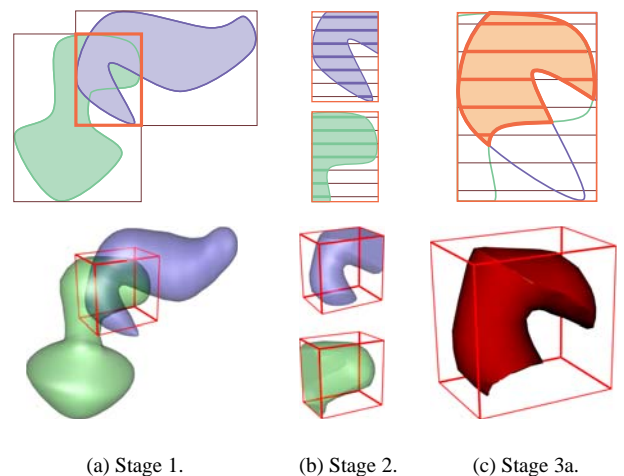


Figure 2: Algorithm overview in 2-D and 3-D. (a) AABB intersection. (b) LDI generation within the Vol. (c) Computation of the intersection volume.

Algorithm overview. Our approach proceeds in three stages.

Stage 1 computes the AABB intersection for a pair of objects. If the intersection is empty, the two objects do not overlap. If it is not, stages 2 and 3 are applied to the AABB intersection volume. We will refer to it as Volume-of-Interest (Vol).

Stage 2 computes two LDIs, one for each object. LDI generation is restricted to the Vol. The depth values of the LDI can be interpreted as the intersections of parallel rays or 3-D scan lines entering or leaving the object. Thus, an LDI classifies the Vol into *inside* and *outside regions* with respect to an object. Likewise, we classify the corresponding intersections into *entry points* and *leaving points*. The concept is similar in spirit to well-known scan-conversion algorithms to fill concave polygons [Foley et al. 1990], where intersections of a scan line with a polygon represent transitions between interior and exterior.

Stage 3 performs the actual collision detection. Currently, we distinguish two different types of queries: *a)* Both LDIs are combined using Boolean intersection. If the intersection of all inside regions is not empty, a collision is detected. This operation also provides an explicit representation of the intersection volume. *b)* Individual vertices within the Vol of one object are transformed into the LDI of the second object. If such a vertex lies in an inside region,

a collision is detected. This query can also be used to check atomic or point-sampled objects for intersection.

Fig. 2 illustrates the three stages of our algorithm. For a detailed description of all stages, refer to Sections 2.1–2.3.

2.1 AABB Intersection

AABB intersections are computed for pairs of objects. If two AABBs do not overlap, the corresponding objects cannot interfere. Otherwise, the intersection volume provides a VoI, which is considered for further processing. Although AABBs do not provide an optimal, i. e. smallest, bounding volume, they can be computed very efficiently. Furthermore, the intersection of two AABBs is again an AABB which keeps subsequent stages of the algorithm simple.

Alternative BVs, such as object-oriented bounding boxes [Gottschalk et al. 1996] or discrete-oriented polytopes [Klosowski et al. 1996], provide tighter fitting object approximations. However, the computation of these representations is more involved and the resulting intersection volume can have a more complex shape making it much more difficult for further processing. In addition, such structures are impractical for deforming objects, which require a dynamic adaptation of the bounding box.

In order to simplify and accelerate further computations, the VoI has to meet the following requirements: Firstly, to guarantee proper boundary conditions, we have to render the LDI for each object from the outside. We select one of the six faces of the VoI, the so-called *outside face*, as the near rendering plane for each object (see Fig. 3). If a sorted LDI is generated with respect to this face, entry points and leaving points alternate. For closed objects, the first layer is assumed to contain entry points, the second layer leaving points and so on. Secondly, outside faces for both objects should correspond to opposite sides of the VoI (see Fig. 3(a)). This simplifies the combination of both LDIs for the computation of the intersection volume.

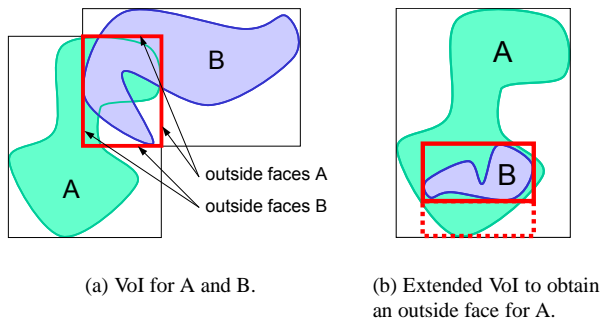


Figure 3: The intersection of the AABBs of two objects A and B provides the VoI.

The intersection volume of two AABBs, i. e. the VoI, is bounded by parts of their faces. Fig. 3(a) shows a VoI with corresponding outside faces. In some cases, for instance when one box is entirely within another box, appropriate outside faces for both objects cannot be found. This problem is solved by extending the VoI. If the outside face of one object is fixed, the opposite face of the VoI can be scaled to touch the bounding box of the other object (see Fig. 3(b)).

If several eligible pairs of outside faces exist, we choose the ones with the smallest distance between them. This criterion is a fast heuristic to optimize depth complexity for LDI generation assuming that depth complexity scales with the distance between the two outside faces.

2.2 LDI Generation

LDIs are computed within the VoI for both objects. This provides an explicit volumetric representation discretized at a predefined LDI resolution. The LDIs will be used for subsequent collision queries. The concept of a volumetric representation using LDIs is exemplified in Fig. 4.




Figure 4: LDIs are computed for objects within the VoI. LDI values are *entry* or *leaving points* and classify the VoI into *inside* and *outside regions*.

In order to generate an LDI, the object is rendered multiple times. The viewing parameters for the rendering process are determined by the selected outside faces of the VoI defining the near planes, and their opposite faces defining the far planes. The remaining faces define top, bottom, left, and right of the rendering volume. Orthographic projection is used for rendering. The resolution of the LDI, which defines the accuracy of the object approximation, is specified by the user. Objects are rendered n_{max} times for LDI generation, where n_{max} denotes the depth complexity of the relevant part of the object within the VoI. Thus, the computational complexity of the LDI generation is $O(n_{max})$. For details regarding the OpenGL implementation, refer to Sec. 3.

The first rendering pass generates a single LDI layer and computes n_{max} . Depth testing and face culling are disabled. Only the stencil test is employed to discard fragments. The stencil test configuration allows the first fragment per pixel to pass, while the corresponding stencil value is incremented. Subsequent fragments are discarded by the stencil test, but still increment the stencil buffer. Hence, after the first rendering pass the depth buffer contains the first object layer per pixel and the stencil buffer contains a map representing the depth complexity per pixel. Thus, n_{max} is found by searching the maximum stencil value. Note, that the $max()$ computation constitutes a very small computational burden, given the typical LDI resolution.

If $n_{max} > 1$, additional rendering passes 2 to n_{max} generate the remaining layers. The rendering setup is similar to the first pass. However, during the n -th rendering pass, the first n fragments per pixel pass the stencil test and, as a consequence, the resulting depth buffer contains the n -th LDI layer. During these passes, the stencil buffer is not incremented, if the stencil test fails.

Since fragments are rendered in arbitrary order, the algorithm generates an unsorted LDI. For further processing, the LDI is sorted per pixel. Only the first n_p layers per pixel are considered, where n_p is the depth complexity of this pixel. n_p is taken from the stencil buffer as computed in the first rendering pass. If n_p is smaller than n_{max} , layers $n_p + 1$ to n_{max} do not contain valid LDI values for this pixel and are discarded (see Fig. 5).

Although the order in which the fragments of an object are rendered is arbitrary, our algorithm relies on consistency across the individual passes. Our LDIs store normalized values. For advanced collision queries, additional information, such as viewing direction and viewing volume are stored as well.

The accuracy of the method can be adjusted. In (x, y) -direction, accuracy corresponds with the user-defined LDI resolution. The accuracy in z -direction is given by the depth-buffer resolution.





Figure 5: LDI for two pixels, unsorted and sorted with respect to depth values.

Compared to alternative data structures, such as bounding-volume hierarchies or distance fields, the LDI representation is very memory efficient. Both the memory consumption and the performance of the LDI generation depend on the geometric complexity and the depth complexity of the objects. However, our results presented in Sec. 4 suggest, that the number of LDI layers does not exceed reasonable values even for complex objects. Further, collision queries performed on the generated LDIs are very efficient even in the case of a large number of LDI layers.

2.3 Collision Detection

The computed LDIs can be utilized to process a variety of collision queries in an efficient way. Our current implementation comprises two different types of collision queries:

Intersection volume. Two LDIs can be combined to compute an intersection volume. The LDIs for two colliding objects are discretized with the same resolution on corresponding sampling grids, but with opposite viewing directions. Hence, pixels in both LDIs represent corresponding volume spans. Therefore, intersection volumes can be computed by a pixelwise intersection of the inside regions of both LDIs. If the resulting intersection is non-empty, a collision is detected. The sum of all pixelwise intersection regions constitutes a discrete representation of the intersection volume (see Fig. 6(a)).



(a) Intersection volume.

(b) Vertex-in-volume.

Figure 6: Two types of collision queries. Pixelwise intersection and testing of the LDIs.

Vertex-in-volume. Individual vertices can also be tested against an LDI. To this end, the vertex is transformed into the local coordinate system of the LDI. If the transformed vertex intersects with an inside region, a collision is detected (see Fig. 6(b)). As demonstrated in Fig. 11 and Fig. 12, this query can be used to compute collisions with atomic objects, such as particles.

3 Implementation

The implementation of VoI computation (stage 1) and collision queries (stage 3) can be easily accomplished by following the descriptions in Sec. 2.1 and Sec. 2.3, respectively. This section describes some details regarding the efficient implementation of the LDI generation (see Sec. 2.2) using core OpenGL 1.4 functionality.

The LDI generation is performed in a multi-pass rendering process following the rendering setup as described in Sec. 2.2. To improve performance, the color buffer is disabled. Each LDI layer is generated in one rendering pass. Except for the first pass, which computes the depth complexities per pixel, all passes are identical. Our actual OpenGL implementation is very similar to the following pseudo code.

```

rendering setup;
// pass 1 computes  $n_{max}$  and the first LDI layer
glClear(GL_DEPTH_BUFFER_BIT |
        GL_STENCIL_BUFFER_BIT);
glStencilFunc(GL_GREATER, 1, 0xff);
glStencilOp(GL_INCR, GL_INCR, GL_INCR);
render object;
depth complexities  $\leftarrow$  stencil buffer;
 $n_{max} \leftarrow \max(\text{depth complexities});$ 
layer [1]  $\leftarrow$  depth buffer;
// passes 2 to  $n_{max}$  compute the remaining LDI layers
 $n \leftarrow 2;$ 
while  $n \leq n_{max}$  begin
    glClear(GL_DEPTH_BUFFER_BIT |
            GL_STENCIL_BUFFER_BIT);
    glStencilFunc(GL_GREATER, n, 0xff);
    glStencilOp(GL_KEEP, GL_INCR, GL_INCR);
    render object;
    layer [ $n$ ]  $\leftarrow$  depth buffer;
     $n \leftarrow n + 1;$ 
end;

```

Each rendering pass actually requires a buffer read-back. Reading back buffers from the GPU, however, can be expensive depending on the actual architecture. We propose two methods for optimizing the data transfer using OpenGL extensions. Firstly, depth and stencil values of a pixel are usually stored in the same 32-bit word in the framebuffer. This allows to read-back both buffers in a single pass using an OpenGL extension, such as `GL_NV_packed_depth_stencil`. Secondly, all rendering passes can be performed independently from each other except for the first pass. We exploit this fact to render individual layers into different regions of the framebuffer. Once finished, the whole framebuffer is read back in a single pass. This optimization reduces the number of read-backs to a maximum of two, assuming that the framebuffer memory is sufficiently large. It reduces stalls in the rendering pipeline and substantially improves the performance of the algorithm.

4 Results

All of the experiments described in this section have been performed on a PC with 2.8 GHz Intel Pentium IV, 2 GB main memory, and NVidia GeForce 4 Ti 4600 graphics card with 128 MB on-board memory. Results for the two types of collision queries are presented using five different test objects (see Fig. 7).

Intersection volume. In Tab. 1, computing times for six cases are given. Since the performance of our algorithm depends on the geometrical complexity of the object and on the depth complexities of both objects within the VoI, we have carried out experiments as

shown in Fig. 8 and Fig. 9. In the worst case, when both AABBs coincide, the VoI and the depth complexity are maximal. Hence, both LDIs are computed for the entire object. The performance of these cases can be derived from Tab. 2, where timings for the LDI generation for entire objects are given.

Objects	Faces	Depth Complexity	Stage 1+2 [ms]	Stage 3a [ms]	Fig.
Knot / Rabbit	12k / 50k	3 / 1	19	1	8
Knot / Rabbit	12k / 50k	4 / 6	56	1	9
Rabbit / Reptile	50k / 110k	7 / 5	36	1	8
Rabbit / Reptile	50k / 110k	4 / 2	48	1	9
Reptile / Santa	110k / 150k	1 / 3	25	1	8
Reptile / Santa	110k / 150k	8 / 6	114	1	9

Table 1: Computation of the intersection volume using various object pairs as shown in Fig. 8 and Fig. 9. The resolution of the LDI is 128x128. The depth complexities are given for both objects within the VoI. Stage 1 and 2 compute the AABBs, the VoI, and both LDIs. Stage 3a generates the intersection volume. In the case of rigid objects, stages 1 and 2 can be pre-computed.

An additional scenario with a very complex intersection volume is illustrated in Fig. 10. The Knot and the Dragon have 12k faces and 520k faces, respectively. Depth complexities within the VoI are 7 and 8. With an LDI resolution of 128x128 the intersection volume can be computed in 256ms.

Vertex-in-volume. Computing times for the vertex-in-volume test are given in Tab. 2. For each object, 100,000 arbitrarily positioned vertices within the AABB of an object are tested. Fig. 11 illustrates the experiment for one of the test objects.

Object	Faces	Depth Complexity	Stage 1+2 [ms]	Stage 3b [ms]
Knot	12k	8	11	26
Rabbit	50k	8	29	27
Reptile	110k	12	83	28
Santa	150k	8	79	26
Dragon	520k	10	307	27

Table 2: Vertex-in-volume test for various objects. The VoI computed in stage 1 coincides with the AABB of the object. Stage 2 computes the LDI representation for the entire object with a resolution of 128x128. In stage 3, 100,000 vertices with arbitrary positions within the AABB of the object are tested. In the case of rigid objects, stages 1 and 2 can be pre-computed. Fig. 11 illustrates the test for one object.

An application of the vertex-in-volume test is shown in Fig. 12. The simulation environment includes 20,000 particles and the deforming Santa with 10k faces. The object deformation, the particle dynamics, the collision detection and response can be computed and rendered in real time with 29 fps.

5 Conclusions

5.1 Discussion

We have presented a fast, hardware-accelerated method for volumetric collision detection. As a central acceleration structure our algorithm computes an LDI that approximates the intersection volume. Unlike many existing collision detection schemes, our method

does not require expensive setup or pre-processing and is, hence, specifically suited for deformable objects and dynamic simulation. Currently, we have implemented two different types of collision queries and demonstrated that the method performs at interactive rates up to very complex geometry.

As a limitation, our collision detection approach is restricted to objects with closed, watertight surfaces, since it relies on the notion of inside and outside.

5.2 Future Work

Currently, collision queries are computed on the CPU. The performance of the algorithm could be improved by using the GPU. Employing pixel shaders with multiple render targets for LDI generation might support GPU-based queries.

While we compute an explicit representation of the intersection volume, additional queries, such as penetration depth or closest point, might be needed for advanced collision response. This is also subject to future research.

At this point, self-collisions are not detected. Actually, our algorithm could be extended towards self-collisions by treating the front and back faces separately during the rendering process. This would allow us to efficiently detect an invalid ordering of front and back faces.

6 Acknowledgements

This project is funded by the National Center of Competence in Research ‘‘Computational Medicine’’ (NCCR Co-Me).

References

- BACIU, G., WONG, W. S.-K., AND SUN, H. 1999. RECODE: an image-based collision detection algorithm. *The Journal of Visualization and Computer Animation* 10, 181–192.
- BARAFF, D. 2001. Collision and contact. *SIGGRAPH 2001 Course Notes*.
- BRIDSON, R., FEDKIW, R., AND ANDERSON, J. 2002. Robust treatment of collisions, contact and friction for cloth animation. In *Proc. of the 29th annual conference on Computer graphics and interactive techniques*, ACM Press, 594–603.
- DEBUNNE, G., DESBRUN, M., CANI, M.-P., AND BARR, A. H. 2001. Dynamic real-time deformations using space & time adaptive sampling. In *Proc. of the 28th annual conf. on Computer graphics and interactive techniques*, ACM Press, 31–36.
- EVERITT, C. 2001. Interactive order-independent transparency. Technical report, NVIDIA Corporation.
- FOLEY, J. D., VAN DAM, A., FEINER, S. K., AND HUGHES, J. F. 1990. *Computer Graphics: Principles and Practice*. Addison-Wesley Publishing Company.
- GOLDFEATHER, J., HULTQUIST, J. P. M., AND FUCHS, H. 1986. Fast constructive-solid geometry display in the pixel-powers graphics system. In *Proc. of the 13th annual conference on Computer graphics and interactive techniques*, ACM Press, 107–116.
- GOTTSCHALK, S., LIN, M. C., AND MANOCHA, D. 1996. Obintree: a hierarchical structure for rapid interference detection. In *Proc. of the 23rd annual conference on Computer graphics and interactive techniques*, ACM Press, 171–180.
- HIROTA, G., FISHER, S., AND LIN, M. C. 2000. Simulation of non-penetrating elastic bodies using distance fields. Technical Report TR00-018, University of North Carolina at Chapel Hill.
- HOFF III, K. E., ZAFERAKIS, A., LIN, M. C., AND MANOCHA, D. 2001. Fast and simple 2d geometric proximity queries using graphics hardware. In *Proc. of 2001 Symposium on Interactive 3D Graphics*, 145–148.
- HUBBARD, P. M. 1993. Interactive collision detection. In *Proc. IEEE Symposium on Research Frontiers in Virtual Reality*, 24–31.
- HUGHES, M., DIMATTIA, C., LIN, M. C., AND MANOCHA, D. 1996. Efficient and accurate interference detection for polynomial deformation and soft object animation. In *Proc. of Computer Animation*, 155–166.

KIM, Y. J., OTADUY, M. A., LIN, M. C., AND MANOCHA, D. 2002. Fast penetration depth computation for physically-based animation. In *Proc. of the ACM SIGGRAPH symposium on Computer animation*, ACM Press, 23–31.

KIM, Y. J., HOFF III, K. E., LIN, M. C., AND MANOCHA, D. 2003. Closest point query among the union of convex polytopes using rasterization hardware. *Journal of Graphics Tools*, to appear.

KLOSOWSKI, J. T., HELD, M., MITCHELL, J. S. B., SOWIZRAL, H., AND ZIKAN, K. 1996. Efficient collision detection using bounding volume hierarchies of k-DOPs. In *Proc. of the 23rd annual conference on Computer graphics and interactive techniques*, ACM Press, 171–180.

LARSSON, T., AND AKENINE-MOELLER, T. 2001. Collision detection for continuously deforming bodies. In *Proc. of Eurographics 2001*, 325–333.

LOMBARDO, J. C., CANI, M.-P., AND NEYRET, F. 1999. Real-time collision detection for virtual surgery. In *Proc. of Computer Animation '99*, 33–39.

QUINLAN, S. 1994. Efficient distance computation between non-convex objects. In *Proc. IEEE Int. Conference on Robotics and Automation*, IEEE, 3324–3329.

RAPPOPORT, A., AND SPITZ, S. 1997. Interactive boolean operations for conceptual design of 3-D solids. In *Proc. of the 24th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., 269–278.

SHADE, J., GORTLER, S., WEI HE, L., AND SZELISKI, R. 1998. Layered depth images. In *Proc. of the 25th annual conference on Computer graphics and interactive techniques*, ACM Press, 231–242.

VAN DEN BERGEN, G. 1997. Efficient collision detection of complex deformable models using aabb trees. *Journal of Graphics Tools* 2, 4, 1–13.




Figure 7: Test objects: Dragon, Santa, Rabbit, Reptile, Knot.




Figure 8: Computation of the intersection volume. AABBs, VoI, and intersection volume are shown for three examples with simple intersections.




Figure 9: Computation of the intersection volume. AABBs, VoI, and intersection volume are shown for three examples with complex intersections.




Figure 10: Intersection volume for Knot and Dragon.




Figure 11: Vertex-in-volume test. Left: The LDI representation is computed for the entire object. VoI and ABB coincide. Right: The vertex-in-volume test is processed for 100,000 arbitrarily positioned vertices.




Figure 12: Deforming Santa and particles. Red color on the right-hand side indicates actual or recent contact.