Efficient and Reliable Self-Collision Culling Using Unprojected Normal Cones

Tongtong Wang[1], Zhihua Liu[1], Min Tang[1], Ruofeng Tong[1] and Dinesh Manocha[2,1]

[1]Zhejiang University
{11421047, liuzhihua, tang_m, trf}@zju.edu.cn
[2]University of North Carolina at Chapel Hill
dm@cs.unc.edu

Abstract

We present an efficient and accurate algorithm for self-collision detection in deformable models. Our approach can perform discrete and continuous collision queries on triangulated meshes. We present a simple and linear time algorithm to perform the normal cone test using the unprojected 3D vertices, which reduces to a sequence point-plane classification tests. Moreover, we present a hierarchical traversal scheme that can significantly reduce the number of normal cone tests and the memory overhead using front-based normal cone culling. The overall algorithm can reliably detect all (self) collisions in models composed of hundreds of thousands of triangles. We observe considerable performance improvement over prior continuous collision detection algorithms.

Keywords: collision detection, animation

1. Introduction

Fast and accurate collision detection is important for generating realistic deformations. At a broad level, prior work can be classified into discrete collision detection (DCD) methods that check for collisions at an instant of time, and continuous collision detection (CCD) techniques that check for colliding regions within a time interval. The latter are used to avoid missing any collisions and to perform reliable simulations by maintaining intersection-free meshes.

Complex mesh models composed of hundreds of thousands of triangle primitives are frequently used in cloth or finite element method (FEM) simulation. Collision detection is regarded as one of the major bottlenecks in these applications and it is important to accurately check for collisions between all the primitives [BFA02, BEB12, Wan14, TTWM14].

Some of the commonly used algorithms use bounding volume hierarchies (BVHs) to accelerate collision detection. These techniques work well in terms of inter-object collision detection. However, self-collision checking in deformable models can be challenging

as many adjacent or nearby primitives of a deforming mesh are in close proximity and not culled by bounding volume tests (BVTs). Even if a mesh is intersection-free, checking all the primitives for self-collisions can be expensive.

The most commonly used methods for self-collisions are based on normal cones (NCs) for DCD [VT94, Pro97, SPO10]. However, the computational overhead of NC tests during BVH traversal can be high. This approach has been extended to CCD [TCYM09], but the additional cost of performing continuous NC culling slows down the overall CCD algorithm due to the quadratic complexity of the continuous contour test. Other techniques include energy-based methods [BJ10, ZJ12] and radial-based CCD culling for skeletal models [WLH*13]. However, self-collision culling continues to remain a bottleneck, especially for CCD between arbitrary deformable models.

Main Results: We present a fast and accurate approach for self-collision detection in triangulated models. Our formulation is based on NCs and is applicable to DCD as well as CCD. In order to reduce the overhead of NC tests during BVH traversal, we present two novel algorithms:

Figure 1: *Improved collision culling. We demonstrate the benefits of our self-collision algorithm on the Funnel benchmark (64K triangles). The three curves highlight the culling efficiencies of various algorithms. These include prior schemes that do not perform self-collision culling (i.e. AABB-only) versus two variants of our algorithm. Our novel CCD algorithm can significantly reduce the number of false positives.*

1. **Linear time NC test using unprojected contours:** We present a novel contour test algorithm that can directly perform the queries on the 3D vertices without computing any projections on a plane. Our formulation reduces to performing point-plane orientation tests on the 3D vertices. The resulting computation is simple and accurate, and reduces to evaluating the sign of algebraic expressions for DCD as well as CCD (Section 3).

2. **Front-based NC culling:** We present an incremental BVH traversal algorithm that combines the NC test with the BVT tree (BVTT) front computation. This reduces the hierarchy traversal overhead based on spatial and temporal coherence, and also the memory overhead (Section 4).

The resulting algorithms are simple to implement and can accurately detect all the collisions. It has no pre-processing overhead and can perform fast collision queries on complex benchmarks on a single CPU core (Section 5). We observe an order of magnitude improvement in the performance of the NC tests over prior approaches. We also highlight the overall speedups in cloth simulation due to our novel CCD algorithm.

2. Related Work

In this section, we give a brief overview of prior work on collision detection between deformable models. The simplest culling algorithms use BVHs that are based on k-DOPs or AABBs. These can be combined with self-collision culling techniques [Pro97, MKE03, TCYM09, SPO10, ZJ12].

Low-level culling: Many low-level culling techniques have been proposed to reduce the number of elementary tests between the triangle pairs for CCD, such as removing redundant elementary tests [GKJ*05, CTM08, WB06, TMT10a, TMY*11] or using bounding volumes of the primitives [HF07]. Our approach can be combined with these low-level acceleration techniques.

Clustering: Many mesh decomposition and clustering strategies have been proposed to compute tighter fitting hierarchies to improve the culling efficiency. Most of these techniques are used as a

pre-process [EL01, WB14]. Schvartzman et al. [SPO10] presented a self-collision test tree (SCTT) that is pre-computed to accelerate the self-collision queries for general deformable models. By executing hierarchically, their cost can be reduced to $O(1)$. Wong et al. [WLH*13] presented a continuous self-collision detection algorithm for skeletal models and extended it to check for collisions between a deformable surface and a simple solid model [WC14]. He et al. [HOEM15] presented a dynamic clustering algorithm for topology changing models. Our approach can be easily combined with these methods to obtain the fastest CCD query performance.

Reliable collision queries: Most earlier methods for DCD and CCD are implemented using floating point computations and numerical tolerances. However, numerical errors in arithmetic operations, along with the tolerances, can impact their accuracy, especially for elementary tests. Recently, many reliable algorithms have been proposed for elementary tests based on exact arithmetic operations [BEB12, TTWM14] or conservative float-point computation with tight bounds [Wan14, WTTM15]. It is also important to ensure that the self-collision culling tests are reliable.

Front-based traversal: BVTT front tracking has been used to accelerate collision detection [LC98]. By performing tests using the BVTT front generated from the last time step, these methods reduce the runtime overhead and make it easier to parallelize on CPUs and GPUs [TMLT11, ZK14]. However, the memory overhead for storing the BVTT front can be high. We combine our NC test with BVTT-based front to significantly reduce the time and storage complexity.

3. Unprojected Normal Cone Test

In this section, we introduce the notation and present our unprojected NC test for self-collision culling.

3.1. Background and notation

In this paper, we consider DCD as well as CCD problems on models represented as 2-manifold triangle meshes, possibly with boundaries. Our approach uses the mesh connectivity information to perform collision culling. DCD deals with checking whether any distinct triangle primitives overlap at a given time instant. On the other hand, CCD algorithms model the motion of each object or triangle using a continuous trajectory between the successive instants of a simulator and check for collisions along the trajectory. Our CCD culling algorithm also uses linearly interpolating trajectories [Pro97, BFA02].

We use the following notation in the rest of the paper: Lower case letters in normal fonts (e.g. a, b and a_i) represent scalar variables and upper case letters (e.g. L and $J(t)$) represent scalar functions. Lower case letters in bold face fonts (e.g. $\mathbf{a}$ and $\mathbf{b_t}$) represent vector quantities and points, and upper case letters in bold face fonts (e.g. $\mathbf{L}$ and $\mathbf{J(t)}$) represent vector-valued functions. The operators '$*$', '$\cdot$' and '$\times$' denote the usual scalar multiplication, dot product and cross product, respectively. We use following acronyms in the rest of the paper: *PPC*: stands for point-plane classification; *BV, BVH, BVTT*: stand for bounding volume, BVH and BVTT, respectively.

Figure 2: *Normal cone. For an input triangle mesh (a), its contour edges are defined by an order list of vertices $\mathbf{v}'_1, \mathbf{v}'_2, \ldots, \mathbf{v}'_n$ (the red arrows). Its corresponding normal cone $C_N(\alpha, \mathbf{l})$ contains all the normal vectors of the triangles (b), where l is the axis and α is the apex angle.*

3.2. Normal cone test

Some of the widely used algorithms for self-collision detection are based on the NC test [VT94]. Given a continuous surface, S, bounded by a contour, C, a sufficient criterion for no self-collision is based on the following two conditions:

1. **Bounds on the normals:** There is a vector, $\mathbf{v}$, such that $(\mathbf{N} \cdot \mathbf{v}) > 0$ for every point of the surface, S, where $\mathbf{N}(S)$ is the normal vector at the point S on the surface.
2. **No boundary self-intersections:** The projection of the contour C along the vector $\mathbf{v}$ does not have any self-intersections on the projected plane.

The first condition is called the *surface normal test* and the second condition is also known as the *contour test*. Provot [Pro97] presented an efficient method to evaluate the first condition based on NCs, which corresponds to a bound on the Gauss map of S. The NC for a mesh of triangles can be computed by combining the normal vectors of individual triangles. Given an input NC $C_N(\alpha, \mathbf{l})$ bounds, all the normal vectors of the triangles of a given triangle patch, where α is the apex angle and $\mathbf{l}$ is the middle axis of the cone (see Figure 2). The NC test has been used for DCD [Pro97, SPO10], and has been extended to CCD by Tang et al. [TCYM09].

Complexity of contour tests: In practice, the contour test tends to be more expensive as compared to surface normal test and can take up to 60% of the running time [SPO10] for DCD computations. This test typically involves computing a projection of the contour of S and checking for self-intersections on the resulting plane. Recently, Schvartzman *et al.* [SPO10] used a line-search star-shaped test to detect star-shaped projection, which guarantees that there will be no self-intersections. Furthermore, they used a pre-computed SCTT to accelerate the computation. As a result, the resulting algorithm has linear time complexity in the number of contour edges. However, the projection computations can be expensive and take up to 37% of total running time [SPO10]. Furthermore, they are prone to floating-point errors and can result in accuracy
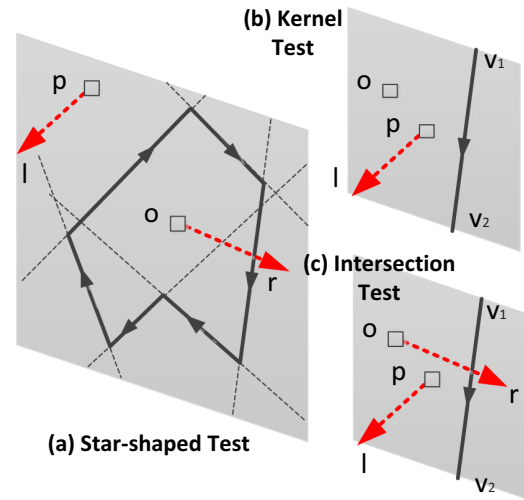
Figure 3: *Star-shaped test. The projection plane is defined by a point $\mathbf{p}$ and a normal vector $\mathbf{l}$. A star-shaped contour polygon in the plane $\{\mathbf{p}, \mathbf{l}\}$ has: (1) at least one kernel point $\mathbf{o}$ located inside the half-spaces of all the contour edges and (2) any ray starting from $\mathbf{o}$ along the direction $\mathbf{r}$ can intersect the boundary only once. Both these conditions can be tested by computing the orientation of a point with respect to different lines.*

issues during the contour tests. Some prior self-collision algorithms either omit the contour test computation for certain cases [VT94] or use some approximations [Pro97]. Heo et al. [HSK*10] described an approximate scheme to avoid the projection computation.

Continuous queries: It is not clear whether the linear time algorithm for contour test in [SPO10] extends to CCD. The resulting formulation based on kernel tests reduces to checking for overlaps between intervals, as opposed to line intersections for DCD. Instead, the only known algorithm for contour tests for CCD reduces to performing $O(n^2)$ EE elementary tests, where n is the number of contour edges [TCYM09]. This is based on projecting the continuously deforming edges to a plane and checking each pair for overlap by solving a cubic equation. In such cases, the contour test becomes a major bottleneck in the overall algorithm.

3.3. Contour test using point-plane classification

We present a novel algorithm for the contour test that does not involve any projection computations. The resulting contour test can be performed in $O(n)$ time for DCD as well as CCD, where n is the number of contour edges in C. Furthermore, our algorithm only needs to perform sign evaluation in terms of 3D point-plane side tests, i.e. determining the orientation of a point with respect to a plane in 3D.

We initially present our contour test for DCD, and later extend it to CCD computations. We first describe the notion of a PPC-based contour test on a 2D projection plane, and then extend it to the unprojected 3D contour points.

Figure 4: *Intersection test. The two different cases (a) and (b), in which there is an intersection between the ray $\{\mathbf{o}, \mathbf{r}\}$ and a contour edge $\{\mathbf{v}_1, \mathbf{v}_2\}$, can be tested by computing the orientations of points with respect to different lines.*

Given a contour, we first project the points onto a projection plane, which is defined by a point $\mathbf{p}$ and a normal vector $\mathbf{l}$ and represented as $\{\mathbf{p}, \mathbf{l}\}$ (see Figure 3a). If the projected contour is a star-shaped polygon, it guarantees no self-intersections [SPO10]. The star-shaped test is performed in two steps:

1. **Kernel test:** There is at least one kernel point (e.g. point $\mathbf{o}$) located at the inside of the half-spaces of all the contour edges (Figure 3b).
2. **Intersection test:** A ray starting from the kernel point in the direction $\mathbf{r}$ can intersect the boundary only once (Figure 3c).

We assume that all the contour edges are oriented in a clockwise or counter-clockwise direction with respect to $\mathbf{l}$.

We define a scalar function $\textbf{Side}(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{l})$ and a sign evaluation function $\textbf{SideSign}(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{l})$, where $\mathbf{a}, \mathbf{b}, \mathbf{c}$ are three points that lie on the projection plane $\{\mathbf{p}, \mathbf{l}\}$:

$$\textbf{Side}(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{l}) = ((\mathbf{a} - \mathbf{c}) \times (\mathbf{b} - \mathbf{c})) \cdot \mathbf{l},$$

$$\textbf{SideSign}(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{l}) = \begin{cases} 1, & \textbf{Side}(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{l}) > 0, \\ -1, & \textbf{Side}(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{l}) < 0, \\ 0, & \textbf{Side}(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{l}) = 0. \end{cases}$$

Note that the choice of $\mathbf{p}$ does not affect the values of $\textbf{Side}()$ and $\textbf{SideSign}()$.

The kernel test on the projected plane is performed using PPC between $\mathbf{o}$ and all the projected contour edges. For an edge defined by the two vertices $\mathbf{v}_1$ and $\mathbf{v}_2$, the orientation between them is computed using $\textbf{SideSign}(\mathbf{o}, \mathbf{v}_1, \mathbf{v}_2, \mathbf{l})$. Therefore, the kernel test can be expressed as checking the signs of n algebraic expressions[1]:

$$\textbf{SideSign}(\mathbf{o}, \mathbf{v}_1, \mathbf{v}_2, \mathbf{l}) == \textbf{SideSign}(\mathbf{o}, \mathbf{v}_2, \mathbf{v}_3, \mathbf{l})$$

$$== \ldots == \textbf{SideSign}(\mathbf{o}, \mathbf{v}_n, \mathbf{v}_1, \mathbf{l}),$$

where $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_n$ are the projected vertices of the contour.

---

[1]We ignore the case that $\textbf{SideSign}(\ldots) = 0$, as the test returns an answer of false.

The intersection test is also performed using sign evaluations. Let us consider the cases (a) and (b) (Figure 4), in which there is an intersection. For both cases, $\mathbf{v}_1$ and $\mathbf{v}_2$ lie on the different sides of the ray $\{\mathbf{o}, \mathbf{r}\}$, and side of $\mathbf{o}$ with respect to $\overrightarrow{\mathbf{v}_1 \mathbf{v}_2}$ is the same side of $\mathbf{v}_2$ with respect to the ray. We can check for these conditions based on evaluating the following expressions:

$$\textbf{SideSign}(\mathbf{o}, \mathbf{v}_1, \mathbf{v}_2, \mathbf{l}) = \textbf{SideSign}(\mathbf{v}_2, \mathbf{o}, \mathbf{o} + \mathbf{r}, \mathbf{l}) \quad \text{and}$$

$$\textbf{SideSign}(\mathbf{v}_1, \mathbf{o}, \mathbf{o} + \mathbf{r}, \mathbf{l}) \neq \textbf{SideSign}(\mathbf{v}_2, \mathbf{o}, \mathbf{o} + \mathbf{r}, \mathbf{l}).$$

Only when both conditions are fulfilled, we count an intersection between the ray and the contour edge. The idea behind the two equations is that the first equation is to ensure that there is an intersection between a line (defined by $\mathbf{o}$ and $\mathbf{r}$) and a line segment (defined by $\mathbf{v}_1$ and $\mathbf{v}_2$). The second equation is to ensure that the intersection is on the right side of the line (i.e. on the ray).

We compute the intersection number between the ray and each contour edge. The intersection test holds only when the intersection number equals to 1. For the case that the intersection number is greater than 1, the test can directly return a false answer.

Overall, the kernel test and the intersection test are reduced to computing dot and cross products, and these tests perform PPC using SideSign() functions.

3.4. Unprojected contour test

The contour test for the 2D projected vertices can be extended to directly operate on the unprojected 3D contour vertices, i.e. the original vertices of the input mesh contours. Let $\mathbf{a}', \mathbf{b}', \mathbf{c}'$ be the original (unprojected) 3D contour vertices, and $\mathbf{a}, \mathbf{b}, \mathbf{c}$ be the corresponding projections on the projection plane $\{\mathbf{p}, \mathbf{l}\}$. We use the following theorem to perform conservative contour tests on the original vertices:

Unprojected orientation test theorem: *Given three 3D points $\mathbf{a}', \mathbf{b}', \mathbf{c}'$, and a projection direction $\mathbf{l}$. Let their projections on a plane perpendicular to $\mathbf{l}$ be $\mathbf{a}, \mathbf{b}, \mathbf{c}$, respectively. In this case, $\textbf{SideSign}(\mathbf{a}', \mathbf{b}', \mathbf{c}', \mathbf{l}) = \textbf{SideSign}(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{l})$.*

Proof. By the definition of $\mathbf{a}, \mathbf{b}, \mathbf{c}$, we have:

$$\mathbf{a} = \mathbf{a}' + u * \mathbf{l}; \quad \mathbf{b} = \mathbf{b}' + v * \mathbf{l}; \quad \mathbf{c} = \mathbf{c}' + w * \mathbf{l},$$

where u, v and w are three scalars. By substituting them into the definition of $\textbf{Side}(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{l})$, we obtain:

$$\textbf{Side}(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{l})$$
$$= ((\mathbf{a} - \mathbf{c}) \times (\mathbf{b} - \mathbf{c})) \cdot \mathbf{l}$$
$$= (((\mathbf{a}' - \mathbf{c}') + (u - w) * \mathbf{l})$$
$$\times ((\mathbf{b}' - \mathbf{c}') + (v - w) * \mathbf{l})) \cdot \mathbf{l}$$
$$= ((\mathbf{a}' - \mathbf{c}') \times (\mathbf{b}' - \mathbf{c}')) \cdot \mathbf{l}$$
$$= \textbf{Side}(\mathbf{a}', \mathbf{b}', \mathbf{c}', \mathbf{l}).$$

$\square$

Algorithm 1 UnprojectedContourTest(C_N): Perform unprojected contour test on the boundary contour of a given normal cone C_N.
Input: A normal cone $C_N(\alpha, \mathbf{l})$, where α is the apex angle, and $\mathbf{l}$ is the axis of the cone, and its boundary contour is defined by an ordered list of vertices, i.e., $\mathbf{v'_1}, \mathbf{v'_2}, ..., \mathbf{v'_n}$.
Output: true for no self-intersection on the projected contour, **false** otherwise.

```
 1: // Prepare parameters for kernel test and intersection test.
 2: o = Σⁿ₁v'ᵢ/n;  l is middle axis of input normal cone C_N
 3: if l parallel to vector {0,1,0} then
 4:     r = {1,0,0}
 5: else
 6:     r = l × {0,1,0}
 7: end if
 8: // Initialize intersection number to zero.
 9: intsNum = 0
10: // Get side sign at the first contour segment.
11: s_0 = SideSign(o, v'_1, v'_2, l)
12: if s_0 == 0 then
13:     return false;
14: end if
15: // Perform kernel test and intersection test on each contour segment.
16: for each edge v'_i, v'_{i+1} do
17:     // Perform side test.
18:     if s_0 ≠ SideSign(o, v'_i, v'_{i+1}, l) then
19:         return false;
20:     end if
21:     // Perform intersection test.
22:     s_1 = SideSign(v'_i, o, o + r, l)
23:     s_2 = SideSign(v'_{i+1}, o, o + r, l)
24:     if s_1 == 0 s_2 == 0 then
25:         return false; // Can't determine the intersection.
26:     end if
27:     if s_2 = s_0 s_1 ≠ s_2  then
28:         intsNum++
29:         if intsNum > 1 then
30:             return false; // More than one intersection.
31:         end if
32:     end if
33: end for
34: return true;
```

Based on this theorem, the kernel test and intersection test, which are based on **Side**(), can be directly evaluated using the original 3D coordinates of the contour vertices. As a result, there is no need to perform projection computations for the conservative star-shape-based contour test.

The pseudo-code for the unprojected contour test is given in Algorithm 1. Given an input NC $C_N(\alpha, \mathbf{l})$ bounds, all the normal vectors of the triangles of a given triangle patch, where α is the apex angle and $\mathbf{l}$ is the middle axis of the cone. The boundary contour of $C_N(\alpha, \mathbf{l})$ is defined by a ordered list of vertices, i.e. $\mathbf{v'_1}, \mathbf{v'_2}, \ldots, \mathbf{v'_n}$. Our algorithm for the contour test will return **true** if there are no self-intersections on the projected contour, but will return **false** otherwise. We first define the parameters to perform the kernel test and intersection test (Lines 1–9, Algorithm 1).

- **Projection direction l**: We use the axis $\mathbf{l}$ of the input NC as the projection direction.

- **Kernel point o**: We use the average of the input contour vertices as the kernel point, i.e. $\mathbf{o} = \frac{\sum_1^n \mathbf{v'_i}}{n}$.

- **Ray direction r**: In theory, it can be any vector perpendicular to $\mathbf{l}$. In practice, we choose it to be $\mathbf{l} \times \{0, 1, 0\}$ if $\mathbf{l}$ is not parallel to vector $\{0, 1, 0\}$; otherwise, we set $\mathbf{r}$ to be $\{1, 0, 0\}$ (Lines 3–7, Algorithm 1).

After specifying all the parameters, we perform the kernel tests and intersection tests over all the contour edges (Lines 9–34). The intersection test is performed at Lines 21–32. If the intersection number exceeds 1, the algorithm returns **false**. The contour test returns true only if the kernel test and the intersection test return true.

Query performance: The worst case for the unprojected contour tests occurs when it returns a true answer and there is no early exit. In that case, it perform $3n$ **SideSign**() evaluations for DCD, where n is the number of contour edges. Each **SideSign**() evaluation requires one cross product and one dot product computation between two 3D vectors. So, the total operation count is bounded by $3n$ cross products and $3n$ dot products for each contour test.

Algorithm 2 SelfCollide(N): Normal Cone Culling using Unprojected Contour Test (NC). By traversing the BVH recursively, meshes that satisfy normal cone test are culled.
Input: A node N on the BVH. C_N is the normal cone associated with N.
Output: No return value.

```
 1: if IsLeaf(N) then
 2:     return; // Traversal terminated.
 3: end if
 4:
 5: if ApexAngle(C_N) < π then
 6:     if UnprojectedContourTest(C_N) = true then
 7:         return; // The corresponding mesh has no self-collisions.
 8:     end if
 9: end if
10: // Check the descendants.
11: SelfCollide(N → LeftChild)
12: SelfCollide(N → RightChild)
13: Collide(N → LeftChild, N → RightChild)
```

3.5. Normal cone culling

Based on Algorithm 1, the overall algorithm for NC culling is shown in Algorithm 2. Our algorithm involves no pre-computation and updates a BVH to perform hierarchical computations. The self-collision checking starts at the root node of the BVH, and traverses in a top-down manner. For a node N of the BVH, and its associated NC C_N, we check whether the apex angle of C_N is less than π and also perform the unprojected contour test. If these two tests are satisfied, then we do not need to traverse to the children of N to check for self-collisions.

3.6. Normal cone culling for CCD

The contour test and NC culling algorithms described above can be extended to CCD. One of the main issues is to compute a bound on the surface normals that varies during the time interval based on linearly interpolating triangle vertices. Moreover, we need to

perform the unprojected contour test on the continuously varying vertices of the contour. We compute a conservative normal bound for each deforming triangle, and merge these bounds in a bottom-up manner on the BVH to update the apex angles of all NCs. In order to perform a continuous unprojected contour test, we assume that a vertex $\mathbf{p_i}$ is moving with the constant velocity $\mathbf{v_i}$ during a given time interval $[0, 1]$.

We first define the function **CSide1**() as

$$\mathbf{CSide1}(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{l}) = ((\mathbf{a} - \mathbf{c}) \times (\mathbf{b} - \mathbf{c})) \cdot \mathbf{l},$$

where $\mathbf{b} = \mathbf{b_0} + \mathbf{v_b} * t$, $\mathbf{c} = \mathbf{c_0} + \mathbf{v_c} * t$. However, $\mathbf{a}$ and $\mathbf{l}$ are unchanged during the time interval $[0, 1]$. $\mathbf{b_0}$ and $\mathbf{c_0}$ are the position of two vertices at $t = 0$. $\mathbf{v_b}$ and $\mathbf{v_c}$ are their moving velocities during the time interval ($t \in [0, 1]$). Therefore:

$$\mathbf{CSide1}(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{l})$$
$$= ((\mathbf{b_0} + \mathbf{v_b} * t - \mathbf{a}) \times (\mathbf{c_0} + \mathbf{v_c} * t - \mathbf{a})) \cdot \mathbf{l}$$
$$= k_0 * B_0^2(t) + k_1 * B_1^2(t) + k_2 * B_2^2(t),$$

where $B_i^2(t)'s$ are second-order polynomials in Bernstein basis.

The detailed derivation of k_i is as follows:

$$k_0 = ((\mathbf{b_0} - \mathbf{a}) \times (\mathbf{c_0} - \mathbf{a})) \cdot \mathbf{l},$$
$$k_1 = \frac{((\mathbf{v_b} \times (\mathbf{c_0} - \mathbf{a}) + (\mathbf{b_0} - \mathbf{a}) \times \mathbf{v_c}) \cdot \mathbf{l}}{2}$$
$$+ ((\mathbf{b_0} - \mathbf{a}) \times (\mathbf{c_0} - \mathbf{a})) \cdot \mathbf{l},$$
$$k_2 = ((\mathbf{b_0} - \mathbf{a}) \times (\mathbf{c_0} - \mathbf{a})) \cdot \mathbf{l}$$
$$+ (\mathbf{v_b} \times (\mathbf{c_0} - \mathbf{a}) + \mathbf{v_c} \times (\mathbf{b_0} - \mathbf{a})) \cdot \mathbf{l}$$
$$+ (\mathbf{v_b} \times \mathbf{v_c}) \cdot \mathbf{l}.$$

With this representation, the value of **CSide** is bounded based on the coefficients k_0, k_1 and k_2 for $0 \le t \le 1$. Based on the formulation of **CSide1**(), we define:

$$\mathbf{CSideSign1}(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{l}) = \begin{cases} 1, & k_0 > 0, k_1 > 0, k_2 > 0, \\ -1, & k_0 < 0, k_1 < 0, k_2 < 0, \\ 0, & \text{otherwise.} \end{cases}$$

This formulation is conservative, but provides a sufficient condition to perform the kernel and intersection tests for CCD.

For the case $\mathbf{a}$ that is moving under constant velocity, i.e. $\mathbf{a} = \mathbf{a_0} + \mathbf{v_a} * t$, and $\mathbf{b}, \mathbf{c}$ are unchanged during the time interval, we define the function:

$$\mathbf{CSide2}(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{l})$$
$$= ((\mathbf{a_0} + \mathbf{v_a} * t - \mathbf{c}) \times (\mathbf{b} - \mathbf{c})) \cdot \mathbf{l}.$$

Similarly, the value of $\mathbf{CSide2}(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{l})$ is bounded by two constants for all $t \in [0, 1]$:

$$k_0' = ((\mathbf{a_0} - \mathbf{c}) \times (\mathbf{b} - \mathbf{c})) \cdot \mathbf{l}$$
$$k_1' = ((\mathbf{a_0} + \mathbf{v_a} - \mathbf{c}) \times (\mathbf{b} - \mathbf{c})) \cdot \mathbf{l}.$$

So, we define:

$$\mathbf{CSideSign2}(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{l}) = \begin{cases} 1, & k_0' > 0, k_1' > 0, \\ -1, & k_0' < 0, k_1' < 0, \\ 0, & \text{otherwise.} \end{cases}$$

By replacing the functions **SideSign**() in Algorithm 1 with new sign functions **CSideSign1**() (Lines 12 and 19) and **CSideSign2**() (Lines 23 and 24), respectively, the contour test algorithm for DCD can be easily extended to perform conservative contour tests during the time interval $[0, 1]$ for CCD (as shown in Algorithm 3).

Algorithm 3 UnprojectedContourTestForCCD(C_N): Perform unprojected contour test on the boundary contour of a given normal cone C_N.

Input: A normal cone $C_N(\alpha, \mathbf{l})$ where α is the apex angle, and $\mathbf{l}$ is the axis of the cone, and its boundary contour which is defined by a ordered list of vertices, i.e., $\mathbf{v_1'}, \mathbf{v_2'}, ..., \mathbf{v_n'}$.

Output: true for no self-intersection on the projected contour, **false** for undeterminable.

1: // Prepare parameters for kernel test and intersection test.
2: $\mathbf{o} = \frac{\sum_1^n \mathbf{v_{i_{t=0}}'}}{n}$
3: $\mathbf{l}$ is middle axis of input normal cone C_N
4: **if** $\mathbf{l}$ parallel to vector $\{\mathbf{0, 1, 0}\}$ **then**
5:    $\mathbf{r} = \{\mathbf{1, 0, 0}\}$
6: **else**
7:    $\mathbf{r} = \mathbf{l} \times \{\mathbf{0, 1, 0}\}$
8: **end if**
9: // Initialize the intersection number to zero.
10: $intsNum = 0$
11: // Get side sign at the first contour segment.
12: $s_0 = \mathbf{CSideSign1}(\mathbf{o}, \mathbf{v_1'}, \mathbf{v_2'}, \mathbf{l})$
13: **if** $s_0 == 0$ **then**
14:    return **false**;
15: **end if**
16: // Perform kernel test and intersection test on each contour segment.
17: **for** each edge $\mathbf{v_i'}, \mathbf{v_{i+1}'}$ **do**
18:    // Perform side test.
19:    **if** $s_0 \ne \mathbf{CSideSign1}(\mathbf{o}, \mathbf{v_i'}, \mathbf{v_{i+1}'}, \mathbf{l})$ **then**
20:       return **false**;
21:    **end if**
22:    // Perform intersection test.
23:    $s_1 = \mathbf{CSideSign2}(\mathbf{v_i'}, \mathbf{o}, \mathbf{o} + \mathbf{r}, \mathbf{l})$
24:    $s_2 = \mathbf{CSideSign2}(\mathbf{v_{i+1}'}, \mathbf{o}, \mathbf{o} + \mathbf{r}, \mathbf{l})$
25:    **if** $s_1 == 0$ $s_2 == 0$ **then**
26:       return **false**; // Can't determine the intersection.
27:    **end if**
28:    **if** $s_2 = s_0$ $s_1 \ne s_2$ **then**
29:       $intsNum + +$
30:       **if** $intsNum > 1$ **then**
31:          return **false**; // More than one intersection.
32:       **end if**
33:    **end if**
34: **end for**
35: return **true**;

Query performance: In the worst case, the unprojected contour test for CCD performs n **CSideSign1**() and $2n$ **CSideSign2**()

evaluations, where *n* is the number of contour edges. Each **CSideSign1**() evaluation requires eight cross products and six dot products computation between two 3D vectors. Each **CSideSign2**() evaluation requires two cross products and two dot products. So, the totally operation count is bounded by 12*n* cross products and 10*n* dot products for each contour test.

4. Front-Based Normal Cone Culling

The hierarchical traversal algorithm based on NC tests presented above can accelerate the performance of each NC test during the hierarchical traversal. However, the resulting algorithm may perform a large number of NC tests while handling complex benchmarks represented using hundreds of thousands of triangles. In this section, we present a novel front-based NC culling algorithm that can significantly reduce the number of NC tests and traversal overhead. Furthermore, it reduces the memory requirements of storing the front.

The basic idea behind our algorithm is shown in Figure 5. For a scene with deformable objects, we compute its BVH as shown in the upper right corner of Figure 5. The self-collision detection algorithm corresponds to the traversal of its BVTT, as shown in Figure 5. The main goal of using BVTT is to compute the front during the traversal from the previous frame [TMT10b, TMLT11]. By starting the traversal of the current frame directly from the BVTT front, we exploit the spatial and temporal coherence to reduce the traversal cost. However, the memory overhead of this front can be high. For example, for a mesh with 4 K triangles, it takes about 17 MB memory to store the BVTT front between successive frames [TMLT11].

A node *A* ∗ *B* in the BVTT represents collision checking between the nodes *A* and *B* of the given BVH. In the worst case, the BVTT can have $O(m^2)$ nodes, where *m* is the number of nodes in the BVH. The front-based algorithms keep track of a subset of BVTT that corresponds to overlapping nodes during the current frame. Instead of storing a global BVTT front (i.e. the thick line in Figure 5), we decompose it into many sub-segments, as shown by the thin line in Figure 5. Each sub-segment of the BVTT is associated with an *X* ∗ *X* node, which corresponds to checking for self-collisions among all the nodes beneath the internal node *X* of BVH. Instead of traversing the children of this node hierarchically, we use our NC culling algorithm to check for self-collisions. This is shown as the segmented areas in Figure 5. In particular, we associate each BVTT front sub-segment for *X* ∗ *X* with the internal node *X* on the BVH.

We use the NC guided BVTT front tracking algorithm (Algorithm 4), to significantly reduce the traversal overhead as well as size of the front. By traversing the BVH recursively, if a node *X* satisfies the NC test, its sub-segment can be directly skipped for testing. Without the NC tests, we will need to track the nodes in this sub-segment to perform self-collision detection. During the BVH tracking, all the visited BVTT nodes are stored into the sub-segment and reused for subsequent frames [TMLT11].

Based on the front-based NC algorithm, we obtain the benefits of BVTT front-based culling as well as NC -based culling. With BVTT front tracking, we can perform collision detection in an efficient manner by utilizing temporal and spatial coherence. Furthermore,

Figure 5: *Normal cone-guided BVTT front tracking. We associate each BVTT front sub-segment for X ∗ X with the internal node X on the BVH. In order to perform self-collision culling, we use a front-based normal cone traversal approach that uses the BVTT front (Algorithm 4). By traversing the BVH recursively, if a node X passes the normal cone test, its associated BVTT front sub-segment is culled. Our approach can reduce the traversal overhead and runtime memory footprint.*

we combine them with the NC test, and this can significantly reduce the size of the BVTT front, as we do not need to store the BVTT front segments for the areas that are culled by our NC tests. Figure 6 highlights the reduced size of the NC-guided BVTT front for the benchmark FlowingCloth. In this case, the size of our NC guided BVTT front (the upper curve) is only about 8% of the conventional BVTT front in [TMLT11] (the lower curve).

Algorithm 4 SelfCollideWithGuidedFrontTracking(*N*): a normal cone guided BVTT front tracking algorithm for self-collision detection (NC + Front).

Input: A node *N* on the BVH. C_N is the normal cone associated with *N*. $Front_N$ is the BVTT front sub-segment associated with node *N*.

Output: No return value.

1: **if** IsLeaf(N) **then**
2:     return; // Traversal terminated.
3: **end if**
4:
5: **if** ApexAngle(C_N) < π **then**
6:     **if** UnprojectedContourTest(C_N) = true **then**
7:         return; // The region is self-collision free.
8:     **end if**
9: **end if**
10:   // Check the descendants.
11: SelfCollide(N → LeftChild)
12: SelfCollide(N → RightChild)
13: FrontTracking($Front_N$)

4.1. Reliable computation

Our overall collision detection algorithm has three main steps: (1) update BVs and NCs; (2) perform front-based BVH traversal and NC

Figure 6: *Front-based normal cone culling. We highlight the size of our NCT-guided BVTT front (the lower curve), which is about 8% of the conventional BVTT front in [TMLT11] (the upper curve) for the benchmark FlowingCloth. Our algorithm can save considerable memory and runtime overhead.*

culling (Algorithm 4) and (3) perform elementary tests between triangle pairs and may use lower level culling methods for acceleration. For stage (1), we compute a conservative bound on the NCs during the merging step. Stage (2) of our algorithm only needs to evaluate the sign of algebraic expressions corresponding to dot products and cross products, as well as other expressions; Stage (3) can be performed with geometric exact elementary test algorithms [BEB12, TTWM14], or conservative CCD algorithms [Wan14, WTTM15], and these computations can be accelerated using floating-point filters. As a result, the overall collision detection algorithm is accurate and not susceptible to errors.

5. Implementation and Results

In this section, we describe our implementation and highlight the performance of our algorithm on several benchmarks.

5.1. Implementation

We have implemented our algorithms on a standard PC (Intel i7-3770K CPU @3.5GHz, 4GB RAM, 64-bits Window 7 OS) with C++ and all the results are generated using a single CPU core. The various components of our overall collision detection system are shown in Figure 7.

Our current implementation is limited to deformable models without topology changes. We first construct a BVH for the entire scene in a top-down manner [TCYM09]. During each time step, we update the bounding volumes and NCs, and perform high-level culling based on these data structures. After these high-level culling operations, we perform low-level culling operations that reduce the number of primitive tests using non-penetration filters (NPFs). Finally, we perform triangle–triangle intersection tests for DCD and exact elementary tests for CCD [TTWM14].

We use AABBs as the underlying bounding volume in the hierarchy. It is possible to use tighter bounding volumes such as k-DOPs, but that increases the overhead of hierarchy update. Instead, we use NPFs [TMT10a] along with AABBs and they provide similar

Figure 7: *Collision detection system. For every time step, we update the bounding volumes and normal cones and perform high-level culling that includes bounding volume culling and front-based normal cone culling. We perform low-level culling to eliminate duplicate elementary tests (for CCD) and reliable primitive tests for each step.*

Figure 8: *Benchmarks. We use seven challenging benchmarks arising from deformable and cloth simulations. We compare the performance of our DCD and CCD algorithms with prior methods.*

culling efficiency as compared to k-DOPs or tight fitting bounding volumes.

5.2. Benchmarks

In order to evaluate the performance of our DCD and CCD algorithms, we used seven different benchmarks that came from different deformable simulation scenarios and have been used by other researchers.

- **Twisting:** It corresponds to complex cloth simulation with twists, as the ball rotates. It has $64K$ triangles. This benchmark has a high number of self-collisions (Figure 8a).
- **Funnel:** A cloth with $64K$ triangles falls into a funnel and folds to fit into the funnel with many self-collisions (Figure 8b).
- **Bishop:** A swing dancer wearing three pieces of cloth (with $124K$ triangles) with self-collisions (Figure 8c).

Bench-marks	CT Time (DCD)		CT Time (CCD)	
	Our	SCT	Our	CBC
Twisting	0.41	1.58	1.09	5.39
Funnel	0.63	2.29	1.17	7.56
Falling	1.4	8.07	17.8	55.6
Bishop	1.12	6.23	5.85	21.53
Flamenco	1.87	7.53	4.22	20.49
SquishyBall	141.1	196.5	454.3	668.5
FlowingCloth	0.033	0.14	0.8	9.37

Figure 9: *Performance of contour tests. We highlight the average running time (in ms) of our unprojected contour test algorithm by comparing with SCT [SPO10] and CBC [TCYM09]. The contour test can take a significant fraction of the overall collision query time as shown in Figure 10.*

Bench-marks	Qry Time (DCD)			Qry Time (CCD)				
				Prior Methods			Our Method	
	Our	AABB Only	SCT	AABB Only	CBC	DC	NC+ Front	NC Only
Twisting	56.79	80.08	60.15	105.55	95.56	96.45	82.6	92.19
Funnel	53.37	78.52	58.76	82.27	79.46	75.94	61.5	64.96
Falling	177.3	278.3	206.2	783.1	643.1	702.4	590.5	639.3
Bishop	75.24	109.23	90.93	124.98	111.9	119.32	97.76	109.5
Flamenco	83.1	102.29	95.4	272.18	267.4	252.31	210.6	215.5
SquishyBall	1842	2484	1921	6553	5749	5697	5142	5609
FlowingCloth	21.27	46.64	23.43	55.22	40.1	34.3	25.03	27.01

Figure 10: *Performance and comparison. We compare the performance of our algorithm with prior techniques for DCD and CCD queries. We report the average time taken by these collision queries (in ms) for each benchmark. Each of these methods performs the same low-level culling operations and different high-level culling operations. For our method, we also highlight the relative benefit of combining BVTT front with normal cone culling (NC). We observe considerable speedups with (NC + Front) over prior CCD algorithms.*

- **Falling:** A man wearing a robe (with $172K$ triangles) falls down rapidly, which introduces wrinkles and self-collisions in the cloth (Figure 8d).
- **Flamenco:** A fiery Flamenco dancer wearing a colourful skirt with ruffles. This benchmark ($49K$ triangles) has a high number of self-collisions (Figure 8e).
- **SquishyBall:** A squishy ball with 820 tentacles and over 1M triangles squishes and bounces on the ground, inducing numerous small inter-penetrations [ZJ12]. (Figure 8f).
- **FlowingCloth:** A flowing cloth with $50K$ triangles [ZJ12] (Figure 8g) is hanging with two corners fixed.

Four of these benchmarks (Twisting, Bishop, Falling and Funnel) are generated using a cloth simulation system. The input for the Flamenco is given as discrete keyframes. We use the linearly interpolating motion of the vertices between successive keyframes and

Bench-marks	Twisting	Funnel	Falling	Bishop	Flamenco	Squishy-Ball	Flowing-Cloth
DCD	34.48%	28.71%	25.5%	37.39%	26.23%	12.21%	8.01%
CCD	37.10%	31.40%	27.4%	39.50%	29.30%	11.40%	9.20%

Figure 11: *Reduction in front size. We highlight the reduction in the memory overhead of the BVTT front due to our novel traversal algorithm based on normal cones (see Section 4). Due to normal cone culling, the size of the BVTT front reduces by 8–37% in our benchmarks.*

check for inter-object and self-collisions. Figure 10 highlights the performance of our algorithm for DCD and CCD queries on these benchmarks.

5.3. Cloth simulation

We integrated our collision detection algorithm into a cloth simulation system. The underlying simulator performs implicit integration and use repulsion forces along with CCD computations to avoid interpenetrations. We compared the overall performance of the cloth simulation by using our collision detection algorithm versus prior method that is based on AABB-only and performs no self-collision culling. We used the simulator to generate the entire simulations corresponding to these four benchmarks: Funnel, Twisting, Falling and Bishop. Our faster CCD algorithm results in $1.2\times$ speedup in the overall performance of the cloth simulator.

6. Comparison and Analysis

6.1. Comparison

We compare the performance of our algorithm with the prior techniques that use hierarchical methods or NCs for DCD and CCD.

1. **SCT:** This corresponds to the implementation of star-contours-based NC algorithm of [SPO10] for DCD. This is based on line-search star-shaped contour test and pre-computing the SCTT.
2. **AABB-only:** These algorithms use AABB as the underlying bounding volumes and perform no self-collision culling. They use low-level culling algorithms to eliminate duplicate elementary tests and perform reliable elementary tests using BSC [TTWM14].
3. **CBC:** This is the continuous NC algorithm for CCD [TCYM09] along with AABB culling and BSC elementary tests and performs the $O(n^2)$ exact contour test.
4. **DC:** This corresponds to the dynamic clustering algorithm [HOEM15] for CCD. It computes new clusters at each frame as the objects deform that is guaranteed to be collision free using observation point [WLH*13]. It is combined with the AABB hierarchy and BSC elementary tests.
5. **NPF:** The main contribution of the algorithm in [TMT10a] is the use of an NPF, which is a light-weight filter (or culling scheme) that can remove most of the false positives before performing exact elementary tests. The use of that filter resulted in $1.7\times$–$3.5\times$ speedups. In this paper, we also used the NPF to reduce the number of false positives (as mentioned at the end of
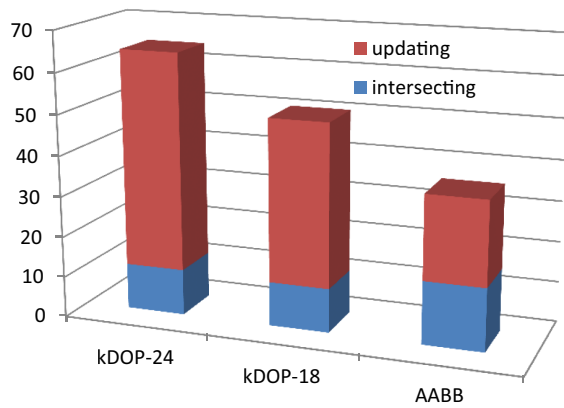
Figure 12: *AABB versus kDOP. We compare the running times (in seconds) of the Funnel benchmark (64K triangles) by using 24-DOPs, 18-DOPs and AABBs as bounding volumes, respectively. As shown by the figure, the time spent in the intersection computations is reduced with the use of tighter bounding volumes, but the updating time increases considerably.*

Section 5.1). In other words, the self-collision culling algorithm presented in this paper is complementary to the use of NPF.

6. **VolCCD and BSC:** The main contribution of Vol-CCD [TMY*11] is in terms of low-level collision culling between volumetric elements; it is orthogonal to the use of NCs. The main contribution of BSC [TTWM14] is related to performing reliable and exact elementary tests for CCD computations. The improved NC algorithm presented in this paper is orthogonal to these methods and can be combined with them. Some of the algebraic formulations used in our NC test is similar to that used in these prior papers, but the overall goal is different.

We also compare the running times of the Funnel benchmark (64K triangles) by using 24-DOPs, 18-DOPs and AABBs as bounding volumes, respectively. As shown in Figure 12, the time spent in the intersection tests is reduced with tighter bounding volumes, the time to update the hierarchy (updating time) increases considerably more. Overall, the algorithms based on AABBs provide the best overall performance on the collision queries.

6.2. Analysis

As compared to the contour test presented in [SPO10], our algorithm is simpler and has lower overhead because we do not perform any projection computations. We observe up to $5.8\times$ improvement in the performance of the contour test for DCD (Figure 9) and up to $1.2\times$ improvement (Figure 10) in the performance of overall DCD algorithm, as compared to [SPO10].

We observe considerable performance improvement for CCD computations, over techniques that either use no self-collisions (AABB-only) or exact continuous NC tests [TCYM09]. This is due to the fact that our NC test has linear complexity, whereas prior methods had quadratic complexity with a higher constant [TCYM09]. As a result, we observe up to $11.7\times$ improvement in the performance

Figure 13: *CCD accelerations with/without low-level culling. We compare the acceleration rates for CCD between our method versus AABB-only with/without performing low-level culling. We observe much higher speedups without low-level culling.*

of the contour test for CCD (Figure 9). In terms of the overall CCD algorithm, the improved NC test results in up to $2.2\times$ improvement over AABB-only. We obtain higher speedups for CCD as compared to DCD, because the NC tests take a larger fraction of the average frame time.

The combination of front-based culling with our improved NC test leads to considerable speedups. We observe additional performance improvement by combining front-based culling with NC (Front + NC) over using only NC in Figure 10. The front-based traversal reduces the cost of hierarchy traversal and the number of SCTs. It does not reduce the number of false positives in terms of elementary tests. As a result, the relative benefits of NC tests and front-based culling are somewhat complementary and we obtain high speedups by combining them. Figure 1 highlights the improved culling with unprojected NC tests. Furthermore, we observe considerable improvements in the size and memory overhead of the BVTT front, when it is combined with NC culling (see Figure 11).

In all the implementations (SCT, AABB-only, CBC and DC), we used low-level culling methods (orphan sets [TCYM09] and NPFs [TMT10a]) for CCD to remove redundant elementary tests. So, the overall performance improvement in CCD query is moderate. If we do not perform low-level culling and only compare the relative performance with BVH culling only, we observe significantly higher speedups, as shown in Figure 13. Many other approaches only seem to compare the relative speedups with only BVH-culling [SPO10, ZJ12]. Basically, these low-level culling methods can significantly reduce the number of exact elementary tests being performed between the primitives. The fastest CCD algorithms use a combination of low-level and high-level culling schemes, as shown in Figure 10 .

7. Limitations, Conclusions and Future Work

We present a fast and reliable algorithm to perform self-collision culling on complex deformable models. There is a general

perception that the overhead of NC tests is high and its applications has mostly been limited to DCD. We presented a novel unprojected contour test that provides $10 - 30\times$ improvement over prior continuous contour tests for CCD. Furthermore, we described a novel traversal scheme using front-based NC culling that reduces the time and space overhead. The combination of these two methods can accelerate the performance of CCD queries by an order of magnitude.

Our approach has some limitations. The unprojected contour test tends to be more conservative than prior methods. The current formulation is limited to linearly interpolating triangles for CCD. Self-collision culling based on NCs works well when the resulting meshes do not have high variations in curvature. As a result, their performance depends on how a large mesh is decomposed into sub-meshes at different nodes of the BVH or the underlying clustering criteria. Our current implementation is limited to deformable models that do not undergo topology changes.

There are many avenues for future work. It would be useful to combine our approach with fast, dynamic clustering schemes that can improve the culling efficiency of NCs and also used for adaptive meshes. We would like to parallelize the approach on multi-core CPUs and GPUs, similar to [GLM05, SGG*06, TMLT11, ZK14, TWT*16], and our reduced-size BVTT front should improve the parallel performance. Finally, we would like to integrate our algorithm with FEM and hair simulation systems.

Acknowledgements

References

[BEB12] BROCHU T., EDWARDS E., BRIDSON R.: Efficient geometrically exact continuous collision detection. *ACM Transactions on Graphics 31*, 4 (July 2012), 96:1–96:7.

[BFA02] BRIDSON R., FEDKIW R., ANDERSON J.: Robust treatment of collisions, contact and friction for cloth animation. *ACM Transactions on Graphics 21*, 3 (July 2002), 594–603.

[BJ10] BARBIČ J., JAMES D. L.: Subspace self-collision culling. *ACM Transactions on Graphics (SIGGRAPH 2010) 29*, 4 (2010), 81:1–81:9.

[CTM08] CURTIS S., TAMSTORF R., MANOCHA D.: Fast collision detection for deformable models using representative-triangles. In *SI3D '08: Proceedings of the 2008 Symposium on Interactive 3D Graphics and Games* (2008), pp. 61–69.

[EL01] EHMANN S. A., LIN M. C.: Accurate and fast proximity queries between polyhedra using convex surface decomposition. In *Computer Graphics Forum* (2001), pp. 500–510.

[GKJ*05] GOVINDARAJU N. K., KNOTT D., JAIN N., KABUL I., TAMSTORF R., GAYLE R., LIN M. C., MANOCHA D.: Interactive collision detection between deformable models using chromatic decomposition. *ACM Transactions on Graphics 24*, 3 (July 2005), 991–999.

[GLM05] GOVINDARAJU N. K., LIN M. C., MANOCHA D.: Quick-cullide: Fast inter-and intra-object collision culling using graphics hardware. In *Virtual Reality, Proceedings IEEE VR* (2005), pp. 59–66.

[HF07] HUTTER M., FUHRMANN A.: Optimized continuous collision detection for deformable triangle meshes. In *Proceedings of WSCG '07* (2007), pp. 25–32.

[HOEM15] HE L., ORTIZ R., ENQUOBAHRIE A., MANOCHA D.: Interactive continuous collision detection for topology changing models using dynamic clustering. In *Proceedings of the 19th Symposium on Interactive 3D Graphics and Games (i3D '15)* (2015), pp. 47–54.

[HSK*10] HEO J.-P., SEONG J.-K., KIM D., OTADUY M. A., HONG J.-M., TANG M., YOON S.-E.: FASTCD: Fracturing-aware stable collision detection. In *Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2010), pp. 149–158.

[LC98] LI T.-Y., CHEN J.-S.: Incremental 3D collision detection with hierarchical data structures. In *Proceedings of the ACM VRST (VRST '98)* (1998), pp. 139–144.

[MKE03] MEZGER J., KIMMERLE S., ETZMUβ O.: Hierarchical techniques in cloth detection for cloth animation. *Journal of WSCG 11*, 1 (2003), 322–329.

[Pro97] PROVOT X.: Collision and self-collision handling in cloth model dedicated to design garments. In *Graphics Interface* (1997), pp. 177–189.

[SGG*06] SUD A., GOVINDARAJU N., GAYLE R., KABUL I., MANOCHA D.: Fast proximity computation among deformable models using discrete voronoi diagrams. *ACM Transactions on Graphics 25*, 3 (July 2006), 1144–1153.

[SPO10] SCHVARTZMAN S. C., PÉREZ A. G., OTADUY M. A.: Star-contours for efficient hierarchical self-collision detection. *ACM Transactions on Graphics 29*, 4 (July 2010), 80:1–80:8.

[TCYM09] TANG M., CURTIS S., YOON S.-E., MANOCHA D.: ICCD: Interactive continuous collision detection between deformable models using connectivity-based culling. *IEEE Transactions on Visualization and Computer Graphics 15* (2009), 544–557.

[TMLT11] TANG M., MANOCHA D., LIN J., TONG R.: Collision-streams: Fast GPU-based collision detection for deformable models. In *Proceedings of the 2011 ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (2011), pp. 63–70.

[TMT10a] TANG M., MANOCHA D., TONG R.: Fast continuous collision detection using deforming non-penetration filters. In *Proceedings of ACM Symposium on Interactive 3D Graphics and Games* (New York, NY, USA, 2010), ACM, pp. 7–13.

[TMT10b] TANG M., MANOCHA D., TONG R.: MCCD: Multi-core collision detection between deformable models using front-based decomposition. *Graphical Models 72*, 2 (2010), 7–23.

[TMY*11] TANG M., MANOCHA D., YOON S.-E., DU P., HEO J.-P., TONG R.: VolCCD: Fast continuous collision culling between deforming volume meshes. *ACM Transactions on Graphics 30* (May 2011), 111:1–111:15.

[TTWM14] TANG M., TONG R., WANG Z., MANOCHA D.: Fast and exact continuous collision detection with Bernstein sign classification. *ACM Transactions on Graphics 33* (November 2014), 186:1–186:8.

[TWT*16] TANG M., WANG H., TANG L., TONG R., MANOCHA D.: CAMA: Contact-aware matrix assembly with unified collision handling for GPU-based cloth simulation. *Computer Graphics Forum (Proceedings of Eurographics 2016) 35*, 2 (2016), 511–521.

[VT94] VOLINO P., THALMANN N. M.: Efficient self-collision detection on smoothly discretized surface animations using geometrical shape regularity. *Computer Graphics Forum 13*, 3 (1994), 155–166.

[Wan14] WANG H.: Defending continuous collision detection against errors. *ACM Transactions on Graphics 33*, 4 (July 2014), 122:1–122:10.

[WB06] WONG W. S.-K., BACIU G.: A randomized marking scheme for continuous collision detection in simulation of deformable surfaces. In *Proceedings of ACM VRCIA* (2006), 181–188.

[WB14] WONG S.-K., BACIU G.: Continuous collision detection for deformable objects using permissible clusters. *The Visual Computer 31*, 4 (2015), 377–389.

[WC14] WONG S.-K., CHENG Y.-C.: Continuous self-collision detection for deformable surfaces interacting with solid models. *Computer Graphics Forum 33*, 6 (2014), 143–153.

[WLH*13] WONG S.-K., LIN W.-C., HUNG C.-H., HUANG Y.-J., LII S.-Y.: Radial view based culling for continuous self-collision detection of skeletal models. *ACM Transactions on Graphics (ACM SIGGRAPH) 32*, 4 (2013), 114:1–114:10.

[WTTM15] WANG Z., TANG M., TONG R., MANOCHA D.: TightCCD: Efficient and robust continuous collision detection using tight error bounds. *Computer Graphics Forum 34*, (2015), 289–298.

[ZJ12] ZHENG C., JAMES D. L.: Energy-based self-collision culling for arbitrary mesh deformations. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2012) 31*, 4 (August 2012), 98:1–98:12.

[ZK14] ZHANG X., KIM Y.: Scalable collision detection using p-partition fronts on many-core processors. *IEEE Transactions on Visualization and Computer Graphics 20*, 3 (March 2014), 447–456.