

A Large Scale Benchmark and an Inclusion-Based Algorithm for Continuous Collision Detection

BOLUN WANG*, Beihang University and New York University

ZACHARY FERGUSON*, New York University

TESEO SCHNEIDER, New York University and University of Victoria

XIN JIANG, Beihang University and Peng Cheng Laboratory in Shenzhen

MARCO ATTENE, IMATI - CNR

DANIELE PANOZZO, New York University




Fig. 1. An overview of the results of our study of different CCD methods run on 60 million queries (both vertex-face and edge-edge). For each method, we show the number of false positives (i.e., the method detects a collision where there is none), the number of false negatives (i.e., the method misses a collision), and the average run time. Each plot reports results in a logarithmic scale. False positives and negatives are computed with respect to the ground truth computed using Mathematica [Wolfram Research Inc. 2020]. Acronyms are defined in Section 4.2.

We introduce a large scale benchmark for continuous collision detection (CCD) algorithms, composed of queries manually constructed to highlight challenging degenerate cases and automatically generated using existing simulators to cover common cases. We use the benchmark to evaluate the accuracy, correctness, and efficiency of state-of-the-art continuous collision detection algorithms, both with and without minimal separation.

We discover that, despite the widespread use of CCD algorithms, existing algorithms are either: (1) correct but impractically slow, (2) efficient but incorrect, introducing false negatives which will lead to interpenetration, or (3) correct but over conservative, reporting a large number of false positives which might lead to inaccuracies when integrated in a simulator.

By combining the seminal interval root finding algorithm introduced by Snyder in 1992 with modern predicate design techniques, we propose a simple and efficient CCD algorithm. This algorithm is competitive with state of the art methods in terms of runtime while conservatively reporting the

time of impact and allowing explicit trade off between runtime efficiency and number of false positives reported.

CCS Concepts: • **Computing methodologies** → **Collision detection**; **Physical simulation**.

Additional Key Words and Phrases: continuous collision detection, computational geometry, physically based animation

ACM Reference Format:

Bolun Wang, Zachary Ferguson, Teseo Schneider, Xin Jiang, Marco Attene, and Daniele Panozzo. 2021. A Large Scale Benchmark and an Inclusion-Based Algorithm for Continuous Collision Detection. *ACM Trans. Graph.* 1, 1, Article 1 (January 2021), 16 pages. <https://doi.org/10.1145/3460775>

1 INTRODUCTION

Collision detection and response are two separate, yet interconnected, problems in computer graphics and scientific computing. Collision detection specializes in finding when and if two objects collide, while collision response uses this information to deform the objects following physical laws. A large research effort has been invested in the latter problem, assuming that collision detection can be solved reliably and efficiently. In this study we focus on the former, using an experimental approach based on large scale testing. We use existing collision response methods to generate collision detection queries to investigate the pros and cons of existing collision detection algorithms.

Static collision detection is popular in interactive applications due to its efficiency, its inability to detect collisions between fast moving objects passing through each other (tunneling) hinders its applicability. To address this limitation, continuous collision detection (CCD) methods have been introduced: by solving a more computationally

*Joint first authors

Authors' addresses: Bolun Wang, LMIB & NLSDE & School of Mathematical Sciences & Shenyuan Honor College, Beihang University, New York University, wangbolun@buaa.edu.cn; Zachary Ferguson, New York University, zfergus@nyu.edu, zfergus@nyu.edu; Teseo Schneider, New York University, University of Victoria; Xin Jiang, LMIB & NLSDE & School of Mathematical Sciences, Beihang University, Peng Cheng Laboratory in Shenzhen; Marco Attene, IMATI - CNR; Daniele Panozzo, New York University, panozzo@nyu.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM. 0730-0301/2021/1-ART1 \$15.00 <https://doi.org/10.1145/3460775>

intensive problem, usually involving finding roots of a low-degree polynomial, these algorithms can detect any collision happening in a time step, often assuming linear trajectories.

The added robustness makes this family of algorithms popular, but they can still fail due to floating-point rounding errors. Floating point failures are of two types: false negatives, i.e., missed collisions, which lead to interpenetration, and false positives, i.e., detecting collisions when there are none.

Most collision response algorithms can tolerate minor imperfections, using heuristics to recover from physically invalid states (in reality, objects cannot inter-penetrate). However, these heuristics have parameters that need to be tuned for every scene to ensure stability and faithfulness in the simulation [Li et al. 2020]. Recently, the collision response problem has been reformulated to avoid the use of heuristics, and the corresponding parameter tuning, by disallowing physically invalid configurations [Li et al. 2020]. For instance, in the attached video, the method in [Li et al. 2020] cannot recover from interpenetration after the CCD misses a collision leading to an unnatural “sticking” and eventual failure of the simulation. This comes with a heavier burden on the CCD algorithm used, which should never report false negatives.

We introduce a large benchmark of CCD queries with ground truth computed using the *exact, symbolic solver* of Mathematica [Wolfram Research Inc. 2020], and evaluate the correctness (lack of false negatives), conservativeness (false positive count), and runtime efficiency of existing state of the art algorithms. The benchmark is composed of both manually designed queries to identify degenerate cases (building upon [Erleben 2018]) and a large collection of real-world queries extracted from simulation sequences. On the algorithmic side, we select representative algorithms from the three main approaches existing in the literature for CCD root-finding: *inclusion-based bisection methods* [Redon et al. 2002; Snyder et al. 1993], *numerical methods* [Vouga et al. 2010; Wang et al. 2015], and *exact methods* [Brochu et al. 2012; Tang et al. 2014]. Thanks to our benchmark, we identified missing cases that were not handled by previous methods, and we did a best effort to fix the corresponding algorithms and implementations to account for these cases.

The surprising conclusion of this study (Section 4.2) is that the majority of the existing CCD algorithms produce false negatives, except three: (1) symbolic solution of the system and evaluation with exact arithmetic computed using Mathematica [Wolfram Research Inc. 2020], (2) Bernstein sign classification (BSC) with conservative error analysis [Wang et al. 2015], and (3) inclusion-based bisection root finding [Redon et al. 2002; Snyder et al. 1993]. (1) is extremely expensive and, while it can be used for generating the ground truth, it is impractical in simulation applications. (2) is efficient but generates many false positives and the number of false positives depends on the geometric configuration and velocities involved. (3) is one of the oldest methods proposed for CCD. It is slow compared to state of the art algorithms, but it is correct and allows precise control of the trade-off between false positives and computational cost.

This extensive analysis and benchmark inspired us to introduce a specialization of the classical inclusion-based bisection algorithm proposed in [Snyder 1992] to the specific case of CCD for triangular meshes (Section 5). The major changes are: a novel inclusion function, an efficient strategy to perform bisection, and the ability

to find CCD roots with minimal separation (Section 6). Our novel inclusion function:

- (1) is tighter leading to smaller boxes on average thus making our method more accurate (i.e., less false positive);
- (2) reduces the root-finding problem into the iterative evaluation of a Boolean function, which allows replacing explicit interval arithmetic with a more efficient floating point filtering;
- (3) can be vectorized with AVX2 instructions.

With these modifications, our inclusion-based bisection algorithm is only $3\times$ slower on average than the fastest inaccurate CCD algorithm. At the same time it is provably conservative, provides a controllable ratio of false positives (within reasonable numerical limits), supports minimal separation, and reports the time of impact. We also discuss how to integrate minimal separation CCD in algorithms employing a line search to ensure the lack of intersections, which are common in locally injective mesh parametrization and have been recently introduced in physical simulation by Li et al. [2020].

Our dataset is available at the NYU Faculty Digital Archive, while the implementation of all the algorithms compared in the benchmark, a reference implementation of our novel inclusion-based bisection algorithm, and scripts to reproduce all results (Section 4) are available at <https://github.com/Continuous-Collision-Detection>. We believe this dataset will be an important element to support research in efficient and correct CCD algorithms, while our novel inclusion-based bisection algorithm is a practical solution that will allow researchers and practitioners to robustly check for collisions in applications where a $3\times$ slowdown in the CCD (which is usually only one of the expensive steps of a simulation pipeline) will be preferable over the risk of false negatives or the need to tune CCD parameters.

2 RELATED WORK

We present a brief overview of the previous works on continuous collision detection for triangle meshes. Our work focuses only on CCD for deformable triangle meshes and we thus exclude discussing methods approximating collisions using proxies (e.g., Hubbard [1995]; Mirtich [1996]).

Inclusion-Based Root-Finding. The generic algorithm in the seminal work of Snyder [1992] on interval arithmetic for computer graphics is a conservative way to find collisions [Redon et al. 2002; Snyder et al. 1993; Von Herzen et al. 1990]. This approach uses inclusion functions to certify the existence of roots within a domain, using a bisection partitioning strategy. Surprisingly, this approach is not used in recent algorithms despite being provably conservative and simple. Our algorithm is based on this approach, but with two major extensions to improve its efficiency (Section 5).

Numerical Root-Finding. The majority of CCD research focuses on efficient and accurate ways of computing roots of special cubic polynomials. Among these, a most popular cubic solver approach is introduced by Provot [1997], in which a cubic equation is solved to check for coplanarity, and then the overlapping occurrence is validated to determine whether a collision actually occurs. Refined constructions based on this idea have been introduced for rigid [Kim

and Rossignac 2003; Redon et al. 2002] and deformable [Hutter and Fuhrmann 2007; Tang et al. 2011] bodies. However, all of these algorithms are based on floating-point arithmetic, requiring numerical thresholds to account for the unavoidable rounding errors in the iterative root-finding procedure. In fact, even if the cubic polynomial is represented exactly, its roots are generally irrational and thus not representable with floating-point numbers. Unfortunately, the numerical thresholds make these algorithms robust only for specific scenarios, and they can in general introduce false negatives. Our approach has a moderately higher runtime than these algorithms, but it is guaranteed to avoid false negatives without parameter tuning. We benchmark Provot [1997] using the implementation of Vouga et al. [2010] in Section 4.

For most applications, false positives are less problematic than false negatives since a false negative will miss a collision, leading to interpenetration and potentially breaking the simulation. Tang et al. [2010] propose a simple and effective filter which can reduce both the number of false positives and the elementary tests between the primitives. Wang [2014] and Wang et al. [2015] improve its reliability by introducing forward error analysis, in which error bounds for floating-point computation are used to eliminate false positives. We benchmark the representative method of Wang et al. [2015] in Section 4.

Exact Root-Finding. Brochu et al. [2012] and Tang et al. [2014] introduce algorithms relying on exact arithmetic to provide exact continuous collision detection. However, after experimenting with their implementations and carefully studying their algorithms, we discovered that they cannot always provide the exact answer (Section 4). Brochu et al. [2012] rephrase the collision problem as counting the number of intersections between a ray and the boundary of a subset of \mathbb{R}^3 bounded by bilinear faces. The ray casting and polygonal construction can be done using rational numbers (or more efficiently with floating point expansions) to avoid floating-point rounding errors. In [Tang et al. 2014] the CCD queries are reduced to the evaluation of the signs of Bernstein polynomials and algebraic expressions, using a custom root finding algorithm. Our algorithm uses the geometric formulation proposed in [Brochu et al. 2012], but uses a bisection strategy instead of ray casting to find the roots. We benchmark both [Brochu et al. 2012] and [Tang et al. 2014] in Section 4.

Minimal Separation. Minimal separation CCD (MSCCD) [Harmon et al. 2011; Lu et al. 2019; Provot 1997; Stam 2009] reports collisions when two objects are at a (usually small) user-specified distance. These approaches have two main applications: (1) a minimal separation is useful in fabrication settings to ensure that the fabrication errors will not lead to penetrations, and (2) a minimal separation can ensure that, after floating-point rounding, two objects are still not intersecting, an invariant which must be preserved by certain simulation codes [Harmon et al. 2011; Li et al. 2020]. We benchmark [Harmon et al. 2011] in Section 6.2. Our algorithm supports a novel version of minimal separation, where we use the L^∞ norm instead of L^2 (Section 6.1).

Collision Culling. An orthogonal problem is efficient high-level collision culling to quickly filter out primitive pairs that do not

collide in a time step. Since in this case it is tolerable to have many false positives, it is easy to find conservative approaches that are guaranteed to not discard potentially intersecting pairs [Curtis et al. 2008; Govindaraju et al. 2005; Mezger et al. 2003; Pabst et al. 2010; Provot 1997; Schwartzman et al. 2010; Tang et al. 2009a, 2008; Volino and Thalmann 1994; Wong and Baciú 2006; Zhang et al. 2007; Zheng and James 2012]. Any of these approaches can be used as a preprocessing step to any of the CCD methods considered in this study to improve performance.

Generalized Trajectories. The linearization of trajectories commonly used in collision detection is a well-established, practical approximation, ubiquitous in existing codes. There are, however, methods that can directly detect collisions between objects following polynomial trajectories [Pan et al. 2012] or rigid motions [Canny 1986; Redon et al. 2002; Tang et al. 2009b; Zhang et al. 2007], and avoid the approximation errors due to the linearization. Our algorithm currently does not support curved trajectories and we believe this is an important direction for future work.

3 PRELIMINARIES AND NOTATION

Assuming that the objects are represented using triangular meshes and that every vertex moves in a linear trajectory in each time step, the first collision between moving triangles can happen either when a vertex hits a triangle, or when an edge hits another edge.

Thus a continuous collision detection algorithm is a procedure that, given a vertex-face or edge-edge pair, equipped with their *linear trajectories*, determines if and when they will touch. Formally, for the vertex-face CCD, given a vertex p and a face with vertices v_1, v_2, v_3 at two distinct time steps t^0 and t^1 (we use the superscript notation to denote the time, i.e., p^0 is the position of p at t^0), the goal is to determine if at any point in time between t^0 and t^1 the vertex is contained in the moving face. Similarly for the edge-edge CCD the algorithm aims to find if there exists a $t \in [t^0, t^1]$ where the two moving edges (p_1^t, p_2^t) and (p_3^t, p_4^t) intersect. We will briefly overview and discuss the pros and cons of the two major formulations present in the literature to address the CCD problem: *multi-variate* and *univariate*.

Multivariate CCD Formulation. The most direct way of solving this problem is to parametrize the trajectories with a parameter $t \in [0, 1]$ (i.e., $p_i(t) = (1-t)p_i^0 + tp_i^1$ and $v_i(t) = (1-t)v_i^0 + tv_i^1$) and write a multivariate polynomial whose roots correspond to intersections. That is finding the roots of

$$F_{vf}: \Omega_{vf} = [0, 1] \times \{u, v \geq 0 | u + v \leq 1\} \rightarrow \mathbb{R}^3$$

with

$$F_{vf}(t, u, v) = p(t) - ((1-u-v)v_1(t) + uv_2(t) + vv_3(t)), \quad (1)$$

for the vertex-face case. Similarly for the edge-edge case the goal is to find the roots of

$$F_{ee}: \Omega_{ee} = [0, 1] \times [0, 1]^2 \rightarrow \mathbb{R}^3$$

with

$$F_{ee}(t, u, v) = ((1-u)p_1(t) + up_2(t)) - ((1-v)p_3(t) + vp_4(t)). \quad (2)$$

In other words, the CCD problem reduces to determining if F has a root in Ω (i.e., there is a combination of valid t, u, v for which the




Fig. 2. Scenes from Erleben [2018] that are used to generate a large part of the handcrafted dataset.

vector between the point and the triangle is zero) [Brochu et al. 2012]. The main advantage of this formulation is that it is direct and purely algebraic: there are no degenerate or corner cases to handle. The intersection point is parameterized in time and local coordinates and the CCD problem reduces to multivariate root-finding. However, finding roots of a system of quadratic polynomials is difficult and expensive, which led to the introduction of the univariate formulation.

Univariate CCD Formulation. An alternative way of addressing the CCD problem is to rely on a *geometric* observation: two primitives intersect if the four points (i.e., one vertex and the three triangle’s vertices or the two pairs of edge’s endpoints) are coplanar [Provot 1997]. This observation has the major advantage of only depending on time, thus the problem becomes finding roots in a univariate cubic polynomial:

$$f(t) = \langle n(t), q(t) \rangle = 0, \quad (3)$$

with

$$n(t) = (v_2(t) - v_1(t)) \times (v_3(t) - v_1(t)) \text{ and } q(t) = p(t) - v_1(t)$$

for the vertex-face case and

$$n(t) = (p_2(t) - p_1(t)) \times (p_4(t) - p_3(t)) \text{ and } q(t) = p_3(t) - p_1(t)$$

for the edge-edge case. Once the roots t^* of f are identified, they need to be filtered, as not all roots correspond to actual collisions. While filtering is straightforward when the roots are finite, special care is needed when there is an infinite number of roots, such as when the two primitives are moving on the same plane. Handling these cases, especially while accounting for floating point rounding, is very challenging.

4 BENCHMARK

4.1 Dataset

We crafted two datasets to compare the performance and correctness of CCD algorithms: (1) a *handcrafted* dataset that contains over 12 thousand point-triangle and 15 thousand edge-edge queries, and (2) a *simulation* dataset that contains over 18 million point-triangle and 41 million edge-edge queries. To foster replicability, we describe the format of the dataset in Appendix A.

The handcrafted queries are the union of queries simulated with [Li et al. 2020] from the scenes in [Erleben 2018] (Figure 2) and a set of




Fig. 3. The scenes used to generate the simulation dataset of queries. We use two simulation methods: (top) a sequential quadratic programming (SQP) method with constraints and active set update from Verschuur and Jalba [2019] and (bottom) the method proposed by Li et al. [2020].

handcrafted pairs for degenerate geometric configurations. These include: point-point degeneracies, near collisions (within a floating-point epsilon from collision), coplanar vertex-face and edge-edge motion (where the function f (3) has infinite roots), degenerated function F_{vf} and F_{ee} , and CCD queries with two or three roots.

The simulation queries were generated by running four nonlinear elasticity simulations. The first two simulations (Figure 3 top row) use the constraints of [Verschuur and Jalba 2019] to simulate two cow heads colliding and a chain of rings falling. The second two simulations (Figure 3 bottom row) use the method of [Li et al. 2020] to simulate a coarse mat twisting and the high speed impact of a golf ball hitting a planar wall.

4.2 Comparison

We compare seven state-of-the-art methods: (1) the interval root-finder (IRF) [Snyder 1992], (2) the univariate interval root-finder (UIRF) (a special case of the rigid-body CCD from [Redon et al. 2002]), (3) the floating-point time-of-impact root finder [Provot 1997] (FPRF) implemented in [Vouga et al. 2010], (4) TightCCD (TCCD) [Wang et al. 2015], (5) Root Parity (RP) [Brochu et al. 2012], (6) a rational implementation of Root Parity (RRP) with the degenerate cases properly handled, and (7) Bernstein Sign Classification (BSC) [Tang et al. 2014]. For each method we collect the average query time, the number of false positives (i.e., there is no collision but the method detects one), and the number of false negatives (i.e., there is a collision but the method misses it). To obtain the ground truth we solve the multivariate CCD formulation (equations (1) and (2)) symbolically using Mathematica [Wolfram Research Inc. 2020] which takes multiple seconds per query. Table 1 summarizes the results. Note that “Ours” corresponds to our new method that will

Table 1. Summary of the average runtime in μs (t), number of false positive (FP), and number of false negative (FN) for the six competing methods.

Handcrafted Dataset (12K) – Vertex-Face CCD									
	IRF	UIRF	FPRF	TCCD	RP	RRP	BSC	MSRF	Ours
t	14942.40	124242.00	2.18	0.38	1.41	928.08	176.17	12.90	1532.54
FP	87	146	9	903	3	0	11	16	108
FN	0	0	70	0	5	5	13	386	0

Handcrafted Dataset (15K) – Edge-Edge CCD									
	IRF	UIRF	FPRF	TCCD	RP	RRP	BSC	MSRF	Ours
t	12452.60	18755.80	0.48	0.33	2.33	1271.32	121.80	2.72	3029.83
FP	141	268	5	404	3	0	28	14	214
FN	0	0	147	0	8	8	47	335	0

Simulation Dataset (18M) – Vertex-Face CCD									
	IRF	UIRF	FPRF	TCCD	RP	RRP	BSC	MSRF	Ours
t	115.89	6191.98	7.53	0.24	0.25	1085.13	34.21	51.07	0.74
FP	2	18	0	95638	0	0	23015	75	2
FN	0	0	5184	0	0	0	0	0	0

Simulation Dataset (41M) – Edge-Edge CCD									
	IRF	UIRF	FPRF	TCCD	RP	RRP	BSC	MSRF	Ours
t	215.80	846.57	0.23	0.23	0.37	1468.70	12.87	10.39	0.78
FP	71	16781	0	82277	0	0	4593	228	17
FN	0	0	2317	0	7	7	27	1	0

be introduced and discussed in Section 5 and MSRF is a minimum separation CCD discussed in Section 6.2.

IRF. The inclusion-based root-finding described in [Snyder 1992] can be applied to both the multivariate and univariate CCD. For the multivariate case we can simply initialize the parameters of F (i.e., t, u, v) with the size of the domain Ω , evaluate F and check if the origin is contained in the output interval [Snyder et al. 1993]. If it is, we sequentially subdivide the parameters (thus shrinking the size of the intervals of F) until a user-tolerance δ is reached. In our comparison we use $\delta = 10^{-6}$. The major advantage of this approach is that it is guaranteed to be conservative: it is impossible to shrink the interval of F to zero. A second advantage is that a user can easily trade accuracy (number of false positives) for efficiency by simply increasing the tolerance δ (Appendix D). The main drawback is that bisecting Ω in the three dimensions makes the algorithm slow, and the use of interval arithmetic further increases the computational cost and prevents the use of certain compiler optimization techniques (such as instruction reordering). We implement this approach using the numerical type provided by the Boost interval library [Schling 2011].

UIRF. [Snyder 1992] can also be applied to the univariate function in Equation (3) by using the same subdivision technique on the single variable t (as in [Redon et al. 2002] but for linear trajectories). The result of this step is an interval containing the earliest root in t which is then plugged inside a geometric predicate to check if the primitives intersect in that interval. While finding the roots with this approach might, at a first glance, seem easier than in the multi-variate case and thus more efficient, this is not the case in our experiments. If the polynomial has infinite roots, this algorithm will have to refine the entire domain to the maximal allowed resolution, and check the validity of each interval, making it correct but very

slow on degenerate cases (Appendix D). This results in a longer average runtime than its multivariate counterpart. Additionally, it is impossible to control the accuracy of the other two parameters (i.e., u, v), thus introducing more false positives.

FPRF. Vouga et al. [2010] aim to solve the univariate CCD problem using only floating-point computation. To mitigate false negatives, the method uses a numerical tolerance η (Appendix E) shows how η affects running time, the false positive, and negative). The major limitations are that the number of false positives cannot be directly controlled as it depends on the relative position of the input primitives and that false negatives can appear if the parameter is not tuned accordingly to the objects velocity and scale. Additionally, the reference implementation does not handle the edge-edge CCD when the two edges are parallel. This method is one of the fastest, which makes it a very popular choice in many simulation codes.

TCCD. TightCCD is a conservative floating-based implementation of Tang et al. [2014]. It uses the univariate formulation coupled with three inequality constraints (two for the edge-edge case) to ensure that the univariate root is a CCD root. The algorithm expresses the cubic polynomial f as a product and sum of three low order polynomials in Bernstein form. With this reformulation the CCD problem becomes checking if univariate Bernstein polynomials are positive, which can be done by checking some specific points. This algorithm is extremely fast but introduces many false positives which are impossible to control. In our benchmark, this is the only non-interval method without false negatives. The major limitation of this algorithm is that it *always* detects collision if the primitives are moving in the same plane, independently from their relative position.

RP and RRP. These two methods use the multivariate formulation F (equations (1) and (2)). The main idea is that the *parity* of the roots of F can be reduced to a ray casting problem. Let $\partial\Omega$ be the boundary of Ω , the algorithm shoots a ray from the origin and counts the parity of the intersection between the ray and $F(\partial\Omega)$ which corresponds to the parity of the roots of F . Parity is however insufficient for CCD: these algorithms cannot differentiate between zero roots (no collision) and two roots (collision), since they have the same parity. We note that this is a rare case happening only with sufficiently large time-steps and/or velocities: we found 13 (handcrafted dataset) and 7 (simulation dataset) queries where these methods report a false negative.

We note that the algorithm described in [Brochu et al. 2012] (and its reference implementation) does not handle some degenerate cases leading to both false negatives and positives. For instance, in Appendix B, we show an example of a “hourglass” configuration where RP misses the collision, generating a false negative. To overcome this limitations and provide a fair comparison to these techniques, we implemented a naïve version of this algorithm that handles all the degenerate cases using rational numbers to simplify the coding (see the additional materials). We opted for this rational implementation since properly handling the degeneracies using floating-point requires designing custom higher precision predicates for all cases. The main advantage of this method is that it is exact (when the degenerate cases are handled) as it does not

contain any tolerance and thus has zero false positives. We note that the runtime of our rational implementation is extremely high and not representative of the runtime of a proper floating point implementation of this algorithm.

BSC. This efficient and exact method uses the univariate formulation coupled with inequality constraints to ensure that the coplanar primitives intersects. The coplanarity problem reduces to checking if f in Bernstein form has a root. Tang et al. [2014] explain how this can be done *exactly* by classifying the signs of the four coefficients of the cubic Bernstein polynomial. The classification holds only if the cubic polynomial has monotone curvature; which can be achieved by splitting the curve at the inflection point. This splitting, however, cannot be computed exactly as it requires divisions (Appendix C). In our comparison, we modified the reference implementation to fix a minor typo in the code and to handle f with inflection points by *conservatively* reporting collision. This change introduces potential false positives, and we refer to the additional material for more details and for the patch we applied to the code.

Discussion and Conclusions. From our extensive benchmark of CCD algorithms, we observe that most algorithms using the univariate formulation have false negatives. While the reduction to univariate root findings provides a performance boost, filtering the roots (without introducing false positives) is a challenging problem for which a robust solution is still elusive.

Surprisingly, only the oldest method, IRF, is at the same time reasonably efficient (e.g., it does not take multiple seconds per query as Mathematica), correct (i.e., no false negatives), and returns a small number of false positives (which can be controlled by changing the tolerance δ). It is however slower than other state of the art methods, which is likely the reason why it is currently not widely used. In the next section we show that it is possible to change the inclusion function used by this algorithm to keep its favorable properties, while decreasing its runtime by ~ 250 times, making its performance competitive with state of the art methods.

5 METHOD

We describe the seminal bisection root-finding algorithm introduced in [Snyder 1992] (Section 5.1) and then introduce our novel Boolean inclusion function and how to evaluate it exactly and efficiently using floating point filters (Section 5.2).

5.1 SOLVE Algorithm [Snyder 1992]

An interval $i = [a, b]$ is defined as

$$i = [a, b] = \{x | a \leq x \leq b, x, a, b \in \mathbb{R}\},$$

and, similarly, an n -dimensional interval is defined as

$$I = i_1 \times \cdots \times i_n,$$

where i_k are intervals. We use $\mathcal{L}(i)$ and $\mathcal{R}(i)$ to refer to the left and right parts of an unidimensional interval i . The width of an interval, written as $w(i) = w([\mathcal{L}(i), \mathcal{R}(i)])$, is defined by

$$w(i) = \mathcal{L}(i) - \mathcal{R}(i)$$




Fig. 4. 1D illustration of the first three levels of the inclusion based root-finder in [Snyder 1992].

and similarly, the width of an n -dimensional interval

$$w(I) = \max_{k=\{1,\dots,n\}} w(i_k).$$

An interval can be used to define an inclusion function. Formally, given an m -dimensional interval D and a continuous function $g: \mathbb{R}^m \rightarrow \mathbb{R}^n$, an inclusion function for g , written $\Box g$, is a function such that

$$\forall x \in D \quad g(x) \in \Box g(D).$$

In other words, $\Box g(D)$ is a n -dimensional interval bounding the range of g evaluated over an m -dimensional interval D bounding its domain. We call the inclusion function $\Box g$ of a continuous function g convergent if for an interval X

$$w(X) \rightarrow 0 \implies w(\Box g(X)) \rightarrow 0.$$

A *convergent* inclusion function can be used to find a root of a function g over a domain bounded by the interval $I_0 = [\mathcal{L}(x_1), \mathcal{R}(x_1)] \times \cdots \times [\mathcal{L}(x_m), \mathcal{R}(x_m)]$. To find the roots of g , we sequentially bisect the initial m -dimensional interval I_0 , until it becomes sufficiently small (Algorithm 1). Figure 4 shows a 1D example (i.e., $g: \mathbb{R} \rightarrow \mathbb{R}$) of a bisection algorithm. The algorithm starts by initializing a stack S of intervals to be checked with I_0 (line 3). At every level ℓ (line 5), the algorithm retrieves an interval I from S and evaluates the inclusion function to obtain the interval I_g (line 7). Then it checks if the root is included in I_g (line 8). If not I can be safely discarded since I_g bounds the range of g over the domain bounded by I . Otherwise ($0 \in I_g$), it checks if $w(I)$ is smaller than a user-defined threshold δ . If so it appends I to the result (line 10). If I is too large, the algorithm splits one of its dimensions (e.g., $[\mathcal{L}(x_1), \mathcal{R}(x_1)]$ is split in $[\mathcal{L}(x_1), \tilde{x}_1]$ and $[\tilde{x}_1, \mathcal{R}(x_1)]$ with $\tilde{x}_1 = (\mathcal{L}(x_1) + \mathcal{R}(x_1))/2$) and appends the two new intervals I_1, I_2 to the stack S (line 13).

Generic Construction of Inclusion Functions. Snyder [1992] proposes the use of interval arithmetic as a universal and automatic way to build inclusion functions for arbitrary expressions. However, interval arithmetic adds a performance overhead to the computation. For example, the product between two intervals is

$$[a, b] \cdot [c, d] = [\min(ac, ad, bc, bd), \max(ac, ad, bc, bd)],$$

which requires four multiplications and two min/max instead of one multiplication. In addition, the compiler cannot optimize composite expressions, since the rounding modes need to be correctly set up and the operation needs to be executed in order to avoid rounding errors [Schling 2011].

Algorithm 1 Inclusion-based root-finder

```

1: function SOLVE( $I_0, g, \delta$ )
2:    $\text{res} \leftarrow \emptyset$ 
3:    $S \leftarrow \{I_0\}$ 
4:    $\ell \leftarrow 0$ 
5:   while  $L \neq \emptyset$  do
6:      $I \leftarrow \text{POP}(L)$ 
7:      $I_g \leftarrow \Box g(I)$  ▷ Compute the inclusion function
8:     if  $0 \in I_g$  then
9:       if  $w(I) < \delta$  then ▷  $I$  is small enough
10:         $\text{res} \leftarrow R \cup \{I\}$ 
11:       else
12:          $I_1, I_2 \leftarrow \text{SPLIT}(I)$ 
13:          $S \leftarrow S \cup \{I_1, I_2\}$ 
14:        $\ell \leftarrow \ell + 1$ 
15:   return  $\text{res}$ 

```

5.2 Predicate-Based Bisection Root Finding

Instead of using interval arithmetic to construct the inclusion function $\Box F$ for the interval $I_\Omega = I_t \times I_u \times I_v = [0, 1] \times [0, 1] \times [0, 1]$ around the domain Ω , we propose to define an inclusion function tailored for F (both for (1) and (2)) as the box

$$B_F(I_\Omega) = [m^x, M^x] \times [m^y, M^y] \times [m^z, M^z] \quad (4)$$

with

$$m^c = \min_{i=1,\dots,8} (v_i^c), \quad M^c = \max_{i=1,\dots,8} (v_i^c), \quad c = \{x, y, z\}$$

$$v_i = F(t_m, u_n, v_l), \quad t_m, u_n, v_l \in \{0, 1\}, \text{ and } m, n, l \in \{1, 2\}.$$

PROPOSITION 5.1. *The inclusion function B_F defined in (4) is the tightest axis-aligned inclusion function of F .*

PROOF. We note that for any given \tilde{u} the function $F(t, \tilde{u}, v)$ is bilinear; we call this function function $F_{\tilde{u}}(t, v)$. Thus, F can be regarded as a bilinear function whose four control points move along linear trajectories $\mathcal{T}(u)_i, i = 1, 2, 3, 4$. The range of $F_{\tilde{u}}$ is a bilinear surface which is bounded by the tetrahedron constructed by the four vertices forming the bilinear surface, which are moving on \mathcal{T}_i . Thus, F is bounded by every tetrahedron formed by $\mathcal{T}(u)_i$, implying that F is bounded by the convex hull of the trajectories' vertices, which are the vertices $v_i, i = 1, \dots, 8$ defining F . Finally, since B_F is the axis-aligned bounding box of the convex-hull of $v_i, i = 1, \dots, 8$, B_F is an inclusion function for F .

Since the vertices of the convex hull belong to F and the convex hull is the tightest convex hull, the bounding box B_F of the convex hull is the tightest inclusion function. \square

THEOREM 5.2. *The inclusion function B_F defined in (4) is convergent.*

PROOF. We first note that F is trivially continuous, second that the standard interval-based inclusion function $\Box F$ constructed with intervals is axis-aligned. Therefore, from Proposition 5.1, it follows that $B_F(I) \subseteq \Box F(I)$ for any interval I . Finally, since $\Box F$ is convergent [Snyder 1992], then also B_F is. \square

The inclusion function B_F turns out to be ideal for constructing a predicate: to use this inclusion function in the SOLVE algorithm (Algorithm 1), we only need to check if, for a given interval I , $B_F(I)$ contains the origin (line 8). Such a Boolean predicate can be conservatively evaluated using floating point filtering.

Conservative Predicate Evaluation. Checking if the origin is contained in an axis-aligned box is trivial and it reduces to checking if the zero is contained in the three intervals defining the sides of the box. In our case, this requires us to evaluate the sign of F at the eight box corners. However, the vertices of the co-domain are computed using floating point arithmetic and can thus be inaccurate. We use forward error analysis to conservatively account for these errors as follows.

Without loss of generality, we focus only on the x -axis. Let $\{v_i^x\}, i = 1, \dots, 8$ be the set of x -coordinates of the 8 vertices of the box represented in double precision floating-point numbers. The error bound for F (on the x -axis) is

$$\begin{aligned} \varepsilon_{ee}^x &= 6.217248937900877 \times 10^{-15} \gamma_x^3 \\ \varepsilon_{vf}^x &= 6.661338147750939 \times 10^{-15} \gamma_x^3 \end{aligned} \quad (5)$$

with

$$\gamma_x = \max(x_{\max}, 1) \quad \text{and} \quad x_{\max} = \max_{i=1,\dots,8} (|v_i^x|).$$

That is, the sign of F_{ee}^x computed using floating-point arithmetic is guaranteed to be correct if $|F_{ee}^x| > \varepsilon_{ee}^x$, and similarly for the vertex face case. If this condition does not hold, we conservatively assume that the zero is contained in the interval, thus leading to a possible false positive. The two constants ε_{ee}^x and ε_{vf}^x are floating point filters for F_{ee}^x and F_{vf}^x respectively, and were derived using [Attene 2020].

Efficient Evaluation. The x, y, z predicates defined above depend only on a subset of the coordinates of the eight corners of $B_F(I)$. We can optimally vectorize the evaluation of the eight corners using AVX2 instructions ($\sim 4\times$ improvement in performance), since it needs to be evaluated on eight points and all the computation is standard floating-point arithmetic. Note that we used AVX2 instructions because newer versions still have spotty support on current processors. After the eight points are evaluated in parallel, applying the floating-point filter involves only a few comparisons. To further reduce computation, we check one axis at a time and immediately return if any of the intervals do not contain the origin.

Algorithm. We describe our complete algorithm in pseudocode in Algorithm 2. The input to our algorithm are the eight points representing two primitives (either vertex-face or edge-edge), a user-controlled numerical tolerance $\delta > 0$ (if not specified otherwise, in the experiment we use the default value $\delta = 10^{-6}$), and the maximum number of checks $m_I > 0$ (we use the default value $m_I = 10^6$). These choice are based on our empirical results (figures 8 and 9). The output is a conservative estimate of the earliest time of impact or infinity if the two primitives do not collide in the time intervals coupled with the reached tolerance.

Our algorithm iteratively checks the box $B = B_F(I)$, with $I = I_t \times I_u \times I_v = [t_1, t_2] \times [u_1, u_2] \times [v_1, v_2] \subset I_\Omega$ (initialized with $[0, 1]^3$). To guarantee a uniform box size while allowing early termination of the algorithm, we explore the space in a breadth-first manner and record the current explored level ℓ (line 6). Since our algorithm

Algorithm 2 Complete overview of our CCD algorithm.

```

1: function SOLVE( $F, \delta, m_I$ )
2:    $n \leftarrow 0$                                  $\triangleright$  Number of check counter
3:    $Q \leftarrow \{[0, 1]^3, 0\}$                  $\triangleright$  Push first interval and level 0 in  $Q$ 
4:    $\ell_p \leftarrow -1$                            $\triangleright$  Previous checked level is -1
5:   while  $Q \neq \emptyset$  do
6:      $I, \ell \leftarrow \text{POP}(Q)$                  $\triangleright$  Retrieve level and interval
7:      $B \leftarrow B_F(I)$                      $\triangleright$  Compute the box inclusion function
8:      $n \leftarrow n + 1$                          $\triangleright$  Increase check number
9:     if  $B \cap C_\varepsilon \neq \emptyset$  then
10:      if  $\ell \neq \ell_p$  then  $\triangleright I$  is the first colliding interval of  $\ell$ 
11:         $I_f \leftarrow I_t$                      $\triangleright$  Save  $t$ -component of  $I$ 
12:      if  $n \geq m_I$  then  $\triangleright$  Reached max number of checks
13:        return  $\mathcal{L}(I_f), w(I_t)$              $\triangleright$  Return left side of  $I_f$ 
14:
15:      if  $w(B) < \delta$  or  $B \subseteq C_\varepsilon$  then
16:        if  $\ell \neq \ell_p$  then
17:          return  $\mathcal{L}(I_f), w(I_t)$              $\triangleright$  Root found
18:        else
19:           $I_1, I_2 \leftarrow \text{SPLIT}(I)$ 
20:           $Q \leftarrow Q \cup \{I_1, \ell + 1\}, \{I_2, \ell + 1\}$ 
21:           $\text{SORT}(Q, \text{ORDER})$ 
22:           $\ell_p = \ell$                          $\triangleright$  Update the previous colliding level
23:        return  $\infty, 0$                      $\triangleright Q$  is empty and no roots were found
24:
25: function SPLIT( $I = I_t \times I_u \times I_v$ )
26:   Compute  $\kappa_t, \kappa_u, \kappa_v$  according to (7)
27:    $c_t \leftarrow w(I_t)\kappa_t, c_u \leftarrow w(I_u)\kappa_u, c_v \leftarrow w(I_v)\kappa_v$ 
28:    $c \leftarrow \max(c_t, c_u, c_v)$ 
29:   if  $c_t = c$  then                         $\triangleright c_t$  is the largest
30:      $I_1 \leftarrow [\mathcal{L}(I_t), (\mathcal{L}(I_t) + \mathcal{R}(I_t))/2] \times I_u \times I_v$ 
31:      $I_2 \leftarrow [(\mathcal{L}(I_t) + \mathcal{R}(I_t))/2, \mathcal{R}(I_t)] \times I_u \times I_v$ 
32:   else if  $c_u = c$  then                     $\triangleright c_u$  is the largest
33:      $I_1 \leftarrow I_t \times [\mathcal{L}(I_u), (\mathcal{L}(I_u) + \mathcal{R}(I_u))/2] \times I_v$ 
34:      $I_2 \leftarrow I_t \times [(\mathcal{L}(I_u) + \mathcal{R}(I_u))/2, \mathcal{R}(I_u)] \times I_v$ 
35:   else                                     $\triangleright c_v$  is the largest
36:      $I_1 \leftarrow I_t \times I_u \times [\mathcal{L}(I_v), (\mathcal{L}(I_v) + \mathcal{R}(I_v))/2]$ 
37:      $I_2 \leftarrow I_t \times I_u \times [(\mathcal{L}(I_v) + \mathcal{R}(I_v))/2, \mathcal{R}(I_v)]$ 
38:   return  $I_1, I_2$ 
39:
40: function ORDER( $\{I_1, \ell_1\}, \{I_2, \ell_2\}$ )
41:   if  $\ell_1 = \ell_2$  then
42:     return  $I_1^t < I_2^t$ 
43:   else
44:     return  $\ell_1 < \ell_2$ 

```

is designed to find the earliest time of impact, we sort the visiting queue Q with respect to time (line 21).

At every iteration we check if B intersects the cube $C_\varepsilon = [-\varepsilon^x, \varepsilon^x] \times [-\varepsilon^y, \varepsilon^y] \times [-\varepsilon^z, \varepsilon^z]$ (line 9); if it does not, we can safely ignore I since there are no collisions.

If $B \cap C_\varepsilon \neq \emptyset$, we first check if $w(B) < \delta$ or if B is contained inside the ε -box (line 15). In this case, it is unnecessary to refine the




Fig. 5. A 2D example of root finding (left) and its corresponding diagram (right). A small colliding (red) box b that is not the earliest, since another box a exists in the same level (a did not trigger the termination of the algorithm since it is too big).




Fig. 6. A 2D example of root finding (left) and its corresponding diagram (right). Our algorithm stops when the number of checks n reaches m_I after checking the box s , which is a non-colliding box (green). The algorithm will return the first colliding box (f) of the same level, right.

interval I more since it is either already small enough (if $w(B) < \delta$) or any refinement will lead to collisions (if $B \subseteq C_\varepsilon$). We return I_t^l (i.e., the left hand-side of the t interval of I) only if I was the first intersecting interval of this current level (line 16). If I is not the first intersecting in the current level, there is an intersecting box (which is larger than δ) with an earlier time since the queue is sorted according to time (Figure 5).

If B is too big we split the interval I in two sub-intervals and push them to the priority queue Q (line 19). Note that, differently from Algorithm 1, we use a *priority queue* Q instead of the stack S . For the vertex-triangle CCD, the domain Ω is a prism, thus, after spitting the interval (line 19), we append I_1, I_2 to Q only if they intersect with Ω . To ensure that B shrinks uniformly (since the termination criteria, Line 15, is $w(B) < \delta$) we *conservatively* estimate the width of B (in the codomain) from the widths of the domain's (i.e., where the algorithm is acting) intervals I_t, I_u, I_v :

$$\alpha > 0, w(I_t) < \frac{\alpha}{\kappa_t}, w(I_u) < \frac{\alpha}{\kappa_u}, w(I_v) < \frac{\alpha}{\kappa_v} \implies w(B_F(I)) < \alpha \quad (6)$$

with α a given constant and

$$\begin{aligned} \kappa_t &= 3 \max_{i,j=1,2} \|F(0, u_i, v_j) - F(1, u_i, v_j)\|_\infty, \\ \kappa_u &= 3 \max_{i,j=1,2} \|F(t_i, 0, v_j) - F(t_i, 1, v_j)\|_\infty, \\ \kappa_v &= 3 \max_{i,j=1,2} \|F(t_i, u_j, 0) - F(t_i, u_j, 1)\|_\infty. \end{aligned} \quad (7)$$

PROPOSITION 5.3. Equation 6 holds for any positive constant α .

PROOF. While $B_F(I)$ is an interval, for the purpose of the proof we equivalently define it as an axis-aligned bounding box whose eight vertices are b_i . We will use the super-script notation to refer to the x, y, z component of a 3D point (e.g., b_i^x is the x -component of b_i) and define the set $I = \{1, \dots, 8\}$. By using the box definition

the width of $B_F(I)$ can be written as

$$w(B_F(I)) = \|b_M - b_m\|_\infty$$

with

$$b_M^k = \max_{i \in I} (b_i^k) \quad \text{and} \quad b_m^k = \min_{i \in I} (b_i^k).$$

Since $B_F(I)$ is the tightest axis-aligned inclusion function (Proposition 5.1)

$$b_M^k \leq \max_{i \in I} v_i^k, \quad b_m^k \leq \min_{i \in I} v_i^k,$$

where $v_i = F(I_t^j, I_u^k, I_v^l)$, with $j, k, l \in \{l, r\}$, thus for any coordinate k we bound

$$b_M^k - b_m^k = \max_{i, j \in I} (v_i^k - v_j^k) \leq \max_{i, j \in I} \|v_i - v_j\|_\infty.$$

For any pair of v_i and v_j we have

$$v_i - v_j = s_1 \alpha_{l,m} + s_2 \beta_{n,p} + s_3 \gamma_{p,q},$$

for some indices $l, m, n, o, p, q \in \{1, 2\}$ and constant $s_1, s_2, s_3 \in \{-1, 0, 1\}$ with

$$\alpha_{i,j} = w(I_t)(F(0, u_i, v_j) - F(1, u_i, v_j)),$$

$$\beta_{i,j} = w(I_u)(F(t_i, 0, v_j) - F(t_i, 1, v_j)),$$

$$\gamma_{i,j} = w(I_v)(F(t_i, u_j, 0) - F(t_i, u_j, 1)),$$

since F is linear on the edges. We note that $\alpha_{i,j}$, $\beta_{i,j}$, and $\gamma_{i,j}$ are the 12 edges of the box B_F . We now define

$$e_t^k = \max_{i,j \in \{1,2\}} |\alpha_{i,j}^k|, \quad e_u^k = \max_{i,j \in \{1,2\}} |\beta_{i,j}^k|, \quad e_v^k = \max_{i,j \in \{1,2\}} |\gamma_{i,j}^k|$$

which allows us to bound

$$\max_{i,j \in I} \|v_i - v_j\|_\infty \leq \|e_t + e_u + e_v\|_\infty \leq \|e_t\|_\infty + \|e_u\|_\infty + \|e_v\|_\infty.$$

Since

$$\|e_t\|_\infty \leq w(I_t) \max_{i,j=1,2} \|F(t_1, u_i, v_j) - F(t_2, u_i, v_j)\|_\infty = w(I_t) \kappa_t / 3,$$

and similarly $\|e_u\|_\infty \leq \kappa_u / 3$, $\|e_v\|_\infty \leq \kappa_v / 3$, we have

$$\|e_t\|_\infty + \|e_u\|_\infty + \|e_v\|_\infty \leq \frac{w(I_t) \kappa_t + w(I_u) \kappa_u + w(I_v) \kappa_v}{3}$$

Finally, from the assumption (6) it follows that

$$w(B_F(I)) \leq \max_{i,j \in I} \|v_i - v_j\|_\infty \leq \|e_t\|_\infty + \|e_u\|_\infty + \|e_v\|_\infty < \alpha.$$

□

Using the estimate of the width of I_t, I_u, I_v we split the dimension that leads to the largest estimated dimension in the range of F (line 28).

Fixed Runtime or Fixed Accuracy. To ensure a bounded runtime, which is very useful in many simulation applications, we stop the algorithm after an user-controlled number of checks m_I . To ensure that our algorithm always returns a *conservative* time of impact we record the first colliding interval I_f of every level (line 11). When the maximum number of check is reached we can safely return the latest recorded interval I_f (line 13) (Figure 6). We note that our algorithm will not respect the user specified accuracy when it terminates early: if a constant accuracy is required by applications, this additional termination criteria could be disabled, obtaining an algorithm with guaranteed accuracy but sacrificing the bound on the maximal running time. Note that without the termination criteria m_I , it is possible (while rare in our experiments) that the algorithm will take a long time to terminate, or run out of memory due to storing the potentially large list of candidate intervals L .

5.3 Results

Our algorithm is implemented in C++ and uses Eigen [Guennebaud et al. 2010] for the linear algebra routines (with the `-avx2` flag). We run our experiments on a 2.35 GHz AMD EPYC™ 7452. We attach the reference implementation and the data used for our experiments, which will be released publicly.

The running time of our method is comparable to the floating-point methods, while being provably correct, for any choice of parameters. For this comparison we use a default tolerance $\delta = 10^{-6}$ and default number of iterations $m_I = 10^6$. All queries in the simulation dataset terminate within 10^6 checks, while for the handcrafted dataset only 0.25% and 0.55% of the vertex-face and edge-edge queries required more than 10^6 checks, reaching an actual maximal tolerance δ of 2.14×10^{-5} and 6.41×10^{-5} for vertex-face and edge-edge respectively. We note that, despite the percentages begin small, by removing m_I the handcrafted queries take 0.015774 and 0.042477 seconds on average for vertex-face and edge-edge respectively. This is due to the large number of degenerate queries, as can be seen from the long tail in the histogram of the run-times (Figure 7). We did not observe any noticeable change of running time for the simulation dataset.

Our algorithm has two user-controlled parameters (δ and m_I) to control the accuracy and running time. The tolerance δ provides a direct control on the achieved accuracy and provides an indirect effect on the running time (Figure 8). The other parameter, m_I , directly controls the maximal running time of each query: for small m_I our algorithm will terminate earlier, resulting in a lower accuracy and thus more chances of false positives (Figure 9 top). We remark that, in practice, very few queries require so many subdivisions: by reducing m_I to the very low value of 100, our algorithm early-terminates only on $\sim 0.07\%$ of the 60 million queries in the simulation dataset.

6 MINIMUM SEPARATION CCD

An additional feature of some CCD algorithms is *minimal separation*, that is, the option to report collision at a controlled distance from an object, which is used to ensure that objects are never too close. This is useful to avoid possible inter-penetrations introduced by numerical rounding after the collision response, or for modeling

Table 2. Summary of the average runtime in μs (t), number of false positive (FP), and number of false negative (FN) for MSRF and our method.

d	Handcrafted – Vertex-Face MSCCD						Handcrafted – Edge-Edge MSCCD						Simulation – Vertex-Face MSCCD						Simulation – Edge-Edge MSCCD					
	MSRF			Ours			MSRF			Ours			MSRF			Ours			MSRF			Ours		
	t	FP	FN	t	FP	FN	t	FP	FN	t	FP	FN	t	FP	FN	t	FP	FN	t	FP	FN	t	FP	FN
10^{-2}	12.89	854	114	18.86K	2.6K	0	3.84	774	189	9.64K	4.8K	0	55.47	156.8K	18.3K	12.04	8.1M	0	14.42	354.1K	7.0K	19.12	8.3M	0
10^{-8}	15.05	216	2	1.60K	159	0	2.89	230	18	3.42K	309	0	55.26	75	0	0.72	8	0	11.12	228	1	0.73	40	0
10^{-16}	13.90	151	35	1.51K	108	0	2.90	231	21	2.92K	214	0	54.83	4	3.8K	0.71	2	0	10.70	10	4	0.72	17	0
10^{-30}	13.59	87	141	1.39K	108	0	2.89	118	157	2.79K	214	0	53.73	0	10.2K	0.66	2	0	10.68	0	1.7K	0.67	17	0
10^{-100}	14.45	16	384	1.43K	108	0	3.05	14	335	2.82K	214	0	53.53	0	18.6K	0.66	2	0	10.59	0	5.0K	0.68	17	0

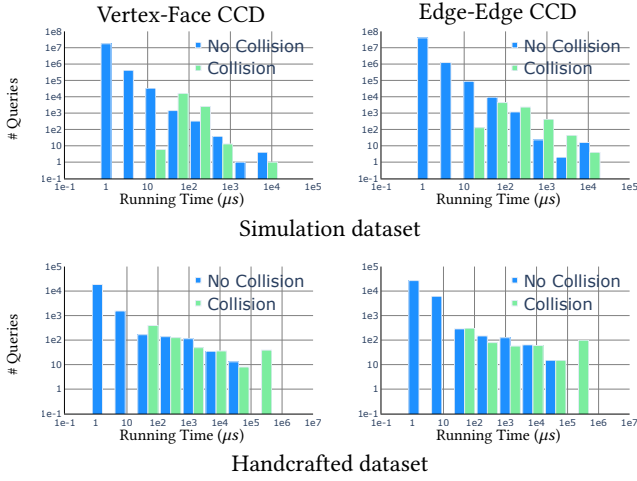




Fig. 7. Log histograms of the running time of positive queries and negative queries on both dataset.

Fig. 8. Top, average runtime of our algorithm for different tolerances δ for the simulation dataset. The shaded area shows the range of the distribution (min and max). Bottom, distribution of running times of our algorithm for three different tolerances $\delta = 10^{-8}$, 10^{-4} , and 1 over the simulation dataset.

fabrication tolerances for additive or subtractive manufacturing. A minimum separation CCD (MSCCD) query is similar to a standard query: instead of checking if a point and a triangle (or two edges) are exactly overlapping, we want to ensure that they are always

Fig. 9. The percentage of early-termination and maximum value of the tolerance δ for different m_I for the simulation dataset.

separated by a user-defined distance d during the entire linear trajectory. Similarly to the standard CCD (Section 3) MSCCD can be expressed using a multivariate or a univariate formulation, usually measuring distances using the Euclidean distance. We focus on the multivariate formulation since it does not require to filter spurious roots, we refer to Section 4.2 for a more detailed justification of this choice.

Multivariate Formulation. We observed that using the Euclidean distance leads to a challenging problem, which can be geometrically visualized as follows: the primitives will not be closer than d if $F(\Omega)$ does not intersect a sphere of radius d centered on the origin. This is a hard problem, since it requires checking conservatively the intersection between a sphere (which is a rational polynomial when explicitly parametrized) and $F(\Omega)$.

Studying the applications currently using minimal separation, we realized that they are not affected by a different choice of the distance function. Therefore, we propose to change the distance definition from Euclidean to Chebyshev distance (i.e., from the L^2 to the L^∞ distance). With this minor change the problem dramatically simplifies: instead of solving for $F = 0$ (Section 5), we need to solve for $|F| \leq d$. The corresponding geometric problem becomes checking if $F(\Omega)$ intersects a cube of side $2d$ centered on the origin.

Univariate Formulation. The univariate formulation is more complex since it requires to redefine the notion of co-planarity for minimum separation. We remark that the function f in (3) measures the length of the projection of $q(t)$ along the normal, thus to find point at distance d the equation becomes $f(t) \leq \langle n(t), q(t) \rangle = d \|n(t)\|$.




Fig. 10. 1D illustration of the first three levels of our MSCCD inclusion based root-finder. Instead of checking if I_g intersects with the origin, we check if it intersects the interval $[-d, d]$ marked in light green.

To keep the equation polynomial, remove the inequality, and avoid square roots, the univariate MSCCD root finder becomes

$$\langle n(t), q(t) \rangle^2 - d^2 \|n(t)\|^2.$$

We note that this polynomial becomes sextic, and not cubic as in the zero-distance version. To account for replacing the inequality with an equality, we also need to check for distance between q and the edges and vertices of the triangle [Harmon et al. 2011]. In addition to finding the roots of several high-order polynomials, this formulation, similarly to the standard CCD, suffers from infinite roots when the two primitives are moving on a plane at distance d from each other.

6.1 Method

The input to our MSCCD algorithm are the same as the standard CCD (eight coordinates, δ , and m_I) and the minimum separation distance $d \geq 0$. Our algorithm returns the earliest time of impact indicating if two primitives become closer than d as measured by the L^∞ norm.

We wish to check whether the box $B_F(\Omega)$ intersects a cube of side $2d$ centered on the origin (Figure 10). Equivalently, we can construct another box $B'_F(\Omega)$ by displacing the six faces of $B_F(\Omega)$ outward at a distance d , and then check whether this enlarged box contains the origin. This check can be done as for the standard CCD (Section 5), but the floating point filters must be recalculated to account for the additional sum (indeed, we add/subtract d to/from all the coordinates). Hence, the filters for F' are:

$$\begin{aligned} \epsilon_{ee}^x &= 7.105427357601002 \times 10^{-15} \gamma_x^3 \\ \epsilon_{vf}^x &= 7.549516567451064 \times 10^{-15} \gamma_x^3 \end{aligned} \quad (8)$$

As before, the filters are calculated as described in [Attene 2020] and they additionally assume that $d < \gamma_x$.

To account for minimum separations, the only change in our algorithm is at line 7 where we need to enlarge B by d and in lines 9 and 15 since C_ϵ needs to be replaced with $C_\epsilon = [-\epsilon^x, \epsilon^x] \times [-\epsilon^y, \epsilon^y] \times [-\epsilon^z, \epsilon^z]$.

6.2 Results

To the best of our knowledge, the minimum separation floating-point time-of-impact root finder [Harmon et al. 2011] (MSRF) implemented in [Lu et al. 2019], is the only public code supporting minimal separation queries. While not explicitly constructed for




Fig. 11. Top, average runtime of our algorithm for varying minimum separation d in the simulation dataset. The shaded area depicts the range of the values. Bottom, distribution of running time for three different minimum separation distances $d = 10^{-50}$, 10^{-8} , and 1 over the simulation dataset.

MSCCD, FPRF uses a distance tolerance to limit false negatives, similarly to an explicit minimum separation. We compare the results and performance in Appendix E.

MSRF uses the univariate formulation, which requires to find the roots of a high-order polynomial, and it is thus unstable when implemented using floating-point arithmetic.

Table 2 reports timings, false positive, and false negatives for different separation distances d . As d shrinks (around 10^{-16}) the results of our method with MSCDD coincide with the ones with $d = 0$ since the separation is small. For these small tolerances, MSRF runs into numerical problems and the number of false negatives increases. Figure 11 shows the average query time versus the separation distance d for the simulation dataset, since our method only requires to check the intersection between boxes, the running time largely depends on the number of detected collision, and the average is only mildly affected by the choice of d .

7 INTEGRATION IN EXISTING SIMULATORS

In a typical simulation the objects are represented using triangular meshes and the vertices are moving along a linear trajectory in a timestep. At each timestep, collisions might happen when a vertex hits a triangle, or when an edge hits another edge. A CCD algorithm is then used to prevent interpenetration; this can be done in different ways. In an active set construction method (Section 7.1) the CCD is used to compute contact forces to avoid penetration assuming linearized contact behaviour. For a line-search based method (Section 7.2), CCD and time of impact are used to prevent the Newton trajectory from causing penetration by limiting the step length. Note that, the latter approach requires a conservative CCD, while the former can tolerate false negatives.

The integration of a CCD algorithm with collision response algorithms is a challenging problem on its own, which is beyond the scope of this paper. As a preliminary study, to show that our method

can be integrated in existing response algorithm, we examine two use cases in elastodynamic simulations:

- (1) constructing an active set of collision constraints [Harmon et al. 2008; Verschoor and Jalba 2019; Wriggers 1995], Section 7.1;
- (2) during a line search to prevent intersections [Li et al. 2020], Section 7.2.

We leave as future work a more comprehensive study including how to use our CCD to further improve the physical fidelity of existing simulators or how to deal with challenging cases such as sliding contact response.

To keep consistency across queries, we compute the numerical tolerances (5) and (8) for the whole scene. That is, x_{\max} , y_{\max} , and z_{\max} are computed as the maximum over all the vertices in the simulation. In algorithms 3 and 4 we utilize a broad phase method (e.g., spatial hash) to reduce the number of candidates C that need to be evaluated with out narrow phase CCD algorithm.

7.1 Active Set Construction

Algorithm 3 Active Set Construction Using Exact CCD

```

1: function CONSTRUCTACTIVSET( $x_0, x_1, \delta, m_I$ )
2:    $C \leftarrow \text{BROADPHASE}(x_0, x_1)$ 
3:    $C_A \leftarrow \emptyset$ 
4:   for  $c \in C$  do                                ▶ Iterate over the collision candidates
5:      $t \leftarrow \text{CCD}(x_0 \cap c, x_1 \cap c, \delta, m_I)$ 
6:     if  $0 \leq t \leq 1$  then
7:        $C_A \leftarrow C_A \cup \{(c, t)\}$ 
8:   return  $C_A$ 
9:
10: function CCD( $c_0, c_1, \delta, m_I$ )
11:   if  $c_0$  and  $c_1$  are edges then
12:      $F \leftarrow \text{build } F_{ee} \text{ from } c_0 \text{ and } c_1$            ▶ Equation (2)
13:   else
14:      $F \leftarrow \text{build } F_{vf} \text{ from } c_0 \text{ and } c_1$            ▶ Equation (1)
15:   return SOLVE( $F, \delta, m_I$ )

```

In the traditional constraint based collision handling (such as that of Verschoor and Jalba [2019]), collision response is handled by performing an implicit timestep as a constrained optimization. The goal is to minimize a elastic potential while avoiding interpenetration through gap constraints. To avoid handling all possible collisions during a simulation, a subset of *active* collisions constraints C_A is usually constructed. This set not only avoids infeasibilities, but also improves performance by having fewer constraints. There are many activation strategies, but for the sake of brevity we focus here on the strategies used by Verschoor and Jalba [2019].

Algorithm 3 shows how CCD is used to compute the active set C_A . Given the starting and ending vertex positions, x_0 and x_1 , we compute the time of impact for each collision candidate $c \in C$. We use the notation $x_i \cap c$ to indicate selecting the constrained vertices from x_i . If the candidate c is an actual collision, that is $0 \leq t \leq 1$, then we add this constraint and the time of impact, t , to the active set, C_A .




Fig. 12. An elastic simulation using the constraints and active set method of Verschoor and Jalba [2019]. From an initial configuration (left) we simulate an elastic torus falling on a fixed cone using three values of δ (from left to right: 10^{-1} , 10^{-3} , 10^{-6}). The total runtime of the simulation is affected little by the change in δ (24.7, 25.2, and 26.2 seconds from left to right compared to 32.3 seconds when using FPRF). For $\delta = 10^{-1}$, inaccuracies in the time-of-impact lead to inaccurate contact points in the constraints and, ultimately, intersections (inset).

From the active constraint set the constraints of Verschoor and Jalba [2019] are computed as

$$\langle n, p_c^1 - p_c^2 \rangle \geq 0,$$

where n is the contact normal (i.e., for a point-triangle the triangle normal at the time of impact and for edge-edge the edge-edge cross product at the time of impact), p_c^1 is the point (or the contact point on the first edge), and p_c^2 is the point of contact on the triangle (or on the second edge) at the end of the timestep. Note that, this constraint requires to compute the point of contact, which depends on the time-of-impact which can be obtained directly from our method.

Because of the difficulty for a simulation solver to maintain and not violate constraints, it is common to offset the constraints such that

$$\langle n, p_c^1 - p_c^2 \rangle \geq \eta > 0.$$

In such a way, even if the η constraint is violated, the real constraint is still satisfied. This common trick, implies that the constraints need to be activated early (i.e., when the distance between two objects is smaller than η) which is exactly what our MSCCD can compute when using $d = \eta$. In Figure 12, we use a value of $\eta = 0.001$ m. When using large values of η , the constraint of Verschoor and Jalba [2019] can lead to infeasibilities because all triangles are extended to planes and edges to lines.

Figure 12 shows example of simulations run with different numerical tolerance δ . Changing δ has little effect on the simulation in terms of run-time, but for large values of δ , it can affect accuracy. We observe that for a $\delta \geq 10^{-2}$ the simulation is more likely to contain intersections. This is most likely due to the inaccuracies in the contact points used in the constraints.

7.2 Line Search

A line search is used in a optimization to ensure that every update decreases the energy E . That is, given an update, Δx , to the optimization variable x , we want to find a step size α such that $E(x + \alpha \Delta x) < E(x)$. This ensure that we make progress towards a minimum.

Algorithm 4 Line Search with Exact CCD

```

1: function LINESEARCH( $E, x_0, \Delta x, p, \delta, m_I$ )
2:    $x_1 \leftarrow x_0 + \Delta x$ 
3:    $C \leftarrow \text{BROADPHASE}(x_0, x_1)$  ▷ Collision candidates
4:    $\alpha \leftarrow 1$ 
5:    $d_i, \rho_i \leftarrow \text{DISTANCE}(C)$ 
6:   Compute  $\epsilon_i$  from (8)
7:    $d \leftarrow \max(p d_i, \delta)$ 
8:   while  $p < (d - \delta - \epsilon_i - \rho_i)/d$  do
9:      $p \leftarrow p/2$ 
10:     $d \leftarrow p d_i$ 
11:     $\delta_i \leftarrow \delta$ 
12:    for  $c \in C$  do ▷  $\alpha$  is bounded by earliest time-of-impact
13:       $t, \tilde{\delta}_i \leftarrow \text{MSCCD}(x_0 \cap c, x_1 \cap c, d, \alpha, \delta, m_I)$ 
14:       $\alpha \leftarrow \min(t, \alpha)$ 
15:       $\delta_i \leftarrow \max(\tilde{\delta}_i, \delta_i)$ 
16:    if  $p < (d - \delta_i - \epsilon_i - \rho_i)/d$  then
17:       $\delta \leftarrow \delta_i$  ▷ Repeat with  $p$  validated from  $\delta_i$ 
18:      Go to line 8.
19:
20:   while  $\alpha > \alpha_{\min}$  do ▷ Backtracking line-search
21:      $x_1 \leftarrow x_0 + \alpha \Delta x$ 
22:     if  $E(x_1) < E(x_0)$  then ▷ Objective energy decrease
23:       break
24:      $\alpha \leftarrow \alpha/2$ 
25:   return  $\alpha$ 
26:
27: function MSCCD( $c_0, c_1, d, t, \delta, m_I$ )
28:   if  $c_0$  and  $c_1$  are edges then
29:      $F \leftarrow \text{build } F_{ee} \text{ from } c_0 \text{ and } c_1$  ▷ Equation (2)
30:   else
31:      $F \leftarrow \text{build } F_{vf} \text{ from } c_0 \text{ and } c_1$  ▷ Equation (1)
32:   return SOLVEMSCCD( $F, t, \delta, m_I, d$ )

```

When used in a line search algorithm, CCD can be used to prevent intersections and tunneling. This requires modifying the maximum step length to the time of impact. As observed by Li et al. [2020], the standard CCD formulation without minimal separation cannot be used directly in a line search algorithm. Let t^* the earliest time of impact (i.e., $F(t^*, \tilde{u}, \tilde{v}) = 0$ for some \tilde{u}, \tilde{v} and there is no collision between 0 and t^*) and assume that the energy at $E(x_0 + t^* \Delta x) < E(x_0)$ (Algorithm 4, line 22). In this case the step $\alpha = t^*$ is a valid descent step which will be used to update the position x in outer iteration (e.g., Newton optimization loop). In the next iteration, the line search will be called with the updated position and the earliest time of impact will be zero since we selected t^* in the previous iteration. This prevents the optimization from making progress because any direction Δx will lead to a time of impact $t = 0$. To avoid this problem we need the line search to find an appropriate step-size α along the update direction that leaves “sufficient space” for the next iteration, so that the barrier in [Li et al. 2020] will be active and steer the optimization away from the contact position.

Formally, we aim at finding a *valid CCD sequence* $\{t_i\}$ such that

$$t_i < t_{i+1}, \quad \lim_{i \rightarrow \infty} t_i = t^*, \quad \text{and} \quad t_i/t_{i+1} \approx 1.$$

The first requirement ensures that successive CCD checks will report an increasing time, the second one ensures that we will converge to the true minimum, and the last one aims at having a “slowly” convergent sequence (necessary for numerical stability). Li et al. [2020] exploit a feature of FPRF to simulate a minimal separation CCD: in this work we propose to directly use our MSCCD algorithm (Section 6).

Constructing a Sequence. Let $0 < p < 1$ be a user-defined tolerance (p close to 1 will produce a sequence $\{t_i\}$ converging faster) and d_i be the distance between two primitives. We propose to set $d = p d_i$, and ensure that no primitive are closer than d . Without loss of generality, we assume that $F(x + \Delta x) = 0$, that is, taking the full step will lead to contact. By taking successive steps in the same direction, d_i will shrink to zero ensuring t_i to converge to t^* . Similarly we will obtain a growing sequence t_i since d decreases as we proceed with the iterations. Finally, it is easy to see that $p = t_i/t_{i+1}$ which can be close to one.

To account for the aforementioned problem, we propose to use our MSCCD algorithm to return a valid CCD sequence when employed in a line search scenario. For a step i , we define δ^i as the tolerance, ϵ_i the numerical error (8), and ρ_i as the maximum numerical error in computing the distances d_i from the candidates set C (line 5). ρ_i should be computed using forward error analysis on the application-specific distance computation: since the applications are not the focus of our paper, we used a fixed $\rho_i = 10^{-9}$, and we leave the complete forward analysis as a future work. (We note that our approximation might thus introduce zero length steps, this however did not happen in our experiments.) If $d_i - (\delta_i + \epsilon_i + \rho_i) > d$, our MSCCD is guaranteed to find a time of impact larger than zero. Thus if we set $d = p d_i$ (line 7), we are guaranteed to find a positive time of impact if

$$d_i > \frac{\delta_i + \epsilon_i + \rho_i}{1 - p}.$$

To ensure that this inequality holds, we propose to validate p before using the MSCCD with δ (line 8), find the time of impact and the actual δ_i (line 12), and check if the used p is valid (line 16). In case p is too large, we divide it by two until it is small enough. Note that, it might be that

$$d_i < \delta_i + \epsilon_i + \rho_i,$$

in this case we can still enforce the inequality by increasing the number of iterations, decreasing δ , or using multi-precision in the MSCCD to reduce ϵ_i . However, this was never necessary in any of our queries, and we thus leave a proper study of these options as a future work.

As visible from Table 2, our MSCCD slows down as d grows. Since the actual minimum distance is not relevant in the line search algorithm, our experiments suggest to cap it at δ (line 7). To avoid unnecessary computations and speedup the MSCCD computations, our algorithm, as suggested by Redon et al. [2002], can be easily modified to accept a shorter time interval (line 13): it only requires to change the initialization of I (Algorithm 2 line 3). These two modifications lead to a 8× speedup in our experiments. We refer




Fig. 13. An example of an elastic simulation using our line search (Section 7.2) and the method of Li et al. [2020] to keep the bodies from intersecting. An octocat is falling under gravity onto a triangulated plane. From left to right: the initial configuration, the final frame with $\delta = 10^{-3}$, $\delta = 10^{-4.5}$, $\delta = 10^{-6}$ all with a maximum of 10^6 iterations. There are no noticeable differences in the results, and the entire simulations takes 63.3, 67.9, and 67.0 seconds from left to right (a speed up compared to using FPRF which takes 102 seconds). ©Brian Enigma under CC BY-SA 3.0.

to this algorithm with MSCCD (i.e., Algorithm 2 with MSCDD, Section 6.1, and modified initialization of I) as SolveMSCCD.

Figure 13 shows a simulation using our MSCCD in line search to keep the bodies from intersecting for different δ . As illustrated in the previous section, the effect of δ is negligible as long as $\delta \leq 10^{-3}$. Timings vary depending on the maximum number of iterations. Because the distance d varies throughout the simulation, some steps take longer than others (as seen in Figure 11). We note that, if we use the standard CCD formulation $F = 0$, the line search gets stuck in all our experiments, and we were not able to find a solution. Note that for a line search based method it is crucial to have a conservative CCD/MSCCD algorithm: the videos in the additional material shows that a false negative leads to an artefact in the simulation.

8 LIMITATIONS AND CONCLUDING REMARKS

We constructed a benchmark of CCD queries and used it to study the properties of existing CCD algorithms. The study highlighted that the multivariate formulation is more amenable to robust implementations, as it avoids a challenging filtering of spurious roots. This formulation, paired with an interval root finder and modern predicate construction techniques leads to a novel simple, robust, and efficient algorithm, supporting minimal separation queries with runtime comparable to state of the art, non conservative, methods.

While we believe that it is practically acceptable, our algorithm still suffers from false positive and it will be interesting to see if the multivariate root finding could be done exactly with reasonable performances, for example employing expansion arithmetic in the predicates. Our definition of minimal separation distance is slightly different from the classical definition, and it would be interesting to study how to extend out method to directly support Euclidean distances. Another interesting venue for future work is the extension of our inclusion function to non-linear trajectories and their efficient evaluation using static filters or exact arithmetic.

Our benchmark focuses only on CPU implementations: reimplementing our algorithm on a GPU with our current guarantees is a major challenge. It will require to control the floating-point rounding on the GPU (and compliant with the IEEE floating-point standard), to ensure that the compiler does not reorder the operations or skip the computation of temporaries. Additionally it would require to recompute the ground truth and the numerical constants for single precision arithmetic, as most GPUs do not yet support

double computation. This is an exciting direction for future work to further improve the performance of our approach.

We will release an open-source reference implementation of our technique with an MIT license to foster adoption of our technique by existing commercial and academic simulators. We will also release the dataset and the code for all the algorithms in our benchmark to allow researchers working on CCD to easily compare the performance and correctness of future CCD algorithms.

ACKNOWLEDGMENTS

We thank Danny Kaufman for valuable discussions and NYU IT High Performance Computing for resources, services, and staff expertise. This work was partially supported by the NSF CAREER award under Grant No. 1652515, the NSF grants OAC-1835712, OIA-1937043, CHS-1908767, CHS-1901091, National Key Research and Development Program of China No. 2020YFA0713700, EU ERC Advanced Grant CHANGE No. 694515, a Sloan Fellowship, a gift from Adobe Research, a gift from nTopology, and a gift from Advanced Micro Devices, Inc.

REFERENCES

- Marco Attene. 2020. Indirect Predicates for Geometric Constructions. *Computer-Aided Design* 126 (2020).
- Tyson Brochu, Essex Edwards, and Robert Bridson. 2012. Efficient Geometrically Exact Continuous Collision Detection. *ACM Transactions on Graphics* 31, 4, Article 96 (July 2012), 7 pages.
- John Canny. 1986. Collision Detection for Moving Polyhedra. *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI-8, 2 (1986), 200–209.
- Sean Curtis, Rasmus Tamstorf, and Dinesh Manocha. 2008. Fast Collision Detection for Deformable Models using Representative-Triangles. *Proceedings of the Symposium on Interactive 3D Graphics and Games*, 61–69.
- Kenny Erleben. 2018. Methodology for Assessing Mesh-Based Contact Point Methods. *ACM Transactions on Graphics* 37, 3, Article 39 (July 2018), 30 pages.
- Naga Govindaraju, David Knott, Nitin Jain, Ilknur Kabul, Rasmus Tamstorf, Russell Gayle, Ming Lin, and Dinesh Manocha. 2005. Collision Detection Between Deformable Models using Chromatic Decomposition. *ACM Transactions on Graphics* 24 (July 2005), 991–999.
- Torbjörn Granlund and the GMP Development Team. 2012. *GNU MP: The GNU Multiple Precision Arithmetic Library* (5.0.5 ed.). <http://gmplib.org/>
- Gaël Guennebaud, Benoît Jacob, et al. 2010. Eigen v3. <http://eigen.tuxfamily.org>
- David Harmon, Daniele Panozzo, Olga Sorkine, and Denis Zorin. 2011. Interference-Aware Geometric Modeling. *ACM Transactions on Graphics* 30, 6 (Dec. 2011), 1–10.
- David Harmon, Etienne Vouga, Rasmus Tamstorf, and Eitan Grinspun. 2008. Robust Treatment of Simultaneous Collisions. *ACM Transactions on Graphics* 27 (Aug. 2008).
- Philip Martyn Hubbard. 1995. Collision Detection for Interactive Graphics Applications. *IEEE Transactions on Visualization and Computer Graphics* 1, 3 (1995), 218–230.
- Marco Hutter and Arnulph Fuhrmann. 2007. Optimized Continuous Collision Detection for Deformable Triangle Meshes. *Journal of WSCG* 15 (July 2007), 25–32.
- Byungmoon Kim and Jarek Rossignac. 2003. Collision Prediction for Polyhedra Under Screw Motions. *Proceedings of the Eighth ACM Symposium on Solid Modeling and Applications*, 4–10.
- Minchen Li, Zachary Ferguson, Teseo Schneider, Timothy Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M. Kaufman. 2020. Incremental Potential Contact: Intersection-and Inversion-Free, Large-Deformation Dynamics. *ACM Transactions on Graphics* 39, 4, Article 49 (July 2020), 20 pages.
- Libin Lu, Matthew J. Morse, Abtin Rahimian, Georg Stadler, and Denis Zorin. 2019. Scalable Simulation of Realistic Volume Fraction Red Blood Cell Flows through Vascular Networks. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC '19)*. Association for Computing Machinery, New York, NY, USA, Article 6, 30 pages. <https://github.com/Continuous-Collision-Detection/Minimum-Separation-Root-Finder>
- Johannes Mezger, Stefan Kimmerle, and Olaf Eitzmuss. 2003. Hierarchical Techniques in Collision Detection for Cloth Animation. *Journal of WSCG* 11, 322–329.
- Brian Vincent Mirtich. 1996. *Impulse-Based Dynamic Simulation of Rigid Body Systems*. Ph.D. Dissertation.
- Simon Pabst, Artur Koch, and Wolfgang Straßer. 2010. Fast and Scalable CPU/GPU Collision Detection for Rigid and Deformable Surfaces. *Computer Graphics Forum*

- 29 (July 2010), 1605–1612.
- Jia Pan, Liangjun Zhang, and Dinesh Manocha. 2012. Collision-Free and Smooth Trajectory Computation in Cluttered Environments. *The International Journal of Robotics Research* 31, 10 (2012), 1155–1175.
- Xavier Provot. 1997. Collision and Self-Collision Handling in Cloth Model Dedicated to Design Garments. In *Computer Animation and Simulation*. Springer, 177–189.
- Stephane Redon, Abderrahmane Kheddar, and Sabine Coquillart. 2002. Fast Continuous Collision Detection between Rigid Bodies. *Computer Graphics Forum* 21 (May 2002).
- Boris Schling. 2011. *The Boost C++ Libraries*. XML Press.
- Sara Schwartzman, Álvaro Pérez, and Miguel Otaduy. 2010. Star-Contours for Efficient Hierarchical Self-Collision Detection. *ACM Transactions on Graphics* 29 (July 2010).
- John M. Snyder. 1992. Interval Analysis for Computer Graphics. In *Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '92)*. Association for Computing Machinery, New York, NY, USA, 121–130.
- John M. Snyder, Adam R. Woodbury, Kurt Fleischer, Bena Currin, and Alan H. Barr. 1993. Interval Methods for Multi-Point Collisions between Time-Dependent Curved Surfaces. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '93)*. Association for Computing Machinery, New York, NY, USA, 321–334.
- Jos Stam. 2009. Nucleus: Towards a Unified Dynamics Solver for Computer Graphics. *IEEE International Conference on Computer-Aided Design and Computer Graphics* (2009), 1–11.
- Min Tang, Sean Curtis, Sung-eui Yoon, and Dinesh Manocha. 2009a. ICCD: Interactive Continuous Collision Detection between Deformable Models Using Connectivity-Based Culling. *IEEE Transactions on Visualization and Computer Graphics* 15 (July 2009), 544–57.
- Min Tang, Young Kim, and Dinesh Manocha. 2009b. C²A: Controlled Conservative Advancement for Continuous Collision Detection of Polygonal Models. *Proceedings of International Conference on Robotics and Automation* (June 2009), 849–854.
- Min Tang, Dinesh Manocha, and Ruofeng Tong. 2010. Fast Continuous Collision Detection using Deforming Non-Penetration Filters. *Proceedings of the Symposium on Interactive 3D Graphics and Games*, 7–13.
- Min Tang, Dinesh Manocha, Sung-eui Yoon, Peng du, Jae-Pil Heo, and Ruofeng Tong. 2011. VolCCD: Fast Continuous Collision Culling Between Deforming Volume Meshes. *ACM Transactions on Graphics* 30 (Jan. 2011).
- Min Tang, Ruofeng Tong, Zhendong Wang, and Dinesh Manocha. 2014. Fast and Exact Continuous Collision Detection with Bernstein Sign Classification. *ACM Transactions on Graphics* 33 (Nov. 2014), 186:1–186:8. Issue 6.
- Min Tang, Sung-eui Yoon, and Dinesh Manocha. 2008. Adjacency-Based Culling for Continuous Collision Detection. *The Visual Computer* 24 (July 2008), 545–553.
- Mickael Verschoor and Andrei C. Jalba. 2019. Efficient and Accurate Collision Response for Elastically Deformable Models. *ACM Transactions on Graphics* 38, 2, Article 17 (March 2019), 20 pages.
- Pascal Volino and Nadia Magnenat Thalmann. 1994. Efficient Self-Collision Detection on Smoothly Discretized Surface Animations using Geometrical Shape Regularity. *Computer Graphics Forum* 13, 3 (1994), 155–166.
- Brian Von Herzen, Alan H. Barr, and Harold R. Zatz. 1990. Geometric Collisions for Time-Dependent Parametric Surfaces. *Computer Graphics (Proceedings of SIGGRAPH)* 24, 4 (Sept. 1990), 39–48.
- Etienne Vouga, David Harmon, Rasmus Tamstorf, and Eitan Grinspun. 2010. Asynchronous Variational Contact Mechanics. *Computer Methods in Applied Mechanics and Engineering* 200 (July 2010), 2181–2194. <https://github.com/evouga/collisiondetection>
- Huamin Wang. 2014. Defending Continuous Collision Detection against Errors. *ACM Transactions on Graphics* 33 (July 2014), 1–10.
- Zhendong Wang, Min Tang, Ruofeng Tong, and Dinesh Manocha. 2015. TightCCD: Efficient and Robust Continuous Collision Detection using Tight Error Bounds. *Computer Graphics Forum* 34 (Sept. 2015), 289–298. Issue 7.
- Wolfram Research Inc. 2020. *Mathematica 12.0*. <http://www.wolfram.com>
- Wingo Sai-Keung Wong and George Baci. 2006. A Randomized Marking Scheme for Continuous Collision Detection in Simulation of Deformable Surfaces. In *Proceedings of the 2006 ACM International Conference on Virtual Reality Continuum and Its Applications (VRCA '06)*. Association for Computing Machinery, New York, NY, USA, 181–188.
- Peter Wriggers. 1995. Finite Element Algorithms for Contact Problems. *Archives of Computational Methods in Engineering* 2 (Dec. 1995), 1–49.
- Xinyu Zhang, Stephane Redon, Minkyung Lee, and Young J. Kim. 2007. Continuous Collision Detection for Articulated Models Using Taylor Models and Temporal Culling. *ACM Transactions on Graphics* 26, 3 (July 2007), 15–25.
- Changxi Zheng and Doug James. 2012. Energy-based Self-Collision Culling for Arbitrary Mesh Deformations. *ACM Transactions on Graphics* 31 (July 2012).




Fig. 14. Prism resulting from the input points and triangle in (9). The origin is marked by the red dot.

A DATASET FORMAT

To avoid any loss of precision we convert every input floating-point coordinate in rationals using GMP [Granlund and the GMP Development Team 2012]. This conversion is exact since every floating point can be converted in a rational number, as long as the numerator and denominator are arbitrarily large integers. We then store the numerator and denominator as a string since the numerator and denominator can be larger than a long number. To retrieve the floating point number we allocate a GMP rational number with the two strings and convert it to double.

In summary, one CCD query is represented by a 8×7 matrix where every row is one of the 8 CCD input points, and the columns are the interleaved x, y, z coordinates of the point, represented as numerator and denominator. For convenience, we appended several such matrices in a common CSV file. The last column represents the result of the ground truth. For instance a CC query between $p_1^0, p_2^0, p_3^0, p_4^0$ and $p_1^1, p_2^1, p_3^1, p_4^1$ is represented as

$$\begin{array}{cccccc}
 p_{1n}^0 & p_{1d}^0 & p_{1n}^0 & p_{1d}^0 & p_{1n}^0 & p_{1d}^0 & T \\
 p_{2n}^0 & p_{2d}^0 & p_{2n}^0 & p_{2d}^0 & p_{2n}^0 & p_{2d}^0 & T \\
 p_{3n}^0 & p_{3d}^0 & p_{3n}^0 & p_{3d}^0 & p_{3n}^0 & p_{3d}^0 & T \\
 p_{4n}^0 & p_{4d}^0 & p_{4n}^0 & p_{4d}^0 & p_{4n}^0 & p_{4d}^0 & T \\
 p_{1n}^1 & p_{1d}^1 & p_{1n}^1 & p_{1d}^1 & p_{1n}^1 & p_{1d}^1 & T \\
 p_{2n}^1 & p_{2d}^1 & p_{2n}^1 & p_{2d}^1 & p_{2n}^1 & p_{2d}^1 & T \\
 p_{3n}^1 & p_{3d}^1 & p_{3n}^1 & p_{3d}^1 & p_{3n}^1 & p_{3d}^1 & T \\
 p_{4n}^1 & p_{4d}^1 & p_{4n}^1 & p_{4d}^1 & p_{4n}^1 & p_{4d}^1 & T,
 \end{array}$$

where p_{in}^t and p_{id}^t are respectively the numerator and denominator of the x -coordinate of p , and T is the same ground truth. The dataset and a query viewer can be downloaded from the NYU Faculty Digital Archive.

B EXAMPLE OF DEGENERATE CASE NOT PROPERLY HANDLED BY [Brochu et al. 2012]

Let

$$\begin{array}{l}
 p^0 = [0.1, 0.1, 0.1], \quad v_1^0 = [0, 0, 1], \quad v_2^0 = [1, 0, 1], \quad v_3^0 = [0, 1, 1], \\
 p^1 = [0.1, 0.1, 0.1], \quad v_1^1 = [0, 0, 0], \quad v_2^1 = [0, 1, 0], \quad v_3^1 = [1, 0, 0]
 \end{array} \quad (9)$$

be the input point and triangle. Checking if the point intersects the triangle is equivalent to check if the prism shown in Figure 14 contains the origin. However, the prism contains a bilinear face that is degenerate (it looks like a “hourglass”). The algorithm proposed in [Brochu et al. 2012] does not consider this degenerate case and erroneously reports no collision.

Table 3. Summary of the average runtime in μs (t), number of false positive (FP), and number of false negative (FN) for FPRF and our method.

d	Handcrafted – Vertex-Face MSCCD						Handcrafted – Edge-Edge MSCCD						Simulation – Vertex-Face MSCCD						Simulation – Edge-Edge MSCCD					
	FPRF			Ours			FPRF			Ours			FPRF			Ours			FPRF			Ours		
d	t	FP	FN	t	FP	FN	t	FP	FN	t	FP	FN	t	FP	FN	t	FP	FN	t	FP	FN	t	FP	FN
10^{-2}	2.41	1.8K	4	18.86K	2.6K	0	1.16	3.3K	19	9.64K	4.8K	0	8.04	869.1K	1	12.04	8.1M	0	8.01	1.1M	0	19.12	8.3M	0
10^{-8}	4.53	83	3	1.60K	159	0	0.60	160	28	3.42K	309	0	8.00	4	2	0.72	8	0	0.77	16	0	0.73	40	0
10^{-16}	2.23	29	69	1.51K	108	0	0.55	45	145	2.92K	214	0	7.78	0	5.2K	0.71	2	0	0.25	0	2.3K	0.72	17	0
10^{-30}	2.24	9	70	1.39K	108	0	0.58	5	147	2.79K	214	0	7.77	0	5.2K	0.66	2	0	0.25	0	2.3K	0.67	17	0
10^{-100}	2.31	9	70	1.43K	108	0	0.80	5	147	2.82K	214	0	7.75	0	5.2K	0.66	2	0	0.25	0	2.3K	0.68	17	0

C EXAMPLE OF INFLECTION POINT NOT PROPERLY HANDLED BY [Tang et al. 2014]

Let

$$\begin{aligned} p^0 &= [1, 1, 0], & v_1^0 &= [0, 0, 5], & v_2^0 &= [2, 0, 2], & v_3^0 &= [0, 1, 0], \\ p^1 &= [1, 1, 0], & v_1^1 &= [0, 0, -1], & v_2^1 &= [0, 0, -2], & v_3^1 &= [0, 7, 0] \end{aligned}$$

be the input point and triangle. Checking if they intersect at time t is equivalent to finding the roots of

$$-72t^3 + 120t^2 - 44t + 3.$$

To apply the method in [Tang et al. 2014] we need to rewrite the polynomial in form of Tang et al. [2014, Equation (1)]:

$$1B_0^3(t) - \frac{35}{3}B_1^3(t) + \frac{82}{3}B_2^3(t) + 14B_3^3(t).$$

Their algorithm assumes no inflection points in the Bezier curve. Thus it proposes to split the curve at the eventual inflection point (as in the case above). The formula proposed in [Tang et al. 2014, Section 4.1] contains a typo, by fixing it we obtain the inflection point at:

$$t = \frac{6k_0 - 4k_1 + k_2}{6k_0 - 6k_1 + 3k_2 - k_3} = \frac{5}{9}.$$

By using the incorrect formula we obtain $t = 155/312$, which is not an inflection point. In both cases, t cannot be computed exactly since it contains a division, and computing it approximately breaks the assumption of not having inflection points in the Bezier form. In the reference code, the authors detect the presence of an inflection point using predicates, but do not split the curve (the case is not handled). We modified the code (patch attached in the additional material) to conservatively return a collision in these cases.

Independently from this problem, their reference implementation returns false negative (i.e. misses collisions) for certain configurations, such as the following degenerate configuration:

$$\begin{aligned} p^0 &= [1, 0.5, 1], & v_1^0 &= [0, 0.57, 1], & v_2^0 &= [1, 0.57, 1], & v_3^0 &= [1, 1.57, 1], \\ p^1 &= [1, 0.5, 1], & v_1^1 &= [0, 0.28, 1], & v_2^1 &= [1, 0.28, 1], & v_3^1 &= [1, 1.28, 1]. \end{aligned}$$

We could not find out why this is happening, and we do not know if this is a theoretical or numerical problem, or a bug in the implementation.

D EFFECT OF δ ON THE INTERVAL-BASED METHODS

UIRF, IRF, and our method have a single parameter δ to control the size of the interval. Increasing δ will introduce more false positive, while making the algorithms faster (Figure 15). Note that we limit the total running time to 24h, thus UIRF does not have result for




Fig. 15. Log plot of the effect of the tolerance δ on the running time (top) and false positives (bottom) for the three (Ours, UIRF, and IRF) interval based methods on the simulation dataset.

$\delta > 10^{-6}$ (for $\delta = 10^{-6}$ it takes 1ms per query in average). δ has a similar effect on the number of false positives for the three interval based methods, while it has a more significant impact on the running time for UIRF and IRF.

E MINIMUM SEPARATION WITH FPRF

In Table 3 we compare our method with FPRF by changing the parameter η that mimics minimum separation.