

Codimensional Incremental Potential Contact

MINCHEN LI, University of California, Los Angeles, University of Pennsylvania, & Adobe Research

DANNY M. KAUFMAN, Adobe Research

CHENFANFU JIANG, University of California, Los Angeles & University of Pennsylvania




Fig. 1. **Table cloth pull:** Codimensional Incremental Potential Contact (C-IPC) enables high-rate time stepping of codimensional models with intersection-free, accurate, controllable strain-limiting, thickness modeling, and frictional contact. Here, stepping shell and volumetric FE models at $h = 0.01s$, a square cloth is laid on a table with heavy, stiff volumetric dinnerware (a). C-IPC stably responds to rapid pulling at varying speeds (b-f) without stretching beyond limits while accurately resolving tight contacts between table edges and dinnerware. With a slower pull in (b) of $0.5m/s$ nearly all objects pull off due to friction. As pull speed increase at $1m/s$ (c) and $2m/s$ (d), less dinnerware fall off with some just pulled a bit. In (e) and (f) pulling at higher speeds $4m/s$ or $8m/s$ leaves the setting on table with rapid acceleration overcoming friction. Here too, combined strain-limiting and resolution of elastodynamics generates detailed wrinkling behaviors in the cloth for these faster pulls as it flies off the table.

We extend the incremental potential contact (IPC) model [Li et al. 2020a] for contacting elastodynamics to resolve systems composed of codimensional degrees-of-freedom in arbitrary combination. This enables a unified, interpenetration-free, robust, and stable simulation framework that couples codimension-0,1,2, and 3 geometries seamlessly with frictional contact. Extending the IPC model to thin structures poses new challenges in computing strain, modeling thickness and determining collisions. To address these challenges we propose three corresponding contributions. First, we introduce a C^2 constitutive barrier model that directly enforces strain limiting as an energy potential while preserving rest state. This provides energetically-consistent strain limiting models (both isotropic and anisotropic) for cloth that enable strict satisfaction of strain-limit inequalities with direct coupling

Authors' addresses: Minchen Li, University of California, Los Angeles, University of Pennsylvania, & Adobe Research, minchernl@gmail.com; Danny M. Kaufman, Adobe Research, danny.kaufman@gmail.com; Chenfanfu Jiang, University of California, Los Angeles & University of Pennsylvania, chenfanfu.jiang@gmail.com.

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

© 2021 Association for Computing Machinery.

0730-0301/2021/8-ART170 \$15.00

<https://doi.org/10.1145/3450626.3459767>

to both elastodynamics and contact via minimization of the incremental potential. Second, to capture the geometric thickness of codimensional domains we extend the IPC model to directly enforce distance offsets. Our treatment imposes a strict guarantee that mid-surfaces (respectively mid-lines) of shells (respectively rods) will not move closer than applied thickness values, even as these thicknesses become characteristically small. This enables us to account for thickness in the contact behavior of codimensional structures and so robustly capture challenging contacting geometries; a number of which, to our knowledge, have not been simulated before. Third, codimensional models, especially with modeled thickness, mandate strict accuracy requirements that pose a severe challenge to all existing continuous collision detection (CCD) methods. To address these limitations we develop a new, efficient, simple-to-implement *additive CCD* (ACCD) method that applies conservative advancement [Mirtich 1996; Zhang et al. 2006] to iteratively refine a lower bound for deforming primitives, converging to time of impact. In combination these contributions enable *codimensional IPC* (C-IPC). We perform extensive benchmark experiments to validate the efficacy of our method in capturing intricate behaviors of thin-structure contact and resulting bulk effects. In our experiments C-IPC obtains feasible, convergent, and so artifact-free solutions for all time steps, across all tested examples – producing robust simulations. We test C-IPC across extreme deformations, large time steps, and exceedingly close contact over all possible pairings of codimensional domains. Finally, with our strain-limit model, we confirm C-IPC guarantees non-intersection and strain-limit satisfaction

for all reasonable (and well below – verified down to 0.1%) strain limits throughout all time steps.

CCS Concepts: • **Computing methodologies** → **Physical simulation**.

Additional Key Words and Phrases: Strain Limiting, Contact Mechanics, Mixed-Dimensional Elastodynamics, Constrained Optimization

ACM Reference Format:

Minchen Li, Danny M. Kaufman, and Chenfanfu Jiang. 2021. Codimensional Incremental Potential Contact. *ACM Trans. Graph.* 40, 4, Article 170 (August 2021), 24 pages. <https://doi.org/10.1145/3450626.3459767>

1 INTRODUCTION

Thin materials like cloth, hair and sand are everywhere. Simulating them has long remained a critical task in computational modeling and animation. Each such solid is generally best modeled by a corresponding reduced codimensional formulation; respectively by shells, rods and particles. Doing so enables improved efficiency with many less simulated degrees-of-freedom (DOF). Likewise it mitigates the numerical conditioning and severe convergence challenges posed by *shear locking* that we would otherwise encounter if simulating thin materials volumetrically [Yang et al. 2000]. But while we can well-capture the *elastic* behavior of individual thin materials with codimensional models, accurately and consistently simulating their collective and contacting behavior on practical size meshes, with real-world materials and conditions remains challenged. We enumerate these challenges here.

First and foremost simulations should remain *intersection free* at all times. Steps that permit even slight intersections of a codimensional geometry can and will produce tangles that simulations can not recover from.

Second, we require energetically consistent, controllable and accurate strain-limiting. While codimensional discretizations mitigate shear locking they can still suffer from severe *membrane locking*. For example, shell models continue to suffer severe numerical stiffening in bending modes for cloth materials unless midsurface meshes apply extremely high (and generally impractical size) resolutions [Quaglino 2012]. Here higher-order elements can also help further reduce locking artifacts but they add prohibitive computational cost. Strain-limiting allows softer cloth materials coupled with tight limits to avoid these artifacts while still recovering proper cloth response. However, if strain-limiting is not accurately and consistently applied, and carefully coupled to elastodynamics, it often produces more numerical artifacts than it fixes.

Third, geometric thickness must be captured in contact. While thin structures typically have small relative thicknesses, whose *elastic* behavior can be indirectly captured by codimensional DOFs, correctly and directly modeling their thickness in contact is crucial for capturing accurate geometries and bulk behaviors. Consider, for example, a deck of cards. Each card has a nearly imperceptible thickness when considered individually. However, when stacked, their combined thickness is large and clear, and can not be ignored. While contact-processing strategies often introduce thickness parameters, they are generally applied




Fig. 2. **Twisted cylinder.** Testing thickness modeling, contact resolution and buckling under extreme and increasing stress we start twisting a 1m-wide cloth cylinder with 88K nodes, a thickness offset of 1.5mm and \hat{d} of 1mm. Time-stepped at $h = 0.04s$, C-IPC quickly obtains interesting folds and soon after forms a thick central cylinder of wound cloth supported by C-IPC’s finite thickness offset (upper left). In contrast, (upper right) without C-IPC’s thickness offset the correct geometry can not form (nor can it later capture the material’s buckling behavior) further emphasizing the importance of consistent thickness modeling for shells. After 32.96s of further twisting, our offset-thickened cloth continues to support complex contact-driven behaviors including the final buckled geometry (bottom); see our supplemental video for the detailed trajectory.

as a nonphysical failsafe to mitigate collision-processing inaccuracies. As analyzed in Li et al. [2020a], these thicknesses can not be consistently enforced and moreover must be changed heuristically per scene and example (e.g., based on collision speeds) to avoid simulation failures.

Fourth, all codimensions (including volumes) should be seamlessly simulated in a common framework without distinction nor special casing. Accurate frictional contact forces should directly couple interactions between all geometric types irrespective of how close the contact or how thin the modeled thicknesses are.

Fifth, contact modeling between codimensional models with small *but finite* thicknesses challenges all collision-detection routines. Most critically we see unacceptable failures and/or inefficiencies in all available CCD methods and codes. We see this both for standard floating point root-finding CCD methods as well as more recent developments in exact CCD methods. These issues are attributable to the well-known numerical sensitivity of the challenging underlying root-finding problems posed. Here, most specifically, we see that IPC is reliant on a bounded guarantee of non-intersection for all CCD evaluations. Large enough conservative bounds can help [Li et al. 2020a] account for these inaccuracies, ensuring non-intersection, but they do so at the cost of progress and so can even stall convergence altogether when it comes to codimensional DOF.

Finally, simulations should not artificially snag in contact on sharp and codimensional geometries and should be able to robustly solve time-step problems to convergence across variations in step size, scene, materials, contact and boundary conditions.

Here we propose a method that, to our knowledge, is the first to directly resolve all six above challenges with a consistent and reliable solution for frictionally contacting dynamics. Our method guarantees strict satisfaction of fully coupled strain-limits, with intersection-free trajectories, and controllable geometric thickness

resolution (at thin material scales) independent of time step size and severity of contact conditions.

1.1 Contributions

Specifically, to address these challenges we extend the incremental potential contact model (IPC) [Li et al. 2020a] for contacting elastodynamics to resolve systems composed of arbitrary combinations of codimensional DOF with three critical components.

Constitutive strain limiting. We introduce a C^2 constitutive barrier model that directly enforces strain limiting as an energy potential while preserving rest state. This provides energetically-consistent strain limiting models (both isotropic and anisotropic) for cloth that enable, for the first time, strict satisfaction of strain-limit inequalities for all iterations and so for all time steps (verified down to 0.1%), while fully and directly coupling to both elastodynamics and frictional contact via minimization of the incremental potential. Thus, as demonstrated in Section 6.3, we avoid artifacts generated by force-splitting errors in traditional strain-limiting methods.

IPC thickness model. To capture the geometric thickness of codimensional domains in contact we extend the IPC model to directly enforce distance offsets. Our treatment imposes a strict guarantee that mid-surfaces (respectively mid-lines and points) of shells (respectively rods and particles) will not move closer than applied thickness values, even as these thicknesses become characteristically small; see Sec. 5. This enables us to account for thickness in the contacting behavior of codimensional structures and so capture challenging contacting effects; a number of which, to our knowledge, have not been simulated before.

Additive CCD. To provide the strict accuracy required of CCD to resolve codimensional models with thickness, we develop a new, efficient and simple-to-implement *additive CCD* (ACCD) method utilizing conservative advancement [Tang et al. 2009; Zhang et al. 2006, 2007]. ACCD iteratively accumulates a lower bound converging to time of impact and is stable for the exceedingly challenging evaluations required. While we most immediately focus here on ACCD’s application within our C-IPC framework, as we show in the following sections, ACCD is exceedingly simple and so easy to implement when compared to all prior CCD routines (exact and floating point) with both improved performance, robustness and guarantees, and so is suitable for easy replacement in all applications where CCD modules are employed.

Together these form the core of codimensional IPC (C-IPC). C-IPC enables unified simulation of all codimensions including elastic volumetric bodies, shells, rods and particles all coupled together via accurately solved contact and friction. Along with these core contributions we test, compare and analyze C-IPC across many new and pre-existing benchmarks. We confirm C-IPC provides tightly controllable geometric thickness behaviors and guarantees feasibility at all time steps. Finally, C-IPC remains intersection-free and strictly satisfies all set strain limits across all computed trajectories, likewise verified in our extensive benchmark testing.

2 RELATED WORK

2.1 Shells and Rods

Beginning with the pioneering work of Terzopoulos et al. [1987] the simulation of codimensional models, especially shells and rods, has remained a focus in computer graphics. Efficiently modeling the complex behaviors of cloth [Baraff and Witkin 1998; Bender et al. 2013; Bridson et al. 2002; Grinspun et al. 2003; Harmon et al. 2009; Li et al. 2020b; Narain et al. 2012; Volino and Thalmann 2000] and hair [Choe et al. 2005; Daviet et al. 2011; Deul et al. 2018; Kaufman et al. 2014; Kugelstadt and Schömer 2016; McAdams et al. 2009; Müller et al. 2012; Selle et al. 2008; Ward and Lin 2003] are now particularly critical tasks across applications.

Pipelines for their simulation most commonly adopt implicit or else linearly implicit time-integration methods [Baraff and Witkin 1998; Bridson et al. 2002; Harmon et al. 2008; Kim 2020; Li et al. 2020b, 2018a; Narain et al. 2012; Otaduy et al. 2009; Tang et al. 2016, 2018] with a diverse array of collision filters and penalties applied to help resolve contact processing.

To accelerate performance, extensions to the GPU [Li et al. 2020b; Schmitt et al. 2013; Tang et al. 2013b, 2016, 2018], fast projection [English and Bridson 2008; Goldenthal et al. 2007] and multigrid [Tamstorf et al. 2015; Wang et al. 2018; Xian et al. 2019] are all being actively explored. While, to improve the fidelity of models, data-driven material estimation [Bhat et al. 2003; Clyde et al. 2017; Miguel et al. 2012], alternate spatial discretizations [Guo et al. 2018; Jiang et al. 2017; Weidner et al. 2018] and even hybrid models combining yarn and shell simulation [Casafranca et al. 2020; Sperl et al. 2020] are being investigated. Likewise, the differentiability of cloth simulation is now also becoming critical [Liang et al. 2019] for neural training applications.

A key challenge then has been to reliably simulate shells with guaranteed results. Harmon et al. [2009] introduce an explicit time-stepping method for simulating shells that offers a guarantee of non-intersection across all timesteps. This method then requires small timesteps and simplified friction. Here we provide a complementary, fully implicit, differentiable simulation method, unified for all codimensional models, with accurate frictional contact, guaranteeing non-intersection, as well non-inversion and strain-limit satisfaction (when desired), that can be stepped at all reasonable time step sizes.

2.2 Strain Limiting

Strain-limiting methods seek to impose bounds on membrane deformation. A wide range of models and algorithms have long been proposed to provide strain limiting [Bridson et al. 2002; Provot et al. 1995]. To understand recent methods’ behaviors and limitations we categorize them by: 1) constraint choice (limits imposed as either equality or inequality); 2) type of DOFs constrained (these are most commonly edge-based or singular values); 3) splitting model; and 4) solver applied to enforce these constraints.

With this breakdown we see that equality constraints on length-based measures remain a consistent constraint choice. Goldenthal et al. [2007] apply bilateral constraints on quadmesh edge lengths and propose a fast projection method to correct predicted displacements.



Fig. 3. **Hairs.** Left: two hair clusters are twisted to form a tightly wound braid without intersection (processing up to 1.5M contacts per time step); the bottom end is then released and the braid unwinds. Right: a hair cluster simulation example inspired by McAdams [2009] where one end of a hair cluster is fixed, and so falls under gravity upon another hair cluster with two ends fixed and then across a sphere.




Fig. 4. **Reef Knot.** To exercise extreme compression, contact and friction with large stresses, we extend the reef knot test from Harmon et al. [2009] by simulating with a larger system 100K nodes (10× that of the original), adding a strain limit of 1.134, friction with $\mu = 0.02$, and tighter pulling. Two ribbons are initially intertwined in (a) and then stretched in (b)-(c) to form a final tight knot as the ribbons are pulled to nearly their strain limit in (d).

English and Bridson [2008] adopt a non-conforming strategy, applying equality constraints on triangle-edge midpoint distances, and extend fast projection to support a BDF-2 integrator. Thomaszewski et al. [2009] apply lagged, corotational small strain on triangles to define equality constraints and enforce limits by post-projection via Gauss-Seidel or Jacobi iterations. Chen and Tang [2010] likewise define equality constraints on triangle edge lengths.

Improvement over edge-based constraints can often be obtained by constraining singular values of the deformation gradient tensor (generally per triangle) to a finite range. However, irrespective of details, equality constraints always remain active and so simple DOF counting generally explains the membrane locking artifacts often encountered when these strain constraints are enforced. Here, alternate constraint DOF choices [Chen et al. 2019; English and Bridson 2008] can help alleviate this issue by reducing the ratio of constraints to DOFs but also can introduce new challenges, e.g. via non-conforming meshes. As an alternative to bilateral enforcement inequality constraints on strain, in the form of upper and lower bounds [Hernandez et al. 2013; Narain et al. 2012; Wang et al. 2010, 2016] have long been applied; while, even more recently, applying

just upper bounds [Jin et al. 2017] has also been considered. Irrespective of bounds, unilateral constraints are only activated when their strain measures are at their limits and so the potential for an over-constrained system is reduced.

While approaches thus clearly vary we observe that all methods, with the exception of Chen and Tang’s [2010] least-squares approximation for frictionless elastostatics, currently employ time-step splitting. Here we see that two-step splits are often applied to first solve elastodynamics before applying a constraint projection to simultaneously resolve contact constraints and strain limits. Alternately, introducing even more errors, we also see three-step splits where sequential solves of elastodynamics, contact and strain limiting are each decoupled and so resolved independently per time step [Narain et al. 2012]. Thus, for the latter three-step approach strain limits and contact constraints cannot be simultaneously satisfied, while for the former two-step methods errors introduced by splitting elasticity and constraints also generate unacceptable artifacts; see Section 6.

Finally, we note that, to our knowledge, no method (irrespective of constraint-type or splitting choice) employs a constraint solver that can guarantee enforcement of strain limits. As constraint enforcement errors then vary across simulation scenes, and even individual time steps, this leads to uncontrollable and inconsistent variations in material behavior per scene and step. To enable consistent, effective strain-limiting C-IPC introduces a fully coupled, inequality based model solved with a strict guarantee of strain-limit satisfaction.

2.3 Modeling Thickness

Thin bodies modeled by codimensional geometries generally have thicknesses orders-of-magnitude smaller than other dimensions. It is thus tempting (and common) to ignore thickness in contact. However, in order to correctly capture thin material interactions we must account for the geometric effects of thickness in contact. We see this everywhere, for example when playing cards stack in a pile (Figure 8) or when noodles fill a bowl (Figure 10).

A direct strategy for capturing thickness is then the solid shell method [Hauptmann and Schweizerhof 1998] which models thin materials volumetrically with full translational DOF. However, in




Fig. 5. **Garments.** A yellow dress with challenging folded and seamed knife pleats (produced by FoldSketch [Li et al. 2018b]) is staged (a) and then draped on a mannequin to static equilibrium (b) with stitches seamed together. It is then simulated on a Rhumba dancing sequence (c). A multilayer skirt is draped in the same way (d), and then used to simulate an input character animation sequence with large and fast motions (e,f,g). As the mannequin rapidly kicks, intricate garment details and high speed collisions are captured by C-IPC simulated with large time steps ($h = 0.04s$).

doing so linear elements suffer from well-known shear-locking artifacts. And while higher-order elements can alleviate these convergence issues, shear-locking artifacts are still not easily avoided altogether, while computational costs increase significantly. To address these challenges solid shell methods commonly employ reduced integration with hourglass mode stabilization [Reese 2007] or assumed natural strain methods [Cardoso et al. 2008] to mitigate shear locking with linear elements. However, this added complexity generally does not compete with codimensional modeling.

To resolve contact with codimensional models, contact-processing strategies often introduce thickness-like parameters to offset constraints and so reduce collision-processing inaccuracies; see e.g. Narain et al. [2012] and Li et al. [2018a] for recent examples. However, as analyzed in Li et al. [2020a], these thicknesses can not be consistently enforced and moreover must be changed heuristically per scene and example (e.g., based on collision speeds) to avoid simulation failures from intersections and instabilities. In turn, while they are sometimes helpful to improve contact resolution, these tuned parameters can not reliably be applied to model consistent and changing thickness behaviors; see Section 6.5.

We extend the IPC model to capture the geometric thickness of codimensional domains with offset barriers that guarantee a requested small minimal separation from mid-surface as well mid-line and point geometries. In turn this enables reliable and consistent modeling of thickness behaviors in the interactions between thin contacting materials.

2.4 Continuous Collision Detection

Continuous Collision Detection (CCD) methods have long been employed as a check against intersecting mesh boundaries. Provat [1997] formulated finding first times of impacts for linearly displacing point-triangle and edge-edge pairs by first testing for coplanarity via solution of a cubic equation and then performing overlap checks to detect collision.

This by now standard formulation of floating point CCD has been broadly applied and fine tuned over time [Harmon et al. 2009; Tang et al. 2011]. However, while root-finding with floating-point arithmetic can be made efficient, significant numerical errors will still generate unacceptable, false results (both negative and positive). This is most commonly encountered when distances are small and/or

configurations are unavoidably degenerate. Here, while switching from floating-point to rationals can help avoid round-off issues, cost can unacceptably increase if special care is not taken.

Recent methods address CCD accuracy [Brochu et al. 2012; Tang et al. 2014; Wang et al. 2015] by application of exact arithmetic for improved robustness with efficiency. Alternately, others [Harmon et al. 2011; Lu et al. 2019] perform conservative CCD with a requested small conservative separation distance to better avoid interpenetration. This latter strategy is employed in volumetric IPC [Li et al. 2020a] using a robust, floating-point CCD implementation¹ as a base solver.

Here, for modeling thickness in C-IPC, we require CCD queries that can maintain finite separation distances between mid-surface, mid-line and point elements. In turn, we rely on accurately preserving orders-of-magnitude smaller distances between the offset surfaces for computing accurate contact forces. This challenges the accuracy and robustness of CCD queries and we see unacceptable errors, resulting in failures, in all available existing CCD methods and codes. We see this both for standard floating-point root-finding CCD methods as well as more recent exact CCD methods. In turn we see that these failures can generate intersections, slow convergence, and often stop simulation progress altogether; see Section 6.6.

As an alternative to root-finding, conservative advancement (CA) methods have been developed to instead iteratively advance rigid and/or articulated bodies until they are closer than some pre-defined small distance. Starting with Mirtich’s work on convex rigid bodies [1996] this is performed by repeated calculation of a lower-bound on time of impact (TOI) followed by a conservative step taken with the bound [Lu et al. 2014; Tang et al. 2009, 2013a; Zhang et al. 2006, 2007].

For robust CCD evaluation of deformable body trajectories we derive lower bounds on time of impact for deforming mesh primitive pairs with arbitrary displacements and apply them in the CA framework to develop a new, simple to compute, numerically robust, floating-point, additive CCD (ACCD) algorithm. Under the CA framework, ACCD monotonically approaches times of impact without error-prone, direct root-finding. In Section 6 we confirm

¹<https://github.com/evouga/collisiondetection>

ACCD efficiently and accurately succeeds on a wide range of challenging cases where all other methods fail and find similar and often better performance in the cases where floating-point CCD methods can succeed. Finally we verify that ACCD is also suitable for replacement in CCD modules outside of the C-IPC framework with improved efficiency and robustness.

2.5 Unified Codimensions

Unified simulation of all codimensions in a common framework is critical for simulation efficiency and accuracy. Martin et al. [2010] focus on a unified elasticity model. They derive elastons – a higher-order integration rule to measure stretch, shear, bending and twisting along all axes without distinction between codimensions. Elastons accurately capture a wide range of elastoplastic behaviors, while contact forces are determined by point-wise penalties, with objects represented by a set of spheres. Chang et al. [2019] address unification for mixed-dimensional elastic bodies by defining all connections between domains of varying codimension via equality constraints. Here, Bridson et al.'s [2002] collision processing algorithm is then applied to resolve contact. The material point method (MPM) [Jiang et al. 2017] also offers general-purpose modeling of all codomains and contact between them. MPM discretizes elasticity on Lagrangian particles while solving momentum balance on the DOFs of an Eulerian grid. Contact between codimensional objects is then directly resolved via the Eulerian grid as a velocity flow. However, sticking and gap errors in contact are well known artifacts in MPM and can be unacceptable if grid resolutions are too low. These artifacts are investigated and mitigated with Lagrangian DOF by Han et al. [2019].

Position-based dynamics (PBD) [Bender et al. 2014, 2015; Müller et al. 2007; Stam 2009] also enables seamless, unified coupling of bodies with varying codimensions. Here both constitutive model and contacts are resolved as constraints that are iteratively processed for efficient time-integration at the cost of controllable accuracy as constraint complexity increases [Li et al. 2020a]. Extensions of PBD with XPBD [Macklin et al. 2016], Projective Dynamics [Bouaziz et al. 2014], and its further generalization via ADMM [Overby et al. 2017] all similarly offer platforms for co-simulating a diversity of codomains. Recent enhancements now provide improved friction [Ly et al. 2020] and increased efficiency [Daviet 2020]. However, these methods both lack convergence guarantees and apply fixed upper iteration caps so that the fundamental trade-off of accuracy and robustness (in the resolution of elasticity and contact) for efficiency in computation remains [Li et al. 2019; Ly et al. 2020]. In practice this means that numerical instabilities and explosions can and will be encountered, especially in challenging scenarios, e.g. with large time step size, stiffness, deformation, or velocities, as demonstrated in Li et al. [2020a; 2019]. Instead, targeting on guaranteed convergence and stability, C-IPC enables the direct simulation of all codimensions simultaneously. Coupling is provided by the interaction between any and all codimensional pairings via accurate, intersection-free contact.

3 FORMULATION

We focus on the combined solution of meshed, codimensional models simulated jointly and so coupled arbitrarily via contact. To do

so we generalize the IPC model to codimensional DOF and so must address the challenges introduced by thin models both coupled by, and stressed via contact and large, imposed boundary conditions. Here we first cover our generalization of IPC to mixed codimensional models and then, in the following sections, construct the key components of our method that enable their simulation.

Elastodynamics with contact. For elastodynamics, we perform implicit time stepping with optimization time integration [Brown et al. 2018; Gast et al. 2015; Kaufman and Pai 2012; Li et al. 2019; Liu et al. 2017; Overby et al. 2017; Wang et al. 2020] to minimize an Incremental Potential (IP) [Kane et al. 2000] with line search, ensuring stability and global convergence². For a wide range of time steppers and assuming hyperelasticity, the IP to update from time step n to $n + 1$ is defined as

$$E(x) = \frac{1}{2} \|x - \hat{x}^n\|_M^2 + \alpha h^2 \Psi(\beta x + \gamma x^n), \quad (1)$$

with the timestep update given by $x^{n+1} = \operatorname{argmin}_x E(x)$. Here h is time step size, M is the mass matrix, and Ψ is an elastic energy potential. Scaling factors $\alpha, \beta, \gamma \in [0, 1]$ and explicit predictor \hat{x}^n then determine the time step method applied. For example, here we focus primarily on the graphics-standard implicit Euler with $\alpha, \beta = 1, \gamma = 0$ and $\hat{x}^n = x^n + h v^n + h^2 M^{-1} f_{\text{ext}}$. Alternate implicit time stepping, e.g., with implicit midpoint or Newmark, are similarly matched by applying alternate scalings and predictors [Kaufman and Pai 2012].

Incremental Potential Contact (IPC) then augments the IP with contact and dissipative potentials [Li et al. 2020a]:

$$E(x) = \frac{1}{2} \|x - \hat{x}^n\|_M^2 + \alpha h^2 \Psi(\beta x + \gamma x^n) + B(x) + D(x). \quad (2)$$

Here $B(x)$ and $D(x)$ are respectively the IPC contact barrier and friction potentials. The contact barrier, B , enforces strictly positive

²Global convergence indicates convergence to a local optimum for arbitrary (feasible) initial configurations [Nocedal and Wright 2006].




Fig. 6. **Cloth on rotating sphere.** A square cloth (85K-node 1st row, 246K-node 2nd row) is dropped upon a sphere set on the ground with friction ($\mu = 0.4$ for both). As the sphere starts to rotate the cloth follows and winds tightly while preserving strain limits throughout. Differing moduli are used for the 85K- and 246K-node cloth respectively to produce intricate wrinkling behaviors with differing frequency. Here from left to right we show the 25th (right before rotation), 50th, and 75th frames.

distances between all primitive pairs while D provides corresponding friction forces. Following Li et al. [2020a] we apply a custom Newton-type solver to minimize each IP with Continuous Collision Detection (CCD) filtering executed in each line search of every Newton iteration to ensure intersection-free trajectories throughout. Please see Li et al. [2020a] for details on the base IPC algorithm and solver implementation.

Mixed-dimensional hyperelasticity. To enable a unified simulation framework for coupled volumetric bodies, shells, rods and particles, we compute mass and volume for all codimensional elements by treating them as continuum regions with respect to standard discretizations and so construct the total elasticity energy $\Psi(x)$ for the IP as

$$\Psi(x) = \Psi_{\text{vol}}(x) + \Psi_{\text{shell}}(x) + \Psi_{\text{rod}}(x). \quad (3)$$

Here, without loss of generality, as representative examples, we apply fixed Corotated elasticity [Stomakhin et al. 2012] for volumes (Ψ_{vol}); the Discrete Shells hinge bending energy [Grinspun et al. 2003; Tamstorf and Grinspun 2013] combined with either isotropic or orthotropic StVK [Chen et al. 2018; Clyde et al. 2017], or neo-Hookean membrane models for shells (Ψ_{shell}); and the Discrete Rods stretch and bending model [Bergou et al. 2008] for rods (Ψ_{rod}). We select these models as best for comparison and evaluation given models standard in existing codes. More generally, the C-IPC framework is agnostic to a broad range of elasticity and time stepper choices as demonstrated in Section 6 and [Li et al. 2020a]. Along with properly integrating all energies with an accurate volume weighting per element, we further parameterize rod bending moduli via Kirchhoff rod theory following Bergou et al. [2010] for direct material settings (see our supplemental document for details). With per-domain elasticity summed in (3) our IPC model now couples objects of arbitrary codimensions directly (without splitting) via frictional contact. Specifically all codimensional DOF are now associated with their respective discrete inertial and potential energies and so are free to move by time stepping. In turn they are coupled by IPC-type barriers. C-IPC barriers now include all point-triangle and edge-edge pairings from all surfaces (both volumes' and shells'), rods and particles; point-edge pairs from all rod nodes and particle-to-rod segments; and finally point-point pairs between all particles.

4 CONSTITUTIVE STRAIN LIMITING

Here we begin by constructing a new constitutive barrier model that directly enforces strain limiting while maintaining rest-state consistency. We start with a formulation that augments existing membrane energies with an added strain-limiting potential for the general, isotropic case (Section 4.1). We then demonstrate its extension to directly augment orthotropic STVK membranes with strain limits in Section 4.2.

4.1 Isotropic Constitutive Strain Limiting

We define isotropic strain-limiting constraints per element t as

$$\sigma_i^t < s, \quad \forall i, \quad (4)$$

with singular value decomposition of each triangle t 's deformation gradient, $F^t = U^t \Sigma^t V^{tT}$, giving $\Sigma^t = \text{diag}(\sigma_1^t, \sigma_2^t)$. The imposed



Fig. 7. **Funnel.** Three cloth panels ($\mu = 0.4$, 26K nodes each, strain upper bound at 1.0608) are dropped on a funnel. Under friction they rest stably on top until a sharp, scripted collision object rapidly pulls the panels down.

constraint bound, s , is then the requested strain limit with practical bounds generally selected for cloth with $s \in [1.01, 1.1]$ [Bridson et al. 2002; Provot et al. 1995].

Strain-limiting requires local support. It should only exert constraint forcing during stretch when strain is close to an applied upper bound. Otherwise, it should leave an underlying membrane model unchanged. This is analogous to the application of contact forces which should only exert when boundaries are nearly touching. As such, we begin with IPC's contact barrier to construct a comparable C^2 , smoothly clamped barrier for strain limiting. Each barrier per σ_i^t is then

$$b_i^{st}(x) = \begin{cases} -\left(\frac{\hat{s}-\sigma_i^t}{s-\hat{s}}\right)^2 \ln\left(\frac{s-\sigma_i^t}{s-\hat{s}}\right) & \sigma_i^t > \hat{s} \\ 0 & \sigma_i^t \leq \hat{s} \end{cases}, \quad (5)$$

and so is only activated when strain exceeds an imposed, small-strain threshold \hat{s} . E.g., see inset for the barrier energy with strain limit $s = 1.1$ and threshold $\hat{s} = 1$.

Now, with barriers in hand, we could potentially next consider treating them directly as constraints (as in primal barrier and interior point methods). However, doing so would then simply sum these barriers over elements and so would obtain inconsistent behavior as we change meshes. Instead, to provide consistent behavior we impose strain limiting constitutively as an energy density integrated over the volume of cloth to obtain the potential

$$\Psi_{\text{SL}}(x) = \sum_i \kappa_s \int_{\Omega} b_i^{st}(x) dV \approx \kappa_s \sum_{t,i} V^t b_i^{st}(x), \quad (6)$$

where κ_s is the barrier stiffness in Pa . For our approximation, we use $V^t = A^t \xi^t$, with A^t and ξ^t respectively the area and thickness of triangle t . Our final isotropic strain-limiting potential, Ψ_{SL} , is then C^2 with local support and can simply be added to our total potential in Equation 3. In turn, strain limiting is then directly handled at each time step by optimization of the IPC in Equation 2, while ensuring that strain-limit barriers are not violated in each line-search step. During steps when most triangles remain under the strain limit threshold, little additional computation is then required. While, when stretch increases, the necessary nonzero terms from newly activated barriers are added to the system and so, as we show in Section 6, strictly enforce strain limits while balancing all applied forces.

Strain-Limit Stiffness. With our barrier potential defined we now specify setting its stiffness. The mollified clamped barrier we apply

for our strain-limit model is inspired by IPC's smoothing of contact barrier energies [Li et al. 2020a]. Here, however, notice an extra ingredient for strain-limiting is the applied $1/(s - \hat{s})^2$ factor. This scaling enables us to apply our barrier directly to a limit-normalized strain, measured by $y = (\hat{s} - \sigma_i^t)/(s - \hat{s})$ with

$$\hat{\delta}_i^{st}(y) = \begin{cases} -y^2 \ln(1+y) & y < 0, \\ 0 & y \geq 0. \end{cases} \quad (7)$$

In turn, as our applied strain-limit (s) and/or clamping threshold (\hat{s}) is varied with application, the barrier function w.r.t. y remains unchanged. Only the gradient of the linear map from σ_i^t to y varies. This allows us to apply a single, consistent initial barrier stiffness κ_s (we use $1KPa$ for all examples) across all choices of differing strain limits and clamping thresholds. Here the barrier potential curve is then simply linearly rescaled each time, to a different strain range, and so provides consistent conditioning to the system. Finally, to avoid numerical issues introduced by tiny gaps between a current strain and the imposed strain limit (e.g., when extreme boundary conditions are imposed as in Figures 9 and 1) we adapt barrier stiffness when needed following IPC's barrier stiffness adjustment strategy. Starting with our initial κ_s , we increase it by $2\times$ (up to max bound of $\kappa_s = 0.1MPa$) whenever the strain gap $s - \sigma_i^t$ of a triangles t is less than $10^{-4}(s - \hat{s})$ over two consecutive iterations. Notice that, when applied, this adjustment varies the strain-limit energy. This occurs, however, solely near the actual limit in extreme cases. Here strain-limit forces impose constraint while this adaptivity provides improved numerical conditioning. Once away from the limit the barrier returns to a consistent energy with no forcing at all below the threshold \hat{s} .

4.2 Anisotropic Constitutive Strain Limiting

Above we have constructed strain-limiting as an isotropic constitutive model. We formed a barrier energy that can be integrated and so directly added to potential energy in order to augment existing membrane models with hard strain limits. The key to making this work is the application of our C^2 continuous clamping so that the application of strain-limits does not alter rest-shape gradients nor rest-shape Hessians.

Alternately, we can apply a comparable constitutive strain-limiting strategy to directly modify membrane elasticity models to include strain limits. To do so we simply substitute an anisotropic membrane model with a barrier energy that prevents violation of strain limits, while matching the original membrane energy gradient and Hessian at rest.

This second strategy is particularly motivated by the need to incorporate strain-limiting into available data-driven models that have been constructed to carefully fit measured cloth data. Here we demonstrate specific application to the anisotropic, data-driven model of Clyde and colleagues[Clyde 2017; Clyde et al. 2017].

Their constitutive model energy is

$$\begin{aligned} \tilde{\psi}(\tilde{E}_{11}, \tilde{E}_{12}, \tilde{E}_{22}) &= \frac{a_{11}}{2} \eta_1(\tilde{E}_{11}^2) + \frac{a_{22}}{2} \eta_3(\tilde{E}_{22}^2) \\ &+ a_{12} \eta_2(\tilde{E}_{11} \tilde{E}_{22}) + G_{12} \eta_4(\tilde{E}_{12}^2), \end{aligned} \quad (8)$$

where $\tilde{E} = D^T E D$ is the reduced and aligned Green-Lagrangian strain, $\tilde{E}^{i3} = \tilde{E}^{3i} = 0$, and column vectors of D are the tangent and normal bases of the shell in material space. Functions η are then a sum of polynomial functions with real-valued exponents α_{ji} that satisfy $\eta(0) = 0$ and $\eta'(0) = 1$. Here the first constraint enforces zero-energy, zero-stress rest configurations, while the second constraint allows a natural correspondence between the parameters $a_{\alpha\beta}$ and G_{12} and linear elasticity at infinitesimal strain. Clyde et al. use

$$\eta_j(x) = \sum_{i=1}^{d_j} \frac{\mu_{ji}}{\alpha_{ji}} ((x+1)^{\alpha_{ji}} - 1).$$

Their measured data is then restricted to deformations within the fracture limit or elasticity range of the cloth. Beyond this range meaningful extrapolation is then unlikely, due to possible overfitting inside the range (exponents α_{ji} from the fitting can range up to 10^5). To address this limitation Clyde et al. propose a quadratic extrapolation for simulation. However such extrapolation is not physically meaningful. To usefully apply data-driven modeling either fracture should be captured beyond this regime or else a stable and controllable strain limit imposed to stay in bounds. Here, we focus on controllably imposing strain limit while respecting the underlying model fitting.

Barrier Formulation. Starting from the same constitutive model we redefine the η basis with barriers

$$\eta(E) = -E^{max} \log((E^{max} - E)/E^{max}).$$

For consistency note this basis again satisfies both $\eta(0) = 0$ and $\eta'(0) = 1$ and so is valid for elasticity. This barrier also diverges at E^{max} and so ensures that E never exceeds measured strain limit bounds. It thus captures data-driven strain limits and at the same time avoids extrapolation. Here, as the underlying model is anisotropic, warp, weft, and shearing directions can all impose different, measured E^{max} bounds, enabling preservation of measured anisotropy for all strain limits.

Free model parameters $a_{\alpha\beta}$ and G_{12} , are then matched at $\nabla^2 \psi(0, 0, 0)$ with the underlying model's fit; see our supplemental document for details. Note this formulation frees us from computing the sensitive (and potentially expensive) polynomials with large real-valued exponents. See Section 6.2 for experiments on the proposed model. Finally we note that, in contrast to the strategy we demonstrated in our isotropic case from the last section (where we worked directly on an upper bound w.r.t. the deformation gradient), here this second energy is symmetric for stretch and compression.

4.3 Line-Search Filtering for Strain Limits

As discussed above, CCD-based filtering is critical to ensure line-search is consistent with our contact barriers. Now that we add strain-limiting barriers we must additionally ensure that every line search will also not violate imposed strain limits. For isotropic limits, although 2×2 SVDs have a closed form solution, it does not provide a polynomial equation for feasible step size computation. Thus, in C-IPC line search, after we compute an intersection-free starting step size, we simply half the step length if we detect a strain-limit violation during backtracking. We repeat until we obtain a step size satisfying energy decrease and strain limits. In practice,

because our Newton step includes higher-order information of our strain-limits from Hessian and gradients, we observe IPC search directions are effective at naturally avoiding strain-limits. In all cases, when required, we so far observe at most three backtracking steps for strain-limiting are required. For our anisotropic model we also currently apply the same backtracking strategy. However, while currently not needed for efficiency, we do note that here we can formulate quadratic equations amenable to directly computing largest feasible strain-limit satisfying step sizes.

5 MODELING THICKNESS

In its original form IPC maintains intersection-free paths for volumetric models by enforcing positivity of unsigned distances d_k , as an invariant between all non-adjacent and non-incident surface primitive pairs k [Li et al. 2020a]. This is suitable for volumetric contact where this constraint permits arbitrarily close, but never intersecting surfaces. For codimensional models, however, this constraint is no longer sufficient. When the 3D deformation of thin materials is reduced to the deformation of a 2D surface or 1D curve, elasticity can be well-resolved on surfaces and curves, but contact can not. Neglecting to account for finite thickness in codimensional contact generates unacceptable artifacts (see e.g. Figure 2) and clearly fails to capture geometries formed by thin-structure interactions (see e.g. Figures 18 and 8).


We begin by observing that applying a larger \hat{d} (the threshold distance at which IPC contact force application begins), to codimensional geometries in IPC, models a thin responsive layer that resists compression in the normal direction and so forms an elastic thickness. In concert with this elastic layer we also generally require a core thickness beyond which further compression is not allowed. This is needed to guarantee a minimum finite (and so e.g. visible) thickness even when deformation is extreme (see e.g. Figure 2 bottom).

To model thickness in contact for codimensional objects we build an inelastic thickness model that combines the purely elastic layer with a hard, and so inelastic offset to the contact barrier. For a nonzero offset, ξ , geometries are then guaranteed to be separated from each other by ξ . Elastic contact forces exert when distances are below $\xi + \hat{d}$ and diverge when distances are at ξ . Here larger \hat{d} generates greater elastic response, while nonzero ξ guarantees a minimum thickness. This holds even under extreme compression.

We equip each surface element i with a finite thickness ξ_i . The distance constraint for primitive pairs k on the surface, formed between element primitives i and j , are then

$$d_k(x) > \xi_k = \frac{\xi_i + \xi_j}{2}.$$

Boundaries of codimensional materials are then approximated with a rounded cross-section, while for interaction between zero-thickness materials our distance constraints reduce to that of the original IPC constraint (e.g. for volume-to-volume contact). Finally for volume-to-codimensional contact, volumes can continue to maintain zero-thickness boundaries while interacting with finite thickness codimensional boundaries.



For scenarios where modeling thickness matters (see e.g. Figures 2, 3-23), we then set ξ_k to the true thicknesses of the material(s) or slightly smaller, and then set \hat{d} near ξ_k to compensate for the remaining thickness depending on how much compression, and so elastic response, is reasonable or required per application.

5.1 Solving IPC with Thickened Boundaries

For a numerically robust and efficient implementation, IPC applies squared distances (with appropriate rescalings) to compute with equivalent barrier functions. This change swaps the model's derived contact barrier $b(d_k(x), \hat{d})$ to an equivalent rescaled implementation using $b(\hat{d}_k^2(x), \hat{d}^2)$ [Li et al. 2020a]. Here we directly apply our thickened contact barrier via squared distances

$$b_\xi(d_k(x), \hat{d}) = b(\hat{d}_k^2(x) - \xi_k^2, 2\xi_k\hat{d} + \hat{d}^2),$$

so that contact forces correctly diverge at $d_k(x) = \xi_k$, and nonzero contact forces are only applied between mid-surface pairs closer than $\hat{d} + \xi_k$. Recall that $b(d^2, \hat{d}^2) = -(d^2 - \hat{d}^2)^2 \ln(d^2/\hat{d}^2)$ if $d^2 < \hat{d}^2$, and is 0 otherwise [Li et al. 2020a].

Here then the barrier function, gradient, and Hessian only need to be evaluated with the modified input distance, while distance gradient and Hessian remain unchanged as

$$\frac{d(\hat{d}_k^2(x) - \xi_k^2)}{dx} = \frac{d\hat{d}_k^2(x)}{dx}.$$

Evaluation of contact force magnitudes for computing friction forces as well as adaptive stiffness (κ) updates for contact barriers follow similarly.

Most constraint set computations then require simple and comparable modifications. For spatial hashing, bounding boxes of all mid-surface primitive i are extended by $\xi_i/2$ in all dimensions on left-bottom and top-right corners before locating voxels. This enables hash queries with \hat{d} to remain unchanged. Next, for broad-phase contact pair detection, query distances \hat{d} are now updated to $\hat{d} + \xi_k$ to check for bounding box overlap. Or else, equivalently, bounding box primitives are extended as in the spatial hash. Finally, to accelerate continuous collision detection (CCD) queries, spatial hash




Fig. 8. “Precision” card shuffle. We divide fifty-four playing cards into two separate piles and bend them in preparation for a “precision” bridge shuffle. Unlike a human-performed bridge finish, where cards begin interleaved, here we increase challenge by keeping the two bridged piles apart and then (a) precisely, one-by-one, bottom up they are rapidly released in turn to shuffle the deck (b). We partially square the deck in (c), obtaining a fully shuffled stack with non-intersecting, interleaved cards that is 15.4 mm high and so well-matching the height of a standard deck (d).




Fig. 9. **Cloth over codimensional needles.** To test contact across sharp boundaries we drop and then drag a cloth panel across a needle bed formed by line segments. With a large time step size of $h = 0.02s$ and strain limit of 1.0608, the cloth comes safely to a static rest on the segments without jittering (a). We then pull the left side of the cloth at $1m/s$ across the segment tips (b) and observe rapid sliding along (c), and then off (d) the top of the bed without snagging nor stretching artifacts (strain limits are satisfied throughout).

construction also requires similar bounding box extension, while for its broad phase, applied gaps are ξ_k (rather than 0).

5.2 Challenges for CCD

While the above initial modifications for thickness are straightforward, finite thickness and codimensional DOF introduce new computational challenges for CCD queries. Here we analyze these challenges and develop a new CCD method to address them.

Standard-form IPC applies position updates in each iteration to obtain minimal separation distances of $s \times$ the current separation distances, d_k^{cur} , at time of CCD evaluations. Here $s \in (0, 1)$ is a conservative rescaling factor (generally set to $s = 0.2$ or 0.1) that allows CCD queries to avoid intersection, even when surfaces are very close [Li et al. 2020a]. To similarly resolve finite thicknesses, our conservative distance bound is now $s(d_k^{cur} - \xi_k) + \xi_k$.

Concretely, the barrier evaluated parameters, $d_k^2(x) - \xi_k^2$, must always remain positive. In turn we require all CCD evaluations, for each displacement p , to provide us with sufficiently accurate times $t \in (0, 1]$ to satisfy positivity. If impact would occur for a pair along p , we require a time t so that the scaled displacement, tp , ensures distance remains greater than ξ_k and is as close as possible to the target separation distance $\xi_k + s(d_k^{cur} - \xi_k)$.

Doing this with finite thickness is then much more challenging than without thickness. In contact $d_{cur} - \xi_k$ can be as small as $10^{-8}m$ (e.g. under large compression) so that for $s \geq 0.1$ absolute errors in the order of $10^{-9}m$ are acceptable. While, in practice, ξ_k is then regularly at the scale of $10^{-4}m$ (e.g. thickness values of $\sim 3 \times 10^{-4}m$ for cloth). Without thickness ($\xi_k = 0$) updated distances are targeting $s d_k^{cur}$ (and are only required to be strictly greater than 0). In other words any step size that prevents primitive-pair intersection is valid and so relative errors less than $100\% \approx 10^{-9}m / (s10^{-8}m)$ are acceptable. On the other hand, with thickness, updated distances are targeting $\xi_k + s(d_k^{cur} - \xi_k)$ and so, for standard values of ξ_k , would require relative errors from the CCD evaluations near $0.01\% \approx 10^{-9}m / (10^{-4}m + s(10^{-8}m))$ in order to avoid interpenetration.

Obtaining CCD evaluations to this accuracy is then extremely challenging for available methods. As a starting example we tested the floating point CCD solver from Li et al. [2020a], requesting a $s(d_{cur} - \xi_k) + \xi_k$ minimal separation for two challenging codimensional examples with thickness (see Figures 10 and 21) with $\xi_k > 0$. Here the CCD solver returns $t = 0$ time of impact (TOI) erroneously even if we remove the conservative scaling factor (i.e. set $s = 0$) altogether (See Figure 19). Note that in doing so this error stops simulation progress altogether.

5.3 CCD Lower Bounding

Next we derive a useful lower bound value for CCD queries that is numerically robust and can be efficiently evaluated in floating point. This lower bound provides a conservative, guaranteed safe step size and also a clear measure to test a CCD query's validity: any CCD-type evaluation with smaller TOI has clearly failed. In Section 5.4 we then apply this bound to derive a simple, effective and numerically accurate explicit CCD solver that is robust and efficient for progress even in the challenging CCD evaluations we require for thin material simulations.

Here, without loss of generality, we will focus on the edge-edge case between edge pairs (x_0, x_1) and (x_2, x_3) with corresponding displacements p_0, p_1, p_2 and p_3 . The distance function between arbitrary points parameterized respectively by γ and β on each edge at any time t is then

$$f(t, \gamma, \beta) = \|d(t, \gamma, \beta)\|, \text{ with}$$

$$d(t, \gamma, \beta) = (1 - \gamma)x_0(t) + \gamma x_1(t) - ((1 - \beta)x_2(t) + \beta x_3(t)), \quad (9)$$

$$x_i(t) = x_i(0) + t p_i, \text{ and } t, \gamma, \beta \in [0, 1].$$

A CCD evaluation then seeks the smallest positive real t satisfying

$$f_1(t) = \min_{\gamma, \beta} f(t, \gamma, \beta) = 0.$$

If such t exists we call it t_a . Parameters $(\gamma_a, \beta_a) = \operatorname{argmin}_{\gamma, \beta} f(t_a, \gamma, \beta)$ in turn give the corresponding points colliding at time t_a on each both edges.

We can then express t_a as

$$t_a = \frac{f(0, \gamma_a, \beta_a)}{\|((1 - \gamma_a)p_0 + \gamma_a p_1 - ((1 - \beta_a)p_2 + \beta_a p_3))\|}.$$

Challenges to CCD evaluations then occur in determining γ_a and β_a , when degeneracies and numerical error make floating-point methods both prone to false positives and false negatives.

We can however, directly and accurately perform a distance query w.r.t. the primitives at start ($t=0$) of evaluation, to find the distance, $f_1(0)$, that certainly satisfies $f_1(0) \leq f(0, \gamma_a, \beta_a)$. Then triangle inequality gives $\|((1 - \gamma_a)p_0 + \gamma_a p_1 - ((1 - \beta_a)p_2 + \beta_a p_3))\| \leq \|(1 - \gamma_a)p_0 + \gamma_a p_1\| + \|(1 - \beta_a)p_2 + \beta_a p_3\|$, and since $\gamma_a, \beta_a \in [0, 1]$ we get $\|(1 - \gamma_a)p_0 + \gamma_a p_1\| \leq \max(\|p_0\|, \|p_1\|)$ and $\|(1 - \beta_a)p_2 + \beta_a p_3\| \leq \max(\|p_2\|, \|p_3\|)$. Put together we then have the practical and directly computable bound on TOI

$$t_a \geq \frac{f_1(0)}{\max(\|p_0\|, \|p_1\|) + \max(\|p_2\|, \|p_3\|)}. \quad (10)$$

More generally, for any query between any two primitives types, our bound on t_a is correspondingly given by the appropriate distance function in the numerator (e.g., $f_1(0) = \min_{\gamma_a, \beta_a, \alpha_a} f(0, \gamma_a, \beta_a, \alpha_a)$




Fig. 10. **Noodles.** (a) We drop 625 40cm-long noodles, modeled by discrete rods with C-IPC offset, into a bowl. (b) We set offset and \hat{d} to model noodle thickness enabling the bowl to fill when they reach equilibrium. (c) We then remove the bowl to show the final intertwined, intersection-free configuration, noting that thickness offsets are satisfied throughout. In (d) we then remove the volume render to highlight the geometry is simulated solely with discrete rods polyline geometry.

for point-triangle and $f_1(0) = \min_{\gamma_a} f(0, \gamma_a)$ for point-edge) and the sum of the max displacement from each of the two primitives in the denominator. We then note that even when there is no smallest positive time satisfying $f_1(t) = 0$ (and so no impact on the interval), our bound remains well-defined as a conservative step size.

We next observe that, perhaps surprisingly, state-of-the-art floating-point CCD solvers can and will return TOI results smaller than our lower bound, and so are clearly in error (See Figure 19).

Finally, to improve our bound, we observe that it holds independently of the choice of the reference frame. Thus we can further tighten the bound on each CCD query independently by picking frames that reduce the norm of displacement vectors p_i . For example by subtracting each p_i by the average $\frac{1}{4} \sum_i p_i$.

5.4 Additive CCD

With our computed lower bound we now apply a CA strategy [Mirtich 1996; Zhang et al. 2006] to build a new CCD algorithm with finite offsets for deformable bodies that iteratively updates and adds our lower bound over successive conservative steps towards TOI. The resulting *additive* CCD (ACCD) method then robustly solves for a bounded TOI, monotonically approaching each CCD solution, while only requiring explicit calls to evaluate distances between updated primitive positions.

At the start of each CCD query, to initialize the ACCD algorithm (see Algorithm 1), we first center the collision stencil's displacement at origin to reduce our bound's denominator, e.g., $l_p = \max(\|p_0\|, \|p_1\|) + \max(\|p_2\|, \|p_3\|)$ for edge-edge pairs, and so increase the step size we can safely take. If there is no relative motion ($l_p = 0$), we of course simply return no collision and so a full unit step is valid. We then compute the requested minimal separation $g = s(\sqrt{d_{sqr}} - \xi)$ to the offset surface based on the current squared distance d_{sqr} and the scaling factor s . For this we use a formula that is more robust to cancellation error; see lines 8-9 in Algorithm 1.

Starting with a most conservative time of impact $t = 0$ (line 10) we create a *local* scratch-pad copy of nodal positions, x_i , and initialize the current lower bound step, t_l , with Equation 10 (line 11).

We then enter our iterative refinement loop (lines 12-21) to monotonically improve our TOI estimate t . At each iteration, we update our local copy of nodal positions x_i with the current step t_l (lines 13-14). If this new position achieves our target distance to the offset

Algorithm 1 Additive CCD

```

1: procedure ADDITIVECCD( $x_i, p_i, \xi, s, t_c, t$ )
2:    $\bar{p} \leftarrow \sum_i p_i / 4$ 
3:   for  $i$  in  $\{0, 1, 2, 3\}$  do
4:      $p_i \leftarrow p_i - \bar{p}$ 
5:    $l_p \leftarrow \max_{i \in 1st \text{ primitive}} (\|p_i\|) + \max_{i \in 2nd \text{ primitive}} (\|p_i\|)$ 
6:   if  $l_p = 0$  then
7:     return false
8:    $d_{sqr} \leftarrow \text{computeSquaredDistance}(x_0, x_1, x_2, x_3)$ 
9:    $g \leftarrow s(d_{sqr} - \xi^2) / (\sqrt{d_{sqr}} + \xi)$ 
10:   $t \leftarrow 0$ 
11:   $t_l \leftarrow (1 - s)(d_{sqr} - \xi^2) / ((\sqrt{d_{sqr}} + \xi)l_p)$ 
12:  while true do
13:    for  $i$  in  $\{0, 1, 2, 3\}$  do
14:       $x_i \leftarrow x_i + t_l p_i$ 
15:     $d_{sqr} \leftarrow \text{computeSquaredDistance}(x_0, x_1, x_2, x_3)$ 
16:    if  $t > 0$  and  $(d_{sqr} - \xi^2) / (\sqrt{d_{sqr}} + \xi) < g$  then
17:      break
18:     $t \leftarrow t + t_l$ 
19:    if  $t > t_c$  then
20:      return false
21:     $t_l \leftarrow 0.9(d_{sqr} - \xi^2) / ((\sqrt{d_{sqr}} + \xi)l_p)$ 
22:  return true

```

surface (becomes smaller than g) we have converged and the previous t is the time of impact that brings distance just up to g (line 17). If not, we update our TOI estimate by adding the current t_l to t (line 18). Note that we always add our first lower bound step to t (line 16) as it is guaranteed to not bring distance closer than g . If our TOI is now larger than t_c , the current minimum first time of impact (or can be simply set to 1), we can return no collision (lines 19-20). Otherwise we compute a new local, lower bound, t_l , (with 0.9 scaling for improved convergence) from the updated configuration (line 21) and begin the next iteration.

ACCD thus provides an exceedingly simple-to-implement, numerically robust CCD evaluation. It requires only explicit calls to

distance evaluations and so no numerically challenging root-finding operations. In turn, ACCD is able to support thickness offsetting and also controllable accuracy and so flexible tuning for performance vs. accuracy trade-offs in CCD applications. In Section 6.6 we compare ACCD with state-of-the-art CCD solvers, there we see that they all fail severely, resulting in intersection or tiny TOI that stalls the optimization in challenging examples. In turn, in easier cases where alternate CCD methods do succeed, we then see that ACCD achieves similar and often improved timing performance. Finally, we point out that ACCD's stopping criteria requires $s > 0$ (we apply $s = 0.1$ in all examples) to provide finite termination. In turn, ACCD does not target computing an exact TOI but rather (as actually required in contact-processing applications) obtains reliable, intersection-free steps towards TOI.

Worst-case performance. As discussed above, we find in practice and particularly throughout all our challenging testing that ACCD converges with solid performance (comparable or faster times when other methods succeed, efficient times when all other methods fail). It is worth analyzing, however, that it should be possible to have unbounded, worst-case performance. Recall we pick our reference frame so that displacements sum to 0. However, if a primitive has a diverging displacement field which cannot be cancelled out, our denominator l_p can remain large. Then, if a primitive also has a small starting distance d and so correspondingly small numerator, iteration count for ACCD on this case can be large. While we see that this scenario does not occur in practice for elastodynamics, extensions of ACCD to accelerate convergence for these possible cases, and so obtain bounded performance guarantees is an interesting future step.

6 EVALUATION

We implement our methods in C++, applying CHOLMOD [Chen et al. 2008], compiled with Intel MKL LAPACK and BLAS for linear solves and Eigen for remaining linear algebra routines [Guennebaud et al. 2010]. Necessary derivatives for C-IPC and our algebraic simplifications applied for efficiency and numerical robustness are detailed in our supplemental document. To enable future applications, development and testing we will release our implementation of C-IPC as an open source project. All our experiments and evaluations are executed on either an Intel 16-Core i9-9900K CPU 3.6GHz \times 8 (32GB memory), an Intel Core i7-8700K CPU @ 3.7GHz \times 6 (64GB memory), or an Intel Core i7-9700K CPU @ 3.6GHz \times 4 (32GB memory) as detailed per experiment below.

Experiments. Below we begin with a study illustrating and analyzing membrane locking behaviors for standard cloth materials and meshes (Section 6.1). We then demonstrate C-IPC's ability to strictly satisfy strain-limits while fully coupling all physical forces in Section 6.2. To our knowledge C-IPC is the first method to both enable strict satisfaction of strain-limits and to support full coupling of strain-limiting, elasticity and contact forces. To analyze the impact of strain-limit satisfaction and full coupling we next consider comparison against prior methods. As discussed in Section 2, existing methods in strain-limiting generate artifacts including locking, jittering and interpenetration. These issues result from two sources:

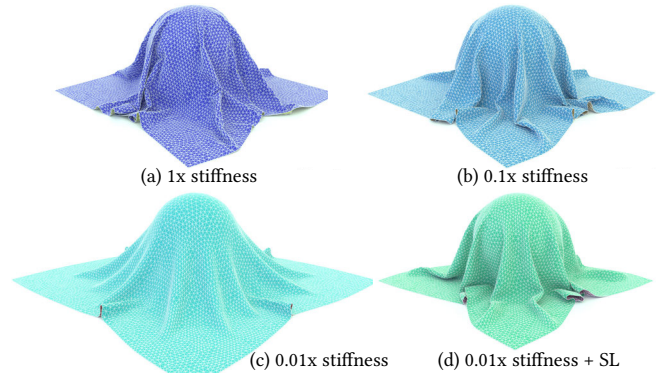


Fig. 11. **Membrane Locking.** A cloth on sphere simulation (8K-node) with (a) $1\times$ stiff membrane of cotton material ($0.8MPa$) suffers from membrane locking where bending behavior is artificially stiffened and creased. Reducing membrane stiffness to (b) $0.1\times$ that of measured cotton values removes artificial bending stiffness but we still observe overly sharp edge creasing. Next an even softer membrane at (c), $0.01\times$ cotton's, then avoids observable membrane locking artifacts altogether, but results in nonphysical stretching inappropriate for cotton. However, (d) applying a softer $0.01\times$ membrane together with our constitutive strain limiting model to enforce measured cotton strain limits prevents membrane locking while obtaining natural stretch limits for the material.

1) inaccuracies from applied splitting models and 2) inability to satisfy strain-limits in computation. Here, for the first time, we first separately analyze the problems that are created by splitting models (Sec. 6.3) independent of solver accuracy errors. Then in Section 6.4 we consider the artifacts and inaccuracies (generated by state-of-the-art cloth code) that jointly result from both splitting errors *and* inability of the method to enforce requested strain-limits. In Section 6.5 and 6.6 we analyze C-IPC's thickness modeling, comparing with both existing cloth simulation codes and prior methods for CCD evaluation. Finally, with all components of C-IPC evaluated, we assess C-IPC's application to garment simulation (Section 6.7) and its resolution of previous stress-test cloth benchmarks (Section 6.8). We then demonstrate C-IPC on simulations of fully coupled systems of arbitrary codimension with consistent thickness modeling (Section 6.9) and finally consider its performance on a new set of cloth simulation stress tests designed to exercise robustness and accuracy (Section 6.10). All experimental setup and statistics are listed in Figure 24. We report the total number of nodes involved in the system, including those that are kinematic (equality constraints). All examples are directly applying IPC's smoothed, semi-implicit friction (1 lagged iteration) and the default Newton tolerance if not otherwise mentioned.

6.1 Cloth Material and Membrane Locking

Here we first illustrate and examine the impact of membrane locking on real-world materials, and strain-limiting's ability to mitigate them. In Penava et al. [2014], density, thickness, and directionally dependent membrane stress-strain curves are measured and validated for a range of cloth materials. However, directly applying these real-world cloth parameters in simulation with standard-resolution meshes can produce severe membrane locking. This is unavoidable

unless exceedingly high and so often impractically large mesh sizes are used. Here we examine this locking behavior, first using our IPC model *without* strain-limiting as a demonstration. We also note that such locking behaviors are independent of algorithm and so, for example, are also easy to demonstrate in other codes like ARC-Sim [Narain et al. 2012] with their real-world captured material parameters [Wang et al. 2011]; see Figure 15d.

We start by considering the behavior of a simple $1m \times 1m$ unstructured mesh dropped on a fixed sphere and ground plane (with friction of $\mu = 0.4$ for all). We apply the measured cotton cloth density ($472.6kg/m^3$) and thickness ($0.318mm$) from Penava et al. [2014] and then consider varying membrane stiffness while keeping bending Young’s modulus at $0.8MPa$ and Poisson ratios for both membrane and bending to the measured value for warp direction as 0.243. Specifically Penava et al. [2014], find membrane Young’s moduli ranging from $0.8MPa$ to $32.6MPa$ for varying in-plane directions.

For our example, we then find that even when applying an isotropic membrane model, with the smallest determined membrane stiffness ($0.8MPa$), we observe severe locking artifacts (see e.g. Figure 11a) where bending is artificially stiffened. If we next lower this smallest measured membrane stiffness by $0.1\times$, we then see that artificial bending artifacts are largely eliminated but we still obtain sharp creasing artifacts forced by the dominating membrane energy (see e.g. Figure 11b). Next, if we try an even smaller $0.01\times$ scaling of membrane stiffness, observable membrane locking effects are now gone, but as expected the resulting material is much too stretched and so not even close to the desired material behavior (Figure 11c). This simple example nicely demonstrates the challenges of membrane locking when simulating stiff real-world cloth materials. We then note that these artifacts are only exacerbated in more challenging simulations with, for example, moving boundary conditions and tight contact.

Next we apply a strict strain-limiting with C-IPC’s isotropic model to constrain strain within the elasticity range measured by Penava et al. [2014]. The broadest range in all directions for cotton allows a stretch factor up to 6.08%. Here, applying this bound, even with $0.01\times$ scaling of membrane stiffness, we now regain a simulation free from membrane locking and, since restricted to measured strain limits, also free of unnatural stretching artifacts (Figure 11d).

Finally, it is important to reiterate that membrane locking is indeed a resolution-dependent artifact whose impact can decrease as mesh size increases. For example, locking artifacts for the cloth draped on sphere at $1\times$ bending stiffness significantly decrease with a mesh of 85K nodes and are imperceptible at 246K nodes. However, this improvement is highly dependent on scene parameters. For example, the same 246K mesh, in the exact same scene, exhibits significant locking artifacts if we just reduce bending stiffness by $0.01\times$ to match the material in the strain-limited simulation in Figure 6, bottom row. Please see our supplemental document for details and illustration of these cases.

6.2 Exact Strain Limits

Above we have demonstrated the well-known importance and impact of applying strain-limiting for simulating cloth materials. Here

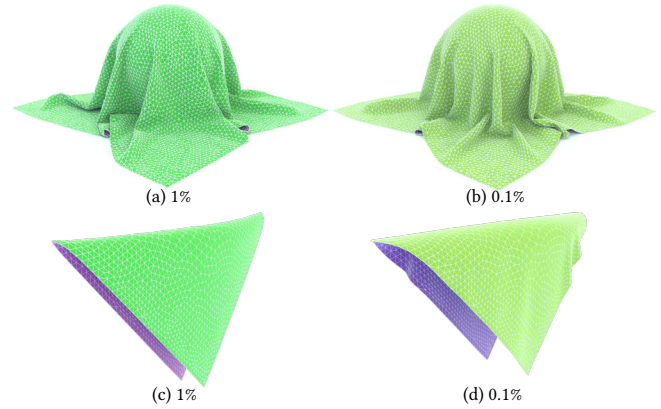


Fig. 12. **Extreme Strain Limits.** We confirm C-IPC’s constitutive strain limiting model provides robust and guaranteed satisfaction down to tightest measured and practically applied limits for cloth (and well below). Here we test C-IPC with strain limits of 1% (already well below most measured limits for cloth) in (a) and (c). To push the challenge further in (b) and (d) we then test C-IPC with an excessively small, largely nonphysical strain limit of 0.1%. Here the applied simulation meshes are low resolution, containing respectively only 8K (for a and b) and 2K (c and d) nodes, and so are challenging to avoid locking artifacts with real-world materials. We confirm that C-IPC preserves requested limits across all elements in all simulation time steps. We also note that, as expected, with the excessively small, nonphysical strain-limits applied in (d) more constraints become active w.r.t. to the DOF count resulting in a small amount of visible membrane locking artifacts; see our discussion below.

we investigate C-IPC’s ability to accurately enforce tight strain limits across both our isotropic and anisotropic models.

Accuracy across decreasing strain limits. To confirm accuracy, controllability and robustness of C-IPC’s constitutive strain limiting we test increasingly severe strain limits of 1% and 0.1% (well below measured limits in standard cloth materials [Clyde et al. 2017; Penava et al. 2014]). We apply them to both the sphere drape example considered in the last section as well as a simple membrane locking test with a two-corner pinned cloth [Chen et al. 2019; Jin et al. 2017]. Here, applying C-IPC, see Figure 12, we observe stable simulation output with stretches on all triangles satisfying their prescribed limit constraints. Closer examination of the draped sphere tests in Figure 12 also confirms that both are free from artificial bending stiffness. For contrast, compare these results with the artificial stiffening in Figure 11a, where the maximum stretch exceeds 4.5% despite applying the measured membrane stiffness directly.

Here, as we are applying our strain-limits *unilaterally*, this agrees with Jin et al.’s [2017] argument that unilateral strain-limits can avoid membrane locking better than bilateral enforcement – at least to a certain degree. Even unilateral enforcement is not a perfect panacea if limits are too tight. For an extreme example we can push strain-limits to a very small and largely unrealistic (0.1%) strain limit. This provides an extreme stress test and we confirm C-IPC does preserve these challenging limits across all time steps. As expected, however, here we can finally observe some slight but distinct sharp edge-creasing artifacts in Figure 12b and similarly visible locking issues in Figure 12d. Although such an extreme limit

is unlikely in practice and generally not encountered for cloth, this experiment does highlight an important point: if most unilateral strain-limit constraints are active then they behave similarly to the bilateral case. Thus, when most strain limit constraints are at the bound active constraint numbers can again approach the number of DOFs, and so, as discussed in Chen et al. [2019], locking can once again be encountered. Here we focus on formulating a controllable and robust constitutive model for strain limiting. We hope this will then enable further increasing the range of locking-free configurations via explorations of alternate discretization for the strain constraints as in Chen et al. [2019]. However, we note that this can best be enabled when the underlying method, as proposed here, can accurately guarantee coupled, constraint satisfaction.

Anisotropic strain limiting. For anisotropy, the story is similar. In Figure 13 we demonstrate C-IPC’s anisotropic model with both the sphere drape and two-corner hang tests. Again we observe that applying measured stiffness values, here from Clyde [2017], generates clear membrane locking artifacts (see Figure 13a). Next in Figures 13b and c we scale down membrane stiffness by $0.1\times$ and $0.01\times$ respectively. Throughout we confirm C-IPC enforces Clyde’s measured strain limits and again obtaining the expected improved results, both without locking artifacts and without non-physical stretching, at $0.01\times$ the measured stiffness. Similarly, in Figure 13d (again at $0.01\times$ reported stiffness) we apply C-IPC to preserve measured strain limits, obtaining artifact-free wrinkling with the anisotropic model.

6.3 Comparison with Splitting Models

Existing strain-limiting methods introduce errors from two primary sources. First, at the modeling level, they split strain limiting from the resolution of elasticity. Second, irrespective of the splitting model employed, algorithms applied to solve them are inaccurate and




Fig. 13. **Anisotropic strain limiting.** We apply the C-IPC anisotropic strain-limiting model to the sphere drape test in (a) through (c), varying only scaling of the measured membrane stiffnesses, and to the two-corner cloth hang test in (d). All are time stepped at $h = 0.04s$ with their final equilibria shown here. In (c) and (d), with a $0.01\times$ scaling of the measured stiffness, we see that C-IPC’s tight enforcement of the measured strain limits [Clyde 2017] obtains artifact-free drapes with differing wrinkle patterns resulting from the anisotropic model.

often unable, as we will see in the following section, to satisfy even moderate requested strain-limits. All prior methods then introduce errors from both of these sources and so it has remained unclear what problems originate from each source. Here, we first separately analyze the problems that are created by splitting models, by solving each step of the split to tight accuracy. This allows us to show that the splitting itself introduces errors that are unavoidable irrespective of how accurately limit constraints could be enforced. Then, next in Section 6.4, we will examine solution accuracy and see that errors in the strain-limit solve itself then produce inconsistent and so often uncontrollable results for simulation.

To examine splitting errors we begin by applying the standard strain-limiting, splitting strategy and so divide each time step solve into two sequential steps. The first solves a predictor step that includes all forces for the whole system, *with the exception of the strain-limit*, to obtain an intermediate configuration \hat{x} . Next, the second step *projects* the predictor to satisfy the strain-limit. To do so we minimize our constitutive strain limiting energy summed with the mass-weighted L^2 distance from the final timestep solution to \hat{x} .

In the simplest case we can consider the effect of this splitting when there are no contact forces. We start with a pinned and pulled square cloth. We fix its two top corners and apply heavy weights to pull its two bottom corners downwards; see Figure 14a and b. For comparison, without strain limiting the cloth is stretched well over $10\times$ (not shown) while with our constitutive strain limits, strains are restricted to the measured elasticity range and, in this fully coupled solution, we see the expected vertically aligned wrinkles from stretching at the two bottom pulled corners; see Figure 14a. On the other hand, in Figure 14b we see that splitting strain limiting from the elasticity solve produces obvious biasing artifacts at the two bottom corners where the wrinkles are aligned in non-physical directions. These errors, resulting from the decoupling of strain-limiting forces, are then increasingly severe as time step sizes increase (here we show with $h = 0.04s$).

Next, to consider splitting errors with contact, we again consider the cloth on sphere example. Here we apply a neo-Hookean membrane elasticity to help the splitting method avoid possible triangle degeneracies. Note that even here our fully coupled C-IPC model does not actually require the neo-Hookean elasticity (as triangle degeneracies are prevented by point-triangle constraints between neighbors). However, for consistent comparison, we apply neo-Hookean models for both. Here in Figure 14d we then see that the two-step splitting with contact now suffers from severe compression artifacts, which again come from resolving elastodynamics in the first step and then applying strain limiting and contact forces in the second. For comparison, consider C-IPC’s corresponding, fully coupled strain-limited solution in Figure 14c. In turn the error in the split solution suggests that an additional lower bound on strain (e.g. as applied in ARCSim) might be helpful to avoid these compression errors. If we then additionally add a strain-limit lower bound (here at 0.7) then we indeed find the resulting splitting solution is now free from severe compression artifacts. Now however, due to the errors in membrane and bending, the cloth in the split solution still remains unnaturally flat against the floor; see Figure 14e.

Of course as with all time-step splitting methods, splitting errors can be reduced by applying increasingly smaller time step sizes.



Fig. 14. **Comparison of C-IPC’s fully coupled strain limiting model with standard splitting models** highlights the differences between the fully coupled solutions in (a) and (c) and the artifacts generated by splitting strain-limiting solves from elastodynamics in (b), (d) and (e). Here we see that with splitting resulting large errors in elasticity (membrane and bending) can generate incorrect wrinkling directions (b), severe compression artifacts (d), or even numerical softening as in (d) and (e).

Here, for example, we find that visually noticeable errors between the split model and our fully coupled solve disappear at $h = 0.004s$. However, as expected, the resultingly large and often unacceptable increase in compute time wipes out any expected gains from splitting (we will see this theme again in the next sections’ comparisons with existing cloth codes). Likewise, the necessary decrease in time step size then varies with example and scene so that robust, controllable time-stepping with splitting, does not appear possible.

6.4 Comparison with ARCSim’s Strain-Limit Solver

Above we analyze the *modeling* errors generated by time-step splitting with strain limiting. To do so we compare models for which each stage is consistently solved to tight accuracy in our common benchmark code. Here we next consider the joint impact that splitting errors combined with inaccurate constraint solves then have in existing cloth simulation codes. To do so we compare with ARCSim [Narain et al. 2012] which is, to our knowledge, the most robust shell simulator currently available to support strain limits with real-world cloth material parameters. See Figure 16 for statistics on the strain limits and timings achieved by ARCSim and C-IPC here.

We again begin by considering the simpler, contact-free case. To do so we consider an even simpler cloth hang example: by dropping the pinned cloth *without* weighted bottom corners. We apply ARCSim’s default cotton material and default strain limit settings which restrict to the range $[0.9, 1.1]$. Here, even with this wide range of permissible strains and at full membrane stiffness, ARCSim results violate the limit bounds significantly at each time step. For visual reference in Figure 15a we illustrate the fully coupled C-IPC solution at equilibrium satisfying the strain-limit bound. Compare with the ARCSim output at equilibrium in Figure 15b where we easily see the artifacts resulting from over-stretching near the fixed corners forming a larger arc. Here both frames are taken at 4s and are time-stepped with $h = 0.04s$.

As we then decrease ARCSim’s time-step size to $h = 0.01s$ these errors still remain significant, and it is not until we reach at step size of $h = 0.001s$ that strain limits are mostly (although still not entirely) satisfied. However, to do so ARCSim then requires 87 minutes to simulate this simple 4s simulation sequence, resulting in a well-over $10\times$ slowdown compared to C-IPC’s fully coupled, strain-limit-satisfying solution stepped at $h = 0.04s$. However, beyond the requisite slowdown, the required decrease in step size changes per scene and material and so it is always unclear, without many time-consuming simulation tests per scene, how to determine the

necessary time step decrease to ensure ARCSim results satisfy a prescribed strain limit. In contrast C-IPC again maintains strain-limit constraint satisfaction across time step, material and scene settings. To investigate ARCSim’s constraint errors further we also observe that the augmented Lagrangian strain-limiting solver [Narain et al. 2012] employed in ARCSim applies a hard-coded max-iteration cap set to 100. Experimentally increasing this limit (and otherwise leaving the ARCSim code unchanged) we confirm that even increasing this cap by an order of magnitude, strain limit satisfaction is still not achieved.

Next we add contact for our ARCSim comparison and consider the 8K-node sphere drape test. Here we again observe that applying strain-limiting to reduce membrane locking does not work and instead see that strain-limit satisfaction worsens and artifacts actually increase as membrane stiffness decreases.

Specifically we start by applying ARCSim’s default cotton material and the same default strain-limit settings as above. For this stiff material we observe the expected membrane locking issues; see Figure 15d. However, we do note that with this stiff membrane, strain limit constraints are certainly reasonably well satisfied except for a few time steps.

Next, as standard, we reduce membrane stiffness in ARCSim with hope of mitigating the locking with the expectation that strain-limiting will compensate. If we *reduce* stretching stiffness by $10\times$, the locking artifacts are indeed less; however, ARCSim’s strain-limiting does not provide the intended compensation for the reduced stiffness. Instead, there are significantly more triangles violating the strain limit throughout the simulation (see Figure 15e) resulting in stretching artifacts. If we then further reduce stiffness to $0.01\times$ the default cotton value (generally necessary in this example to remove all visible locking artifacts; e.g. compare with our comparable stiffness, C-IPC result in Figure 11d and 15c) ARCSim’s results in Figure 15f, show even more significantly violated strain limits and so suffer from the same extreme compression issues exhibited by the split model analysis in the last section; see Figure 14e for comparison. In addition, we now also observe jittering resulting from ARCSim’s further three-way split that further separates the strain limiting solve and the contact forces into two separate, sequential projection steps. We note that this last issue is not new to our analysis here and is discussed in Narain et al. [2012].




Fig. 15. **Comparison of C-IPC and ARCSim results** illustrate the different behaviors of C-IPC’s fully coupled constitutive strain limiting solutions in (a) and (c) and the output of ARCSim’s split model and augmented Lagrangian strain-limiting solver in (b) and (d-f). Here frames from ARCSim near equilibrium exhibit excessive stretch due to lack of limit enforcement in (b), (e) and (f). At default material settings in (d) ARCSim exhibits the expected locking artifacts. Then, as we decrease ARCSim’s membrane stiffness to avoid membrane locking in (e) and (f) its strain-limiting can not preserve bounds, while the split model introduces additional artifacts so that the over-compression and stretching artifacts are now most obvious. For reference compare with C-IPC’s strain-limited, fully coupled solutions in (a) and (c) at comparable stiffness (0.01 \times) and same times and see Figure 16 for summary of timings and limits achieved.

Test	Method	h (s)	Max s_{max}	Avg s_{max}	Min s_{min}	Avg s_{min}	Time (m)
Hang 0.01 \times	C-IPC	0.04	1.0988	1.0884	0.9801	0.9824	6.5
Hang 1 \times	ARCSim	0.04	2.1281	1.4836	0.7686	0.8760	2.0
Hang 1 \times	ARCSim	0.01	1.8959	1.3197	0.8127	0.8879	8.0
Hang 1 \times	ARCSim	0.001	1.1110	1.1009	0.8959	0.9093	87.0
Sphere 0.01 \times	C-IPC	0.04	1.0617	1.0535	0.9548	0.9652	7.5
Sphere 1 \times	ARCSim	0.04	1.4176	1.0875	0.2397	0.9220	1.0
Sphere 0.1 \times	ARCSim	0.04	1.3796	1.1587	0.0568	0.8599	1.5
Sphere 0.01 \times	ARCSim	0.04	1.7944	1.3216	0.6714	0.8445	1.0

Fig. 16. **Strain limiting comparison with ARCSim.** We compare ARCSim and C-IPC strain limiting on two cloth simulation test cases: a pinned hanging cloth (without contact) and fixed sphere drape (with contact). For both methods we apply ARCSim’s default strain limit settings targeting strains to $[0.9, 1.1]$. Above we summarize statistics of the achieved max, s_{max} , (respectively, min s_{min}) strain-limit over all elements in a timestep as the averaged and max (respectively min) over all time steps. We apply C-IPC stepped at $h = 0.04$ with stiffness rescaled by 0.01 \times (necessary to avoid membrane locking artifacts and so requiring accurate strain-limiting to compensate) for both examples. We correspondingly consider ARCSim across a range of time-step sizes and stiffness re-scalings (1 \times , 0.1 \times and 0.01 \times). See Figure 15 for visualization of results. In the hang test, even for the full membrane stiffness, we see ARCSim’s strain limits are far from satisfied. Then, as step size decreases we observe ARCSim results better satisfy strain limits – although never fully. Correspondingly, in the sphere drape, as we reduce membrane stiffness towards a value required to mitigate locking, we ARCSim results increasingly violate strain limits more significantly. In comparison we observe C-IPC satisfies strain limit satisfaction at the necessary soft membrane energy, across all elements and time steps for both tests.

6.5 Thickness Modeling

Next we examine modeling small but finite thicknesses in contact. To illustrate differences between the C-IPC model and existing state-of-the-art shell simulators we examine a simple, yet challenging cloth stacking benchmark. We construct a scene with a stack of ten 8K-node square cloth panels aligned with vertical spacing and rotated orientation, dropped simultaneously onto a fixed square board. See Figure 17a for initial set up. This should form a cloth pile




Fig. 17. **Cloth stack comparisons.** Ten 8K-node square cloth panels are dropped in stack upon onto a square board with friction $\mu = 0.1$ (a). Here ARCSim [Narain et al. 2012] and ARGUS [Li et al. 2018a] are tested for varying cloth thicknesses. Time stepping at $h = 0.001s$ ARCSim fails to resolve contact for a thickness of 5mm (b), and at 10mm there is no longer interpenetration but the dynamics and geometry have noticeable artifacts (c). ARGUS can successfully simulate cloth with 5mm and 10mm thicknesses (e,f) at $h = 0.001s$ with some smaller artifacts; but at 1mm thickness it fails to resolve contact altogether. See the top row Figure 18 for our corresponding IPC results for this example.

where as we vary material thickness the overall pile height and bulk properties should also correspondingly change.

ARCSim [Narain et al. 2012] computes³ collision response applying Harmon et al’s [2008] non-rigid impact zones and a repulsion thickness applied to both model cloth thickness and to stabilize collision-processing constraints. ARCSim’s thickness parameter default is set to 5mm and is exposed as a configuration that can be changed by users. This parameter is then associated in code with additional thickness parameters directly applied for collision projection, analysis and contact force computation. Starting with the default 5mm thickness setting, and testing with increasingly smaller time steps down to $h = 0.001s$, ARCSim consistently reports collision handling failures and demonstrates severe artifacts in the simulation results; see Figure 17b. Next applying a 10mm thickness ARCSim reports success in resolving contacts at $h = 0.001s$, but still generates large artifacts in both geometry and dynamics; see Figure 17c. We did not push ARCSim further to time steps smaller than $h = 0.001s$, as this already makes ARCSim excessively slow for computation; e.g. as compared to C-IPC.

ARGUS [Li et al. 2018a] updates ARCSim with improved contact and friction processing for shell simulation. As in ARCSim it also

³We use the most recent ARCSim v0.3.1 in our testing.

applies and exposes to users the same thickness parameters. With default solver parameters, at $h = 0.001s$ ARGUS is able to simulate the cloth stack with the default $5mm$ material thickness (Figure 17e) and also with a thicker $10mm$ material (Figure 17f). Here we also observe that some height differences in the piles are captured. However, as we decrease thickness towards thinner, more standard garment material thicknesses, e.g. setting thickness to $1mm$, ARGUS produces severe artifacts as shown in Figure 17d. As with ARCSim, for all three thickness settings ARGUS' timing remains much slower than C-IPC.

In summary we observe that state-of-the-art methods require small and often impractical time step sizes to avoid failures as modeled geometric thickness decreases. With smaller thickness and/or increasing collision speeds, the required time step size to avoid failure decreases, correspondingly increasing simulation cost and making it unclear what time step size is required for success per scene and thickness.

C-IPC in contrast, for the same thickness benchmark consistently models the varying thickness by simply setting \hat{d} (recall the distance we start exerting contact forces at) to effective thickness values $1mm$, $5mm$, and $10mm$. In the first row of Figure 18 C-IPC successively models all varying material thicknesses and corresponding heights of the cloth stacks. C-IPC's model then further provides constitutive behavior for thickness in the normal direction. This is demonstrated in middle and bottom rows of Figure 18 where we then drop an elastic ball on the stack. In the middle row we show maximum compression for each stack's collision, observing the increasing indentation with thickness. Also note the wrinkles forming solely in the thickest, $10mm$ case. In the bottom row we show the final frame at equilibrium for each simulation after collision, illustrating the different rest heights of the shell piles weighted by the volumetric FE ball model. This example also illustrates C-IPC's natural coupling of differing elastic models which we explore in detail in Section 6.9. These examples also help illustrate how computational cost for C-IPC can vary with the threshold \hat{d} ; see Figure 24. Generally we see




Fig. 18. **Sphere on cloth stack.** Ten 8K-node square cloth are dropped onto a square board with friction ($\mu = 0.1$). Top: Varying IPC \hat{d} per column, C-IPC is able to provide controllable thickness modeling for cloth. Middle: A soft elastic volumetric sphere ($E = 10KPa$) is dropped onto to the cloth stacks and is illustrated here at maximum compression; here with $10mm$ thickness (far right) we observe the resulting wrinkling enabled by the thicker depression. Bottom: at rest equilibrium heights continue to vary with thickness under the weight of the ball.

Example	ACCD	IPC	Root Parity	MinSep	BSC	T-BSC
Sprinkles	72.8	tiny toi	—	tiny toi	—	—
Noodles	98.0	tiny toi	—	tiny toi	—	—
Hair clusters	855.5	tiny toi	—	tiny toi	—	—
Braids	965.4	1012.8	—	tiny toi	—	—
Card Shuffle	165.2	tiny toi	—	tiny toi	—	—
Needle bed	43.4	55.6	intersect	tiny toi	error	∞
Kick	229.6	262.3	intersect	tiny toi	error	tiny toi
Stack, 1mm	441.5	498.3	448.5	tiny toi	error	∞
Stack, 5mm	189.6	236.8	191.9	tiny toi	error	∞
Stack, 10mm	123.3	135.0	124.4	tiny toi	error	∞

Fig. 19. **CCD comparison statistics.** Here we summarize the per-Newton iteration, total CCD query time (excluding spatial hash) in milliseconds for ten simulation benchmark examples across CCD methods. We compare our ACCD method, IPC's conservative floating point CCD [Li et al. 2020a], MinSep [Harmon et al. 2011; Lu et al. 2019], Root Parity [Brochu et al. 2012], BSC [Tang et al. 2014], and T-BSC [Wang et al. 2015]. First five examples in the benchmark include finite thickness offset, which can not be supported by Root Parity, BSC, and T-BSC methods. When simulation cannot proceed from CCD errors we report type. Here " ∞ " indicates unable to find non-intersecting step (endless line-search) and "tiny toi" too small a step size (no progress).

that at larger threshold values, e.g. $\hat{d} = 10mm$, C-IPC processes more contact pairs per step, making simulation more expensive than at smaller values, with less contacts per solve, e.g. \hat{d} at $1mm$ and $5mm$. At the same time setting \hat{d} to values orders of magnitude smaller, e.g. $1\mu m$, leads to slower convergence, and so longer compute times, due to increased sharpness in the barrier functions.

6.6 CCD Benchmarking

For the last section's comparisons it was sufficient to model thickness effects in C-IPC solely via \hat{d} and so purely as an elastic behavior. More generally, as discussed in Section 5, as thicknesses become smaller and contacts tighter, C-IPC's inelastic thickness model utilizing both \hat{d} and our offset ξ is required.

For simulating codimensional objects these two mechanisms are combined together. Here larger \hat{d} introduces increased elastic response while a nonzero ξ provides guarantees for minimum thickness even under extreme compression; see e.g. Figure 2. For scenes where consistently modeling thickness matters; see e.g. Figures 2, 3, 10, 21, 22, and 23, we then set ξ to the reported material thickness value of the object or slightly smaller and then set \hat{d} near ξ to control the degree of elastic response required.

As discussed in Section 5.3 simulating these examples produces much more degenerate contact pairs and requires higher precision to capture thickness offsets and so makes great demands on CCD accuracy. Here the conservative CCD strategy from IPC [Li et al. 2020a] is no longer sufficient. Employing standard floating point CCD routines can and will simply return 0 TOI for many queries, and so incorrectly stop the IPC optimization progress altogether. Likewise alternative CCD methods with higher reported accuracies similarly fail in most such cases.

This motivates our new ACCD method which, as demonstrated below, remains robust and accurate – obtaining stable, forward progress in IPC-optimization for even our most complex examples where all available CCD alternatives fail. At the same time ACCD continues to provide comparable and generally faster timings on easier examples where alternate CCD methods are able to succeed.

For comparison with ACCD we test a range of CCD methods: Minimal Separation (MinSep) CCD [Harmon et al. 2011; Lu et al. 2019], Root Parity [Brochu et al. 2012], BSC [Tang et al. 2014], T-BSC [Wang et al. 2015] and IPC’s conservative floating point CCD [Li et al. 2020a] on a set of ten benchmark examples with challenging codimensional collisions. The first five examples in the benchmark have thickness offsets and the remaining five do not; see Figure 19.

We are able to directly test MinSep by applying it as a base solver within the same conservative CCD strategy as IPC [Li et al. 2020a]. Here minimal separation is set to $s \times d_{\text{cur}}$ with the current scaling factor s and distance d_{cur} . If the returned TOI is less than 10^{-6} the query is performed again without scaling ($s = 0$) in order to attempt to make the query easier. To test the Root Parity, BSC and T-BSC methods, which only decide whether an interval "has collision" or "has no collision" without computing TOI, we apply them with backtracking for each query until finding a step size without collision. When found, we then return a conservative scaling of the step by $1 - s$. As backtracking cannot directly extend these methods to compute times or step sizes that *first* bring a distance of a contact pair to ξ , we can not test these three methods on our examples with thickness offsets; and so restrict our testing of them to just the last five test examples in Figure 19).

We summarize our comparisons in Figure 19. IPC’s conservative strategy with floating-point CCD works well on simpler scenes, albeit with slightly slower timings than our ACCD method, but then fails on the even more complex ones. We also note that while conservative floating point CCD handles degenerate cases like parallel edge-edge by querying additional point-edge and point-point pairs, ACCD only needs to query general pairs (e.g. point-triangle and edge-edge pairs for surface-only scenes). On the other hand, MinSep always returns tiny values, sometimes even 0, for all ten benchmark examples and so makes no progress. Root Parity performs well with nice efficiency on most benchmark examples it can be applied to (those without thickness offsets) but unfortunately misses some collisions altogether leading to unacceptable interpenetrations in examples like Needle Bed and Garment. For BSC we find it reports runtime errors for all benchmark examples it can be applied to with an internal message "Inflection points are not handled exactly in BSC", while T-BSC returns tiny values for the Garment example, and is unable to determine step sizes without collision (we confirm both BSC variants work well on simple cases) and so becomes trapped in line search in the other five benchmark examples it can be applied to. In contrast we see that ACCD enables all examples to complete to convergence while, for examples where alternate CCD routines are able to complete, it offers fastest times. Finally, as a preliminary investigation of ACCD application outside of the C-IPC framework we apply ACCD in place of ARGUS’s default CCD routines and test with both 1mm and 5mm cloth stack examples. While, as expected, the simulation results are nearly identical (with no visual difference) CCD costs decrease by $\sim 50\%$.

Example	Method	Nodes	h (s)	Time (m)	SL	Contacts (avg., max.)	Cores
Arabesque	C-IPC	31.9K	0.032	0.32	1.260	6.3K, 15.9K	4
	C-IPC	31.9K	0.032	0.23	no	4.8K, 16.5K	4
	Argus	31.9K	0.002	0.87	no	—	4
Clubbing	C-IPC	31.9K	0.032	0.57	1.253	4.5K, 16.8K	4
	C-IPC	31.9K	0.032	0.54	no	6.3K, 16.8K	4
	Argus	31.9K	0.002	0.82	no	—	4
Rhumba Layer	C-IPC	43.3K	0.04	1.2	1.134	31.7K, 43.7K	6
Kick Layer	C-IPC	43.3K	0.04	2.2	1.134	51.5K, 82.0K	6
Rhumba Knife	C-IPC	12.8K	0.04	0.4	1.134	15.4K, 21.3K	6

Fig. 20. **Garment simulation statistics** for a pair of test dance sequences, Arabesque and Clubbing, from ARGUS and three additional tests with garment patterns not handled by ARGUS. Here **SL** denotes the strain limit applied.

6.7 Garment Simulation

We next apply C-IPC to simulate garment drapes and dynamics.

ARGUS Comparisons. We start by considering C-IPC on a pair of test dance sequences, Arabesque and Clubbing, proposed for ARGUS [Li et al. 2018a]. Both examples use the highest resolution input meshes tested in Li et al. [2018a]. Here Clubbing exercises contact processing more extensively, with faster motions and transitions than Arabesque. For comparable results we run C-IPC and re-run ARGUS on the same machine, with fixed mesh resolution for ARGUS. Here we observe robust, smooth simulation results across the entire sequences for C-IPC both with and without strain limits. Notably, even with the addition of strict non-intersection enforcement, convergent, fully implicit time steps and resolving all contact stencils, not solely node-node as in ARGUS’s treatment, we still observe 3.7X and 1.5X speedup for C-IPC on Arabesque and Clubbing over ARGUS results. Likewise, if we also enable C-IPC’s constitutive strain-limiting, speedups remain similarly at 2.7X and 1.4X. See Figure 20 for comparative timings and our supplemental video for results.

Garment Simulation Challenges. Two primary challenges in garment simulation are resolving complex, multilayered designs and simulating clothing worn by rapidly moving characters. To test the robustness of C-IPC for these cases we construct two new examples. Due to their complex stitching patterns these garments can not be simulated in the ARGUS and ARCSim codes.

We simulate the yellow dress (15.7K nodes) produced by FoldSketch [Li et al. 2018b] with challenging-to-resolve knife pleats created where the cloth is folded upon itself and then stitched tight in contact. We start C-IPC with the flat pattern staged around the 12.8K-node mannequin (Figure 5a). We then drape by stepping C-IPC forward with the stitching forces (springs) seaming the design together at large time step. C-IPC obtains the static equilibrium of the draped garment after just a few large time steps obtaining sharp knife pleats while preserving the correct layering order, free of intersection (Figure 5b). We then move the knife pleat through




Fig. 21. **Sprinkles.** We drop 25K 6mm-long and 2mm-thick “sprinkles” into a bowl (a). Setting IPC’s offset and \hat{d} to 1.5mm and 0.5mm respectively to model each sprinkle’s thickness C-IPC fills the bowl’s volume at rest (b). In (c) we remove the bowl to highlight the non-interpenetrating rest configuration while in (d) we remove volume rendering as well to show the underlying edge segments simulated.

the steps of the Rhumba, preserving the flowing motion of the dress with knife pleats remain crisp and properly colliding throughout (Figure 5c).

We then apply C-IPC to similarly initialize a multilayer skirt pattern (30.5K nodes) designed using Sensitive Couture [Umetani et al. 2011]. This obtains the initial pose static drape in Figure 5d. Next we script the mannequin mesh dressed in the layered garment through a fast-moving martial arts sequence⁴ with large motions. In the bottom row of Figure 5 we show three frames of the resulting C-IPC simulation featuring a portion of a rapid kicking sequence. Please also see our supplemental video. C-IPC stably resolves these rapid movements, continuing at a large time step of $h = 0.04s$, while capturing intricate garment details, enforcing strain limits and non-intersection (and so preserving the layering order) throughout. Finally we take the same garment through the slower Rhumba sequence and observe a correspondingly reduced cost as less iterations are required to resolve the time steps with the gentler motion and correspondingly lower contact count.

6.8 Benchmarks for Cloth Simulation

In this section we confirm C-IPC’s performance on challenging benchmark tests from prior work.

Cloth on rotating sphere. In Figure 6 we test spinning sphere-type examples from Bridson et al. [2002] designed to exercise friction, contact processing and tight wrinkling. We start by dropping an 86K node square cloth (unstructured mesh with $\mu = 0.01$ and 1.0608 strain upper bound) onto a sphere and ground (both $\mu = 0.4$). As the sphere rotates, friction catches the cloth and fine wrinkling is captured as the cloth is pulled inward without locking and stretching artifacts. We confirm that as we increase resolution (to 246K nodes) and reduce bending stiffness (by 10 \times) higher frequency wrinkles resolve in the folds.

Funnel. To test layering, friction and tight contact with strain limiting we extend the funnel examples from Tang et al. [2018] and Harmon et al. [2008]. In Figure 7 we drop three cloth panels ($\mu = 0.4$, 26K nodes each, strain upper bound at 1.0608) upon a funnel. Under friction they rest stably on top until a sharp, scripted collision object

⁴We use a motion sequence from Mixamo <https://www.mixamo.com/>, and apply Fang et al. [2021] to obtain an intersection-free mannequin sequence.

(four-node tetrahedra) rapidly thrusts the panels down and through the funnel. C-IPC solve all time steps in the sequence to convergence while preserving an intersection free, strain-limit satisfying trajectory throughout.

Reef knot. To combine extreme compression, contact and friction with large stresses we match the initial set up of the reef knot example from Harmon et al. [2009]. Two ribbons are initially intertwined and then stretched to form a knot in Figure 4. Here we extend the challenge of the original test by simulating with 100K nodes (10 \times that of the original), applying a strain limit of 1.134, friction with $\mu = 0.02$, and pulling much tighter – stretching the ribbons to nearly their strain limit. This captures the frictional dynamics of the trajectory (without jitter) and results in extreme stress at the center, forming a much smaller knot than originally demonstrated.

6.9 Shells, Rods, Particles and Coupling

In this section we demonstrate C-IPC simulation across all codomains. This covers volumetric FE materials, shells, rods and even, when desired, particles. Because these domains interact via IPC barriers we are able to seamlessly and directly couple all such mesh-based models in the same simulation. These tests also highlight the advantages and suitability of our consistent and controllable thickness modeling for these thin material interactions.

Card shuffle. We test accuracy and precision in very thin shell frictional contact during a complex, interleaving shuffle. We begin with fifty-four stiff shell playing cards (63.5mm \times 88.8mm) with IPC offset and \hat{d} set to 0.1mm and 0.2mm respectively to resolve card thickness. In Figure 8a we break the cards into two piles and apply moving boundary conditions on their sides to bend them in preparation for a “precision” bridge shuffle. Unlike a human-performed bridge finish, where cards begin interleaved, here we increase challenge by keeping the two bridged piles apart. We then release the held boundary conditions card-by-card from the bottom up, alternating between left and right side piles (Figure 8b). Released cards quickly regain flat rest shapes and fall on the ground to form a new, shuffled, intersection-free stacked pile. We then partially square the deck with two scripted collision boundaries (Figure 8c) obtaining a final pile 15.4 mm high and so well matching that of a




Fig. 22. **Granules on cloth.** We drop a column of 50K particles onto a cloth with fixed corners. Each particle is represented by a single node, with C-IPC offset and \hat{d} set to $2mm$ and $1mm$ respectively in order to capture the granule’s volume at rest while the C-IPC barrier ensures particles do not tunnel through each other nor the cloth shell. Particles at rest capture tight spherical packings (see zoom in) confirming offsets are preserved throughout. Finally we release one corner of the cloth to let the granules fall onto the ground.

real deck (Figure 8d). Here we applied $0.5\times$ the default tolerance to capture the intricate dynamics.

Hairs. Modeling hair with self-collision is a severe challenge for rod contact processing – in part due to their extremely small cross-sections. We examine C-IPC’s behavior on two challenging hair simulation tests. Hair thickness is small, but its bulk volumetric effects are important and have remained difficult to capture. Here C-IPC captures hair thickness by enforcing an offset of $0.08mm$ for all rods with $\hat{d} = 0.1mm$. In Figure 3(a)-(e) we first test C-IPC with braiding. We form a tight braid by twisting the bottom DOFs of two clusters of hair about each other. Despite undergoing large stress the resulting braid remains intersection free and preserve the thickness offset throughout. We also confirm this in dynamics by releasing the bottom held boundary conditions and letting the braid unwind without snagging artifacts nor instabilities. Next, we test hair self-collision and friction further by dropping a cluster of hairs with one free end across a similar perpendicularly hanging cluster [McAdams et al. 2009] and over a fixed sphere in Figure 3(f)-(k). C-IPC again simulates this scene with convergent, large time step solves throughout.

Noodles. A large bowl of noodles is delicious but also challenging to simulate due to the large numbers of DOF and high contact counts. In Figure 10a we drop a stack of 625 $40cm$ -long cooked noodles into a $13.85cm$ -diameter bowl (fixed geometry). Each noodle is a 200-segment discrete rod with $\mu = 0.4$ and offset and \hat{d} set to $1mm$ and $0.5mm$ respectively to model the noodles’ $1.5mm$ thickness. With C-IPC’s thickness handling the dropped rods come to a rest (their total volume should be $4.42 \times 10^{-4} m^3$) largely filling the bowl (Figure 10b) whose volume is $6.96 \times 10^{-4} m^3$. At equilibrium In Figure 10c we remove the bowl to show the final intertwined, interpenetration-free configuration. Finally in Figure 10d we remove the volume render to highlight that this complex non-intersecting volume is simulated solely with discrete rod polylines.

Sprinkles. Similarly we can do the same for just edge segments. We drop a grid of 25K “sprinkles” each formed by $6mm$ -long, single edge segment into a bowl (Figure 21a). The sprinkles’ $2mm$ thickness are modeled by a C-IPC offset of $1.5mm$ and \hat{d} of $0.5mm$. Making the offset $3\times$ larger than \hat{d} models a strongly inelastic response for their collisions. As with the noodles we can estimate the total expected volume of the sprinkles as $5.76 \times 10^{-4} m^3$ which at equilibrium fill the bowl’s $6.88 \times 10^{-4} m^3$ volume without interpenetration (Figure 21b and c). Again in Figure 21d we remove the volume rendering to highlight the geometry is simulated solely with codimensional line

segments. For both the noodles and sprinkles we render the fitted cylindrical surface $0.2mm$ thinner than $\xi + \hat{d}$.

Spherical Granules on cloth. In the extreme C-IPC can even simulate small spherical granules by treating individual vertices as particles endowed with a strict thickness offset. We test resolution of non-intersecting particle trajectories by dropping a column of 50K particles onto a cloth with fixed corners (Figure 22). Collisions remain non-intersecting even at high-speed and with the cloth. Particles at rest capture tight spherical packings (see zoom in) confirming offsets are preserved. Once particles pile in the center we release a cloth corner letting them flow to the ground. To model each particle we compute mass with $\rho = 1600kg/m^3$ and $V = 4/3\pi r^3 \approx 4.2 \times 10^{-9} m^3$, and set thickness with a C-IPC offset of $2mm$ and \hat{d} of $1mm$. Simulating particles this way then resembles the Discrete Element Method (DEM) [Cundall and Strack 1979; Yue et al. 2018] where particles are “coated” with an interaction zone (analogously $\xi < d < \xi + \hat{d}$ in our setting) for contact forces determination based on the depth, area, or volume of the zones’ intersections [Jiang et al. 2020]. DEM is a critical tool in geomechanics where it provides a reasonable simplification by ignoring particle deformation while accurately capturing momentum and inter-particle contact forces. Here C-IPC relates contact forces directly to distance via our barrier function and so enables seamless coupling with deformable materials as in the cloth here. Unlike DEM, C-IPC then also additionally guarantees intersection-free trajectories for particle volumes defined by offsets.

All-in. Finally we test coupling all codimension together. We drop an FE armadillo volume, a shell cloth and a particle column onto a net formed by interleaved rods (Figure 23a), applying a C-IPC offset of $0.5mm$ and \hat{d} of $1mm$ for all domains. We successively drop these codomains one at a time. First the dropped armadillo is tangled and caught between the net’s rods (Figure 23b). Next the cloth falls upon the armadillo (Figure 23c) and finally the column pours onto cloth and flows into the two resulting folds of cloth supported by the armadillo and the rod net (Figure 23d, e, and f (top view)).

6.10 New Stress Tests for Cloth Simulation

Finally we propose and evaluate C-IPC on a set of three new, challenging cloth simulation examples designed to require extreme robustness and accuracy in modeling elastodynamics with frictional contact, strain limiting and thickness. To our knowledge C-IPC is the first method to achieve simulations of these scenarios and behaviors.




Fig. 23. **All-in.** We simulate objects of all codimensions coupled together via C-IPC contact in a scene with an armadillo volume, a shell cloth, and a column formed by individual particles released in sequence to fall onto a net formed by interleaved rods.

Pulling cloth over codimensional needles. Cloth contacts against sharp geometries notoriously stress contact simulations with snagging and sedropping and then dragging a cloth panel across a needle bed formed by line segments (Figure 9). With a large time step size of $h = 0.02s$ and strain limit of 1.0608, the cloth comes safely to a static rest on the bed of needles without jittering. We pull the left side of the cloth at $1m/s$ across the segment tips and observe rapid sliding along and then off the top of the bed of needles without snagging nor stretching artifacts (strain limits are satisfied throughout). Likewise, if these contact are sliding across asperities with large displacement (i.e., large time step and/or large speed) these challenges only grow. Here we test C-IPC’s handling of extreme cases with both challenges, by In cases like these, were we have extreme stretch from sharp contacts, C-IPC’s adaptive strain-limiting stiffness is especially important to ensure that time step solves remain numerically feasible and efficient.

Tablecloth trick. The classic tablecloth trick attempts to pull a cloth out from under a table setting with the aim of keeping the dinnerware upright and ideally staying mostly in place. Key to making this trick work is a pull with sufficient speed so that sliding acceleration of the cloth can overcome friction forces with little or no pull on the settings⁵. Simulation of this effect then requires tight coupling of elasticity and strain limiting with accurate resolution of friction and contact as well as robust handling of sliding contact with sharp edges. This is needed to correctly model the balances between friction, the large stresses and forces from pulling, and the normal forces from heavy dinnerware. In Figure 1 we set a table for C-IPC with both heavy and stiff dinnerware (volumetric FE models). We simulate pulling the strain-limited (1.0608 upper bound) tablecloth, time-stepped at $h = 0.01s$, with increasing speeds. As we start with lower pulling speeds, as expected, objects do not follow the cloth. As we increase speed we then get less messy results and finally, with a pull at $4m/s$, the full setting remains upright and on the table with just a bit of movement (we also note detailed folds begin to be formed by the strain-limited cloth pulled at high speed). Finally, with a high-speed pull of $8m/s$ the tablecloth pulls out smoothly with almost no change to the setting at all. Here, for the latter two successful pulls we observe extreme forces from the moving boundary conditions pulling the cloth at a high speed and the large friction forces exerted by contact with the heavier dinnerware.

Twisted cylinder. Finally, we test thickness modeling, contact resolution and buckling under extreme and increasing stress. We start

with a $1m$ -wide cloth cylinder ($0.25m$ radius) modeled by a shell with 88K nodes, with a thickness offset of $1.5mm$ and \hat{d} of $1mm$, time-stepped at $h = 0.04s$. We then simultaneously twist the cylinder and slowly move the two sides together at $72^\circ/s$ and $5mm/s$ respectively to introduce wrinkling, folding, and eventual tight buckling. In order to allow this scripted torture to continue for the entire 38 seconds length of this sequence we do not apply strain limiting – so C-IPC must resolve exceedingly high stretching forces here. Likewise, to clarify buckling behavior we do not apply gravity (avoiding droop). In the first second of simulation we immediately obtain interesting global folds (see our supplemental document and video). Soon after a thick central cylinder of wound cloth forms. The volume of this cylinder’s structure is supported by C-IPC’s finite thickness offset (Figure 2 upper left). For contrast consider in Figure 2 upper right a frame captured at the same time step but now simulated *without* C-IPC’s offset. Here we clearly see that correct geometry *can not* be formed without C-IPC’s thickness offset (nor can it capture the material’s later buckling behavior – see below) further clarifying the importance of consistent thickness modeling for codimensional models. Next, as we continue our C-IPC simulation *with offset* for 32.96s of further twisting, our offset-thickened cloth continues to support complex behaviors including the final buckled geometry in Figure 2 bottom.

7 CONCLUSION

C-IPC extends the incremental potential contact framework from volumetric deformables [Li et al. 2020a] to codimensional and mixed-dimensional structures with controllable thickness, fully coupled strain limiting, accurate frictional behavior and non-intersection maintained throughout all time steps. Of course as in the original IPC method, in order to provide these non-intersection guarantees IPC requires an initial non-interpenetrating geometry at start of simulation. Currently, e.g. for garments, this can often be easily achieved by staging. However, loosening these constraints to achieve guarantees from a starting tangled state is an important and exciting direction of future exploration.

As demonstrated in our stress tests, C-IPC consistently converges to optimal solutions even when subject to challenging boundary conditions that push against strain limits and contact barriers. However, (C-)IPC can not progress with boundary conditions that impose infeasible (e.g., intersecting or strain-limit violating) trajectories. Here detecting and/or finding means for IPC to progress through infeasible conditions would be a challenging future extension. For friction, C-IPC directly utilizes the smoothed, semi-implicit discretization of maximal dissipation from IPC with the option of lagged iterations.

⁵www.juggle.org/table-cloth-pull-trick

Examples	# nodes	h (s)	SL	dhat (m), offset (m)	mu, eps_v (m/s)	contacts avg. (max.) per timestep	# cores	avg. t (min), iterations per timestep
Cloth on sphere	10.0K	0.04	1.0608	1e-3, 0	0.4, 1e-3	10.6K (13.7K)	4	0.1, 12.2
Cloth on sphere (aniso SL)	10.0K	0.04	*	1e-3, 0	0.4, 1e-3	8.2K (10.4K)	4	0.1, 12.4
Sphere on cloth stack (1mm)	85.0K	0.02	1.0608	1e-3, 0	0.1, 1e-3	526.5K (875.9K)	8	2.0, 8.5
Sphere on cloth stack (5mm)	85.0K	0.02	1.0608	5e-3, 0	0.1, 1e-3	876.2K (1.6M)	8	2.3, 8.6
Sphere on cloth stack (10mm)	85.0K	0.02	1.0608	1e-2, 0	0.1, 1e-3	1.7M (2.8M)	8	4.7, 12.9
Cloth on rotating sphere	87.1K	0.04	1.0608	1e-3, 0	0.4, 1e-3	176.7K (362.5K)	6	11.6, 76.0
Cloth on rotating sphere	247.5K	0.04	1.0608	1e-3, 0	0.4, 1e-3	539.8K (1.2M)	6	35.5, 70.7
Funnel	79.3K	0.04	1.0608	1e-3, 0	0.4, 1e-3	284.3K (377.3K)	8	9.9, 76.0
Ribbon knot	103.8K	0.04	1.134	1e-3, 0	0.02, 1e-5	1.6M (1.7M)	8	13.4, 62.5
Pulling cloth on needles	26.3K	0.02	1.0608	1e-3, 0	0, N/A	0.3K (6.6K)	6	1.8, 66.6
Table cloth pull (4m/s)	57.1K	0.01	1.0608	1e-3, 0	0.2, 1e-3	63.4K (89.1K)	6	0.7, 13.8
Twisting cylinder (with offset)	87.6K	0.04	No	1e-3, 1.5e-3	0, N/A	1.5M (2.3M)	8	10.2, 33.4
Cards shuffling	6.1K	0.02	No	2e-4, 1e-4	0.02, 3e-5	36.2K (46.5K)	6	0.7, 35.2
Braids	77.8K	0.01	N/A	1e-4, 8e-5	0.1, 1e-4	566.5K (1.5M)	8	3.5, 16.9
Hair clusters	91.6K	0.01	N/A	1e-4, 8e-5	0.2, 1e-4	71.9K (89.2K)	4	2.7, 31.0
Noodles	126.8K	0.02	N/A	5e-4, 1e-3	0.4, 1e-5	321.7K (454.0K)	8	9.8, 41.9
Sprinkles	51.2K	0.02	N/A	5e-4, 1.5e-3	0.15, 1e-5	170.0K (228.8K)	6	3.3, 69.0
Sand on cloth	58.7K	0.01	1.0608	1e-3, 2e-3	0.05, 1e-3	127.6K (185.0K)	8	10.6, 183.2
All in	48.2K	0.02	1.0608	1e-3, 5e-4	0.2, 1e-3	10.2K (27.8K)	8	5.8, 100.8

Fig. 24. **Simulation Statistics** for IPC on a subset of our benchmark examples. Remaining garment examples are summarized in Figure 20. For each simulation we report geometry, time step, strain-limit enforced (SL), thickness settings (\hat{d} and offset), friction parameters (μ and accuracy ϵ_μ), number of contacts processed per time step and machine, as well as average timing and number of Newton iterations per time step solve. (*The tested anisotropic model has strain limits at 0.14 for both diagonal entries and 0.063 for the off-diagonal entry of the Green-Lagrangian strain [Clyde 2017].)

As validated in the original IPC work, we confirm that C-IPC accurately captures correct stick-to-slip thresholds with convergent lagging (see our cloth sliding benchmark). Otherwise, for C-IPC simulation we apply a single lagged iteration and find it sufficient to model complex frictional behaviors (see e.g. the tablecloth pull and rotating sphere) and detailed garment drapes. However, as with the original IPC work (and prior nonsmooth methods) we also note that there is no guarantee of convergence to accurate satisfaction of fully implicit friction relations. An algorithm that achieves this remains an open and important challenge for simulation.

Along with enabling new and improved applications in cloth, hair and many other codimensional simulation tasks we are also excited to extend C-IPC further for geomechanics applications in the simulation of granular flows – especially for complex granule shapes which have so far proven challenging to resolve.

Here we have first focused on providing a reliable simulator for all materials and conditions. As noted in previous works, in many applications the biggest overhead is the many laborious simulation passes required to hand-tune collision parameters to make a scene work at all, and then often many more time-consuming steps to still correct output. As such, reliability leads to overall increased speed of output. However, although performance for C-IPC, as analyzed above, is more than competitive with state of the art cloth codes when they can offer comparable accuracies, much more can be done to improve C-IPC performance. We look forward to both better leveraging parallel architectures and improving underlying

IPC solver methods. Similarly, here we are investigating and extensively testing with *implicit* methods. Another fruitful direction of exploration is *explicit* time stepping. It would be exciting to consider relative behavior of C-IPC and ACM [Harmon et al. 2009] for side-by-side comparison between state of the art implicit and explicit methods providing non-intersection guarantees.⁶ Likewise, we look forward to leveraging the differentiability of our approach. Here we expect that the combined benefits of C-IPC’s reliable performance across sweeps of input and its differentiability, should together enable many design and training tasks for garments and complex structures as demonstrated in Macklin et al. [2020] and Geilinger et al. [2020]. Finally, here we have focused on providing strain-limiting solely with upper bounds and while we have not yet seen cases requiring lower bound limits these are easy to apply and so should be interesting for future exploration.

While proven critical for C-IPC we have also shown that ACCD is an exceedingly simple replacement for pre-existing complex, and often sensitive CCD modules. Here in preliminary testing we have found that in application outside of C-IPC framework ACCD can also provide significant speed-up. We look forward to further applications of ACCD as an extremely easy-to-implement and efficient alternative CCD algorithm for tasks in both simulation and geometry processing.

⁶Towards this end we have been communicating with the ACM authors for over year towards obtaining a representative code for comparison. After hard work on both ends we are hopeful, but licensing concerns required for access to the code have hampered us so far. We hope to have the comparison in future work.

In summary, independent from elasticity model and time step size, C-IPC guarantees out-of-the-box intersection-free simulation with a strict satisfaction of strain limits (confirmed down to 0.1%) and accurately captures geometrically meaningful thickness. This is enabled, in part, by a new, simple, additive CCD (ACCD) method with highly stable and accurate output for challenging, first-time-of-impact tasks.

ACKNOWLEDGMENTS

We thank Jacky (Jiecong) Lu for set-up of ribbon and garment examples; Jie Li, Florence Descoubes, and Laurence Boissieux for assistance with ARGUS; and Rasmus Tamstorf and Etienne Vouga for valuable discussion. This work was supported in part by NSF CAREER IIS-1943199, CCF-1813624, and ECCS-2023780.

REFERENCES

- David Baraff and Andrew Witkin. 1998. Large steps in cloth simulation. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*. 43–54.
- Jan Bender, Dan Koschier, Patrick Charrier, and Daniel Weber. 2014. Position-based simulation of continuous materials. *Computers & Graphics* 44 (2014), 1–10.
- Jan Bender, Matthias Müller, and Miles Macklin. 2015. Position-Based Simulation Methods in Computer Graphics. In *Eurographics (tutorials)*. 8.
- Jan Bender, Daniel Weber, and Raphael Diziol. 2013. Fast and stable cloth simulation based on multi-resolution shape matching. *Computers & Graphics* 37, 8 (2013), 945–954.
- Miklós Bergou, Basile Audoly, Etienne Vouga, Max Wardetzky, and Eitan Grinspun. 2010. Discrete viscous threads. *ACM Trans. Graph. (TOG)* 29, 4 (2010), 1–10.
- Miklós Bergou, Max Wardetzky, Stephen Robinson, Basile Audoly, and Eitan Grinspun. 2008. Discrete elastic rods. In *ACM SIGGRAPH 2008 papers*. 1–12.
- Kiran S Bhat, Christopher D Twigg, Jessica K Hodgins, Pradeep Khosla, Zoran Popovic, and Steven M Seitz. 2003. Estimating cloth simulation parameters from video. (2003).
- Sofien Bouaziz, Sebastian Martin, Tiantian Liu, Ladislav Kavan, and Mark Pauly. 2014. Projective Dynamics: Fusing Constraint Projections for Fast Simulation. *ACM Trans. Graph. (TOG)* 33, 4 (2014).
- Robert Bridson, Ronald Fedkiw, and John Anderson. 2002. Robust treatment of collisions, contact and friction for cloth animation. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*. 594–603.
- Tyson Brochu, Essex Edwards, and Robert Bridson. 2012. Efficient geometrically exact continuous collision detection. *ACM Trans. Graph. (TOG)* 31, 4 (2012), 1–7.
- George E Brown, Matthew Overby, Zahra Forootaninia, and Rahul Narain. 2018. Accurate dissipative forces in optimization integrators. *ACM Trans. Graph. (TOG)* 37, 6 (2018), 1–14.
- Rui PR Cardoso, Jeong Whan Yoon, Made Mahardika, S Choudhry, RJ Alves de Sousa, and RA Fontes Valente. 2008. Enhanced assumed strain (EAS) and assumed natural strain (ANS) methods for one-point quadrature solid-shell elements. *Int. J. Numer. Methods Eng.* 75, 2 (2008), 156–187.
- Juan J Casafra, Gabriel Cirio, Alejandro Rodríguez, Eder Miguel, and Miguel A Otaduy. 2020. Mixing Yarns and Triangles in Cloth Simulation. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 101–110.
- Jumyung Chang, Fang Da, Eitan Grinspun, and Christopher Batty. 2019. A Unified Simplicial Model for Mixed-Dimensional and Non-Manifold Deformable Elastic Objects. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 2, 2 (2019), 1–18.
- Hsiao-yu Chen, Paul Kry, and Etienne Vouga. 2019. Locking-free Simulation of Isometric Thin Plates. *arXiv preprint arXiv:1911.05204* (2019).
- Hsiao-Yu Chen, Arnav Sastry, Wim M van Rees, and Etienne Vouga. 2018. Physical simulation of environmentally induced thin shell deformation. *ACM Trans. Graph. (TOG)* 37, 4 (2018), 1–13.
- Ming Chen and Kai Tang. 2010. A fully geometric approach for developable cloth deformation simulation. *The visual computer* 26, 6-8 (2010), 853–863.
- Yanqing Chen, Timothy A Davis, William W Hager, and Sivasankaran Rajamanickam. 2008. Algorithm 887: CHOLMOD, supernodal sparse Cholesky factorization and update/downdate. *ACM Transactions on Mathematical Software (TOMS)* 35, 3 (2008), 1–14.
- Byoungwon Choe, Min Gyu Choi, and Hyeong-Seok Ko. 2005. Simulating complex hair with robust collision handling. In *Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation*. 153–160.
- David Clyde. 2017. *Numerical Subdivision Surfaces for Simulation and Data Driven Modeling of Woven Cloth*. Ph.D. Dissertation. UCLA.
- David Clyde, Joseph Teran, and Rasmus Tamstorf. 2017. Modeling and data-driven parameter estimation for woven fabrics. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. 1–11.
- Peter A Cundall and Otto DL Strack. 1979. A discrete numerical model for granular assemblies. *geotechnique* 29, 1 (1979), 47–65.
- Gilles Daviet. 2020. Simple and scalable frictional contacts for thin nodal objects. *ACM Trans. Graph. (TOG)* 39, 4 (2020), 61–1.
- Gilles Daviet, Florence Bertails-Descoubes, and Laurence Boissieux. 2011. A hybrid iterative solver for robustly capturing coulomb friction in hair dynamics. In *Proceedings of the 2011 SIGGRAPH Asia Conference*. 1–12.
- Crispin Deul, Tassilo Kugelstadt, Marcel Weiler, and Jan Bender. 2018. Direct position-based solver for stiff rods. In *Computer Graphics Forum*, Vol. 37. Wiley Online Library, 313–324.
- Elliot English and Robert Bridson. 2008. Animating developable surfaces using non-conforming elements. In *ACM SIGGRAPH 2008 papers*. 1–5.
- Yu Fang, Minchen Li, Chenfanfu Jiang, and Danny M Kaufman. 2021. Guaranteed Globally Injective 3D Deformation Processing. *ACM Trans. Graph. (TOG)* 40, 4 (2021).
- Theodore F Gast, Craig Schroeder, Alexey Stomakhin, Chenfanfu Jiang, and Joseph M Teran. 2015. Optimization integrator for large time steps. *IEEE transactions on visualization and computer graphics* 21, 10 (2015), 1103–1115.
- Moritz Geilinger, David Hahn, Jonas Zehnder, Moritz Bächer, Bernhard Thomaszewski, and Stelian Coros. 2020. ADD: analytically differentiable dynamics for multi-body systems with frictional contact. *ACM Trans. Graph. (TOG)* 39, 6 (2020), 1–15.
- Rony Goldenthal, David Harmon, Raanan Fattal, Michel Bercovier, and Eitan Grinspun. 2007. Efficient simulation of inextensible cloth. In *ACM SIGGRAPH 2007 papers*. 49–es.
- Eitan Grinspun, Anil N Hirani, Mathieu Desbrun, and Peter Schröder. 2003. Discrete shells. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*. Citeseer, 62–67.
- Gaël Guennebaud, Benoît Jacob, et al. 2010. Eigen v3. <http://eigen.tuxfamily.org>.
- Qi Guo, Xuchen Han, Chuyuan Fu, Theodore Gast, Rasmus Tamstorf, and Joseph Teran. 2018. A material point method for thin shells with frictional contact. *ACM Trans. Graph. (TOG)* 37, 4 (2018), 1–15.
- Xuchen Han, Theodore F Gast, Qi Guo, Stephanie Wang, Chenfanfu Jiang, and Joseph Teran. 2019. A hybrid material point method for frictional contact with diverse materials. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 2, 2 (2019), 1–24.
- David Harmon, Daniele Panozzo, Olga Sorkine, and Denis Zorin. 2011. Interference-aware geometric modeling. *ACM Trans. Graph. (TOG)* 30, 6 (2011), 1–10.
- David Harmon, Etienne Vouga, Breannan Smith, Rasmus Tamstorf, and Eitan Grinspun. 2009. Asynchronous contact mechanics. In *ACM SIGGRAPH 2009 papers*. 1–12.
- David Harmon, Etienne Vouga, Rasmus Tamstorf, and Eitan Grinspun. 2008. Robust treatment of simultaneous collisions. In *ACM SIGGRAPH 2008 papers*. 1–4.
- R Hauptmann and K Schweizerhof. 1998. A systematic development of ‘solid-shell’ element formulations for linear and non-linear analyses employing only displacement degrees of freedom. *Int. J. Numer. Methods Eng.* 42, 1 (1998), 49–69.
- F Hernandez, G Cirio, AG Perez, and MA Otaduy. 2013. Anisotropic strain limiting. In *Proc. of Congreso Español de Informática Gráfica*, Vol. 2.
- Chenfanfu Jiang, Theodore Gast, and Joseph Teran. 2017. Anisotropic elastoplasticity for cloth, knit and hair frictional contact. *ACM Trans. Graph. (TOG)* 36, 4 (2017).
- Yupeng Jiang, Minchen Li, Chenfanfu Jiang, and Fernando Alonso-Marroquin. 2020. A hybrid material-point spheropolyhedron-element method for solid and granular material interaction. *Int. J. Numer. Methods Eng.* (2020).
- Ning Jin, Wenlong Lu, Zhenglin Geng, and Ronald P Fedkiw. 2017. Inequality cloth. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*.
- Couro Kane, Jerrold E Marsden, Michael Ortiz, and Matthew West. 2000. Variational integrators and the Newmark algorithm for conservative and dissipative mechanical systems. *Int. J. Numer. Methods Eng.* 49, 10 (2000), 1295–1325.
- Danny M Kaufman and Dinesh K Pai. 2012. Geometric numerical integration of inequality constrained, nonsmooth Hamiltonian systems. *SIAM J. on Sci. Comp.* 34, 5 (2012).
- Danny M. Kaufman, Rasmus Tamstorf, Breannan Smith, Jean-Marie Aubry, and Eitan Grinspun. 2014. Adaptive Nonlinearity for Collisions in Complex Rod Assemblies. *ACM Trans. on Graphics (SIGGRAPH 2014)* (2014).
- Ted Kim. 2020. A finite element formulation of baraff-witkin cloth. In *Symposium on Computer Animation*.
- Tassilo Kugelstadt and Elmar Schömer. 2016. Position and orientation based Cosserat rods. In *Symposium on Computer Animation*. 169–178.
- Cheng Li, Min Tang, Ruofeng Tong, Ming Cai, Jieyi Zhao, and Dinesh Manocha. 2020b. P-cloth: interactive complex cloth simulation on multi-GPU systems using dynamic matrix assembly and pipelined implicit integrators. *ACM Trans. Graph. (TOG)* 39, 6 (2020), 1–15.
- Jie Li, Gilles Daviet, Rahul Narain, Florence Bertails-Descoubes, Matthew Overby, George E Brown, and Laurence Boissieux. 2018a. An implicit frictional contact solver for adaptive cloth simulation. *ACM Trans. Graph. (TOG)* 37, 4 (2018), 1–15.

- Minchen Li, Zachary Ferguson, Teso Schneider, Timothy Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M. Kaufman. 2020a. Incremental Potential Contact: Intersection- and Inversion-free Large Deformation Dynamics. *ACM Trans. Graph.* 39, 4 (2020).
- Minchen Li, Ming Gao, Timothy Langlois, Chenfanfu Jiang, and Danny M Kaufman. 2019. Decomposed optimization time integrator for large-step elastodynamics. *ACM Trans. Graph. (TOG)* 38, 4 (2019), 1–10.
- Minchen Li, Alla Sheffer, Eitan Grinspun, and Nicholas Vining. 2018b. FoldSketch: Enriching Garments with Physically Reproducible Folds. *ACM Trans. Graph.* 37, 4 (2018).
- Junbang Liang, Ming Lin, and Vladlen Koltun. 2019. Differentiable cloth simulation for inverse problems. In *Advances in Neural Information Processing Systems*. 772–781.
- Tiantian Liu, Sofien Bouaziz, and Ladislav Kavan. 2017. Quasi-newton methods for real-time simulation of hyperelastic materials. *ACM Trans. Graph. (TOG)* 36, 3 (2017).
- Libin Lu, Matthew J Morse, Abtin Rahimian, Georg Stadler, and Denis Zorin. 2019. Scalable simulation of realistic volume fraction red blood cell flows through vascular networks. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–30.
- Yanyan Lu, Zhonghua Xi, and Jyh-Ming Lien. 2014. Collision Prediction: Conservative Advancement Among Obstacles With Unknown Motion. In *ASME 2014 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. American Society of Mechanical Engineers Digital Collection.
- Mickaël Ly, Jean Jouve, Laurence Boissieux, and Florence Bertails-Descoubes. 2020. Projective Dynamics with Dry Frictional Contact. *ACM Trans. Graph. (TOG)* 39, 4 (2020).
- Miles Macklin, Kenny Erleben, Matthias Müller, Nuttapon Chentanez, Stefan Jeschke, and Tae-Yong Kim. 2020. Primal/dual descent methods for dynamics. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 89–100.
- Miles Macklin, Matthias Müller, and Nuttapon Chentanez. 2016. XPBD: position-based simulation of compliant constrained dynamics. In *Proceedings of the 9th International Conference on Motion in Games*. 49–54.
- Sebastian Martin, Peter Kaufmann, Mario Botsch, Eitan Grinspun, and Markus Gross. 2010. Unified simulation of elastic rods, shells, and solids. *ACM Trans. Graph. (TOG)* 29, 4 (2010), 1–10.
- Aleka McAdams, Andrew Selle, Kelly Ward, Eftychios Sifakis, and Joseph Teran. 2009. Detail preserving continuum simulation of straight hair. *ACM Trans. Graph. (TOG)* 28, 3 (2009), 1–6.
- Eder Miguel, Derek Bradley, Bernhard Thomaszewski, Bernd Bickel, Wojciech Matusik, Miguel A Otaduy, and Steve Marschner. 2012. Data-driven estimation of cloth simulation models. In *Computer Graphics Forum*, Vol. 31. Wiley Online Library.
- Brian Vincent Mirtich. 1996. *Impulse-based dynamic simulation of rigid body systems*. Matthias Müller, Bruno Heidelberger, Marcus Hennix, and John Ratcliff. 2007. Position based dynamics. *Journal of Visual Communication and Image Representation* 18, 2 (2007), 109–118.
- Matthias Müller, Tae-Yong Kim, and Nuttapon Chentanez. 2012. Fast simulation of inextensible hair and fur. *VRIPHYS* 12 (2012), 39–44.
- Rahul Narain, Armin Samii, and James F O'Brien. 2012. Adaptive anisotropic remeshing for cloth simulation. *ACM Trans. Graph. (TOG)* 31, 6 (2012), 1–10.
- Jorge Noddedal and Stephen Wright. 2006. *Numerical optimization*.
- Miguel A Otaduy, Rasmus Tamstorf, Denis Steinemann, and Markus Gross. 2009. Implicit contact handling for deformable objects. In *Computer Graphics Forum*, Vol. 28. Wiley Online Library, 559–568.
- Matthew Overby, George E Brown, Jie Li, and Rahul Narain. 2017. ADMM \supseteq Projective Dynamics: Fast Simulation of Hyperelastic Models with Dynamic Constraints. *IEEE Transactions on Visualization and Computer Graphics* 23, 10 (2017), 2222–2234.
- Željko Penava, Diana Šimić-Penava, and Ž Knezic. 2014. Determination of the elastic constants of plain woven fabrics by a tensile test in various directions. *Fibres & Textiles in Eastern Europe* (2014).
- Xavier Provot. 1997. Collision and self-collision handling in cloth model dedicated to design garments. In *Computer Animation and Simulation '97*. 177–189.
- Xavier Provot et al. 1995. Deformation constraints in a mass-spring model to describe rigid cloth behaviour. In *Graphics interface*. Canadian Information Processing Society.
- Alessio Quaglino. 2012. *Membrane locking in discrete shell theories*. Ph.D. Dissertation. Georg-August-Universität Göttingen.
- Stefanie Reese. 2007. A large deformation solid-shell concept based on reduced integration with hourglass stabilization. *Int. J. Numer. Methods Eng.* 69, 8 (2007).
- Nikolas Schmitt, Martin Knuth, Jan Bender, and Arjan Kuijper. 2013. Multilevel Cloth Simulation using GPU Surface Sampling. *VRIPHYS* 13 (2013), 1–10.
- Andrew Selle, Michael Lentine, and Ronald Fedkiw. 2008. A mass spring model for hair simulation. In *ACM SIGGRAPH 2008 papers*. 1–11.
- Georg Sperl, Rahul Narain, and Chris Wojtan. 2020. Homogenized yarn-level cloth. *ACM Trans. Graph. (TOG)* 39, 4 (2020), 48–1.
- Jos Stam. 2009. Nucleus: Towards a unified dynamics solver for computer graphics. In *2009 11th IEEE International Conference on Computer-Aided Design and Computer Graphics*. IEEE, 1–11.
- Alexey Stomakhin, Russell Howes, Craig Schroeder, and Joseph M Teran. 2012. Energetically consistent invertible elasticity. In *Proceedings of the 11th ACM SIGGRAPH/Eurographics conference on Computer Animation*. 25–32.
- Rasmus Tamstorf and Eitan Grinspun. 2013. Discrete bending forces and their Jacobians. *Graphical models* 75, 6 (2013), 362–370.
- Rasmus Tamstorf, Toby Jones, and Stephen F McCormick. 2015. Smoothed aggregation multigrid for cloth simulation. *ACM Trans. Graph. (TOG)* 34, 6 (2015), 1–13.
- Min Tang, Young J Kim, and Dinesh Manocha. 2009. C 2 A: Controlled conservative advancement for continuous collision detection of polygonal models. In *2009 IEEE International Conference on Robotics and Automation*. IEEE, 849–854.
- Min Tang, Dinesh Manocha, and Young J Kim. 2013a. Hierarchical and controlled advancement for continuous collision detection of rigid and articulated models. *IEEE transactions on visualization and computer graphics* 20, 5 (2013), 755–766.
- Min Tang, Dinesh Manocha, Sung-Eui Yoon, Peng Du, Jae-Pil Heo, and Ruo-Feng Tong. 2011. VolCCD: Fast continuous collision culling between deforming volume meshes. *ACM Trans. Graph. (TOG)* 30, 5 (2011), 1–15.
- Min Tang, Ruofeng Tong, Rahul Narain, Chang Meng, and Dinesh Manocha. 2013b. A GPU-based streaming algorithm for high-resolution cloth simulation. In *Computer Graphics Forum*, Vol. 32. Wiley Online Library, 21–30.
- Min Tang, Ruofeng Tong, Zhendong Wang, and Dinesh Manocha. 2014. Fast and exact continuous collision detection with bernstein sign classification. *ACM Trans. Graph. (TOG)* 33, 6 (2014), 1–8.
- Min Tang, Huamin Wang, Le Tang, Ruofeng Tong, and Dinesh Manocha. 2016. CAMA: Contact-aware matrix assembly with unified collision handling for GPU-based cloth simulation. In *Computer Graphics Forum*, Vol. 35. Wiley Online Library, 511–521.
- Min Tang, Tongtong Wang, Zhongyuan Liu, Ruofeng Tong, and Dinesh Manocha. 2018. I-cloth: incremental collision handling for GPU-based interactive cloth simulation. *ACM Trans. Graph. (TOG)* 37, 6 (2018), 1–10.
- Demetri Terzopoulos, John Platt, Alan Barr, and Kurt Fleischer. 1987. Elastically deformable models. In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*. 205–214.
- Bernhard Thomaszewski, Simon Pabst, and Wolfgang Strasser. 2009. Continuum-based strain limiting. In *Computer Graphics Forum*, Vol. 28. Wiley Online Library, 569–576.
- Nobuyuki Umetani, Danny M Kaufman, Takeo Igarashi, and Eitan Grinspun. 2011. Sensitive couture for interactive garment modeling and editing. *ACM Trans. Graph.* 30, 4 (2011), 90.
- Pascal Volino and N Magnenat Thalmann. 2000. Implementing fast cloth simulation with collision response. In *Proceedings Computer Graphics International 2000*. IEEE.
- Huamin Wang, James O'Brien, and Ravi Ramamoorthi. 2010. Multi-resolution isotropic strain limiting. *ACM Trans. Graph. (TOG)* 29, 6 (2010), 1–10.
- Huamin Wang, James F O'Brien, and Ravi Ramamoorthi. 2011. Data-driven elastic models for cloth: modeling and measurement. *ACM Trans. Graph. (TOG)* 30, 4 (2011).
- Xinlei Wang, Minchen Li, Yu Fang, Xinxin Zhang, Ming Gao, Min Tang, Danny M Kaufman, and Chenfanfu Jiang. 2020. Hierarchical optimization time integration for cf-rate mpm stepping. *ACM Trans. Graph. (TOG)* 39, 3 (2020), 1–16.
- Zhendong Wang, Min Tang, Ruofeng Tong, and Dinesh Manocha. 2015. TightCCD: Efficient and robust continuous collision detection using tight error bounds. In *Computer Graphics Forum*, Vol. 34. Wiley Online Library, 289–298.
- Zhendong Wang, Tongtong Wang, Min Tang, and Ruofeng Tong. 2016. Efficient and robust strain limiting and treatment of simultaneous collisions with semidefinite programming. *Computational Visual Media* 2, 2 (2016), 119–130.
- Zhendong Wang, Longhua Wu, Marco Fratarcangeli, Min Tang, and Huamin Wang. 2018. Parallel Multigrid for Nonlinear Cloth Simulation. *Computer Graphics Forum* (2018).
- Kelly Ward and Ming C Lin. 2003. Adaptive grouping and subdivision for simulating hair dynamics. In *11th Pacific Conference on Computer Graphics and Applications, 2003. Proceedings*. IEEE, 234–243.
- Nicholas J Weidner, Kyle Piddington, David IW Levin, and Shinjiro Sueda. 2018. Eulerian-on-lagrangian cloth simulation. *ACM Trans. Graph. (TOG)* 37, 4 (2018), 1–11.
- Zangyueyang Xian, Xin Tong, and Tiantian Liu. 2019. A Scalable Galerkin Multigrid Method for Real-time Simulation of Deformable Objects. *ACM Trans. Graph. (TOG)* 38, 6 (2019).
- H.T.Y. Yang, S. Saigal, A. Masud, and R.K. Kapania. 2000. A survey of recent shell finite elements. *Int. J. Numer. Methods Eng.* 47 (2000).
- Yonghao Yue, Breannan Smith, Peter Yichen Chen, Maytee Chantharayukhonthorn, Ken Kamrin, and Eitan Grinspun. 2018. Hybrid grains: adaptive coupling of discrete and continuum simulations of granular media. *ACM Trans. Graph. (TOG)* 37, 6 (2018), 1–19.
- Xinyu Zhang, Minkyoung Lee, and Young J Kim. 2006. Interactive continuous collision detection for non-convex polyhedra. *The Visual Computer* 22, 9 (2006), 749–760.
- Xinyu Zhang, Stephane Redon, Minkyoung Lee, and Young J Kim. 2007. Continuous collision detection for articulated models using taylor models and temporal culling. *ACM Trans. Graph. (TOG)* 26, 3 (2007), 15–es.