CLODs: Dual Hierarchies for Multiresolution Collision Detection

Miguel A. Otaduy and Ming C. Lin

Department of Computer Science
University of North Carolina at Chapel Hill

Abstract

We present a novel framework for multiresolution collision detection using "*contact levels of detail*" (CLOD). Given a polyhedral model, our algorithm automatically builds a "dual hierarchy". A dual hierarchy is both a multiresolution representation of the original model and its bounding volume hierarchy for accelerating collision queries. We have proposed various error metrics, including object-space errors, velocity dependent gap, screen-space errors and their combinations. At runtime, our algorithm uses these error metrics to select the appropriate levels of detail *independently* at each potential contact location. Compared to the existing exact collision detection algorithms, we observe significant performance improvement using CLODs on some benchmarks, with little degradation in the visual rendering of simulations.

1 Introduction

Collision detection is an important problem in physically-based modeling, computer animation, virtual environments, electronic prototyping, robotics and many other engineering applications. Collision detection is often one of the computational bottlenecks in achieving interactive simulation for complex environments. The running time of any collision query algorithm depends on both the input size (the combinatorial complexity of the models) and the output size (the number of contact points). Despite the huge body of literature, the existing techniques cannot offer the desired performance for many real-time applications, such as haptic rendering of large structures, real-time interaction with massive CAD models, and interactive dynamic simulations of complex objects.

In this paper, we investigate the use of simplification techniques for collision detection. Model simplification has been an active research area for the past decade. Many efficient and effective simplification algorithms have been proposed to accelerate the visual rendering of highly complex models in real time. Applications of mesh simplification algorithms to the problem of collision detection can potentially accelerate collision queries. However, only relatively simple algorithms have been proposed for convex polytopes. A simple approach would be to generate a series of simplified representations, also known as levels-of-detail (LOD), and use them directly for collision detection. But, collision queries require auxiliary data structures, such as bounding volume hierarchies (BVH) or spatial partitioning, to achieve good runtime performance. These data structures often take a substantial amount of memory, while the LODs themselves can require a large amount of storage as well. To represent the auxiliary data structures for each static LOD can quickly increase the memory requirement to the point where the memory access time may dominate the overall runtime performance. On the other hand, constructing BVHs for on-the-fly simplification can be time consuming.

Main Contribution: We present a general framework for constructing "*contact levels-of-detail*" (CLOD) for accelerating collision queries. The resulting set of simplified representations is a dual hierarchy. It serves both as a BVH for accelerating collision queries and a multiresolution representation of the original model for computing contact information. We have also proposed several error metrics, including object-space errors, velocity dependent gap, screen-space errors and their combinations. At runtime, the algorithm uses these error metrics to select adaptive CLOD at each potential contact location *independently*. In addition, the framework offers the capability to perform time-critical collision queries for real-time applications. We have tested our approach on a diverse set of benchmarks with challenging contact scenarios. We observed noticeable performance improvement with little degradation in the visual rendering of simulation results on some of the benchmarks.

Organization: Section 2 briefly surveys previous work. We present the framework on contact levels-of-detail for collision queries in section 3. Section 4 and 5 describes a specific implementation of hierarchy construction using convex hulls and its use for runtime contact queries respectively.

We discuss the implementation issues and demonstrate the effectiveness of our approach in section 6.

2 Related Work

This research is built upon a large collection of knowledge in mesh simplification and collision detection. We briefly survey related work in this section.

2.1 Model Simplification

Many techniques for mesh decimation and model simplification have been proposed for the last few years. We refer readers to an excellent book on this subject[Luebke et al. 2002]. There has been some growing interest in perception-based simplification for interactive rendering[Luebke and Hallen 2001]. However, these techniques are based on different criteria than those for collision detection. Although we precompute the level of detail (LOD) hierarchy off line, the way we select the appropriate LOD on the fly is "contact dependent" at each potential contact location across the object surfaces. It shares a more common theme with view-dependent simplification[Luebke et al. 2002], which uses higher resolution representations on the silhouette of the object and much coarser approximations on the rest of the object that is not as noticeable to the viewpoint.

El-Sana and Varshney proposed to construct a continuous, multiresolution hierarchy of the model during preprocessing[El-Sana and Varshney 2000]. At run time, a high-detail representation is used for regions around the probe point and a coarser representation farther away using a bell-shaped filter. The proposed approach only applies to haptic rendering using a point probe exploring a 3D model. It uses the commercial haptic rendering toolkit, GHOST, to perform point-object collision query, which has a much lower computational complexity. This approach does not extend naturally to collision queries between two interacting 3D objects, since multiple disjoint contacts can occur simultaneously at widely varying locations without much spatial coherence. The latter problem is the focus of our paper.

2.2 Collision Detection

Data structures based on hierarchical representations have been extensively used in the design of efficient algorithms for collision detection between geometric models (see survey[Klosowski et al. 1998; Lin and Gottschalk 1998]). Examples of typical bounding volumes include spheres and axis-aligned bounding boxes, due to their simplicity in performing overlap tests between two such volumes. Other hierarchies include octrees, k-d trees, K-DOPs, OBBTrees, trees based on S-bounds, R-trees and their variants, etc.[Klosowski et al. 1998; Lin and Gottschalk 1998]. Other spatial representations are BSP's and their extensions to multi-space partitions, space-time bounds or four-dimensional tests (see a brief survey[Redon et al. 2002]), and many more.

The concept of time-critical collision detection was first introduced by Hubbard using sphere-trees[Hubbard 1994]. Collision queries are performed as far down the sphere-trees as time allows, without traversing the entire hierarchy. The contact information was derived from two colliding bounding spheres and could deviate arbitrarily from the actual contact normals and contact locations. This idea can be applied to any type of bounding volume hierarchies (BVH). An error metric is often desirable for interactive applications to formally and rigorously quantify the amount of error introduced. However, no tight error bounds have been provided using such approaches. This can be problematic, especially when contact normals and contact points are required to compute a plausible collision response. Other techniques that exploit hierarchical representations and motion coherence for fast distance computation between only convex polytopes have also been proposed[Ehmann and Lin 2000; Guibas et al. 1999].

O'Sullivan, et al. investigated LOD techniques for collision simulations and studied various factors affecting collision perception, including separation, eccentricity, causality, distractors, and accuracy of simulation results[O'Sullivan and Dingliana 2001]. Their work is based on the model of human visual perception and validated by psychophysical experiments. The feasibility of using these factors for scheduling interruptible collision detection among large numbers of visually homogeneous objects is also demonstrated. Instead of addressing the scheduling of multiple collision events among many objects, we focus primarily on the problem of contact queries between two highly complex objects.

Recently techniques based on GPU acceleration have also been proposed for collision queries[Hoff et al. 2001; Lombardo et al. 1999]. However, these approaches are not necessarily faster for rigid bodies where precomputation can be effectively carried out off-line, since the readback from frame-buffer and depth buffer cannot be done fast enough.

3 Framework for Contact Levels of Detail

In this section, we present our approach for constructing hierarchies of CLODs. First we analyze the advantage of using CLODs in hierarchical collision detection, then we address the problems of building such CLODs and the process for creating them, and finally we explain how CLODs are used in runtime collision queries.

3.1 Hierarchical Collision Detection

Bounding volume hierarchies (BVHs) are commonly used for collision detection between general geometric objects. To perform intersection tests, two models are queried by recursively traversing their BVHs in tandem. Each recursive step tests whether a pair of bounding volumes a and b, one from each hierarchy, overlap. If a and b do not overlap, the recursion branch is terminated. Otherwise, if they do overlap, the algorithm is applied recursively to their children. If a and b are both leaf nodes, the primitives within them are tested directly.

The test between the two BVHs can be described by the bounding volume test tree (BVTT)[Larsen et al. 2000], a tree structure that holds in each node the result of the queries between two BVHs. When temporal coherence is present, collision tests can be accelerated by *generalized front tracking* (GFT)[Ehmann and Lin 2001]. GFT caches the front of the BVTT, where the result of the queries switches from "true" to "false", for collision query in the next time step. The overall cost of a collision test is proportional to the number of nodes in the front of the BVTT.

When large areas of the two objects are in a close proximity, a larger portion of the BVTT front is close to the leaves, and it consists of a larger number of nodes. The size of the front also depends on the resolution at which the objects are modeled; higher resolution implies a BVTT with a greater depth. We can draw the conclusion that the cost of a collision query depends on two key factors:

- The size of the contact area
- The resolution of the models

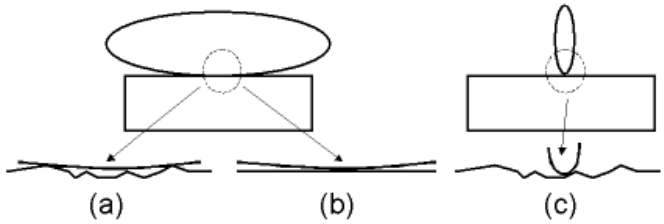

Figure 1: a) Large contact area in high resolution; b) Large contact area in low resolution; c) Small contact area in high resolution

The contact between two real objects typically occurs along a certain contact area. With polygonalized models, this may result in multiple contact points. The finer the resolution of the objects, the larger the number of contact points, as seen in Fig. 1. However, employing a larger resolution may have little effect on the forces computed between the objects, because these forces are computed as a sum of contact forces arising from a net of contact points. We can argue that intuitively a larger contact area allows the objects to be described at a coarser resolution. In this paper we exploit this hypothesis to create multiresolution representations of the objects, and use them at each contact location independently for selecting the appropriate resolution of each approximation. Using CLODs we can achieve nearly constant cost for collision queries by exploiting, among other factors, the relation between contact area and the resolution of local contact features.

3.2 Construction of Dual Hierarchies

In order to perform efficient multiresolution collision detection, we need to achieve two main objectives:

1. Create *accurate multiresolution representations.*
2. Embed the multiresolution representations in *effective bounding volume hierarchies.*

We create multiresolution representations by performing mesh decimation on the given polyhedral models. The difficulty arises when trying to embed these representations in BVHs. If we consider

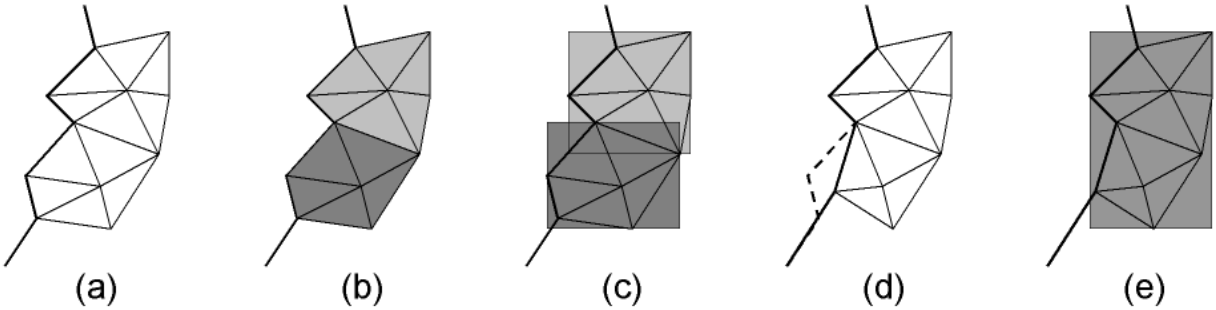

Figure 2: Constructing Dual Hierarchies: (a) Initial surface; (b) Clusters of triangles; (c) BVs for each cluster (in this case, AABBs); (d) Mesh simplification; (e) BV of the union of clusters after some conditions are met.

each LOD of the object as one whole model, each LOD would require a distinct BVH for collision detection. This would result in a very inefficient collision query, because the front of the BVTT would have to be updated for the BVH of each LOD. Instead, we introduce a procedure to create one unique dual hierarchy that serves as both a multiresolution representation and a BVH.

The root of the hierarchy will be the BV of the coarsest LOD. Descending to the next level of the hierarchy will yield the children BVs, whose union encloses the next LOD. At the end of the hierarchy, the leaf BVs will enclose the original surface.

The process of creating the CLODs, depicted in Fig. 2, starts by grouping the triangles of the original surface into clusters. The size and properties of these clusters depend on the type of BV that is used for the BVH, and will be such that the collision query performance between two BVs is optimized. The next step is to compute the BV of each cluster. After this initialization, we start a mesh decimation process with a bottom-up construction of the BVH. This is achieved by merging clusters and computing the BV of their union.

The atomic simplification operations need to satisfy the following:

- **Constraints imposed by the BVH:** The BVs have to remain *valid* after each simplification operation. This may impose topological and/or geometric constraints.

- **Design requirements to achieve better efficiency:** The union of clusters is possible when certain conditions are met. The BVH will be more effective in collision pruning, if these conditions are taken into account when designing the atomic simplification operations.

In Sec. 4, we present a specific implementation of CLODs that uses convex hulls as the BVs.

3.3 Runtime Queries with CLOD

Using CLODs, multiresolution collision detection can be implemented by slightly modifying the typical collision detection procedures using BVHs. For each node x of the BVTT, we perform a collision query. If the query returns "false", we do not need to descend to the children. If the query result is "true", then we perform a test for selective refinement. This test can embed various perceptual error metrics, and it determines if the resolution of x is fine enough to describe the contact information at each query location. If the refinement test returns "true", then we can directly compute contact information for x, otherwise we descend to its children in the BVTT. Descending to the children involves descending to the children BVs, as well as refining the surface representation. This approach handles the selective refinement at each query location independently in a very efficient way.

CLODs can also be used to perform time-critical collision detection[Hubbard 1994]. We need to store the nodes of the BVTT front, and assign priorities to each node based on the refinement test.

In Sec. 5 we describe how we implement the collision queries using CLODs based on convex hulls, as well as various error metrics that we have designed.

4 Dual Hierarchies of Convex Hulls

In this section we describe a particular implementation of CLODs using BVHs of convex hulls. We first address the reasons for choosing convex hulls as the BVs, and then describe the details of constructing dual hierarchies.

4.1 Selection of the BVs

Overlap tests between convex hulls can be executed rapidly in expected constant time with motion coherence[Guibas et al. 1999]. Furthermore, convex hulls provide at least equally good, if not superior, fitting to the underlying geometry as OBBs[Gottschalk et al. 1996] or k-dops[Klosowski et al. 1998].

Most importantly the fitting property is related to the performance of the query to obtain contact information of a certain CLOD. As explained in Sec. 3.3, when the refinement test determines that a CLOD does not need to be refined, we must get contact information at that level. That implies getting contact information from the triangles of that specific CLOD. If the BVs at that level are one of AABB, OBB or k-dops, the efficiency of getting contact information using triangles is related to the number of triangles in each cluster. However, with convex hulls, if we ensure that the clusters are themselves convex surface patches, the contact information at the triangle level is obtained for free when performing the query between BVs[Ehmann and Lin 2001].

The initial clusters will be defined as the surface patches of a convex surface decomposition[Ehmann and Lin 2001]. The definition of convex surface patches imposes two types of convexity constraints on the process of creating dual hierarchies:

- **Local constraints**: the interior edges of convex patches have to remain convex after simplification operations are applied.
- **Global constraints**: the enclosing convex hulls cannot intersect with the surface of the object.

However, it is important to note that convex hulls pose some limitations on the types of models that can be handled. In our implementation, the convex hulls are formed by convex surface decomposition, which requires the input models to be described as 2-manifold oriented surfaces.

4.2 Construction of the Dual Hierarchy

As already mentioned, the process is initialized by performing a convex surface decomposition of the object and computing the convex hulls of the patches. This is followed by a simplification loop represented schematically in Fig. 3.

We apply atomic simplification operations taking into account convexity constraints. After each operation, we test if the union of every pair of neighboring convex patches is a valid convex patch. If so, we merge the patches, construct the convex hull of the union, and set parent children relationships between the BVs. We have chosen to generate a new LOD every time that the number of convex patches is halved.

Due to the convexity constraints, we might reach a point where no more simplification operations are possible. We complete the construction of the hierarchy by unconstrained pairwise merging of patches[Ehmann and Lin 2001]. The levels of the hierarchy created in this manner cannot be used as CLODs in collision queries, but are necessary to complete the BVH.

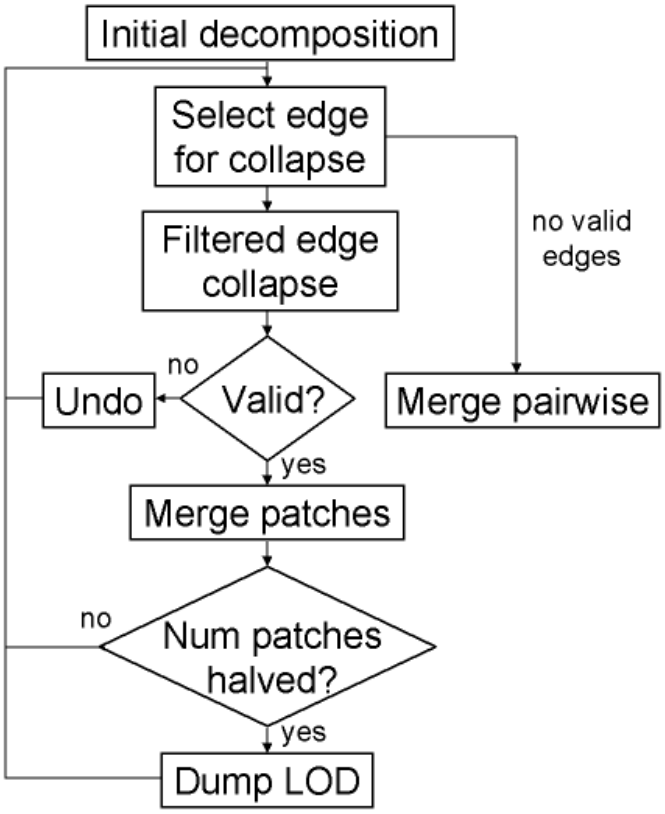

Figure 3: Generation of CLODs

There are various aspects of the simplification process central to our framework of CLODs that need to be defined:

- The type of simplification operation
- The assignment of priorities for simplification
- The local retriangulation after the simplification

We select edge collapse as the atomic simplification operation. There are many options for prioritizing edges and selecting the position of the resulting vertices: minimization of energy functions[Hoppe 1996], optimization approaches[Lindstrom and Turk 1998], quadric error metrics for

measuring surface deviation[Garland and Heckbert 1997], and more. However, none of these approaches meets the convexity constraints or takes into account the factors that maximize the efficiency of CLODs. Our local simplification operations are inspired by multiresolution analysis and signal processing of meshes.

4.2.1 Multiresolution Analysis of Meshes

An LOD M^j of a mesh M at resolution r_j is defined as an approximation of M that stores all the surface detail at resolutions lower than r_j. Following this definition, we have decided to prioritize the edges to be collapsed based on their resolution. We use the same definition of edge resolution as in[Otaduy and Lin 2003], which is inspired by the definition of resolution for functions in the 1D setting. By the same definition of LODs, the collapse of an edge at resolution r_j involves removing the geometric detail at resolutions higher than r_j. We implement this by placing the vertices resulting from edge collapse operations using Guskov's minimization of second order divided differences[Guskov et al. 1999].

We have compared this approach to other techniques for placing the vertices resulting from edge collapses, and it has proved to accelerate the process of merging convex patches. This is probably because unconstrained filtering tends to remove concavities at high resolution, while other techniques fail to remove these concavities in favor of preserving other properties. Note that a more efficient merging of convex patches has beneficial consequences in the performance of CLODs at runtime.

4.2.2 Filtered Edge Collapse

We define a local simplification operation **filtered edge collapse** subject to convexity constraints[Ehmann and Lin 2001; Otaduy and Lin 2003]. This operation is composed of the following steps:

1. A topological edge collapse.

2. An initialization process that sets the position of the resulting vertex using quadric error metrics.

3. Unconstrained filtering of the position of the vertex.

4. Solving an optimization problem to minimize the distance of the vertex to its unconstrained position while taking into account the local convexity constraints.

5. A bisection search to find a location where the global convexity constraints are also met.

The local convexity constraints for each edge are linear in the position of the vertex v, and can be expressed in terms of the neighboring vertices (see Fig. 4) as:

$$((v - a) \times (b - a)) \cdot (c - a) \leq 0 \tag{1}$$

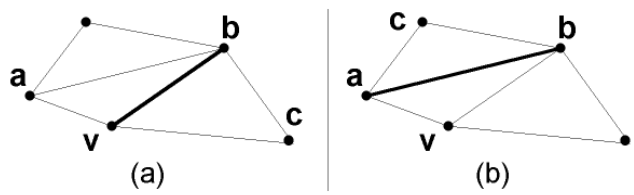

Figure 4: Assignment of vertices for local convexity constraints: a) incident edges; b) opposite edges

In Fig. 5 we show an example of a dual hierarchy of CLODs. The process is applied to the model of a wrinkled torus. Fig. 5-b shows a detailed view of the convex decomposition of the original surface. In the next CLODs, the filtering smoothens the concave areas, and the convex patches grow forming triangle strips on the hyperbolic regions. Likewise, the simplification operations coarsen the triangulation. Also notice that initially the simplification operations are concentrated in the region where the wrinkles merge and is also the area with the finest detail.

5 Runtime Queries

In this section we discuss our implementation of collision queries using CLODs of convex hulls, present various error metrics for selecting the appropriate CLODs, and demonstrate the application to rigid body simulation.

5.1 Collision Queries between CLODs of Convex Hulls

A typical collision query between two BVs a and b determines if a and b are closer than a distance δ. The collision detection algorithms based on BVHs must ensure that if a leaf node x of the BVTT

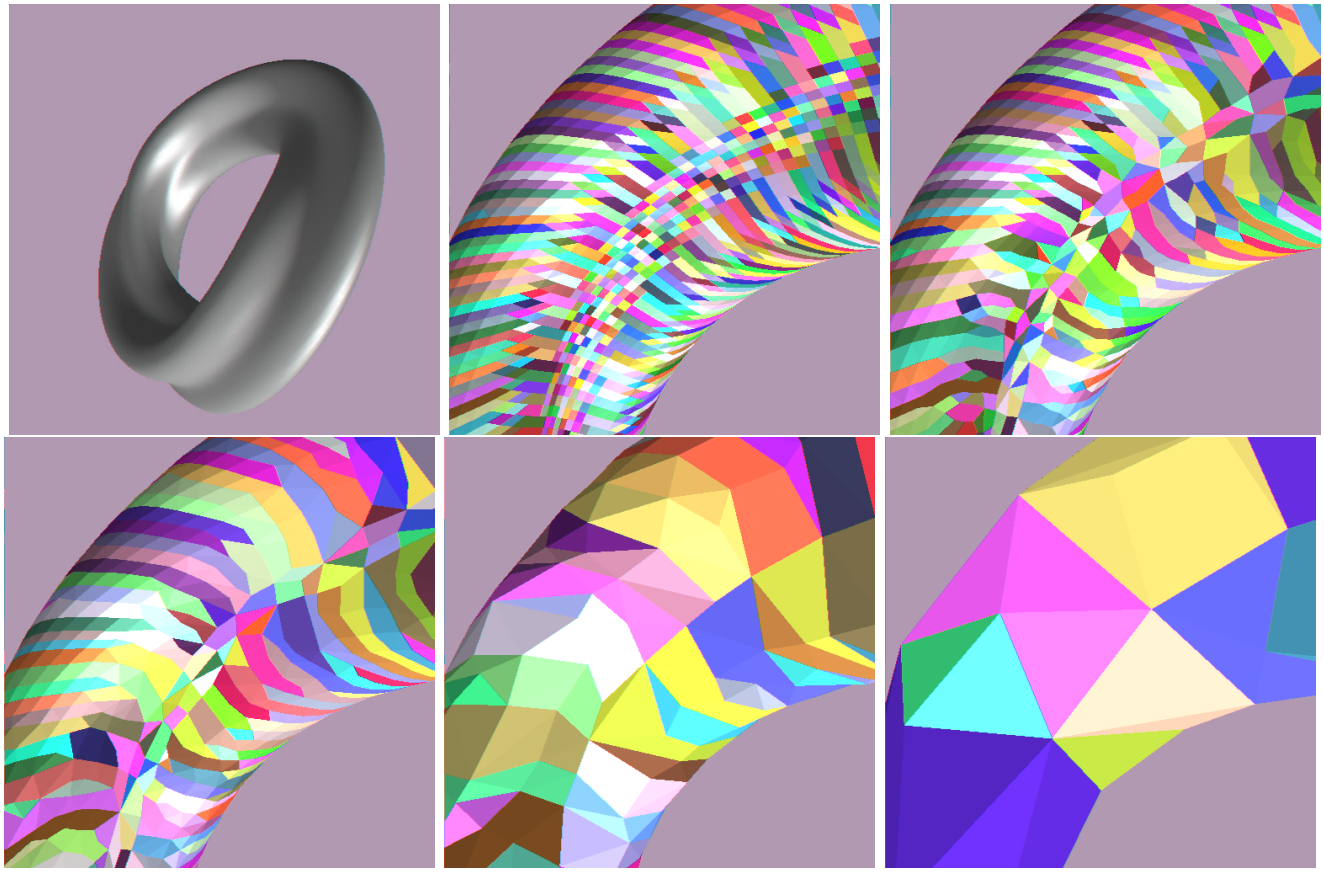

Figure 5: From left to right and top to bottom, original mesh and detail view of $CLOD_0$, $CLOD_1$, $CLOD_2$, $CLOD_4$ and $CLOD_6$.

returns "true" to the collision query, then all its ancestors must return "true" too. This is usually achieved by containing the surface of each object inside the union of the BVs at every level of its BVH. In our implementation of CLODs this containtment property does not hold, but we can ensure the correctness of the collision detection by modifying the collision distance δ_{ab} between two convex hulls a and b:

$$\delta_{ab} = \delta + h(a^i, a) + h(b^j, b) \tag{2}$$

where $h(a^i, a)$ and $h(b^j, b)$ are maximum directed Hausdorff distances from the descendant BVs of a and b to a and b respectively.

5.2 Error Metrics

We have designed a functional ϕ to evaluate the resulting error, when we stop a collision query at a node ab of the BVTT.

$$\phi_a = \frac{s_a}{r_a^2} \tag{3}$$

where s_a is the surface deviation of the convex hull a with respect to the original surface and r_a is its resolution. This functional can be regarded as a measure of the volume of the fictitious features that are filtered out when using a as the CLOD.

As introduced in Sec. 3.1, larger contact areas allow the models to be described at coarser resolution. We take into account this observation by averaging the functional ϕ over an estimated contact area D. Thus, we compute a weighted surface deviation s^* as:

$$s^*_{ab} = \frac{max(\phi_a, \phi_b)}{D} \tag{4}$$

s^* can been considered as the surface deviation errors weighted by a constant that depends on both the contact area and the resolutions of local surface features. Contact area is too expensive to be computed in run time. Instead, we compute support areas of the vertices of the models as a preprocessing, and we use those to estimate the contact area.

The node of the BVTT ab is refined if s^*_{ab} is larger than a threshold value, s_0. This threshold can be determined based upon different perceptual metrics.

Size Dependant Metric

s_0 will be determined as a fixed percentage of the radii of the objects involved in the contact queries.

Velocity Dependant Metric

Set s_0 as a value proportional to the relative velocity of the colliding objects at the contact location. This is based on the fact that the gap between the objects is less noticeable as the objects move faster.

View Dependant Metric

Determine s_0 based on screen-space errors. Given N pixels of error, a distance l from the camera to the contact location, a distance n to the near plane of the view frustum, a size f of the frustum in world coordinates, and a size i of the image plane in pixels:

$$s_0 = \frac{N \cdot l \cdot f}{n \cdot i} \tag{5}$$

Selective refinement using CLODs can be implemented combining any of these error metrics. It can also support other types of perceptual error metrics[O'Sullivan and Dingliana 2001].

Constant Frame-rate: Another important feature is the fact that CLODs can be used for time-critical collision detection. The factor $\frac{s^*}{s_0}$, computed at every potential contact location, can be used to prioritize the refinement. To achieve a guaranteed framerate for real-time applications, the algorithm will perform as many collision queries as possible, within a fixed time interval. The query event queue will be prioritized based on $\frac{s^*}{s_0}$.

5.3 Application to Rigid Body Simulation

In rigid body simulations, many of the factors involved in the selection of CLODs, such as the velocity of the objects, the contact area and the distance to the camera, will vary dynamically. This results in transitions between the refinement tests and thus switching between selected CLODs. Rigid body simulations often require the execution of time-stepping algorithms to search for the time instants prior to collision events. Special treatment is necessary so that switching CLODs do not generate inconsistencies or deadlock situations in the time-stepping algorithms. (The analogy in visual rendering is the "popping" effects often seen during LOD switching.)

Given a node x_i of the BVTT, with collision query "false" at times t_i and t_{i+1} of the simulation, and a node x_j, child of x_i, with collision query "true" at both time instants, if the refinement test of x_i switches from "false" to "true" at $t \in [t_i, t_{i+1}]$, the time stepping method will encounter an inconsistency. It will try to search for an inexistent collision event in the interval $[t_i, t_{i+1}]$. We solve this problem by estimating a collision time t_c interpolating the separation distance of the node x_i at t_i and the penetration depth of the node x_j at t_{i+1}. We apply the contact response and restart the numerical integration from t_c with x_i as the active CLOD.

6 Implementation and Performance

In this section we describe the implementation issues and the benchmarks used to analyze the performance and effectiveness of CLODs.

6.1 Benchmark Models

Fig. 6 shows the set of benchmark models. They are:

- A wrinkled torus falling along a spiral peg
- A spoon sliding inside a cup
- A soup of numbers settling in a bowl

In Table 1 we present the statistics of the dual hierarchies created to represent the benchmark models. Note that in all cases the models are simplified to a coarsest CLOD with a complexity in the order of 1000 triangles and 100 convex pieces. This limitation is imposed by the topology of the objects and the localized influence of the simplification operations. The statistics of the *numbers* represent an average of all the models from *number* 0 to *number* 9.

Models	Cup	Spoon	Spiral	Torus	Bowl	Numbers
Tris_{orig}	64000	86016	32160	32000	12288	1080
BVs_{orig}	15490	16125	6448	5919	4336	161
Tris_{simp}	1532	1074	716	644	634	306
BVs_{simp}	241	61	100	92	63	37
CLODs	7	9	7	7	8	3
LODs	14	14	13	13	13	9

Table 1: **Models and Associated Hierarchies.** *The number of triangles (Tris_{orig}) and convex patches (BVs_{orig}) of the initial mesh of the models; the number of triangles (Tris_{simp}) and convex patches (BVs_{simp}) of the coarsest valid CLOD; and CLODs and total number of LODs in the hierarchies.*

6.2 Runtime Performance

The rigid body simulations captured in the snapshots in Fig. 6 were computed using impulsive methods[Mirtich and Canny 1995]. These simulations were performed on a Pentium-4 2.4Ghz processor PC with 2.0GB of memory, a NVidia GeForce-FX graphics card and Windows2000 OS.

The timing profile of Fig. 7 shows query times and the number of nodes in the BVTT front for a simulation of the spoon sliding inside the cup. We have compared results obtained with a exact collision detection algorithm[Ehmann and Lin 2001] against the results of our CLOD technique using the size dependant error metric. In particular, we have used values of s_0 of 3.5%, 0.35%, and 0.035% of the radius of the cup. As the timings show, CLODs with $s_0 = 3.5\%$ perform at least as good as the exact algorithm for most of the simulation duration. During several time intervals, we observe a performance gain of almost one order of magnitude. In the other two benchmarks, we have not observed so remarkable performance gains using CLODs.

We have also evaluated CLODs with velocity and view dependant error metrics. In Fig. 8 coarse CLODs are selected when the spoon hits the bottom of the cup, and fine resolution CLODs are selected when the spoon slides along the side of the cup. In the first case the polygon count of the representations is roughly 10 times larger than in the second case.

6.3 Limitations and Analysis

We have chosen a diverse set of benchmarks with varying model complexity and observed the performance of CLODs on different simulation scenarios. As a result of the modification to the collision tolerance introduced in Sec. 5.1, the multiresolution collision detection algorithm using CLODs cannot prune as efficiently as the exact algorithm[Ehmann and Lin 2001], when the objects are separated by a distance notably larger than δ. However, when they come close to contact, the temporal coherence in the front of the BVTT is much higher with CLODs than with the exact algorithm. This can be verified in Fig. 7, where the number of nodes does not vary much with CLODs. This also explains the spikes present in the timings for the exact algorithm. These spikes take place at the instants when the objects are about to interpenetrate. Even if the number of nodes in the front of the BVTT is smaller for the exact algorithm than for CLODs with $s_0 = 0.035\%$, the

coherence is lower, because suddenly many nodes in the BVTT come closer than δ, thus yielding longer query times.

We believe that the insignificant performance gain in the simulation of the torus falling along the spiral is due to the similar reasons. Because of the high traveling speed and the impacts of the falling torus, the simulated objects are rarely in close proximity. As for the simulation of the numbers settling in a bowl, the polygonal complexity of the models of the numbers is not high enough to benefit from CLODs.

To summarize, we can conclude that CLODs perform remarkably better than exact collision detection algorithms (up to one order of magnitude) for situations of sliding or resting contacts between complex models. These are in fact some of the most challenging contact scenarios, because large areas of the objects are in parallel close proximity[Gottschalk et al. 1996].

7 Summary and Future Work

We have introduced a novel framework for multiresolution collision queries using "contact levels of detail" (CLOD). CLOD is a dual hierarchy that serves as both a multiresolution representation of the model and its bounding volume hierarchy for collision queries. At run time, the algorithm dynamically selects adaptive CLOD at each potential contact independently using various error metrics. This approach considerably speeds up the overall performance of collision queries in complex environments with challenging contact scenarios.

This research may be extended in several directions:

- A further investigation of the problem of switching CLODs and its implications on the contact response models of various simulation methods;
- Optimization of the simplification process to obtain more effective hierarchies, perhaps incorporating an operation with a more global support than the current ones;
- The application of the general framework of CLODs to other types of BVHs.

References

Ehmann, S., and Lin, M. C. 2000. Accelerated proximity queries between convex polyhedra using multi-level voronoi marching. *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems*.

Ehmann, S., and Lin, M. C. 2001. Accurate and fast proximity queries between polyhedra using convex surface decomposition. *Computer Graphics Forum (Proc. of Eurographics'2001) 20*, 3.

El-Sana, J., and Varshney, A. 2000. Continuously-adaptive haptic rendering. *Virtual Environments 2000*, pp. 135–144.

Garland, M., and Heckbert, P. S. 1997. Surface simplification using quadric error metrics. In *Proc. of ACM SIGGRAPH*, 209–216.

Gottschalk, S., Lin, M., and Manocha, D. 1996. OBB-Tree: A hierarchical structure for rapid interference detection. *Proc. of ACM Siggraph'96*, 171–180.

Guibas, L., Hsu, D., and Zhang, L. 1999. *H-Walk*: Hierarchical distance computation for moving convex bodies. *Proc. of ACM Symposium on Computational Geometry*.

Guskov, I., Sweldens, W., and Schroder, P. 1999. Multiresolution signal processing for meshes. *Proc. of ACM SIGGRAPH*, pp. 325 – 334.

Hoff, K., Zaferakis, A., Lin, M., and Manocha, D. 2001. Fast and simple geometric proximity queries using graphics hardware. *Proc. of ACM Symposium on Interactive 3D Graphics*.

Hoppe, H. 1996. Progressive meshes. In *SIGGRAPH 96 Conference Proceedings*, Addison Wesley, H. Rushmeier, Ed., Annual Conference Series, ACM SIGGRAPH, 99–108. held in New Orleans, Louisiana, 04-09 August 1996.

Hubbard, P. 1994. *Collision Detection for Interactive Graphics Applications*. PhD thesis, Brown University.

Klosowski, J., Held, M., Mitchell, J., Sowizral, H., and Zikan, K. 1998. Efficient collision detection using bounding volume hierarchies of k-dops. *IEEE Trans. on Visualization and Computer Graphics 4*, 1, 21–37.

Larsen, E., Gottschalk, S., Lin, M., and Manocha, D. 2000. Distance queries with rectangular swept sphere volumes. *Proc. of IEEE Int. Conference on Robotics and Automation*.

Lin, M., and Gottschalk, S. 1998. Collision detection between geometric models: A survey. *Proc. of IMA Conference on Mathematics of Surfaces*.

Lindstrom, P., and Turk, G. 1998. Fast and memory efficient polygonal simplification. *Proc. of IEEE Visualization*, 279–286.

Lombardo, J. C., Cani, M.-P., and Neyret, F. 1999. Real-time collision detection for virtual surgery. *Proc. of Computer Animation*.

Luebke, D., and Hallen, B. 2001. Perceptually driven simplification for interactive rendering. *Rendering Techniques; Springer-Verlag*.

Luebke, D., Reddy, M., Cohen, J., Varshney, A., Watson, B., and Huebner, R. 2002. *Level of Detail for 3D Graphics*. Morgan-Kaufmann.

Mirtich, B., and Canny, J. 1995. Impulse-based simulation of rigid bodies. In *Proc. of ACM Interactive 3D Graphics*.

O'Sullivan, C., and Dingliana, C. 2001. Collisions and perception. *ACM Trans. on Graphics 20*, 3, pp. 151–168.

Otaduy, M. A., and Lin, M. C. 2003. Sensation preserving simplification for haptic rendering. *Proc. of ACM SIGGRAPH*. (To appear).

Redon, S., Kheddar, A., and Coquillart, S. 2002. Fast continuous collision detection between rigid bodies. *Proc. of Eurographics (Computer Graphics Forum)*.

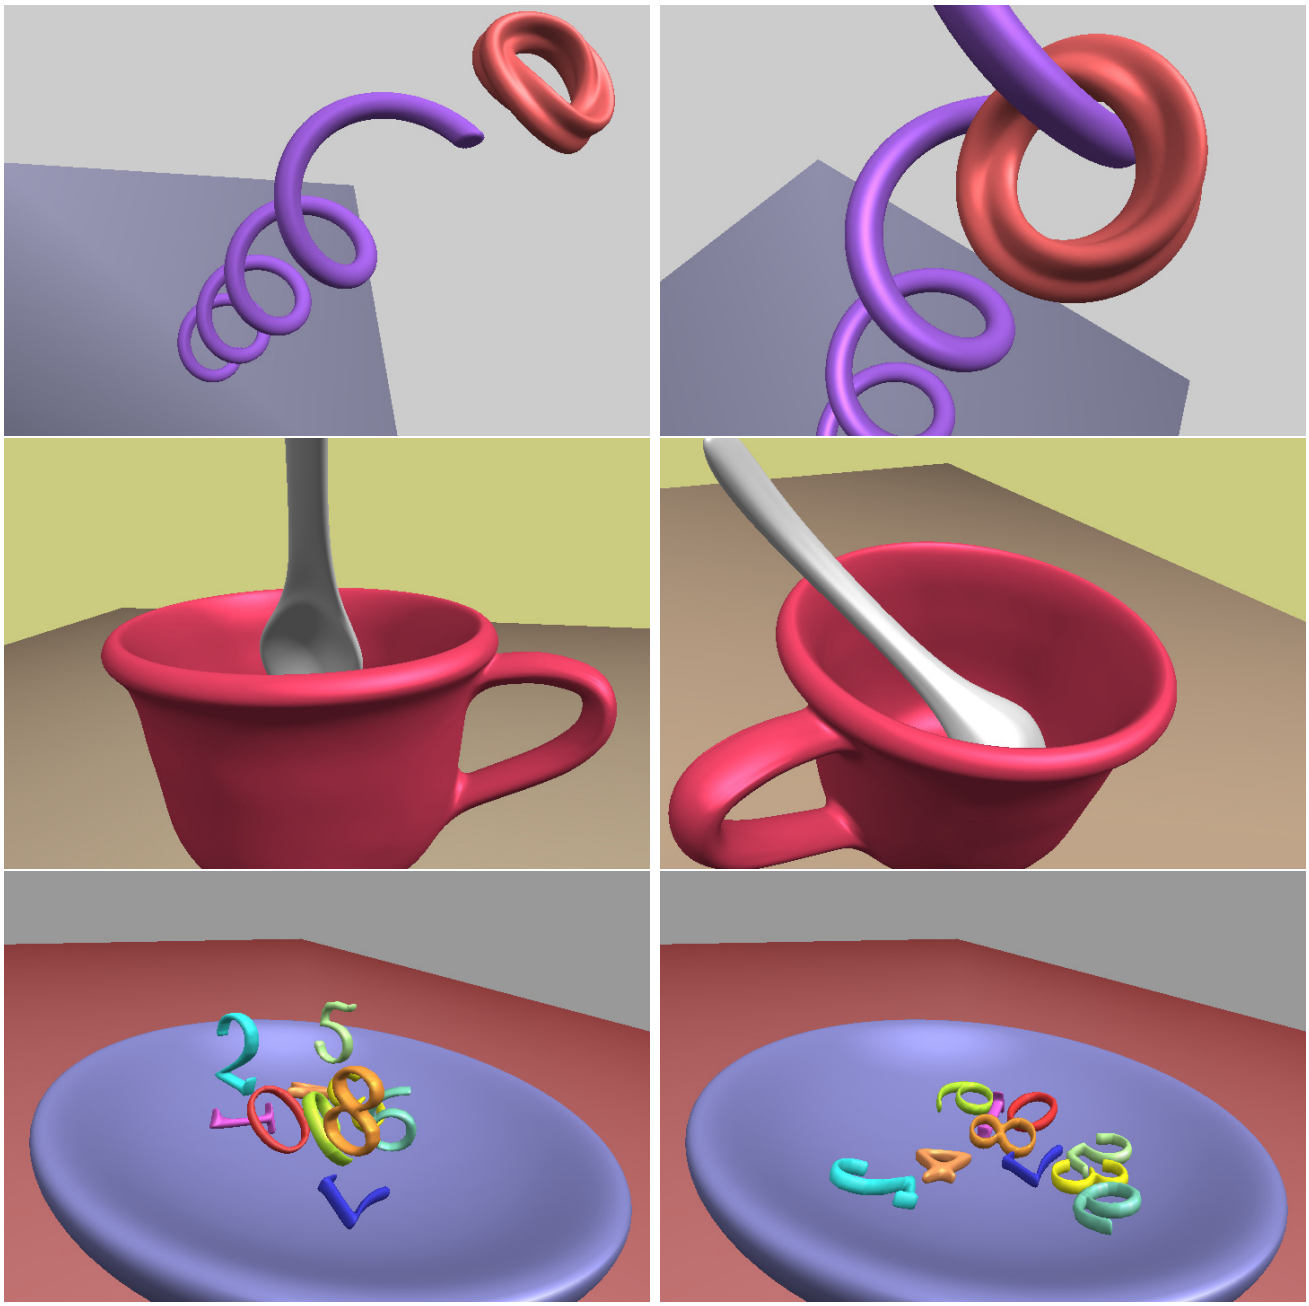

Figure 6: Snapshots of our benchmark simulations. From top to bottom: torus falling along a spiral peg, a spoon in a cup, and a soup of numbers in a bowl.

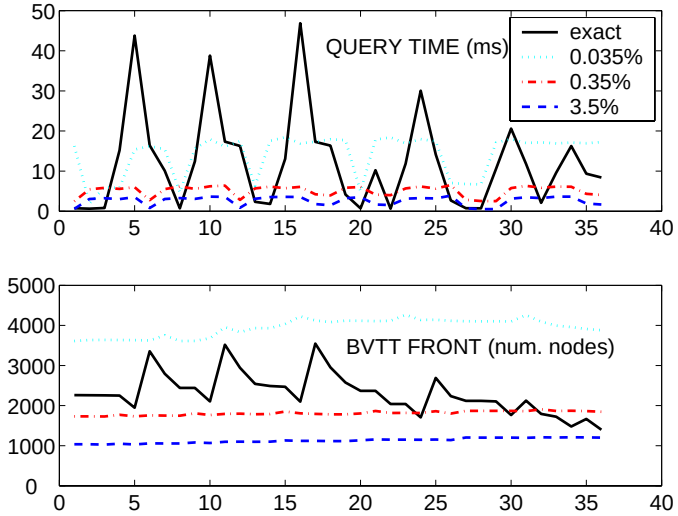

Figure 7: Comparison of query time and size of the BVTT front for exact full res. collision detection and CLODs

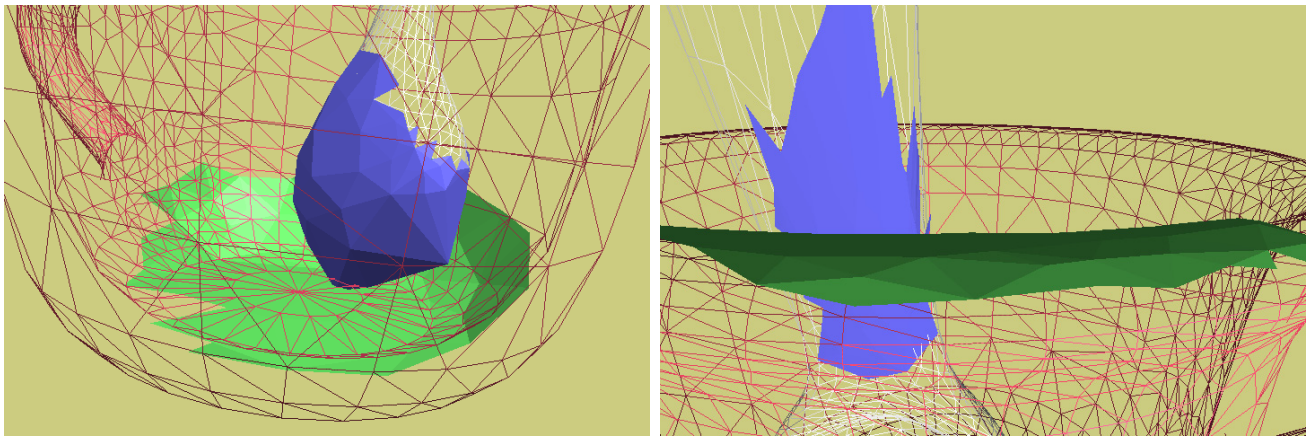

Figure 8: Velocity dependant adaptive selection of CLODs in the contact between a cup and a spoon. In blue and green, the vicinity of the contact locations shown at the resolution of the adaptively selected CLODs.