

Détection continue de collisions entre triangles et fonctions
distance signée en temps réel

par

Joël PELLETIER-GUÉNETTE

MÉMOIRE PAR ARTICLES PRÉSENTÉ À L'ÉCOLE DE TECHNOLOGIE
SUPÉRIEURE
COMME EXIGENCE PARTIELLE À L'OBTENTION DE LA MAÎTRISE
AVEC MÉMOIRE EN GÉNIE DES TECHNOLOGIES DE L'INFORMATION
M. Sc. A.

MONTRÉAL, LE 14 DÉCEMBRE 2025

ÉCOLE DE TECHNOLOGIE SUPÉRIEURE
UNIVERSITÉ DU QUÉBEC



Joël Pelletier-Guénette, 2025



Cette licence Creative Commons signifie qu'il est permis de diffuser, d'imprimer ou de sauvegarder sur un autre support une partie ou la totalité de cette oeuvre à condition de mentionner l'auteur, que ces utilisations soient faites à des fins non commerciales et que le contenu de l'oeuvre n'ait pas été modifié.

PRÉSENTATION DU JURY

CE MÉMOIRE A ÉTÉ ÉVALUÉ

PAR UN JURY COMPOSÉ DE:

M. Sheldon Andrews, directeur de mémoire
Département de génie logiciel et TI à l'École de technologie supérieure

M. Matthew Harker, président du jury
Département de génie des systèmes à l'École de technologie supérieure

M. Eric Paquette, membre du jury
Département de génie logiciel et TI à l'École de technologie supérieure

IL A FAIT L'OBJET D'UNE SOUTENANCE DEVANT JURY ET PUBLIC

LE 28 NOVEMBRE 2025

À L'ÉCOLE DE TECHNOLOGIE SUPÉRIEURE

REMERCIEMENTS

Je souhaite remercier mon directeur de recherche, professeur Sheldon Andrews, pour m'avoir introduit au monde de la recherche en infographie, ainsi que pour son appui, ses conseils et son expertise qui m'ont été d'une valeur inestimable tout au long de ce projet.

Je remercie également le professeur Alexandre Mercier-Aubin de l'université de Sherbrooke, qui nous a rejoints en tant qu'étudiant postdoctoral pour contribuer à la supervision du projet et à la rédaction de notre article.

Je souhaite aussi exprimer ma gratitude envers les membres du jury pour le temps passé à évaluer ce mémoire, ainsi que pour leurs commentaires.

Un grand merci aux membres du laboratoire multimédia de l'ETS, à Thomas Friedlaender de Symgery, à Paul Kry de l'université McGill et à Kenny Erleben de l'université de Copenhague pour de nombreux échanges et pour leurs commentaires précieux.

Je remercie aussi Nicholas Sharp pour Polyscope Sharp *et al.* (2019), Inigo Quilez pour ses ressources sur les SDF, et le laboratoire d'infographie de l'université Stanford pour les modèles 3D du lapin, de l'Armadillo, et du dragon XYZRGB, qui ont été d'une grande aide dans le développement d'outils et de simulations pour nos expériences.

Je remercie le Conseil de recherches en sciences naturelles et en génie du Canada (CRSNG) pour le soutien financier (Bourses d'études supérieures en recherche du Canada au niveau de la maîtrise 585438-2023).

Enfin, un énorme merci à mes parents, Gilles et Micheline, à mes petits frères Hugo et Rémi, et à mes amis pour leur soutien tout au long de mes études. Tout particulièrement, merci à Geoffrey Glaum pour son support sans faille au cours des dernières années, pour les centaines de fois qu'il m'a rappelé qu'il est important de manger, et pour les millions de tasses de café extra-fort qu'il a contribué à ce projet. :)

Détection continue de collisions entre triangles et fonctions distance signée en temps réel

Joël PELLETIER-GUÉNETTE

RÉSUMÉ

De l'entraînement virtuel de pilotes et de chirurgiens à l'optimisation de trajectoire, en passant par l'esquive d'obstacles en temps réel pour l'aérospatial ou la robotique, plusieurs systèmes modernes dépendent de simulations physiques en temps réel. La détection précise et robuste des collisions est au cœur des simulations physiques utiles à ces domaines. Elle est toutefois extrêmement difficile à effectuer dans les applications en temps réel. La bonne représentation des objets simulés a un impact direct sur le réalisme et sur l'efficacité avec laquelle nous pouvons simuler divers phénomènes physiques. Ce mémoire introduit des solutions efficaces pour la détection de collision de haute qualité entre des maillages de triangles et des fonctions distance signée (SDF). Le problème de détection continue des collisions (CCD) est abordé comme un problème d'optimisation spatio-temporel cherchant le moment exact d'impact (TOI) entre un triangle et l'isosurface d'une SDF. Cette méthode est plus robuste que les précédentes, qui se basent sur l'échantillonnage par points, et plus performante que les méthodes de points pour la détection discrète des collisions (DCD) triangle-SDF. Une méthode de raffinement adaptatif est aussi proposée pour fournir efficacement de l'information supplémentaire lors de collisions avec de grands triangles. Cela permet l'utilisation de géométrie efficace pour réduire le coût de la détection de collision. Une comparaison avec des algorithmes de pointe expose les capacités et les avantages de l'approche proposée par la simulation de milliers de scénarios difficiles avec collisions. Les résultats expérimentaux démontrent que ce travail de recherche améliore la qualité des simulations physiques tout en maintenant la performance nécessaire au temps réel. Dans l'ensemble, ce travail fournit la première méthode fonctionnelle pour la CCD triangle-SDF robuste en temps réel, et ouvre la porte à un grand nombre de nouvelles applications potentielles pour la simulation physique, particulièrement quand de petits objets aux mouvements rapides interagissent.

Mots-clés: Infographie, Simulation physique, Détection de collisions, Fonction distance signée, Analyse numérique

Real-time continuous collision detection between triangles and signed distance functions

Joël PELLETIER-GUÉNETTE

ABSTRACT

From virtual training for surgeons and pilots to live obstacle avoidance and trajectory optimization in aerospace and robotics, many modern systems rely on real-time physics simulations. Accurate and robust collision detection is key for useful simulations. Yet, it is excessively difficult to achieve in real-time applications. The proper modelling of objects directly affects the realism of physics-based simulations and how efficiently we can simulate physical phenomena. This thesis introduces efficient solutions for high-quality collision detection between triangle meshes and signed distance functions (SDFs). The continuous collision detection (CCD) problem is tackled with a novel spatio-temporal optimization to seek the exact time of impact (TOI) between a triangle and an SDF isosurface. This method offers improved robustness over existing point sampling methods, and outperforms recent triangle-SDF discrete collision detection (DCD) algorithms. Also proposed, is a novel adaptive refinement method to efficiently provide supplemental contact information over large triangles. Efficient coarse geometry can thus be used to reduce the cost of collision detection involving triangle meshes. A comparison against state-of-the-art algorithms is provided, and exposes the capabilities and advantages of the proposed approach through simulating thousands of difficult collision scenarios. Experimental results demonstrate that this work improves the quality of physics simulations while maintaining real-time performance. Overall, this work provides the first working method for robust triangle-SDF CCD, and opens the door to a wide array of new potential applications for physics-based simulations where small and fast moving objects interact.

Keywords: Computer graphics, Physics simulation, Collision detection, Signed distance function, Numerical analysis

TABLE DES MATIÈRES

	Page
INTRODUCTION	1
0.1 Contexte et problématique	1
0.2 Objectif et méthodologie	4
0.3 Contributions	4
0.4 Organisation du document	5
CHAPITRE 1 REAL-TIME TRIANGLE-SDF CONTINUOUS COLLISION DETECTION	7
1.1 Abstract	7
1.2 Introduction	8
1.3 Related Work	10
1.3.1 Collision detection	10
1.3.2 Geometry representation	12
1.4 Preliminaries	13
1.4.1 Signed distance field (SDF)	13
1.4.2 Triangle-SDF discrete collision detection	14
1.4.3 Continuous collision detection	15
1.4.4 Golden section search	15
1.5 Triangle-SDF CCD	17
1.5.1 Local optimization	18
1.5.2 Frank-Wolfe for CCD	19
1.5.3 Frank-Wolfe with line search	21
1.5.4 Starting iterate	24
1.5.5 Adaptive triangle subdivision	24
1.5.6 Acceleration structure	25
1.6 Results	27
1.6.1 Contact detection	28
1.6.2 Ground truth comparisons	31
1.6.3 Performance	34
1.7 Discussion and limitations	36
1.8 Conclusion	38
1.9 Acknowledgements	39
CONCLUSION ET RECOMMANDATIONS	41
2.1 Futurs travaux et recommandations	41
BIBLIOGRAPHIE	45

LISTE DES TABLEAUX

	Page
Tableau 1.1	Time of impact comparison of our method vs ground truth 33
Tableau 1.2	Travel distance comparison of our method vs ground truth 33
Tableau 1.3	Performance scalabilty with regard to triangle count 34
Tableau 1.4	Performance timings 35

LISTE DES FIGURES

		Page
Figure 0.1	Collision détectée par la DCD	1
Figure 0.2	Collision manquée par la DCD	2
Figure 0.3	Collision détectée par la CCD	2
Figure 1.1	DCD and CCD comparison of a shuriken impacting an armor breastplate at 500 m/s	7
Figure 1.2	Spatio-temporal local minima of an SDF for a sphere passing rightward through its boundary	17
Figure 1.3	Visual representation of the adaptive barycentric triangle subdivision process and search for supplemental contact points	26
Figure 1.4	Visual representation of a bounding sphere and the padding required to captures non-linear motions in linear sphere tracing	27
Figure 1.5	Visual results for a skateboard mesh sliding down a SDF ramp	28
Figure 1.6	Visual results for a coarse cone mesh falling onto a SDF box	29
Figure 1.7	Visual results for a bunny mesh falling onto a detailed SDF object	29
Figure 1.8	Visual results for a hoop mesh falling over a SDF sphere of matching radius	30
Figure 1.9	Visual results for a deformable cloth mesh falling onto a SDF bunny	30
Figure 1.10	Visual results for a thin box mesh falling onto a spiky SDF	31
Figure 1.11	A triangle cannot lay flat on a SDF box with a single contact point, but succeeds with adaptive subdivision	32
Figure 1.12	Our CCD method captures collision from a triangle mesh pyramid rotating 360° in a single time step	32

LISTE DES ALGORITHMES

	Page
Algorithme 1.1 Golden Section Search (GSS)	16
Algorithme 1.2 Frank-Wolfe with GSS (FWGSS)	22

LISTE DES ABRÉVIATIONS, SIGLES ET ACRONYMES

BVH	Hiérarchie de volumes englobants (<i>Bounding Volume Hierarchy</i>)
CCD	Détection continue des collisions (<i>Continuous Collision Detection</i>)
CRSNG	Conseil de Recherches en Sciences Naturelles et en Génie du Canada
DCD	Détection discrète des collisions (<i>Discrete Collision Detection</i>)
ETS	École de Technologie Supérieure
FW	Frank-Wolfe
FWGSS	FW avec méthode du nombre d'or (<i>FW with Golden Section Search</i>)
GSS	Méthode du nombre d'or (<i>Golden Section Search</i>)
GT	Base vérité (<i>Ground Truth</i>)
IPC	<i>Incremental Potential Contact</i>
NSERC	CRSNG (<i>Natural Sciences and Engineering Research Council of Canada</i>)
PGD	Descente de gradient avec projection (<i>Projected Gradient Descent</i>)
SDF	Fonction distance signée (<i>Signed Distance Function</i>)
TOI	Temps de l'impact (<i>Time of Impact</i>)

LISTE DES SYMBOLES ET UNITÉS DE MESURE

UNITÉS DE BASE

Longueur

m mètre

Temps

s seconde

ms milliseconde

μ s microseconde

SYMBOLES MATHÉMATIQUES

Opérateurs

$\vec{\square}$ Vecteur / Position

$\|\square\|$ Norme euclidienne

$|\square|$ Valeur absolue

Fonction distance signée (SDF)

ϕ SDF

$|\phi|$ Fonction distance absolue

$\phi(\vec{x})$ SDF évaluée à la position \vec{x}

$\nabla\phi$ Gradient de la SDF

$\nabla\phi(\vec{x})$ Gradient de la SDF évalué à la position \vec{x}

Statistique

σ Écart type

Optimisation

φ Nombre d'or

α Coefficient d'avancement de l'algorithme de Frank-Wolfe

Optimisation (suite)

i Itération

$i + 1$ Prochaine itération

t_i Instant dans le temps évalué à l'itération i

t_{i+1} Instant dans le temps à évaluer à la prochaine itération $i + 1$

\vec{x}_i Position du point initial évalué à l'itération i

\vec{x}_{i+1} Position du point évalué au temps de la prochaine itération $i + 1$

\vec{s}_i Sommet de support du triangle à l'itération i de l'algorithme de Frank-Wolfe

SYMBOLES MÉCANIQUES

Temps

t Un instant dans le temps

Δt Différence entre deux instants

Positions

\vec{g} Centre de masse / Centre de rotation en chute libre

$\Delta\vec{x}$ Différence de position

Vélocités

\vec{v} Vélocité linéaire

\vec{v}_g \vec{v} du centre de masse

Vélocités (suite)

$\vec{\omega}$ Vitesse angulaire

$\vec{\omega}_g$ $\vec{\omega}$ du centre de masse

Détection de collisions

\vec{x}_c Point de contact

t_{hit} Temps de l'impact (TOI)

t_{start} Début d'un intervalle de temps

t_{end} Fin d'un intervalle de temps

INTRODUCTION

Le Québec est au cœur de l'industrie dans les domaines multimédia et en infographie. Des entreprises de partout font appel à l'expertise locale pour produire des animations et des effets visuels. Grâce à de nouveaux algorithmes et à d'ingénieuses abstractions mathématiques, il nous est possible de simuler des phénomènes physiques d'une complexité telle qu'il aurait été impossible d'imaginer il y a quelques années. L'intérêt envers ce genre de simulations n'est pas limité qu'aux films et aux jeux-vidéo. Par exemple, les simulations en temps réel ont un grand nombre d'applications en ingénierie, en robotique, en aérospatiale et en santé. Elles peuvent assister dans des procédures délicates, prévenir des manœuvres dangereuses ou déclencher des mesures de sécurité en cas de catastrophe imminente. Elles permettent aussi l'entraînement sécuritaire et à faible coût de pilotes, de chirurgiens et d'opérateurs machinistes. Leur potentiel n'est limité que par nos capacités à simuler divers phénomènes physiques de façon rapide et réaliste.

0.1 Contexte et problématique

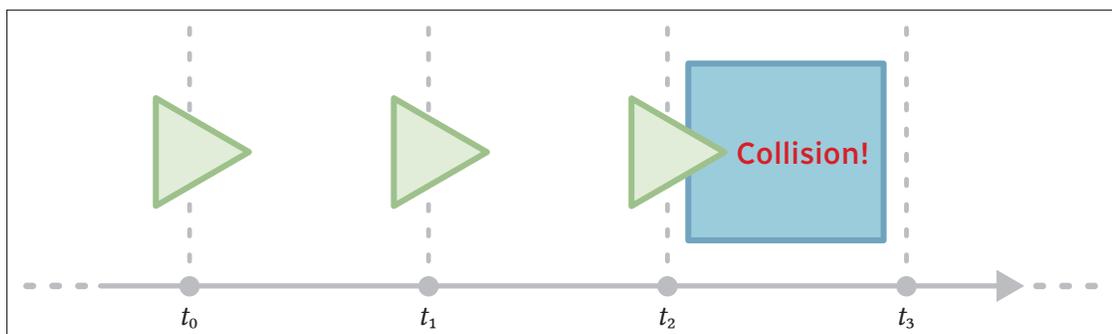


Figure 0.1 Un triangle bouge dans le temps. La DCD est effectuée à des moments précis. Une collision est détectée à l'instant t_2 , mais il y a interpénétration. Les objets requièrent donc séparation

Même avec du matériel et des algorithmes de pointe, la simulation en temps réel nécessite toujours des compromis en termes de réalisme et de portée. Par exemple, la détection de collisions

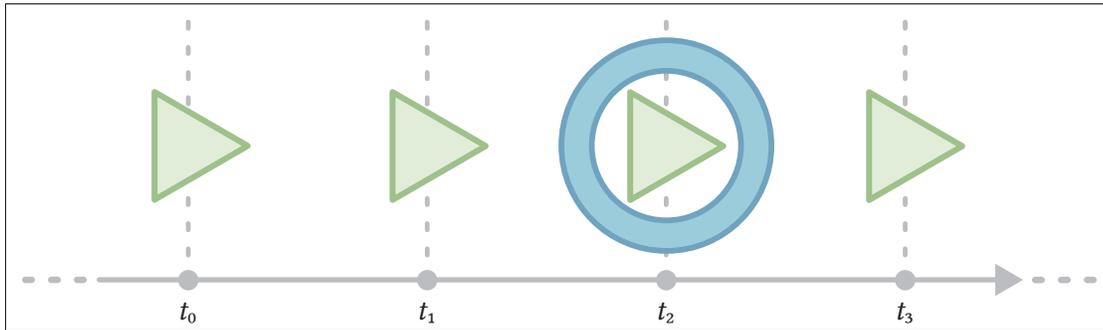


Figure 0.2 Un triangle bouge dans le temps. La DCD est effectuée à des moments précis. La collision entre le triangle et l'obstacle est manquée par la DCD

joue un rôle clé en déterminant si, où et quand les objets entrent en collision. La détection discrète de collisions (DCD) s'exécute à des intervalles de temps fixes et sépare les objets qui s'intersectent, tel que présenté à la Figure 0.1. Elle est rapide, mais elle ne peut pas détecter les collisions qui ne se produisent qu'entre ses exécutions. La Figure 0.2 présente un exemple d'une telle situation. Ceci est particulièrement problématique pour les simulations avec de petits objets ou de grandes vitesses relatives. La détection continue des collisions (CCD) calcule le moment exact des collisions pour prévenir les intersections (Nie, Zhao, Xu & Li, 2020). La Figure 0.3 présente visuellement le principe de CCD. Elle est essentielle aux méthodes robustes

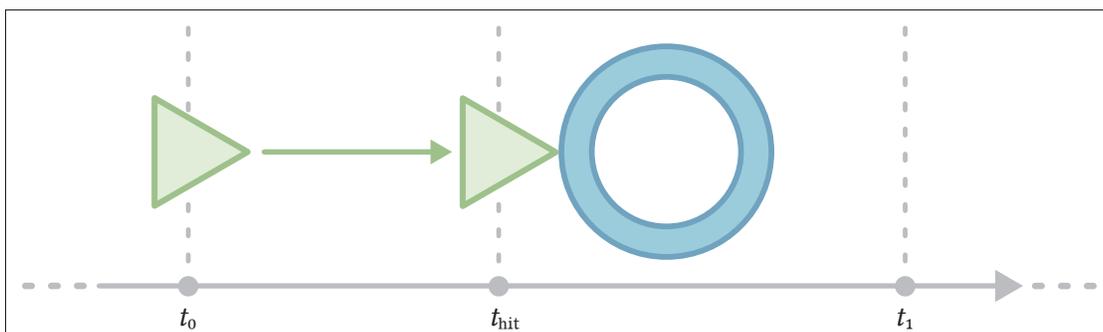


Figure 0.3 Un triangle bouge dans le temps. La CCD détecte qu'une collision a lieu à l'instant t_{hit} . Le triangle est libre d'avancer jusqu'au moment du contact

avec garantie de non-interpénétration comme « incremental potential contact » (IPC) (Li *et al.*, 2020). Les collisions tendent à s'enchaîner rapidement lorsque les objets entrent en contact et

s'immobilisent contre une surface. Quelques millisecondes de simulation peuvent donc entraîner des millions de requêtes à la CCD pour le prochain temps d'impact (TOI). La simulation en temps réel nécessite donc de limiter les requêtes, et donc d'ignorer des collisions, ou d'autres restrictions sévères sur la portée et le réalisme de la simulation (Ericson, 2004).

La représentation des objets simulés affecte directement l'efficacité de la détection de collisions. IPC utilise les maillages de triangles (Li *et al.*, 2020), qui peuvent modéliser presque toutes les formes, mais qui sont coûteux en termes de calcul (Erleben, 2018). Une fonction distance signée (SDF) fournit efficacement l'information nécessaire à la détection de collisions, mais ne s'adapte pas bien aux formes saillantes ou sans intérieur et extérieur distincts (Fuhrmann, Sobotka & Groß, 2003). L'idée d'utiliser ces deux types de représentations dans une même simulation est donc intéressante, car cela permet de faire bon usage des avantages de chacune lorsque approprié.

Traditionnellement, la détection de collisions entre SDFs et maillages de triangles se limitait à la détection des contacts au niveau des sommets du maillage, ou nécessitait un échantillonnage dense de la surface des triangles (Xu & Barbič, 2017; McNeely, Puterbaugh & Troy, 2005). Des méthodes récentes par Macklin *et al.* (2020) permettent de résoudre spécifiquement le problème de DCD entre une SDF et la surface des triangles d'un maillage. Ces méthodes améliorent la qualité des contacts générés, à comparer des méthodes à base de points, mais elles ne supportent qu'un seul point de contact par triangle. Puisque les triangles à grande surface peuvent facilement être traversés par plusieurs éléments d'une SDF en même temps, des maillages denses avec une haute résolution sont souvent nécessaires pour la détection fiable des collisions. Les maillages de triangles plus grossiers à basse résolution sont toutefois préférés par bien des applications pour leur efficacité. Enfin, ces méthodes reposent sur la DCD. Elles effectuent donc des tests d'intersection à des moments précis et n'empêchent pas les formes de passer au travers les

unes des autres, comme la CCD le permettrait. Néanmoins, la CCD triangle-SDF demeure un problème inexploré.

0.2 Objectif et méthodologie

L'objectif principal du travail présenté dans ce mémoire est de générer des algorithmes efficaces et robustes pour la CCD triangle-SDF, afin d'améliorer la performance et la qualité des simulations physiques en temps réel. L'efficacité des algorithmes est mesurée par le temps d'exécution de la détection de collisions. La performance des simulations est mesurée par le temps nécessaire à la simulation physique complète d'intervalles de temps fixes. La qualité est mesurée en comparant le nombre de collisions détectées et l'état des objets simulés par différents algorithmes face aux résultats obtenus par force-brute. La robustesse fait référence au maintien des capacités des algorithmes dans une multitude de scénarios. Elle est évaluée en comparant les autres métriques au travers d'un ensemble de scénarios créés à la main (pour couvrir les cas difficiles) et de scénarios aléatoires générés par ordinateur (pour des données statistiques supplémentaires). Le tout est implémenté en C++ dans un simulateur sur mesure et comparé aux méthodes de pointe existantes.

0.3 Contributions

Cette recherche contribue plusieurs algorithmes qui se complètent. D'abord, une méthode d'optimisation locale spatio-temporelle est proposée pour approcher le problème de la CCD. La résolution du problème d'optimisation fournit le moment du prochain impact et les points les plus près entre le maillage de triangle et l'isosurface de la SDF. Nous introduisons une méthode de résolution itérative innovante, basée sur le gradient de la SDF, qui résout pour le minimum global du problème plutôt que pour un minimum local. Ensuite, une nouvelle méthode de résolution grossière et efficace de la CCD réduit la quantité d'évaluations coûteuses de la CCD triangle-SDF. Cette méthode capture la rotation des triangles dans une hiérarchie de

sphères englobantes et utilise une CCD linéaire efficace pour trouver une borne inférieure au domaine temporel du problème de la CCD triangle-SDF, réduisant ainsi sa complexité. Enfin, nous introduisons une méthode de raffinement barycentrique efficace pour obtenir des points de contacts simultanés supplémentaires sur la surface de grands triangles, ce qui permet l'utilisation de géométrie réduite efficace.

0.4 Organisation du document

Ce mémoire se concentre sur la présentation de notre article (Pelletier-Guénette, Mercier-Aubin & Andrews, 2025), qui couvre le travail de recherche effectué dans son ensemble. Le professeur Mercier-Aubin a participé aux expérimentations en créant des scénarios de tests supplémentaires et en aidant à la capture des données expérimentales. Le professeur Andrews, superviseur officiel du projet, a participé au développement du cadre théorique. Les deux ont joué un rôle de mentors et ont aidé à l'écriture de l'article. Notre article, publié dans PACMCGIT, a obtenu une mention honorable pour meilleur article au « Symposium on computer animation 2025 » de l'ACM. Il couvre une majorité des éléments attendus dans un mémoire de maîtrise par article. Notamment, une revue critique de la littérature existante et une discussion détaillée des résultats obtenus sont déjà inclus dans l'article. Puisqu'il n'y a qu'un seul article dans ce mémoire, et qu'il couvre déjà en détail l'ensemble du travail de recherche, nous n'incluons pas de chapitres redondants face au contenu de l'article.

Ce mémoire est organisé en deux chapitres distincts. Le premier chapitre présente notre article. Celui-ci débute par une revue de littérature détaillée qui couvre les travaux existants (Section 1.3) et les concepts préliminaires nécessaires à la compréhension du travail de recherche (Section 1.4). Les méthodes développées pour la CCD triangle-SDF sont ensuite présentées en détail à la Section 1.5. Les expériences et les résultats sont présentés à la Section 1.6, suivis par une discussion de ces derniers à la Section 1.7. L'article expose enfin nos conclusions sur le sujet ainsi que quelques pistes pour de futurs travaux à la Section 1.8. Le deuxième chapitre résume

les principales conclusions du travail ainsi que les recommandations mentionnées dans l'article. Bien que déjà couvert par l'article, ce dernier chapitre est inclus afin de présenter l'information dans le contexte du mémoire et de cette introduction.

CHAPITRE 1

REAL-TIME TRIANGLE-SDF CONTINUOUS COLLISION DETECTION

Joël Pelletier-Guénette¹, Alexandre Mercier-Aubin^{1,2}, Sheldon Andrews¹

¹ Département de génie logiciel et TI, École de Technologie Supérieure,
1100 Notre-Dame Ouest, Montréal, Québec, Canada H3C 1K3

² Département d'Informatique, Université de Sherbrooke,
2500, boulevard de l'Université, Sherbrooke, Québec, Canada J1K 2R1

Article publié dans « Proceedings of the ACM on Computer Graphics and Interactive Techniques (PACMCGIT) » en août 2025 (Pelletier-Guénette *et al.*, 2025).

© 2025 Copyright held by the authors. Publication rights licensed to ACM.

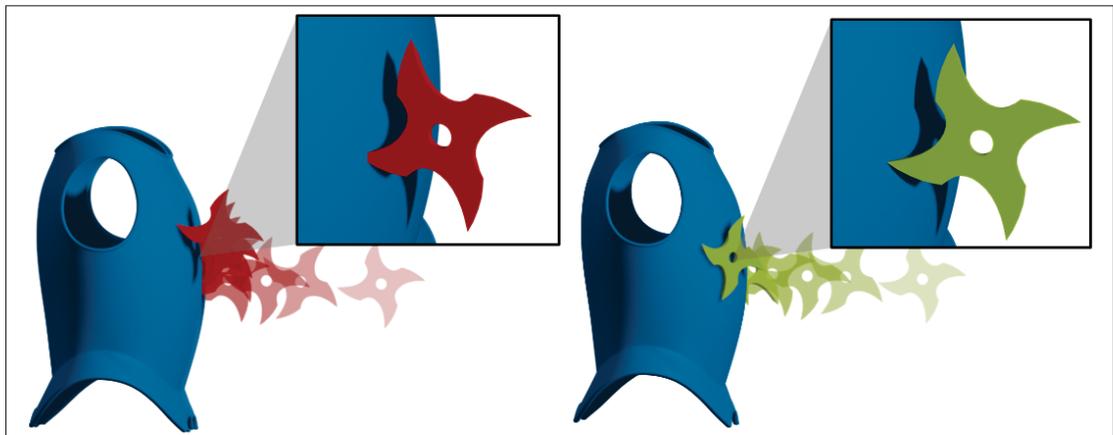


Figure 1.1 A fast-flying shuriken moving at 500 m/s hits an armor breastplate. The shuriken is represented by a triangle mesh, and the armor by an implicit function. Left: Discrete collision detection allows penetration into the armor, resulting in a different trajectory after collision response. Right: Our method gives an interpenetration-free simulation

1.1 Abstract

We introduce an efficient solution to the problem of continuous collision detection (CCD) between triangle geometry and signed distance fields (SDFs). We formulate the triangle-SDF collision problem as a novel spatio-temporal local optimization that solves for the first time of impact between a triangle and an SDF isosurface. Our method offers improved robustness over

point sampling methods, and outperforms recent triangle-SDF discrete collision detection (DCD) algorithms. Furthermore, a novel method for adaptively refining the potential collision points on large triangles is proposed for robust triangle-SDF collision detection with coarse meshes. This enables the use of reduced geometry for efficient simulations. We demonstrate the benefits of our approach by comparing to state-of-the-art algorithms for triangle-SDF collision detection, and showcase its effectiveness through simulations involving complex collision scenarios.

1.2 Introduction

Collision detection is the computational process of determining whether, where, and when objects in a simulated environment are intersecting. It is a key technology for numerous applications in computer graphics, robotics, and virtual reality. Continuous Collision Detection (CCD) is particularly interesting as it permits computing intersection-free trajectories of the simulated objects, whereas discrete methods may miss collisions or require shape separation (Nie *et al.*, 2020).

The accuracy, complexity and performance of collision detection methods are largely dependent on shape representation. For instance, triangle meshes are popular shape representations in computer graphics — due to their ability to model complex geometry — but intersection tests using triangle meshes are expensive, especially as higher resolutions are required to capture fine details. Simulation performance can be drastically improved with simpler shape representations. For example, efficient methods exist for distance, overlap and closest-point queries between spheres, oriented or axis-aligned boxes, and convex hulls (Ericson, 2004). While such primitives cannot accurately represent most objects, they are invaluable when real-time performance is required. They can provide efficient approximate bounding volumes for use in broad-phase collision detection, drastically reducing the number of computationally expensive tests required by more accurate narrow-phase collision detection. Likewise, Signed Distance Fields (SDFs) provide efficient distance and inside-outside information, making them powerful tools for shape approximation in collision detection. They can be stored compactly as analytical functions (Andrews, Erleben & Ferguson, 2022), discrete grids (Koschier, Deul, Brand & Bender,

2017a), or even deep neural networks (Park, Florence, Straub, Newcombe & Lovegrove, 2019). They excel at representing smooth or curved surfaces, and can efficiently represent complex geometry (Koschier *et al.*, 2017a). Common drawbacks of SDFs, however, are that sharp features may be lost due to the smooth nature of implicit surfaces, and they are not well suited to represent non-manifold or self-intersecting surfaces where space cannot be properly partitioned into a well-defined interior, exterior, and surface (Jones, Baerentzen & Sramek, 2006). Triangle meshes, on the other hand, are unaffected by these issues. It is thus of particular interest to combine shape representations and leverage their respective advantages.

Discrete collision detection involving meshes and primitives is well explored already (Andrews *et al.*, 2022; Teschner *et al.*, 2005; Ericson, 2004). Methods involving meshes and SDFs have traditionally been limited to point-based contacts over the mesh vertices or require dense point sampling (Xu & Barbič, 2017; McNeely *et al.*, 2005). The recently proposed method by Macklin *et al.* (2020) specifically addresses the problem of discrete collision detection between SDFs and triangle meshes. They improve the quality of contact generation compared to point-based approaches, but they only support a single contact point per triangle. Larger triangles may intersect multiple features of an SDF, and thus dense and high-resolution meshes are often required for reliable collision detection. However, low-resolution coarse meshes are preferred by many applications for their efficiency. The existing triangle-SDF methods are further limited to intersection tests performed at discrete points in time. This can result in shapes with a high relative velocity passing through each other (tunnelling). Similar problems occur for SDFs with thin features, where it is easy for shapes to pass through regions where the SDF gradient is undefined or changes direction. This can lead to inaccurate or unstable simulations. Triangle-SDF CCD is, to the best of our knowledge, an unexplored problem.

We address the aforementioned problems by proposing a novel triangle-SDF collision detection algorithm that reformulates the collision detection problem as a local spatio-temporal optimization. Solving this optimization problem provides not only the time of impact, but also the closest points between triangles and an SDF isosurface. Our CCD algorithm avoids tunnelling and offers improved robustness over point sampling methods; it also outperforms recent triangle-SDF

discrete collision detection (DCD) algorithms. Our technique does not assume a specific contact model (constraint-based, barrier methods, penalty forces, impulse-based) and can be integrated into existing simulators. Our technique can be used alongside other methods to handle different collision shapes; it does not limit the simulation to triangles and SDFs. Additionally, we introduce a technique to efficiently compute multiple contact points per triangle to support coarse triangle meshes without globally refining the whole mesh. Rather, new contact points are generated only when needed. We evaluate our method for speed, accuracy, and robustness, using a series of hand-crafted extreme collision scenarios involving triangles and SDFs. We demonstrate our approach using several complex rigid body and cloth simulations, with meshes up to hundreds of thousands of triangles and SDFs hand-picked to challenge gradient methods with gradient discontinuities, sharp features, and multiple boundaries along triangle trajectories. Our algorithm is suitable for real-time applications, and interactive speeds are easily achieved, even on a single-threaded CPU-based implementation.

1.3 Related Work

Collision detection is a key technology for many computer graphics applications, including rendering, animation, 3D modeling, and computational design. It is the process of determining if two or more geometric shapes are intersecting or will intersect. In this section, we provide a brief overview of seminal work in computer graphics on collision detection, with a particular emphasis on the CCD regime of methods, as well as collision algorithms involving geometry represented as an SDF.

1.3.1 Collision detection

Previous works in computer graphics have provided overviews of the fundamentals of collision detection and response for physics-based animation (Andrews *et al.*, 2022; Teschner *et al.*, 2005; Ericson, 2004). However, numerous algorithms have been developed by the graphics community over the years to address collision detection for specific phenomena and geometry pairs. For instance, Fuhrmann *et al.* (2003) developed a method for rapid collision detection

between rigid and highly deformable objects using distance fields, and Zhang, Marschner, Solomon & Tamstorf (2023a) use Sum-of-Squares programming to handle higher-order geometry. Alternatively, Guendelman, Bridson & Fedkiw (2003) uses interference detection through testing vertices to the inside/outside function of another model and time step limitation.

Computational complexity and performance are often primary concerns where collision detection is concerned, since the process is often viewed as a bottleneck in the simulation pipeline. Acceleration data structures such as spatial hash grids (Quinlan, 1994), BSP trees (Naylor, 1998), and bounding volume hierarchies (BVH) (Ericson, 2004) are common since they effectively prune the search space of narrow-phase interaction tests, which can be costly. Such acceleration techniques are compatible with our proposed methodology, and we show how a BVH of spheres can be used to effectively reduce the number of triangle-SDF collision tests.

Animations involving fast-moving objects or thin geometry can produce tunnelling and severe penetration artefacts, due mainly to the discrete collision detection used by many applications. Continuous collision detection (CCD) avoids these artifacts by identifying intersections along the entire path of motion, resulting in specialized algorithms for simulation of rigid bodies (Redon, Kheddar & Coquillart, 2002), deformable solids (Govindaraju, Kabul, Lin & Manocha, 2006), and cloth (Bridson, Fedkiw & Anderson, 2002). CCD is also an essential component of robust contact simulation with non-interpenetration guarantees, such as in incremental potential contact (IPC) (Li *et al.*, 2020). Xu & Barbič (2017) developed a continuous algorithm for collisions between point clouds and implicit geometry representations, such as distance functions. Such methods can often use properties like Bernstein basis or Bézier curves to reduce the problem (Tang, Tong, Wang & Manocha, 2014). This is useful to model the continuous nature of curved regions. Alternatively, some approaches assume linear motions to reduce computational complexity (Chen *et al.*, 2024; Wang *et al.*, 2021, 2022). This assumption can significantly streamline calculations, making it easier to predict and manage object trajectories. Similarly, it is possible to use interval calculations to do SDF-SDF collision detection (Liu *et al.*, 2024) and geometrically exact CCD (Brochu, Edwards & Bridson, 2012), though at prohibitive costs for real-time CCD for multiple contact points queries.

Our work focuses on collision detection combining an explicit mesh-based representation of geometry with an implicit SDF. It can be integrated as part of a typical physical simulation pipeline and does not prevent the use of other collision detection methods to handle different collision shapes. Our proposed acceleration structure relies on this flexibility. It uses less accurate shapes for which vastly more efficient SDF CCD methods can provide early approximations and collision pair culling. Haptic simulation has similarly focused on using similar hybrid representations, with collision between a point cloud and a voxelized distance field receiving particular attention (McNeely, Puterbaugh & Troy, 1999; McNeely *et al.*, 2005). Robustness is improved by continuous collision detection versions of these algorithms (Xu & Barbič, 2017).

To the best of our knowledge, our method is unique in addressing CCD between triangles and arbitrary SDFs. IPC implementations use triangle-triangle representations (Li *et al.*, 2020), and related methods e.g. (Xu & Barbič, 2017), require dense point sampling or specific SDF discretizations.

1.3.2 Geometry representation

The efficiency of SDF-based algorithms is directly impacted by the efficiency of their underlying data structures and shape representations. Numerous adaptive methods have been proposed to reduce the storage overhead of sampled signed distance fields (Friskin, Perry, Rockwood & Jones, 2000). Notably, Koschier *et al.* (2017a) proposed a highly accurate grid-based SDF by fitting a polynomial approximation of the local shape to each grid cell, and we use this representation for several of our experiments. Deep neural models have also demonstrated effectiveness for learning implicit shape representations (Park *et al.*, 2019). SDF-based swept volumes have proven to be effective in computing collision-free trajectories for animation (Wang *et al.*, 2024) and robotics planning (Zhang, Wang, Xu, Gao & Gao, 2023b), albeit at computational costs that prohibit real-time applications. Our method treats the SDF as a “black box”, so it is not tied to its representation. Thus, future development of efficient shape representations will remain compatible.

1.4 Preliminaries

We first offer a brief overview of the concepts used in our paper and the technique introduced by Macklin *et al.* (2020) to make our paper self-contained. We note that our technique does not require a specific contact model (constraint-based, barrier methods, penalty forces, impulse-based). Contact handling and the resolution of subsequent potential contacts are out of scope for our paper. For this, we refer the reader to related work on the topic (Andrews *et al.*, 2022). While our work focuses on collision detection between triangles and SDFs, it does not prevent other methods from handling different collision shapes in the same physics pipeline. Our proposed acceleration structure, presented in Section 1.5.6, relies on this flexibility. It uses less accurate shapes for which vastly more efficient SDF CCD methods can provide early approximations and collision pair culling. Other efficient collision detection algorithms involving common primitives, and a primer on using multiple specialized collision detection algorithms together can be found in (Ericson, 2004).

1.4.1 Signed distance field (SDF)

A signed distance field is a function $\phi(\vec{\mathbf{x}}): \mathbb{R}^3 \rightarrow \mathbb{R}$ that gives the signed Euclidean distance from a point $\vec{\mathbf{x}} \in \mathbb{R}^3$ to a shape boundary, where $\phi = 0$, and with $\phi(\vec{\mathbf{x}}) < 0$ for all points inside the shape. Its gradient $\nabla\phi \in \mathbb{R}^3$ satisfies $\|\nabla\phi\| = 1$ and gives the direction in which distance increases the most, so the closest point from $\vec{\mathbf{x}}$ on the boundary is $\vec{\mathbf{x}} - \nabla\phi(\vec{\mathbf{x}})\phi(\vec{\mathbf{x}})$. In the case of physics-based animation, the SDF gradient conveniently provides a contact normal $\vec{\mathbf{n}} = \nabla\phi$, which can otherwise be challenging to compute in triangle-triangle collision scenarios (Erleben, 2018).

Certain types of SDFs have limitations. When the SDF uses a discrete representation, its domain may be limited, e.g., to the region covered by a grid. Furthermore, an SDF is not differentiable at points along the medial axis of a shape, resulting in discontinuities of the gradient. However, we find that such issues rarely occur in practice or can be easily circumvented. The impact of SDF quality on our method is further discussed in Section 1.7.

1.4.2 Triangle-SDF discrete collision detection

A technique for discrete collision detection (DCD) between triangle meshes and SDFs was introduced in Macklin *et al.* (2020) using the *closest point methodology* (Erleben, 2018). Collision points are generated by a minimization of the SDF over the barycentric coordinates u, v, w of a triangle with vertices $\vec{\mathbf{a}}, \vec{\mathbf{b}}, \vec{\mathbf{c}} \in \mathbb{R}^3$:

$$\min_{u,v,w} \phi(u\vec{\mathbf{a}} + v\vec{\mathbf{b}} + w\vec{\mathbf{c}}) \quad (1.1)$$

$$\text{s.t.} \quad u, v, w \geq 0 \quad (1.2)$$

$$u + v + w = 1. \quad (1.3)$$

The resulting barycentric coordinates give the deepest penetrating point inside the triangle. Gradient methods, such as *projected gradient descent* or the *Frank-Wolfe algorithm* (Frank & Wolfe, 1956), can solve for local minima by using the derivative of the SDF with respect to the barycentric coordinates:

$$\vec{\mathbf{d}} = \left[\frac{\partial \phi}{\partial u}, \frac{\partial \phi}{\partial v}, \frac{\partial \phi}{\partial w} \right]^\top = \begin{bmatrix} \nabla \phi(\vec{\mathbf{x}}) \cdot \vec{\mathbf{a}} \\ \nabla \phi(\vec{\mathbf{x}}) \cdot \vec{\mathbf{b}} \\ \nabla \phi(\vec{\mathbf{x}}) \cdot \vec{\mathbf{c}} \end{bmatrix}. \quad (1.4)$$

This approach improves the quality of contact generation compared to point-based methods, but only supports a single contact point per triangle (Macklin *et al.*, 2020). Larger triangles may intersect multiple features of an SDF, and thus dense and high-resolution meshes are often required for reliable collision detection. Mesh subdivision can be applied ahead of time to provide better coverage, but it increases the cost of the optimization accordingly. We later propose an efficient method to solve these issues (see Section 1.5.5).

Since DCD only provides contact information if a contact is occurring at a given instant in time, it cannot effectively prevent penetration. Larger time steps, fast movements, or thin SDFs can cause discrete collision techniques to miss contacts. For instance, SDFs may fail to provide a realistic separation direction for sufficiently deep points, or if they are inside sharp features. This

results in shapes being forced to penetrate or passing through one another instead of colliding as expected. In such cases, continuous methods for collision detection are preferred.

1.4.3 Continuous collision detection

Continuous collision detection (CCD) is generally applied in order to provide accurate collision detection in simulations where discrete time steps may cause fast-moving or thin objects to pass through each other or interpenetrate in a way such that collisions cannot be handled accurately or in a stable fashion. CCD gives the time of the first impact between objects in a given time interval $t_{\text{hit}} \in [t_{\text{start}}, t_{\text{end}}]$, and guarantees that no collision occurs up to t_{hit} .

1.4.4 Golden section search

Golden section search (GSS) is an iterative minimization technique discovered by Kiefer (1953), and is well-suited to local minimization problems. We use it to solve single-dimension sub-problems as part of our method in Section 1.5.3. Although GSS is a well-known algorithm and many public domain implementations are available, we provide algorithmic details in Algorithm 1.1 for completeness. The method seeks the minimizer for a function $f(l)$ over some interval, such that

$$l_{\min} = \arg \min_{l \in [l_{\text{start}}, l_{\text{end}}]} f(l). \quad (1.5)$$

The algorithm maintains function values f_0, f_1, f_2 , and f_4 , for four points l_0, l_1, l_2 , and l_4 , respectively. The points form three contiguous intervals with a width ratio of $\varphi : 1 : \varphi$, where φ is the golden ratio:

$$\varphi = (1 + \sqrt{5})/2, \quad (1.6)$$

$$\varphi^{-1} = (\sqrt{5} - 1)/2. \quad (1.7)$$

The algorithm operates by successively narrowing the range over which f is evaluated. Note that the points l_i can be of any dimensions. In the context of CCD, we use them both as points in

Algorithm 1.1 Golden Section Search algorithm for minimizing function $f(l)$ over the interval $l_{\text{start}} l_{\text{end}}$. Returns the minimum value l_{min} .

```

1 GSSMinimize( $l_{\text{start}}, l_{\text{end}}, f$ )
2    $r \leftarrow \varphi^{-1}, \quad r^{-1} \leftarrow 1 - r$ 
3    $\alpha_0 \leftarrow 0, \quad \alpha_1 \leftarrow r^{-1}, \quad \alpha_2 \leftarrow r, \quad \alpha_3 \leftarrow 1$ 
4    $l_0 \leftarrow l_{\text{start}}, \quad l_1 \leftarrow \text{lerp}(l_{\text{start}}, l_{\text{end}}, \alpha_1), \quad l_2 \leftarrow \text{lerp}(l_{\text{start}}, l_{\text{end}}, \alpha_2), \quad l_3 \leftarrow l_{\text{end}}$ 
5    $f_0 \leftarrow f(l_0), \quad f_1 \leftarrow f(l_1), \quad f_2 \leftarrow f(l_2), \quad f_3 \leftarrow f(l_3)$ 
6   while  $(l_3 - l_0) \leq \text{tol} (l_1 + l_2)$  do
7     if  $\min(f_0, f_1) < \min(f_2, f_3)$  then
8        $\alpha_3 \leftarrow \alpha_2, \quad l_3 \leftarrow l_2, \quad f_3 \leftarrow f_2$ 
9        $\alpha_2 \leftarrow \alpha_1, \quad l_2 \leftarrow l_1, \quad f_2 \leftarrow f_1$ 
10       $\alpha_1 \leftarrow r \alpha_2 + r^{-1} \alpha_0$ 
11       $l_1 \leftarrow \text{lerp}(l_{\text{start}}, l_{\text{end}}, \alpha_1)$ 
12       $f_1 \leftarrow f(l_1)$ 
13    else
14       $\alpha_0 \leftarrow \alpha_1, \quad l_0 \leftarrow l_1, \quad f_0 \leftarrow f_1$ 
15       $\alpha_1 \leftarrow \alpha_2, \quad l_1 \leftarrow l_2, \quad f_1 \leftarrow f_2$ 
16       $\alpha_2 \leftarrow r \alpha_1 + r^{-1} \alpha_3$ 
17       $l_2 \leftarrow \text{lerp}(l_{\text{start}}, l_{\text{end}}, \alpha_2)$ 
18       $f_2 \leftarrow f(l_2)$ 
19    end if
20  end while
21   $l_{\text{mid}} \leftarrow \text{lerp}(l_{\text{start}}, l_{\text{end}}, \frac{1}{2}(\alpha_0 + \alpha_3))$ 
22   $f_{\text{mid}} \leftarrow f(l_{\text{mid}})$ 
23  if  $f_0 < f_{\text{mid}}$  and  $f_0 < f_3$  then
24     $l_{\text{min}} \leftarrow l_0$ 
25  else if  $f_{\text{mid}} < f_3$  then
26     $l_{\text{min}} \leftarrow l_{\text{mid}}$ 
27  else
28     $l_{\text{min}} \leftarrow l_3$ 
29  end if
30  return  $l_{\text{min}}$ 

```

time and as points in space, depending on the specific line search problem. Points are obtained by linear interpolation over the interval $[l_{\text{start}}, l_{\text{end}}]$. The complement φ^{-1} to the golden ratio is used as part of the interval refinement process. Each iteration, we compare the minimum of the left values, i.e., $\min(f_0, f_1)$, to the minimum of the right values, i.e., $\min(f_2, f_3)$. If the left values are smaller, the boundaries of the two leftmost intervals are used to form a new set of

three intervals with the same width ratio; the boundaries of the two rightmost intervals are used otherwise. Only a single new point and function value need to be computed when the interval is reduced. One point stays the same, and the other two simply reuse previously computed points. Once the algorithm terminates, the point with the smallest function evaluation from l_0 , l_3 , and midpoint l_{mid} is returned.

1.5 Triangle-SDF CCD

In this section, we present an approach that extends triangle-SDF collision to the continuous regime. Previous work on triangle-SDF collision by Macklin *et al.* (2020) uses iterative gradient-based methods to solve the triangle-SDF discrete collision detection problem. Likewise, we are interested in Frank-Wolfe (FW) methods, which are iterative first-order optimization algorithms for constrained convex optimization (Frank & Wolfe, 1956; Braun *et al.*, 2022; Jaggi, 2013). However, the landscape of CCD optimization is potentially non-convex, and this can result in tunnelling, as shown in Figure 1.2. Hence, one of our central contributions is a novel Frank-Wolfe-based algorithm that identifies and converges toward a spatio-temporal global minimum.

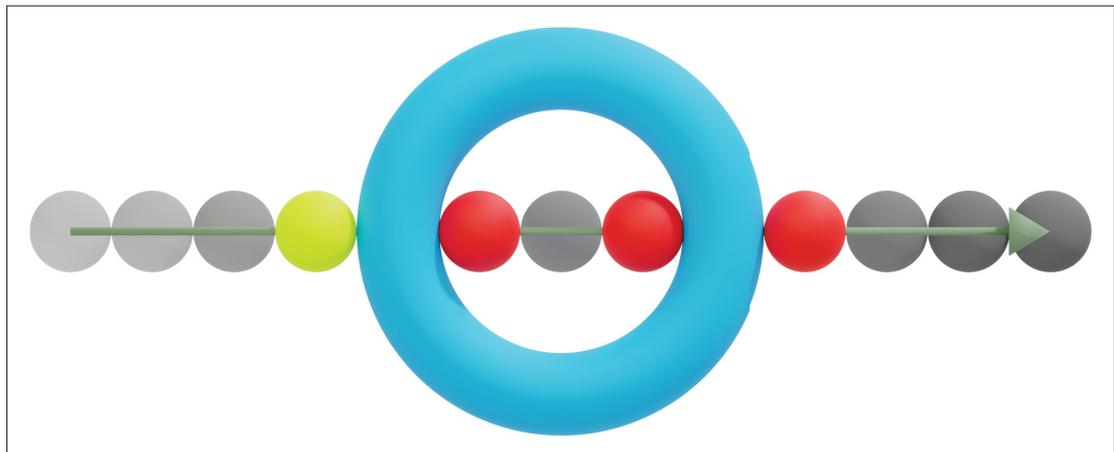


Figure 1.2 A sphere moving rightward passes through the boundary of an SDF isosurface (torus) multiple times. Local minima of the spatio-temporal problem are encountered each time the sphere touches the SDF. CCD requires solving for a global minimum with regard to time (green), but gradient methods may converge toward erroneous local minima (red)

We focus on FW algorithms since preliminary experiments revealed that Projected Gradient Descent (PGD) often fails to converge due to a constant step size. Furthermore, backtracking variations were costly due to the increased number of projections back to the problem space required by PGD, whereas FW algorithms require none (as noted by Macklin *et al.* (2020)).

1.5.1 Local optimization

The goal of the CCD optimization is to provide the spatial coordinates and time of the earliest contact over some time interval $t_{\text{hit}} \in [t_{\text{start}}, t_{\text{end}}]$, such that no collision occurs before t_{hit} . We seek, for each triangle, a point of contact $\vec{\mathbf{x}}_c$ at some future time t where $\phi(\vec{\mathbf{x}}_c) = 0$. Since multiple instances of the the time interval can yield $\phi(\vec{\mathbf{x}}_t) = 0$, solving for the unsigned distance $|\phi(\vec{\mathbf{x}}_t)|$ could converge on the wrong side of the SDF (as shown in Figure 1.2). Therefore, we cannot directly solve for t_{hit} using the unsigned distance $|\phi(\vec{\mathbf{x}}_t)|$ if we intend to find a temporal global minimum. Hence, we minimize the signed distance over the time variable t and the barycentric coordinates u, v, w of a triangle with vertices $\vec{\mathbf{a}}_t, \vec{\mathbf{b}}_t, \vec{\mathbf{c}}_t \in \mathbb{R}^3$:

$$\min_{u,v,w,t} \phi(u\vec{\mathbf{a}}_t + v\vec{\mathbf{b}}_t + w\vec{\mathbf{c}}_t), \quad (1.8)$$

$$\text{s.t.} \quad \phi(u\vec{\mathbf{a}}_t + v\vec{\mathbf{b}}_t + w\vec{\mathbf{c}}_t) \geq 0, \quad (1.9)$$

$$u, v, w \geq 0, \quad (1.10)$$

$$u + v + w = 1, \quad (1.11)$$

$$t \in [t_{\text{start}}, t_{\text{end}}] . \quad (1.12)$$

Here, $\vec{\mathbf{x}}_t = u\vec{\mathbf{a}}_t + v\vec{\mathbf{b}}_t + w\vec{\mathbf{c}}_t$ and the subscript \cdot_t indicates that the vertex positions change over the time interval, and that we seek the global minimum for t . If no t is found where the unsigned distance is zero, then a collision does not occur over the time interval. Otherwise, the simulation cannot advance past the time of collision and remain intersection-free. Violations of the constraint in Equation 1.9 provide new upper bounds to the time interval; see Section 1.5.2 for more details.

We revise the derivative from Equation 1.4 to account for the added time dimension:

$$\vec{\mathbf{d}} = \begin{bmatrix} \nabla\phi(\vec{\mathbf{x}}_t) \cdot \vec{\mathbf{a}}_t \\ \nabla\phi(\vec{\mathbf{x}}_t) \cdot \vec{\mathbf{b}}_t \\ \nabla\phi(\vec{\mathbf{x}}_t) \cdot \vec{\mathbf{c}}_t \\ \nabla\phi(\vec{\mathbf{x}}_t) \cdot \vec{\mathbf{v}}_t \end{bmatrix}. \quad (1.13)$$

Here, $\vec{\mathbf{v}}_t$ is the linear velocity of the potential solution $\vec{\mathbf{x}}_t$ at time t . The velocity $\vec{\mathbf{v}}_t$ is obtained differently depending on the type of simulation or required accuracy. For instance, velocities can be extracted directly from the values stored at the nodal coordinates in cloth and elastic simulation. For standard rigid-body scenarios, we project the angular and linear velocities ($\vec{\omega}_g$ and $\vec{\mathbf{v}}_g$) from the rigid-body's centre of rotation $\vec{\mathbf{g}}$ to the surface point $\vec{\mathbf{x}}_t$, such that

$$\vec{\mathbf{v}}_t = \vec{\mathbf{v}}_g + \vec{\omega}_g \times (\vec{\mathbf{x}}_t - \vec{\mathbf{g}}). \quad (1.14)$$

We note that SDFs are queried in their own local space, as it is more efficient and much easier to transform points than it is to update SDFs at runtime (Andrews *et al.*, 2022). Thus, if the SDF is used as a collision shape for an object which also undergoes rigid motions, the position and velocities should be further projected into the SDF space.

1.5.2 Frank-Wolfe for CCD

Although the standard Frank-Wolfe algorithm cannot guarantee convergence toward a global minimum, it has excellent convergence properties when seeking a local minimum (Frank & Wolfe, 1956; Braun *et al.*, 2022; Jaggi, 2013). The algorithm thus provides an efficient approach to detect if a point violates the lower bound on the SDF value (see Equation 1.9) for the CCD problem. At any point, $\phi(\vec{\mathbf{x}}_t) \leq 0$ automatically guarantees that $t \geq t_{\text{hit}}$, which gives a new upper bound to the time minimization problem. The time interval is reduced to $t_{\text{hit}} \in [t_{\text{start}}, t]$, thus culling later temporal minima from the problem space. Unfortunately, the gradient at the end of the newly truncated time interval may cause the algorithm to get stuck trying to converge toward

the same local minimum as before. We address this issue by splitting the direction-finding sub-problem of the Frank-Wolfe algorithm into two steps.

Our algorithm begins much like Macklin *et al.* (2020). Each iteration i , a support vertex is selected from the triangle in its configuration at time t_i , which provides a direction in space:

$$\min_{\vec{s}_i} \quad \vec{s}_i^\top \nabla \phi(\vec{\mathbf{x}}_{t_i}), \quad (1.15)$$

$$\text{s.t.} \quad \vec{s}_i \in \left\{ \vec{\mathbf{a}}_{t_i}, \vec{\mathbf{b}}_{t_i}, \vec{\mathbf{c}}_{t_i} \right\}. \quad (1.16)$$

Then, we introduce a new single-dimension temporal direction

$$d_i = \begin{cases} -1, & \text{if } \phi(\vec{\mathbf{x}}_{t_i}) \leq 0 \\ -\text{sign}(\nabla \phi(\vec{\mathbf{x}}_{t_i}) \cdot \vec{\mathbf{v}}_{t_i}), & \text{otherwise} \end{cases} \quad (1.17)$$

that moves $\vec{\mathbf{x}}_{t_i}$ toward the zero isosurface of the SDF when outside, and otherwise moves the solution backward in time. The solution is updated using the step $\alpha = 2/(i + 2)$, standard to Frank-Wolfe methods, such that

$$\Delta t = \alpha(t_{\text{end}} - t_{\text{start}}) d_i, \quad (1.18)$$

$$t_{i+1} \leftarrow \max(t_{\text{start}}, \min(t_{\text{end}}, t_i + \Delta t)), \quad (1.19)$$

$$\vec{\mathbf{x}}_{t_{i+1}} \leftarrow \text{proj}(\vec{\mathbf{x}}_{t_i} + \alpha(\vec{\mathbf{s}}_i - \vec{\mathbf{x}}_{t_i})). \quad (1.20)$$

Note that in the last line, the position of the updated solution is projected into the triangle at t_{i+1} to obtain $\vec{\mathbf{x}}_{t_{i+1}}$. This can be achieved by barycentric interpolation or by applying the same transformation used to produce the triangle at t_{i+1} .

Splitting the direction-finding sub-problem this way forces the algorithm to ignore the temporal part of the gradient when $\phi(\vec{\mathbf{x}}_t) \leq 0$. The spatial direction still follows the gradient, so it may find deeper points, but the temporal direction will be able to push them out of the shape by moving back in time to refine the current solution and avoid penetration at the local minimum.

It also allows the algorithm to find earlier local minima, in which case the boundary can be adjusted again until the only remaining minimum is the one at the end of the adjusted time interval. The algorithm terminates by detecting when $t_{i+1} \approx t_i$ and $\vec{\mathbf{x}}_{t_{i+1}} \approx \vec{\mathbf{x}}_{t_i}$, up to a desired precision, or when a hard limit on the number of iterations is reached.

This temporally modified version of the Frank-Wolfe algorithm usually reduces the problem down to a simple convex problem with only a single global temporal minimum, yet offers no guarantees. We found multiple scenarios where this method fails, such as configurations involving repeated tunnelling over a single time step, as shown in Figure 1.2. We address these issues in the following section.

1.5.3 Frank-Wolfe with line search

The domain of the spatial minimization sub-problem changes as the modified Frank-Wolfe method described in Section 1.5.2 steps through time. The triangle at the new time point is effectively a different 3D slice of the four-dimensional spatio-temporal domain, which introduces discontinuities. In addition, sudden modifications of the time domain occur when an evaluated point violates the lower bound constraint $\phi(\vec{\mathbf{x}}_{t_i}) \geq 0$. This may cause multiple local minima to be pushed out of bound, and force the spatial minimization to keep seeking new ones. This is made difficult by the ever-decreasing step size $\alpha = 2/(i + 2)$, as the step size can become too small to find other minima.

It is possible to avoid this issue by restarting the minimization whenever we reach sufficient precision to ensure no penetration at a local minimum, but this quickly becomes costly. The candidate contact point $\vec{\mathbf{x}}_{t_i}$ will often “wobble” around a solution, getting continually closer, but may never quite reach its target. We show an example of this behaviour in the supplemental video, where the iterations over a triangle in a rotating pyramid are compared to those of our method. This behaviour is expected from Frank-Wolfe methods, however repeating this behaviour for multiple local minima slows down the algorithm drastically. Furthermore, because we ignore the direction of one of the gradient components while in contact with the SDF, the new modified

Algorithm 1.2 Frank-Wolfe with GSS (FWGSS)

```

1 FWGSS( $t_{\text{start}}, t_{\text{end}}$ )
2    $t_1 \leftarrow t_{\text{start}}$ 
3   Compute the barycentric coordinates  $u, v, w$  of the starting iterate.  $\triangleright$  Section 1.5.4
4   for  $i \in 1 \dots \text{max iterations}$  do
      // Solve temporal sub-problem
5    $\vec{\mathbf{x}}_{t_i} \leftarrow \text{BarycentricInterpolate}(u, v, w, t_i)$ 
6   Compute  $\vec{\mathbf{v}}_{t_i}$   $\triangleright$  Équation 1.14
7   Query  $\phi(\vec{\mathbf{x}}_{t_i})$  and  $\nabla\phi(\vec{\mathbf{x}}_{t_i})$ 
8   if  $\phi(\vec{\mathbf{x}}_{t_i}) \leq 0$  then
9      $t_{\text{end}} \leftarrow \min(t_i, t_{\text{end}})$   $\quad$  // Update interval
10    def UnsignedDistanceAtTime( $t$ ):
11       $\vec{\mathbf{x}} \leftarrow \text{BarycentricInterpolate}(u, v, w, t)$ 
12      return  $|\phi(\vec{\mathbf{x}})|$ 
13     $t_{i+1} \leftarrow \text{GSSMinimize}(t_{\text{start}}, t_i, \text{UnsignedDistanceAtTime})$ 
       $\triangleright$  Algorithm 1.1
14    else
15      Compute time direction  $d_i$   $\quad$   $\triangleright$  Équation 1.17
16      def SignedDistanceAtTime( $t$ ):
17         $\vec{\mathbf{x}} \leftarrow \text{BarycentricInterpolate}(u, v, w, t)$ 
18        return  $\phi(\vec{\mathbf{x}})$ 
19      if  $d_i < 0$  then  $t_{i+1} \leftarrow \text{GSSMinimize}(t_{\text{start}}, t_i, \text{SignedDistanceAtTime})$ 
20      else  $t_{i+1} \leftarrow \text{GSSMinimize}(t_i, t_{\text{end}}, \text{SignedDistanceAtTime})$ 
21    end if
      // Solve spatial sub-problem
22     $\vec{\mathbf{x}}_{t_{i+1}} \leftarrow \text{BarycentricInterpolate}(u, v, w, t_{i+1})$ 
23    Compute  $\vec{\mathbf{v}}_{t_{i+1}}$   $\quad$   $\triangleright$  Équation 1.14
24    Query  $\phi(\vec{\mathbf{x}}_{t_{i+1}})$  and  $\nabla\phi(\vec{\mathbf{x}}_{t_{i+1}})$ 
25    if  $\phi(\vec{\mathbf{x}}_{t_{i+1}}) \leq 0$  then  $t_{\text{end}} \leftarrow \min(t_{i+1}, t_{\text{end}})$   $\quad$  // Update interval
26
27    Compute support vertex  $\vec{\mathbf{s}}_1$   $\quad$   $\triangleright$  Équation 1.15
28    def SignedDistanceAtPoint( $\vec{\mathbf{x}}$ ):
29      return  $\phi(\vec{\mathbf{x}})$ 
30     $\vec{\mathbf{x}}_{t_{i+1}} \leftarrow \text{GSSMinimize}(\vec{\mathbf{x}}_{t_{i+1}}, \vec{\mathbf{s}}_1, \text{SignedDistanceAtPoint})$ 
31    Update barycentric coordinates  $u, v, w$  using  $\vec{\mathbf{x}}_{t_{i+1}}$ 
32    if  $t_{i+1} \approx t_i$  and  $\vec{\mathbf{x}}_{t_{i+1}} \approx \vec{\mathbf{x}}_{t_i}$  then break
33  end for

```

direction does not have to be the descent direction, so theoretically the method can oscillate. We address the aforementioned issues with two more modifications over the temporally modified Frank-Wolfe method presented in Section 1.5.2.

First, α from Equations 1.18 - 1.20 is replaced by two line searches: one to solve for the optimal spatial step toward \vec{s}_i , and another for the optimal temporal step in the direction obtained from Equation 1.17. Each line search is its own minimization of the SDF over a one-dimensional slice of the problem space – one in time, and the other in space. Solving for the step size with a line search maintains the convergence guarantees of the Frank-Wolfe algorithm (Braun *et al.*, 2022; Jaggi, 2013), but allows the spatial minimization to better adjust to the discontinuities and reductions of the domain. For this, we use the GSS method from Section 1.4.4. Our choice is based on the robustness and efficiency of the method, demonstrated by Macklin *et al.* (2020), when minimizing SDFs over single-dimension problems. The second modification is to use the result of the temporal sub-problem minimization to improve the spatial minimization. Each iteration, we find the time direction, solve for the temporal step size, and compute t_{i+1} as before. However, we additionally project \vec{x}_t to t_{i+1} and reevaluate the SDF and gradient before we proceed with Equation 1.15, solve for the spatial step size, and update \vec{x}_{i+1} . We found this to increase the efficiency of our technique in solving the spatiotemporal problem.

To summarize, our Frank-Wolfe with line search (FWGSS) converges toward the first time of impact by solving different problems depending on the situation. First, it minimizes the current solution by finding barycentric coordinates and time that maximize penetration. When the current solution is inside the SDF, i.e., $\phi(\vec{x}_{t_i}) < 0$, the temporal line search seeks to minimize the unsigned distance, while the spatial line search continues minimizing the SDF. Additionally, $\phi(\vec{x}_{t_i}) \leq 0$ provides a new temporal upper bound, and forces the temporal line search to ignore the gradient and seek solutions backward in time. This effectively causes triangles to exit the SDF, culls local minima past the temporal upper bound, and avoids searching for solutions near the local minimum at the end of the time interval. By repeating until no further local minima are found, we obtain a simpler problem with a single temporal (global) minimum at the first time of impact.

Visualizations of the convergence in our supplemental video compare versions of the modified Frank-Wolfe method alongside projected gradient descent and shows that FWGSS (ours) is a clear winner. Our proposed algorithm is further summarized as pseudocode in Algorithm 1.2.

1.5.4 Starting iterate

Picking a good starting iterate can also improve the convergence of gradient-based algorithms. In (Macklin *et al.*, 2020), the triangle vertex $\vec{s}_i \in \{\vec{\mathbf{a}}, \vec{\mathbf{b}}, \vec{\mathbf{c}}\}$ closest to the SDF is selected. We found this heuristic to hinder convergence in our method, as the closest point to the SDF at the start of the time interval might not necessarily be moving toward the SDF. Instead, the vertex that is most likely to collide with the SDF is selected as the one that minimizes $\vec{\mathbf{v}}_i \cdot \nabla \phi(\vec{\mathbf{s}}_i)$, where $\vec{\mathbf{v}}_i$ is the linear velocity of $\vec{\mathbf{s}}_i$.

1.5.5 Adaptive triangle subdivision

Algorithm 1.2 and the algorithm described by Macklin *et al.* (2020) are triangle-local. That is, they provide a single contact point per triangle. As shown in Figure 1.10 and Figure 1.11, this is not sufficient when large triangles or edges collide with the SDF at many simultaneous contact points. Subdividing the mesh reduces the risk of encountering this situation by splitting contact planes into smaller coplanar triangles. Unfortunately, this drastically increases the number of triangles on the mesh, all of which must be tested for contacts at every step of the simulation. We address this issue by providing an adaptive domain subdivision scheme to generate more contact points, when needed, without the cost of subdividing the whole mesh or modifying the geometry.

Collision detection provides a contact point $\vec{\mathbf{x}}_c$ with barycentric coordinates u, v, w on a triangle with vertices $\vec{\mathbf{a}}, \vec{\mathbf{b}}, \vec{\mathbf{c}}$. The triangle’s centroid $\vec{\mathbf{g}} = (\vec{\mathbf{a}} + \vec{\mathbf{b}} + \vec{\mathbf{c}})/3$ is cheap to compute and allows us to define the three barycentric regions A, B, C shown in Figure 1.3. These regions each form

a sub-triangle with the following vertices:

$$A \leftarrow \vec{\mathbf{b}}, \vec{\mathbf{c}}, \vec{\mathbf{g}}, \quad B \leftarrow \vec{\mathbf{a}}, \vec{\mathbf{g}}, \vec{\mathbf{c}}, \quad C \leftarrow \vec{\mathbf{a}}, \vec{\mathbf{b}}, \vec{\mathbf{g}}. \quad (1.21)$$

The barycentric coordinates can be used to find in which region $\vec{\mathbf{x}}_c$ resides. It is within A if $u \leq v$ and $u \leq w$, within B if $v \leq u$ and $v \leq w$, or within C if $w \leq u$ and $w \leq v$. The discrete collision detection algorithm described in Section 1.4.2 can then search for supplemental simultaneous contact points in sub-triangles which do not include $\vec{\mathbf{x}}_c$. DCD is sufficient here, even in a CCD pipeline, because $\vec{\mathbf{x}}_c$ is already obtained at the earliest time of impact, and we seek simultaneous contacts.

However, it is not guaranteed that contacts exist in sub-regions, or that only a single contact should be generated in the sub-triangle where $\vec{\mathbf{x}}_c$ resides. The subdivision process is therefore applied iteratively to each of the sub-triangles. We show a typical example of one such case in Figure 1.3. It is possible to balance accuracy and performance by using a threshold, based on a minimum sub-triangle area, to limit iterative subdivisions to bigger triangles or sub-triangles. Only two additional non-collinear points are required for rigid face-face contact stability, so iterative subdivision may be interrupted early upon finding two valid points. In practice, we found that testing the two original sub-regions without iterative subdivision is generally sufficient even for extremely coarse meshes, as shown in the supplemental video and Figure 1.10.

1.5.6 Acceleration structure

Narrow-phase intersection tests can effectively be reduced by using bounding volumes around each triangle, or a hierarchy thereof, and evaluating their distance to the SDF zero isosurface. Broad-phase collision tests between a bounding sphere and an SDF are particularly efficient, because a single distance query is enough to determine if collision is possible. If the distance from a sphere's centre is bigger than its radius, nothing within the sphere can collide with the SDF. However, using bound spheres is more difficult with CCD, especially when simulating complex motions. For instance, a swept-sphere test can help cull entire rigid bodies in some

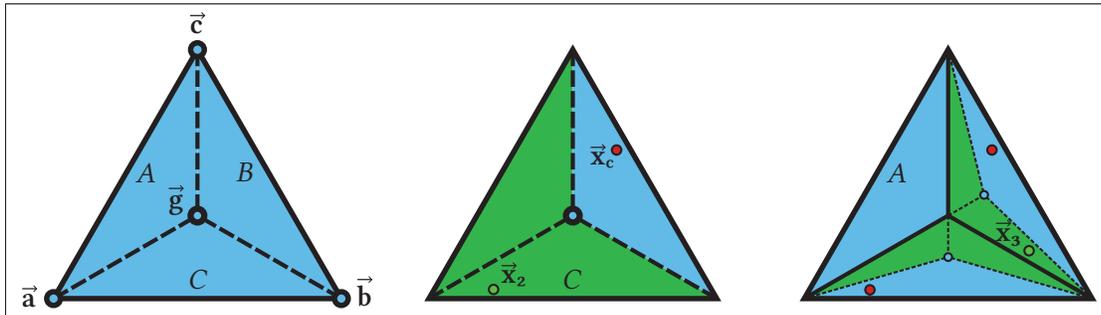


Figure 1.3 A typical example of the subdivision process and recursive search for supplemental contacts. Left: a triangle can be subdivided in three sub-triangles at its centroid \vec{g} . Centre: based on the position of a first contact point \vec{x}_c , two (green) sub-triangles are searched for new contact points. A new contact \vec{x}_2 is found in sub-triangle C . Right: The sub-triangles are further subdivided to search for a third contact point. Sub-triangle A is ignored, because a previous search found no contact in it. Based on the existing contacts, four new (green) sub-triangles are searched, and a third point of contact \vec{x}_3 is found. The three points satisfy the conditions of being coplanar and not collinear, the algorithm is done

scenarios, but cannot capture rotational motions of individual triangles, leading to missed collisions (Ericson, 2004).

We found that a sphere tracing method covers complex motions, including rotations, over time well enough to efficiently cull triangles, as long as a sufficient collision margin is added around each sphere, and that the shapes do not rotate more than 180 degrees around their centre of rotation. This is easily ensured by splitting the sphere tests into multiple intervals based on the rotation threshold. Note that we seek the earliest potential contact. There is no need to test later parts of the full interval if a potential contact is found in an earlier segment. Then, our iterative optimization algorithm moves the “padded” spheres linearly, between their position at the start and at the end of the time interval. A visualization of the bounding sphere and padding is shown in Figure 1.4. The movement at each iteration is defined by the difference between the padded sphere’s radius and the distance to the SDF. If the padded sphere intersects the SDF anywhere along this trajectory, the corresponding triangle is considered for further collision testing. For time integration methods assuming constant velocities over a time step, we found that a good heuristic for the proper padding is to compute the distance between the position of the bounding

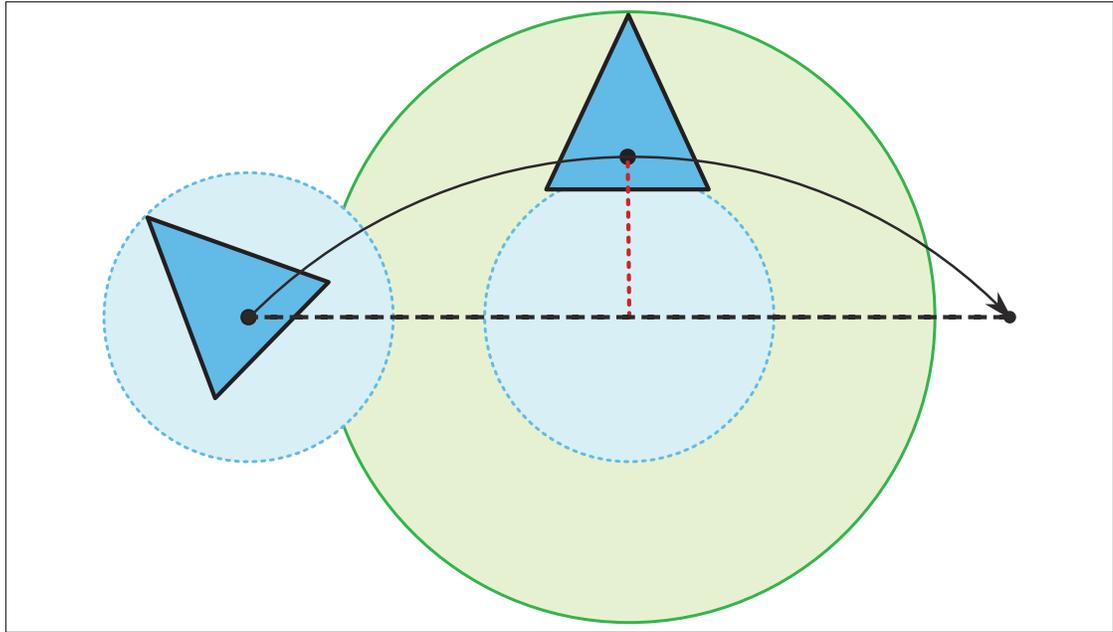


Figure 1.4 A triangle undergoes non-linear motion. Linear sphere tracing, using the triangle’s bounding sphere (blue), cannot capture the difficult motion. The maximum distance (red) between the trajectory and its linear approximation is added to the radius of the bounding sphere. The augmented bounding sphere (green) covers non-linear motions, even along the linear approximation of the trajectory

sphere, on the real trajectory in the middle of the time interval, and the midpoint of the line between the start and end position, as shown in Figure 1.4.

1.6 Results

We evaluate our proposed method for speed, accuracy and robustness using a series of challenging scenarios involving triangles-SDF collision that requires CCD for success. The scenarios use meshes with up to hundreds of thousands of triangles and SDFs that challenge gradient methods with gradient discontinuities, sharp features, and multiple boundaries along triangles trajectories. All scenarios use the meter and second as base units, and no normalization is applied. A measure of 1 cm is 0.01 m in our simulators. We compare our method against the discrete collision detection method proposed by Macklin *et al.* (2020), which gives a baseline for comparing performance and robustness.

A custom rigid body and cloth simulator, written in C++, are used for all experiments. We use Discregrid by Koschier *et al.* (2017b), modified to support both single and double floating point precision, to generate and discretize SDFs from meshes. All examples were run on a consumer-grade laptop computer with an Intel i9 processor and an RTX 4070 laptop GPU, on a single CPU thread. Unless specified otherwise, we run our simulations at 60 time steps per simulated second.

We refer to the Frank-Wolfe with Golden Section Search (FWGSS) as our approach, but we also developed the Frank-Wolfe (FW) method for triangle-SDF CCD, which is included in relevant comparisons.

1.6.1 Contact detection

We first compare the quality of the collision detection to other approaches. In all the examples, our method cheaply offers robust collision detection.

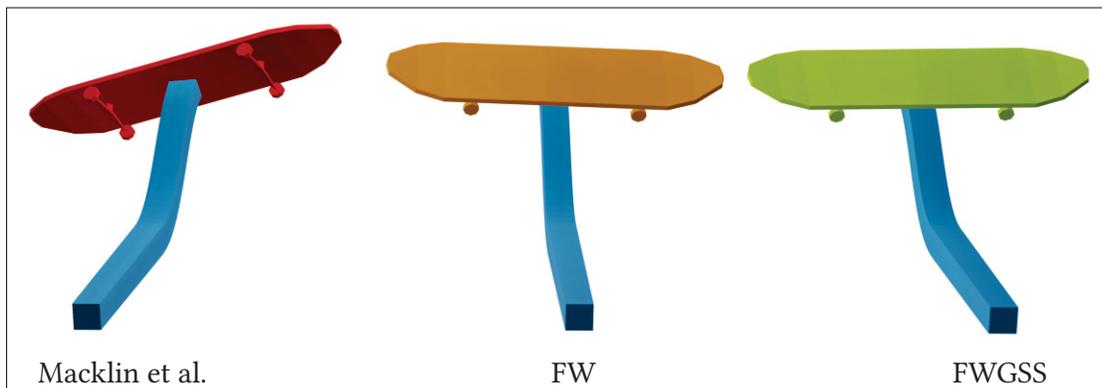


Figure 1.5 A skateboard slides down a ramp using the Macklin *et al.* (2020), FW, and FWGSS methods. The FWGSS method allows the skateboard to settle at the tip of the rail, while the others are unstable

As seen in Figure 1.5, a poorly chosen first time of impact can have drastic influences on the scene, making the discrete collision skateboard in red fall from the ramp.

In Figure 1.6 and Figure 1.7, we compare Macklin *et al.* (2020) and our method. On first impact, discrete collision detection fails to detect sharp contacts like the ears of the bunny, while our

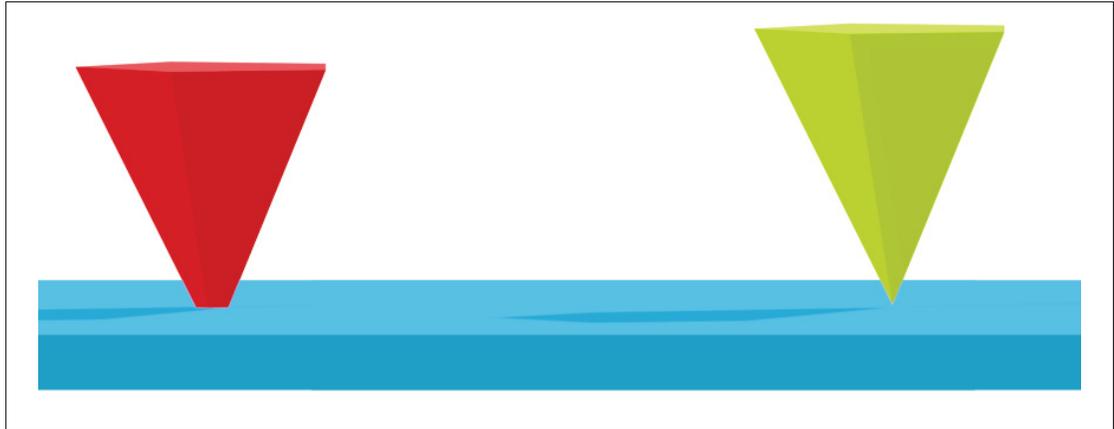


Figure 1.6 Two coarse pyramids fall onto an analytical SDF box. The simulation with discrete collision detection (left) shows clear interpenetration, while our method on the right detects sharp contacts in time

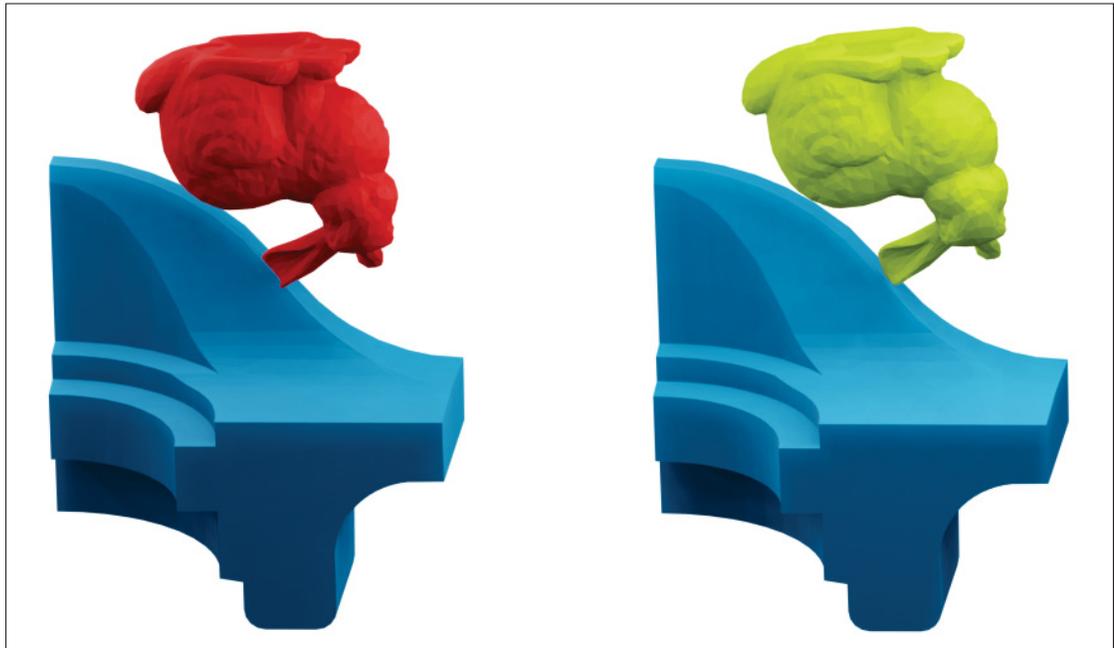


Figure 1.7 A bunny simulated with Macklin *et al.* (2020) and one with FWGSS fall onto an SDF shape. The discrete collision detection bunny clips through the SDF while our method shows no interpenetration

approach accurately captures the fine details of the collision. In Figure 1.8, we present a hoop made with Boolean operators to have exactly a one by one hole matching the radius of one for

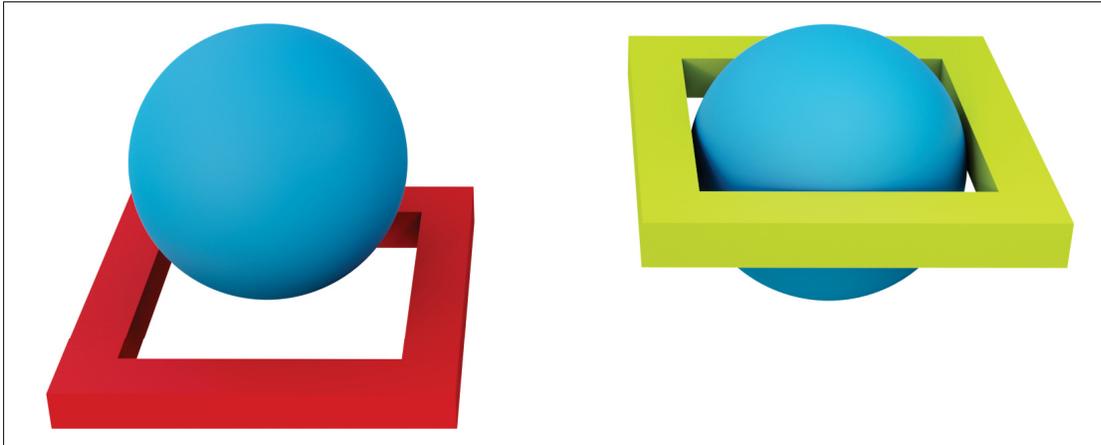


Figure 1.8 A triangle mesh hoop falls over an analytical SDF sphere. The hole width exactly matches the sphere diameter. Discrete collision detection (Macklin *et al.*, 2020) (left) misses the collision event, while our method (right) detects it

the analytical SDF sphere. Our algorithms successfully capture this challenging contact, while the previous state-of-the-art completely misses the collision. This is because the opportunity for contacts is small and can easily be traversed in a single time step.

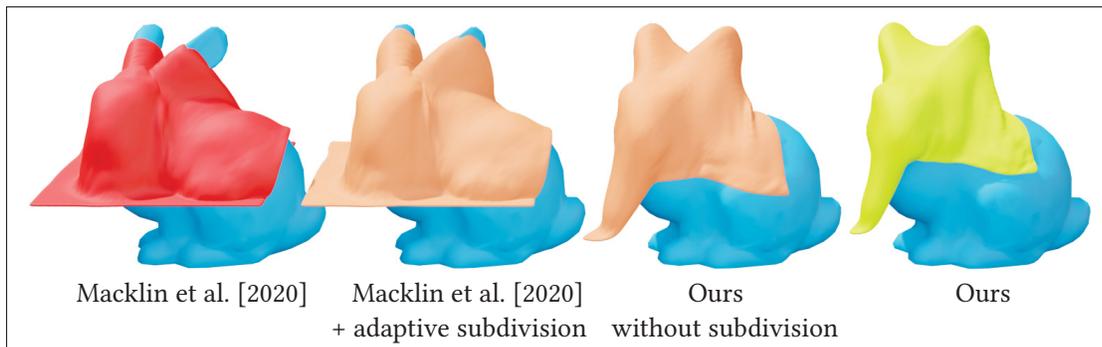


Figure 1.9 A mass-spring cloth with triangle mesh geometry falls on a bunny with SDF geometry. Only our method using adaptive triangle subdivision produces a penetration-free simulation

In Figure 1.9, we compare the effects with and without adaptive triangle subdivision using a cloth simulation, where any interpenetration is instantly apparent due to the 2D nature of the mesh. We simulate two versions of the scene, one with a time step of 1 ms and one with a large

step of 10 ms. The subdivision method already significantly reduces the interpenetration using discrete collision detection.

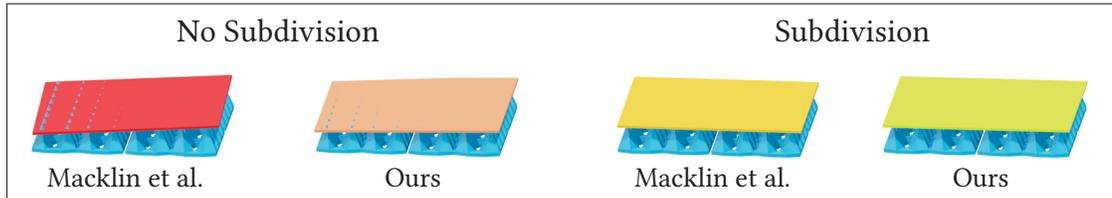


Figure 1.10 A box modeled using a coarse triangle mesh falls on a spiky SDF. Without adaptive triangle subdivision, simulations using discrete and continuous collision both exhibit jitter and severe penetration

Likewise, our method without subdivision features some small level of interpenetration at a triangle level because only a single contact point is generated, while the subdivision adds key contact points. In Figure 1.10, we show the impact of subdivision over time, where both discrete and continuous collision detection sink into the SDF without it. Subdivision allows the low-poly box to remain stable on top of the spikes, leading to more realistic simulations when used in more complex scenes. Figure 1.11 shows the first frame after a triangle falls flat against a box SDF. Without subdivision, a single contact constraint is generated, which is not sufficient to avoid penetration. With adaptive subdivision, three contact constraints are generated, allowing the triangle to remain stable.

As for large time steps, like the one frame 360° spin of Figure 1.12, our method catches them, while discrete collision detection has no way of detecting the contact, staying endlessly in the same position. The collision for a full 360° rotation in a single step is accurately captured with our method at the right time of impact, while discrete collision detection fails to detect the collision entirely, doing a full degree rotation in a time step, and for every time step.

1.6.2 Ground truth comparisons

To further validate the robustness of our method, we generated a data set containing 2048 randomized starting configurations (position and velocity) along with ground truth (GT) collision

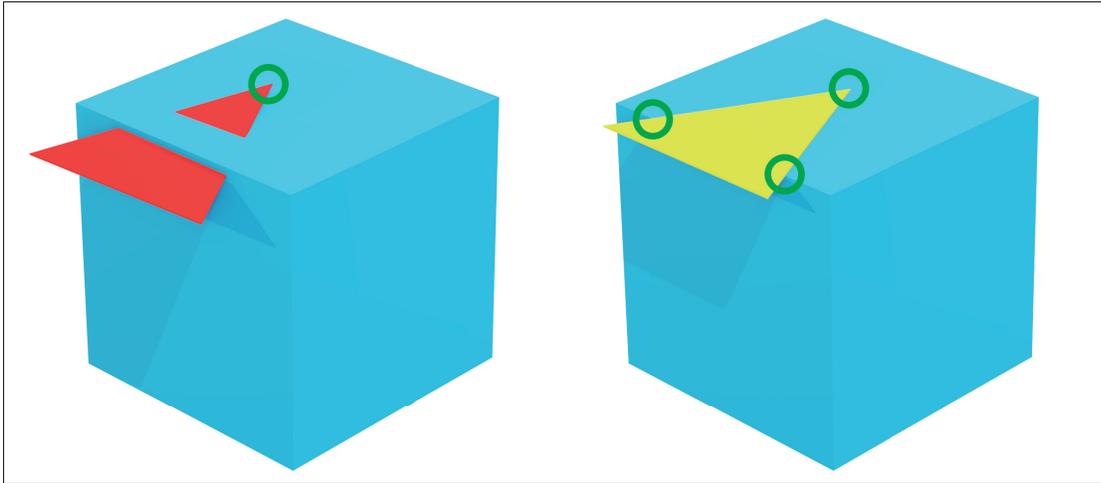


Figure 1.11 Left: a single contact point cannot prevent the triangle to pass through the box. Right: we use adaptive subdivision to generate additional contact points, which provide stable face-face contact

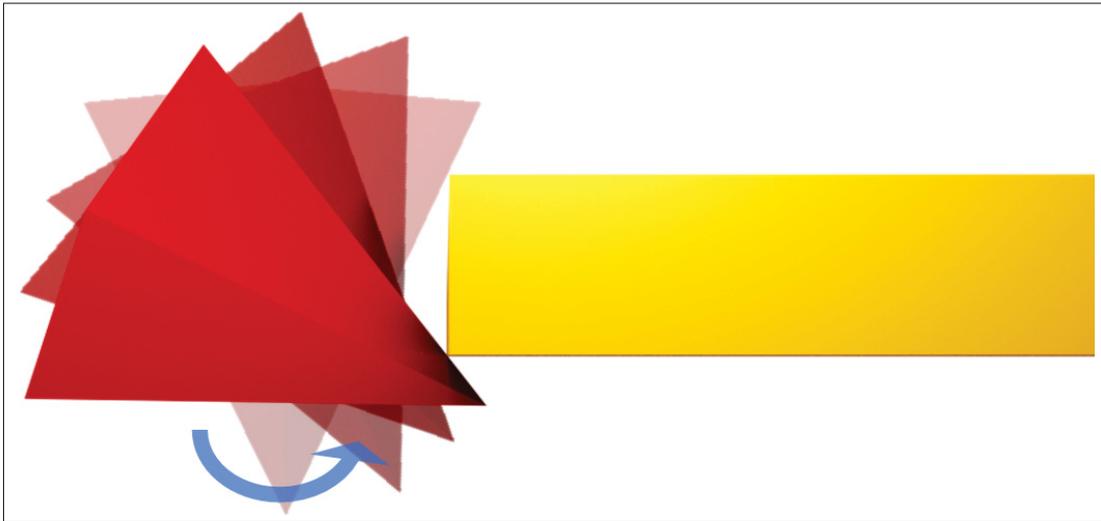


Figure 1.12 A triangle mesh pyramid rotates 360° in a single time step. Our CCD method detects the collision on the analytical SDF box, preventing intersection

solutions. The data set comprises the torus SDF from Figure 1.2, and two meshes, each with 1024 random configurations. The first mesh consists of a single triangle, and the second is the shuriken with 888 triangles from Figure 1.1. The torus SDF is 30 cm wide, and the triangle and shuriken are roughly 10 cm wide. We note that no normalization is applied. The meter and

Tableau 1.1 Time of impact comparison of our method vs ground truth

	$\Delta\text{TOI (s)}$				
	Min	Max	Avg	Median	σ
Triangle	-1.88E-05	+1.18E-04	+6.87E-07	-1.00E-12	+7.25E-06
Shuriken	-6.47E-06	+7.78E-04	+7.29E-07	-3.00E-12	+2.43E-05
All	-1.88E-05	+7.78E-04	+6.09E-07	-1.00E-12	+1.50E-05

Tableau 1.2 Travel distance comparison of our method vs ground truth

	$\Delta\vec{x} \text{ (m)}$				
	Min	Max	Avg	Median	σ
Triangle	-3.64E-04	+1.76E-03	+1.26E-05	-1.79E-11	+1.23E-04
Shuriken	-2.08E-04	+1.90E-03	+1.22E-06	-5.04E-11	+6.23E-05
All	-3.64E-04	+1.90E-03	+6.98E-06	-3.27E-11	+9.25E-05

second are always the base units. The GT solutions are obtained from a brute force approach by advancing through the collision time interval in steps of 1E-06 s. Each step tests each triangle for collision using a thousand iterations of the DCD method by Macklin *et al.* (2020). The method backtracks using a thousand iterations of GSS every time a collision is detected, and repeats the brute force DCD over each triangle until the first time of impact is found.

Our FWGSS method uses a bounding sphere hierarchy (BSH) acceleration structure, but the GT method doesn't. The tests are generated from a combination of random mesh positions, speed, and angular velocity, but we ensure that the mesh is on a collision course with the SDF.

Statistical results of the tests are presented in Tableau 1.1 and Tableau 1.2. They show that our FWGSS method finds the correct time of impact with an average precision of roughly 0.5 μs . As multiple points on the same triangle may be valid contact points, we cannot directly compare the distance between those found by the two methods. Instead, we compare the distance between the stopping point of the mesh between the FWGSS and GT results. We see that FWGSS gives a distance of < 0.5 mm compared to GT, with an average in the order of μm . Of the 2048 tests, FWGSS gave approximately 0.98% false positives for at least one triangle. The false positives are likely due to a slight penetration when FWGSS finds a time of impact a bit later than GT.

FWGSS missed at least one triangle in contact for approximately 0.94% of tests. However, all tests for which FWGSS gave false negatives still found a time of impact very close to the GT method. False positive and false negative collisions occurred only for the shuriken, not the single triangle tests.

Tableau 1.3 Performance scalability with regard to triangle count

Triangles	Method	Total* (ms)	Tri. mean** (μs)	Tri. median** (μs)	Tri. σ ** (μs)
1K	DCD	10.41	25.99	25.37	3.69
10K	DCD	63.94	28.57	28.59	2.52
100K	DCD	306.66	27.34	27.10	3.52
1K	FWGSS	6.47	6.29	0.89	25.25
10K	FWGSS	24.18	3.31	0.86	5.91
100K	FWGSS	86.85	0.96	0.65	3.39

* Total physics computation time up to the time of impact.

** Per triangle statistics.

1.6.3 Performance

We evaluate the wall clock time of our method against other approaches by monitoring the computation time of the collision detection and adaptive subdivision. We only compare each technique up to the first time of impact and ignore subsequent simulation frames, as they are influenced by previous results and the different techniques quickly lose a common comparison basis. Other factors unrelated to the compared techniques, like constraint solve or collision response, would also otherwise begin to influence the results.

Tableau 1.4 Performance timings

Scene	Method	Total (ms)	Runtime %	Tri. Tested	Mean (μ s)	Min (μ s)	Max (μ s)	SD (μ s)
Skate	DCD	2.33	94.72	76	30.69	24.99	80.17	13.94
Skate	FW	1.35	32.22	153	8.85	1.71	42.35	3.7
Skate	FWGSS	3.91	57.42	153	25.57	7.14	94.7	24.72
Bunny	DCD	5.09	98.64	89	57.18	26.29	113.5	29.06
Bunny	FW	0.52	27.81	134	3.9	0.39	25.01	4.7
Bunny	FWGSS	1.06	52.74	134	7.93	0.40	36.29	4.94
Shuriken	DCD	1.43	97.28	21	67.97	28.05	86.74	16.79
Shuriken	FWGSS	0.4	74.07	36	11.13	0.57	39.95	7.71

In Tableau 1.3, we evaluate the scalability of the algorithms as the number of triangle increases. This scenario uses the SDF spike grid from Figure 1.10, and the armadillo mesh at different resolutions. We note that, normally, it would make a lot more sense for the spike grid to be a triangle mesh (coarse and sharp), and for the armadillo to be an SDF (dense and relatively smooth). We deliberately chose this configuration because the spike grid SDF is extremely challenging for gradient methods. Both methods are accelerated using a bounding sphere hierarchy to cull narrow-phase tests. Our method of choice (FWGSS) scales better than the DCD method by Macklin *et al.* (2020). The total run time is drastically improved with CCD preventing penetration, whereas DCD runs for a full extra simulation frame before having to handle a lot more intersections. Furthermore, the time interval for CCD shrinks whenever an earlier contact is found by our method. This makes subsequent triangle tests more efficient, and explains the larger standard deviation of our method. Per-triangle statistics demonstrate that our method becomes more efficient as the number of triangle increases. This is also due to the shrinking time interval, which provides a simpler problem space for subsequent narrow-phase triangle tests. We note that the average time per triangle is measured in microseconds.

In Tableau 1.4, we compare the time in milliseconds taken by the triangle-SDF collision detection algorithms and triangle subdivision, aggregated over the simulation frames for each scene. We include per-triangle timing statistics in microseconds to provide better insight about the efficiency of each algorithm. In all scenarios, our method of choice (FWGSS) is more accurate, and it outperforms the Macklin *et al.* (2020) DCD method. This holds true, even if the number of triangles tested over multiple time steps is higher with our methods, because false positives provided by the broad-phase tests are quickly resolved by the narrow-phase tests. The increased number of tests is due to the relative inaccuracy of our acceleration structure (see Section 1.5.6) compared to structures suited for DCD.

1.7 Discussion and limitations

Knowing the lower bound of the minimized function may allow our method to outperform Macklin *et al.* (2020) while enabling CCD by providing the time of impact, but contact handling

can be (very) expensive. Our collision detection method is fast, and compatible with efficient coarse meshes thanks to the adaptive triangle subdivision. It is also time-integration agnostic, and can be used with any type of contact handling. In real-time interactive simulations, where performance is prioritized over accuracy, our method may only be called once per frame for each potential collision pair requiring CCD. In this case, performance and robustness will be improved by our technique, and real-time performance is easily maintained. If higher accuracy is required, and interpenetration must absolutely not happen, then our method may need to be called a lot more due to the collision response itself causing more contacts to happen within the time step. Our technique may still improve performance and robustness in these simulations but cannot guarantee real-time performance.

While our method can accurately find the first time and point of impact, it cannot guarantee an absolutely interpenetration-free simulation and maintain its real-time performance capacity without at least a small collision margin. We found that in practice, a margin of about one millimetre is sufficient, but accuracy is otherwise limited by floating point precision in the numerical solver and by limiting the number of iterations to maintain real-time performance.

In the discrete case, we get a cheap, relatively tight bound on the solution space, a static bounding sphere around the triangle. In the CCD case, we cannot accurately and cheaply estimate the full volume a triangle would pass through. Our ray-marching approach is cheap and avoids missing contacts, but requires a wide padding. Therefore, more triangles can pass the filter of the broad phase, even if they do not end up colliding.

Our method depends on the quality of the SDF, but inexact or approximate SDFs are common. We even use approximate polynomial forms in multiple of our scenarios. Imperfect SDFs affect the efficiency and guarantees our method can offer. For example, the more discontinuities there are, the harder it is for gradient methods to converge. Floating point error is also a concern for numerical methods, and may be exacerbated by SDFs with low precision or gradients with non-unit norms. High relative velocities and longer time steps also contribute to this issue. In practice, though, we found that our method still fares well on difficult scenarios with

discontinuities, high velocities and rotations, and even difficult SDFs for shapes which would normally be represented by triangle meshes (like the spike grid in Figure 1.10).

Existing benchmarks for collision detection, such as (Wang *et al.*, 2020), compare continuous triangle-triangle collision detection methods against ground truth solution. Although it is technically possible to generate a triangle mesh from an SDF using methods such as marching cubes (Lorensen & Cline, 1987), the resulting geometry only provides a rough discrete approximation of the SDF. Standard triangle-triangle CCD would thus only be able to provide results relative to this approximation, so we exclude them as a baseline for ground-truth comparison. Not to mention that triangle-triangle collision detection brings additional challenges (Erleben, 2018). We therefore opt to compare our results to the ground-truth obtained using the brute force approach described in Section 1.6.2.

1.8 Conclusion

Our method addresses challenges in collision detection by reformulating the problem as a spatio-temporal local optimization. Through the introduction of adaptive triangle subdivision and the development of an enhanced problem-specific Frank-Wolfe method with line search, our method leads to robust and accurate collision detection for both coarse and complex geometries. This reduces issues such as tunnelling and missed collisions. Our method offers improved robustness over point sampling methods and outperforms recent triangle-SDF discrete collision detection (DCD) algorithms.

The golden section search is used for the various line searches required by our method. This approach is relatively efficient, and can often find minima more quickly than the version of our algorithm without it. In the future, we would like to study how different line search methods affect the efficiency and accuracy of our method. For instance, it is possible that a simpler bisection method could provide better efficiency.

The current acceleration structure has large bounds for detection. In the future, we would like to improve its efficiency by reusing spatial acceleration structures often used for the discretization

of SDFs. Likewise, we could potentially improve the convergence of Frank-Wolfe methods using momentum-based approaches similar to Montaut, Le Lidec, Petrik, Sivic & Carpentier (2024). No benchmark exists for triangle-SDF collisions. We evaluated our method with a series of extreme hand-crafted scenarios and random variations. However, it would be beneficial for the research community to develop standard benchmarks to provide rigorous and varied tests for triangle-SDF collisions and related edge cases. We note that our method could potentially be compatible with the incremental potential contacts (IPC) framework (Li *et al.*, 2020), which would completely rid the simulations of the possibility of interpenetration past the first contact. We are excited to see the adoption of fast continuous collision detection in industry applications, and hope that efficient methods using SDFs will enable penetration-free simulations and further advance the field of real-time physics-based animations.

1.9 Acknowledgements

This project was partly funded by, and we acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), under Application Id: 585438-2023. The authors would like to thank Thomas Friedlaender, Paul Kry, and Kenny Erleben for providing valuable feedback and discussions. The authors would also like to thank Nicholas Sharp for Polyscope (Sharp *et al.*, 2019), Inigo Quilez for the invaluable resources on SDFs, and the Stanford University Computer Graphics Laboratory for the Bunny, Armadillo, and XYZRGB dragon models, which helped in developing tools and simulations for our experiments.

CONCLUSION ET RECOMMANDATIONS

Ce travail de recherche avait comme objectif d'améliorer la performance et la qualité atteignables en temps réel dans les simulations basées sur la physique. Ces améliorations sont essentielles, car le potentiel et l'utilité de ces simulations sont présentement limités par notre capacité à simuler divers phénomènes physiques de manière rapide et réaliste. Notre analyse des méthodes de pointe et de la littérature existante souligne que rapidité et qualité sont souvent opposées ; l'une nécessite des compromis au niveau de l'autre. Elle présente aussi l'intérêt d'utiliser les SDFs comme représentation efficace des formes simulées. Les SDFs permettent à la technique de Macklin *et al.* (2020) d'obtenir des performances impressionnantes pour la DCD. C'est donc dans cette optique que s'inscrivent nos contributions.

Une reformulation du problème de détection des collisions en une optimisation spatio-temporelle permet à notre méthode d'adresser les défis qui y sont liés. L'introduction d'une méthode de Frank-Wolfe avec recherche linéaire, améliorée et spécifique au problème, permet la détection précise et robuste des collisions en fournissant efficacement le TOI pour la CCD triangle-SDF. En plus d'offrir une meilleure robustesse que les méthodes à base d'échantillonnages denses, les résultats démontrent que notre méthode de CCD maintient de meilleures performances que les méthodes de pointe pour la DCD triangle-SDF, même lorsque le nombre de triangles augmente drastiquement. De plus, notre méthode de sous-division barycentrique adaptative des triangles permet aussi l'utilisation de géométrie réduite efficace. Bref, notre méthode réduit efficacement les problèmes d'intersections et les collisions manquées, tout en offrant de meilleures performances, et ce, autant avec une géométrie dense que réduite.

2.1 Futurs travaux et recommandations

Notre méthode de détection de collisions améliore la qualité et les performances atteignables en temps réel dans les simulations physiques. Toutefois, la gestion des collisions peut s'avérer très

coûteuse. Les collisions s'enchaînent souvent rapidement lorsque les formes entrent en collision et tentent de se mettre au repos. Notre méthode, bien que très efficace, peut donc être appelée à résoudre des millions de problèmes de CCD par milliseconde de simulation pour résoudre les nouvelles collisions générées continuellement par la gestion des collisions précédentes. La simulation en temps réel nécessite donc de limiter le nombre de résolutions de la CCD par intervalle de simulation. Par conséquent, bien que notre méthode améliore la performance et la robustesse des simulations, il lui est impossible de garantir, à elle seule, l'absence totale d'interpénétration et les performances nécessaires à la simulation en temps réel.

Notre méthode est la première à permettre la CCD robuste entre une SDF et la surface entière d'un triangle. Elle ouvre donc la porte à d'amples nouvelles avenues pour améliorer d'avantage la performance et la qualité des simulations physiques en temps réel. Par exemple, notre méthode d'accélération a besoin de larges bornes pour tenir compte des rotations dans le temps. Une approche potentielle pour améliorer son efficacité serait de réutiliser les structures spatiales communément utilisées pour discrétiser les SDFs. De plus, notre méthode FWGSS utilise la méthode du nombre d'or (GSS) pour diverses recherches linéaires. Cette approche est relativement efficace, mais il serait intéressant d'étudier comment d'autres types de recherches linéaires affecteraient la convergence et l'efficacité de notre méthode. Par exemple, est-ce qu'une méthode par bisection fournirait de meilleures performances, ou est-ce que cela diminuerait la capacité à éviter les minimums locaux ? Autrement, la convergence de notre méthode, basée sur l'algorithme de Frank-Wolfe, serait potentiellement améliorable par une modification tenant compte de la quantité de mouvement, similairement à Montaut *et al.* (2024).

Enfin, aucun banc d'essais libre d'accès n'existe pour la détection de collision entre triangles et SDFs. La détection de collisions triangle-triangle, quant à elle, est bien couverte à cet égard (Wang *et al.*, 2021). Nos méthodes ont été évaluées au travers d'un ensemble de scénarios créés à la main et de scénarios aléatoires générés par ordinateur. Bien que suffisante pour couvrir

les cas difficiles et pour obtenir des données statistiques supplémentaires, cette approche est difficile à mettre en place. Nous sommes donc intéressés par l'étude approfondie des limites des algorithmes existants et des problèmes non résolus en lien avec l'utilisation des SDFs pour la création d'un banc de test standard et libre d'accès. En plus d'accélérer la recherche dans le domaine, cela permettrait l'accumulation de meilleures données, et de meilleures comparaisons entre les différentes techniques dans les futurs travaux.

BIBLIOGRAPHIE

- Andrews, S., Erleben, K. & Ferguson, Z. (2022). Contact and friction simulation for computer graphics. *ACM SIGGRAPH 2022 Courses*, (SIGGRAPH '22), 1-172. doi : 10.1145/3532720.3535640.
- Braun, G., Carderera, A., Combettes, C. W., Hassani, H., Karbasi, A., Mokhtari, A. & Pokutta, S. (2022, Novembre). Conditional Gradient Methods.
- Bridson, R., Fedkiw, R. & Anderson, J. (2002). Robust treatment of collisions, contact and friction for cloth animation. *ACM Transactions on Graphics*, 21(3), 594–603. doi : 10.1145/566654.566623.
- Brochu, T., Edwards, E. & Bridson, R. (2012). Efficient geometrically exact continuous collision detection. *ACM Transactions on Graphics*, 31(4), 1-7. doi : 10.1145/2185520.2185592.
- Chen, X., Yu, C., Ni, X., Chu, M., Wang, B. & Chen, B. (2024). A Time-Dependent Inclusion-Based Method for Continuous Collision Detection between Parametric Surfaces. *ACM Transactions on Graphics*, 43(6), 1-11. doi : 10.1145/3687960.
- Ericson, C. (2004). *Real-time collision detection*. USA : CRC Press.
- Erleben, K. (2018). Methodology for Assessing Mesh-Based Contact Point Methods. *ACM Transactions on Graphics*, 37(3), 39 :1–39 :30. doi : 10.1145/3096239.
- Frank, M. & Wolfe, P. (1956). An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3(1-2), 95-110. doi : 10.1002/nav.3800030109.
- Friskien, S. F., Perry, R. N., Rockwood, A. P. & Jones, T. R. (2000). Adaptively sampled distance fields : a general representation of shape for computer graphics. *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, (SIGGRAPH '00), 249–254. doi : 10.1145/344779.344899.
- Fuhrmann, A., Sobotka, G. & Groß, C. (2003). Distance fields for rapid collision detection in physically based modeling. *Proceedings of GraphiCon*, 2003, 58–65.
- Govindaraju, N. K., Kabul, I., Lin, M. C. & Manocha, D. (2006). Fast continuous collision detection among deformable models using graphics processors. *Proceedings of the 12th Eurographics Conference on Virtual Environments*, (EGVE'06), 19–26.
- Guendelman, E., Bridson, R. & Fedkiw, R. (2003). Nonconvex rigid bodies with stacking. *ACM Transactions on Graphics*, 22(3), 871–878. doi : 10.1145/882262.882358.

- Jaggi, M. (2013, Juin). Revisiting Frank-Wolfe : Projection-Free Sparse Convex Optimization. *Proceedings of the 30th International Conference on Machine Learning*, 28(Proceedings of Machine Learning Research), 427–435.
- Jones, M., Baerentzen, J. & Sramek, M. (2006). 3D Distance Fields : A Survey of Techniques and Applications. *IEEE Transactions on Visualization and Computer Graphics*, 12(4), 581–599. doi : 10.1109/TVCG.2006.56.
- Kiefer, J. (1953). Sequential Minimax Search for a Maximum. *SciSpace - Paper*, 4(3), 502–506. doi : 10.1090/S0002-9939-1953-0055639-3.
- Koschier, D., Deul, C., Brand, M. & Bender, J. (2017a). An hp-adaptive discretization algorithm for signed distance field generation. *IEEE Transactions on Visualization and Computer Graphics*, 23(10), 2208–2221. doi : 10.1109/TVCG.2017.2730202.
- Koschier, D. et al. (2017b). Discregrid Library. Repéré à <https://github.com/InteractiveComputerGraphics/Discregrid>.
- Li, M., Ferguson, Z., Schneider, T., Langlois, T., Zorin, D., Panozzo, D., Jiang, C. & Kaufman, D. M. (2020). Incremental potential contact : intersection-and inversion-free, large-deformation dynamics. *ACM Transactions on Graphics*, 39(4), 49 :1-49 :20. doi : 10.1145/3386569.3392425.
- Liu, P., Zhang, Y., Wang, H., Yip, M. K., Liu, E. S. & Jin, X. (2024). Real-time collision detection between general SDFs. *Comput. Aided Geom. Des.*, 111(C), 102305 :1-102305 :13. doi : 10.1016/j.cagd.2024.102305.
- Lorensen, W. E. & Cline, H. E. (1987). Marching cubes : A high resolution 3D surface construction algorithm. *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*, (SIGGRAPH '87), 163–169. doi : 10.1145/37401.37422.
- Macklin, M., Erleben, K., Müller, M., Chentanez, N., Jeschke, S. & Corse, Z. (2020). Local Optimization for Robust Signed Distance Field Collision. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 3(1), 1-17. doi : 10.1145/3384538.
- McNeely, W. A., Puterbaugh, K. D. & Troy, J. J. (1999). Six degree-of-freedom haptic rendering using voxel sampling. *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, (SIGGRAPH '99), 401–408. doi : 10.1145/311535.311600.
- McNeely, W. A., Puterbaugh, K. D. & Troy, J. J. (2005). Advances in voxel-based 6-DOF haptic rendering. *ACM SIGGRAPH 2005 Courses*, (SIGGRAPH '05), 50–es. doi : 10.1145/1198555.1198606.

- Montaut, L., Le Lidec, Q., Petrik, V., Sivic, J. & Carpentier, J. (2024). GJK++ : Leveraging Acceleration Methods for Faster Collision Detection. *IEEE Transactions on Robotics*, 40, 2564-2581. doi : 10.1109/TRO.2024.3386370.
- Naylor, B. (1998). A Tutorial On Binary Space Partitioning Trees. *Computer Games Developer Conference Proceedings*, pp. 433–457.
- Nie, Q., Zhao, Y., Xu, L. & Li, B. (2020). A Survey of Continuous Collision Detection. *2020 2nd International Conference on Information Technology and Computer Application (ITCA)*, pp. 252-257. doi : 10.1109/ITCA52113.2020.00061.
- Park, J. J., Florence, P., Straub, J., Newcombe, R. & Lovegrove, S. (2019, Juin). DeepSDF : Learning Continuous Signed Distance Functions for Shape Representation. *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 165-174. doi : 10.1109/CVPR.2019.00025.
- Pelletier-Guénette, J., Mercier-Aubin, A. & Andrews, S. (2025). Real-Time Triangle-SDF Continuous Collision Detection. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 8(4), 49 :1–49 :22. doi : 10.1145/3747862.
- Quinlan, S. (1994). Efficient distance computation between non-convex objects. *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, pp. 3324-3329 vol.4. doi : 10.1109/ROBOT.1994.351059.
- Redon, S., Kheddar, A. & Coquillart, S. (2002). Fast Continuous Collision Detection between Rigid Bodies. *Computer Graphics Forum*, 21(3), 279-287. doi : 10.1111/1467-8659.t01-1-00587.
- Sharp, N. et al. [www.polyscope.run]. (2019). Polyscope.
- Tang, M., Tong, R., Wang, Z. & Manocha, D. (2014). Fast and exact continuous collision detection with Bernstein sign classification. *ACM Transactions on Graphics*, 33(6), 1-8. doi : 10.1145/2661229.2661237.
- Teschner, M., Kimmerle, S., Heidelberger, B., Zachmann, G., Raghupathi, L., Fuhrmann, A., Cani, M.-P., Faure, F., Magnenat-Thalmann, N., Strasser, W. & Volino, P. (2005). Collision Detection for Deformable Objects. *Computer Graphics Forum*, 24(1), 61-81. doi : 10.1111/j.1467-8659.2005.00829.x.
- Wang, B., Ferguson, Z., Schneider, T., Jiang, X., Attene, M. & Panozzo, D. (2020). A Large-scale Benchmark and an Inclusion-based Algorithm for Continuous Collision Detection. *ACM Transactions on Graphics*, 40(5), 188 :1–188 :16. doi : 10.1145/3460775.

- Wang, B., Ferguson, Z., Schneider, T., Jiang, X., Attene, M. & Panozzo, D. (2021). A Large-scale Benchmark and an Inclusion-based Algorithm for Continuous Collision Detection. *ACM Transactions on Graphics*, 40(5), 1-16. doi : 10.1145/3460775.
- Wang, B., Ferguson, Z., Jiang, X., Attene, M., Panozzo, D. & Schneider, T. (2022). Fast and Exact Root Parity for Continuous Collision Detection. *Computer Graphics Forum*, 41(2), 355-363. doi : 10.1111/cgf.14479.
- Wang, J., Zhang, T., Zhang, Q., Zeng, C., Yu, J., Xu, C., Xu, L. & Gao, F. (2024). Implicit Swept Volume SDF : Enabling Continuous Collision-Free Trajectory Generation for Arbitrary Shapes. *ACM Transactions on Graphics*, 43(4), 1-14. doi : 10.1145/3658181.
- Xu, H. & Barbič, J. (2017). 6-DoF Haptic Rendering Using Continuous Collision Detection between Points and Signed Distance Fields. *IEEE Transactions on Haptics*, 10(2), 151-161. doi : 10.1109/TOH.2016.2613872.
- Zhang, P., Marschner, Z., Solomon, J. & Tamstorf, R. (2023a). Sum-of-Squares Collision Detection for Curved Shapes and Paths. *ACM SIGGRAPH 2023 Conference Proceedings*, (SIGGRAPH '23), 1-11. doi : 10.1145/3588432.3591507.
- Zhang, T., Wang, J., Xu, C., Gao, A. & Gao, F. (2023b). Continuous Implicit SDF Based Any-Shape Robot Trajectory Optimization. *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 282-289. doi : 10.1109/IROS55552.2023.10342104.