Orientation-aware Incremental Potential Contact

ZIZHOU HUANG, New York University, USA

MAX PAIK, New York University, USA

ZACHARY FERGUSON, Massachusetts Institute of Technology, USA

DANIELE PANOZZO, New York University, USA

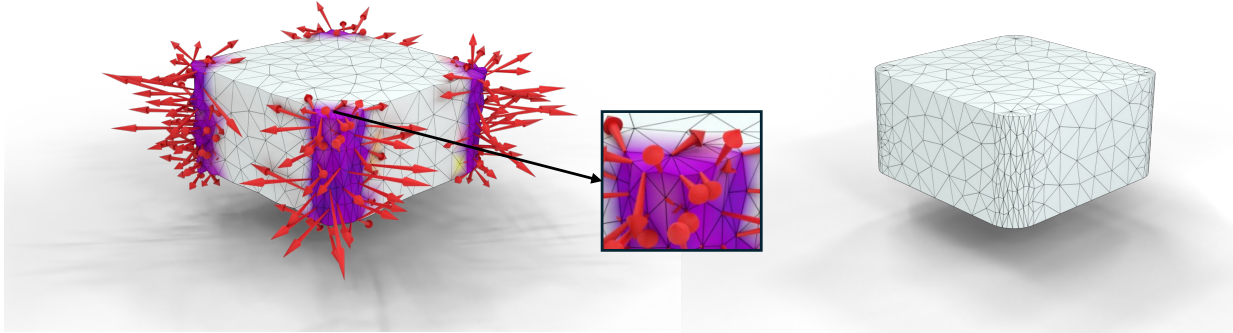DENIS ZORIN, New York University, USA

Fig. 1. **Spurious stresses: nonuniform 3D mesh.** Incremental potential contact [Li et al. 2020] introduces spurious contact forces in many configurations. We show a 3D block with rounded corners, maximal edge length $0.42\,\mathrm{m}$, and minimum edge length $0.03\,\mathrm{m}$ Left: With $\hat{d}$ chosen suitably small for the coarsely discretized regions (in this example $\hat{d} = 0.07\,\mathrm{m}$,), IPC introduces spurious forces at finely meshed corners, leading to jagged deformation artifacts. Right: our method maintains the rest configuration in the absence of other outside forces by deriving a new IPC-like contact formulation from the continous case.

The Incremental Potential Contact (IPC) method enables robust complex simulations of deformable objects with contact and friction. The key to IPC's robustness is its strict adherence to geometric constraints, avoiding intersections, which are a common cause of robustness issues in contact mechanics. A key element of the IPC approach to contact is a geometric barrier function, which is defined directly in the discrete setting. While IPC achieves its main goal of providing guarantees for contact constraints, its parameters need to be chosen carefully to avoid significant simulation artifacts and inaccuracies. We present a systematic derivation of an IPC-like continuum potential defined for smooth and piecewise smooth surfaces, starting from identifying a set of natural requirements for contact potentials, including the barrier property, locality, differentiable dependence of shape, and absence of forces in rest configurations, based on the idea of candidate sets. Our potential is formulated in a way independent of surface discretization.

This new potential is suitable for piecewise-linear surfaces and its efficiency is similar to standard IPC. We demonstrate its behavior and compare it to IPC on a range of challenging contact examples.

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

Additional Key Words and Phrases: Finite element method, Elastodynamics, Contact dynamics

1 INTRODUCTION

Contact modeling is a critical component of simulation tools for many domains, including physical simulation for computer graphics, robotics, mechanical design, and biomedical engineering. Recent

Authors' addresses: Zizhou Huang, zizhou@nyu.edu, New York University, USA; Max Paik, mp6569@nyu.edu, New York University, USA; Zachary Ferguson, zfergus@nyu.edu, Massachusetts Institute of Technology, USA; Daniele Panozzo, panozzo@nyu.edu, New York University, USA; Denis Zorin, dzorin@cs.nyu.edu, New York University, USA.

advances, in particular, Incremental Potential Contact (IPC) [Li et al. 2020], make it possible to perform, in a reliable way, complex, large-deformation, and long-duration simulations of deformable objects with self-contact, with minimal parameter tuning.

The fundamental principle of the IPC solvers is the strict adherence to geometric constraints over the entire trajectory (no intersections including self-intersections), even when approximations are made in other aspects of the problem. This feature is the key to robustness: overwhelmingly, robustness issues in contact mechanics are geometric in nature; once a contact constraint is violated, it may be very difficult for the simulation to recover.

Geometric constraint satisfaction in IPC is accomplished by a combination of using a geometric barrier function, i.e. an interaction potential that approaches infinity for contact points, and conservative continuous collision detection (CCD). IPC formulation excels at its main goal, qualitatively improving the robustness of contact simulations.

At the same time, the approach has limitations, and certain conditions may lead to inaccurate results and simulation artifacts. The primary reason for this behavior is several specific features of the design of the barrier potentials used in the simulation. The potential was designed directly for discretized models, rather than obtained from a discretization of a continuum potential; as we discuss in more detail in Section 2. In particular, there is a direct, and strong limitation on how local the potential is (and, as a consequence, how stiff it is) and how the object is discretized.

In this paper, we undertake a systematic derivation of a continuum barrier potential and its discretization from a set of desiderata,

elaborated in Section 4. Similarly to Convergent IPC [Li et al. 2023], we assume that the total contact potential is obtained by integration over the surface of a pointwise potential $\psi(x)$. The following requirements, which we make more precise in Section 3 include:

- *Generality*: The total potential over the surface is well-defined for any piecewise smooth surface not in contact.
- *Barrier*: The potential $\psi(x)$ grows to infinity as the objects in the simulation approach (self-)contact at x, but is finite for any configuration not in contact.
- *No spurious forces*: In the undeformed configuration the potential is zero.
- *Localization*: The potential has a parameter $\hat{d}$ and vanishes if the objects are further away than $\hat{d}$ from contact.
- *Differentiability*: The potential depends smoothly on the surface configuration (e.g., for piecewise surfaces, on mesh vertex positions).
- *Resolution independence*: The localization of the contact potential can be controlled independently from mesh resolution.

Note that two key requirements (*Localization* and, less directly, *No spurious forces* depend on the notion of *distance to contact*: how we measure this distance for self-contact is critical for choosing the right potential Section 3.

The original IPC formulation [Li et al. 2020] meets these requirements only partially for piecewise linear surfaces. Namely, it can violate *No spurious forces* and *Resolution independence* depending on the choice of $\hat{d}$. To address this we introduce an IPC-like potential designed to satisfy these requirements. The key features of our approach include:

- The concept of *contact candidate sets*, that, in turn, supports a definition of *distance to contact* suitable for contact barrier potentials.
- An approach for defining a contact potential for piecewise smooth surfaces that meets the requirements enumerated above, independently of the choice of discretization.
- A discrete version of this potential for piecewise-linear surfaces retaining the same properties and efficiency of standard IPC.

We demonstrate in a set of benchmarks that satisfying the properties above is crucial for accurate elastodynamic: our formulation avoids spurious, discretization-dependent forces introduced by the original IPC formulation [Li et al. 2020] (Section 5). The decoupling of the locality from mesh resolution further reduces the parameter tuning required for scenes with large deformations (Section 5).

2 RELATED WORK

Traditional methods for contact response mostly formulate non-interpenetration constraints based on the normal displacements of the contacting regions with locally approximated geometries. Then the constrained time integration problem is solved with Lagrange multipliers [Carpenter et al. 1991; Deuflhard et al. 2008; Hiermeier et al. 2018; Kane et al. 1999; Laursen and Love 2002; Paoli and Schatzman 2002a,b; Taylor and Papadopoulos 1993; Vola et al. 1998], penalty methods [Armero and Petőcz 1998; Belytschko and Neal 1991; Hauret and Le Tallec 2006; Hiermeier et al. 2018; Puso and

Laursen 2004], barrier methods [Belytschko et al. 2000; Christensen et al. 1998; Kloosterman et al. 2001], etc. To guarantee global convergence for the solution of frictionless large deformation contact problems for hyperelastic materials, Youett et al. [2019] proposed a filter-trust-region method to solve the constrained minimization problem formulated from implicit time integration.

Mortar methods [Belgacem et al. 1998; Hüeber and Wohlmuth 2006; Puso and Laursen 2004] are commonly used in engineering [Krause and Zulian 2016] and biomechanics [Maas et al. 2012] but require to (a priori) mark all possible pairs of contacting mesh boundary. A clear limitation of these methods is then that they cannot handle, without extensions, self-collisions or collisions in regions with more than two contacting surface regions. Li et al. [2020] provide a didactic comparison of IPC and one such mortar method [Krause and Zulian 2016]. There it can be seen that such methods enforce contact constraints weakly and therefore allow intersections (especially at large timesteps and/or velocities).

Instead of directly treating contact as unilateral constraints, Wriggers and Laursen [2006] and Sauer and De Lorenzis [2013] formulate frictionless contact based on distance-dependent surface interaction potentials, providing a unified framework for various contact formulations including Lagrange multiplier, penalty, and barrier. Starting with the first and second laws of thermodynamics, Duong and Sauer [2019] recently extended [Sauer and De Lorenzis 2013] to account for Coulomb friction under isothermal conditions, and demonstrated the convergence of the method under refinement.

Gap functions. Gap functions compute a *signed distance* for pairs of points or primitives, where negative distance measures the interpenetration depth. Collisions are then resolved by enforcing that all gap functions are non-negative. Commonly used gap functions include projecting between pairwise surface primitives onto a fixed geometric normal [Otaduy et al. 2009; Wriggers 1995] or using CCD to determine the point of contact and compute a projected distance along the contact normal [Harmon et al. 2008; Verschoor and Jalba 2019]. Harmon et al. [2011] introduced the method of space-time interference volumes which was in part motivated by the need to define a version of the gap function in a fully general setting; the complexity of their construction shows that even rigorously formulating the self-contact problem in inequality-constrained is nontrivial. We discuss the connection of our method to one type of gap function in Section 4 in more detail.

Barrier-based methods. The use of barrier functions for handling contact has been utilized in graphics for over a decade now. Most notably, Harmon et al. [2009] and Vouga et al. [2011] utilize a set of layered discrete penalty barriers that grow unbounded as the configuration approaches contact. However, this incremental construction makes it unsuitable for optimization-based implicit time integration and therefore requires small time-steps for stability.

Kaldor et al. [2008] simulate knitted cloth at the yarn level where yarn-yarn collisions are handled through a continuous potential integrating a barrier function over two disjoint spline segments. This approach, however, ignores self-collisions within a single segment and between neighboring segments. To discretize their model, Kaldor et al. [2008] use "Simpson's quadrature at fixed positions in parameter space" which has the limitation of potentially missing

the nearly singular portion of the integral and therefore introducing intersections.

Our method largely follows the IPC method of Li et al. [2020], which transforms discrete gap constraints into a sum of log barrier functions. The method of Li et al. [2020] has been directly utilized and extended by numerous follow-up works to include support for codimensional elements [Li et al. 2021], rigid/affine body dynamics [Ferguson et al. 2021; Lan et al. 2022], medial elastics [Lan et al. 2021], solid-fluid interactions [Xie et al. 2023], higher-order finite element analysis [Ferguson et al. 2023], etc.

The most significant limitation of IPC [Li et al. 2020] is its lack of convergence under mesh refinement due to its purely discrete formulation (see Li et al. [2023] and Du et al. [2023] for details). To address this, Li et al. [2023] form a continuous formulation for the original IPC model which requires minimal adjustments to existing implementations. The limitation of their approach is the requirement for joint refinement of $\hat{d}$ and mesh resolution. We provide further discussion of the limitations below.

Du et al. [2023] show the tendency of "Node-to-Segment" methods (e.g., IPC [Li et al. 2020]) to produce spurious tangential contact forces. To remedy this they propose to use a fully C^1-continuous surface representation; computing distances used in the contact potential between the nodes and smooth surface (extending the method of Larionov et al. [2021]). The main limitation of [Du et al. 2023], however, is the inability to handle self-collisions.

Tangent-point energies. A possible continuous approach to avoiding intersections are the tangent-point energies [Buck and Orloff 1995; Strzelecki and von der Mosel 2013] which repel surface points by integrating over all surface pairs $(x, y) \in \Omega \times \Omega$:

$$\int_{x \in \Omega} \int_{y \in \Omega} \frac{1}{r(f(x), f(y))^\alpha} \, \mathrm{d}x \mathrm{d}y,$$

where r is the radius of the smallest sphere tangent to $f(x)$ and passing through $f(y)$ and α controls the repulsion strength. This avoids having infinite energy for geodesic neighbors as the radius is large, but grows to infinity for points approaching zero distance in Euclidean space. Recently, Yu et al. [2021a,b] utilize tangent-point energies to prohibit intersections while optimizing surface geometry. In pure form, this method is unsuitable for collision response due to its global coupling of all points, but it is closely related to our approach as discussed in Section 4 in the case of smooth surfaces. This not only leads to a dense system that needs to be factorized in the optimization but also spurious contact forces between distant pairs.

3 CONTACT POTENTIAL REQUIREMENTS

We consider a collection of deformable objects defined on a material domain $\mathrm{Int}(\Omega)$, whose boundary is Ω. $\mathrm{Int}(\Omega)$ may have multiple connected components (Figure 2).

An admissible deformation $f : \mathrm{Int}(\Omega) \mapsto \mathbb{R}^n$, $n = 2, 3$ is noninjective on the boundary only, i.e., f is injective in the interior of the domain, but may have *contact points* $x \neq y$ on the boundary Ω, for which $f(x) = f(y)$.

We say that an admissible f is a *contact deformation*, if it has contact points, otherwise, we call it *contact-free*. We assume that for
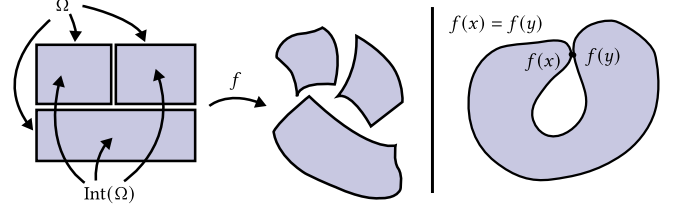

Fig. 2. A collection of objects is transformed using a deformation f (left), which can cause the objects to intersect at a contact point (right).

a point $x \in \Omega$, we have a measure $d_c(x, f)$ of how far it is from being a contact point. Defining this measure is a key aspect of our construction, and is necessary to formulate our contact requirements.

We note that the problem of defining $d_c(x, f)$ is far from trivial. While for rigid objects the distance to the closest point on *another* object is an adequate measure, this is not the case for deformable objects. Due to the possibility of self-contact, we cannot exclude points on the same object and, in this case, there are always points on Ω arbitrarily close to x in Euclidean distance that need to be considered to be far from being in contact with x. At the same time, an arbitrarily small, if measured by Euclidean distance, perturbation of a surface may create an actual contact, so the distance from x to a real contact point y can be arbitrarily small.

Before describing our approach to the distance to contact, we make the properties we stated in the introduction more precise.

Requirement 1 (*Generality*). A total contact potential

$$\Psi(f) = \int_\Omega \psi_\epsilon(x; f) \, \mathrm{d}x$$

is well-defined, i.e., $\psi_\epsilon(x; f)$ is integrable on Ω for a sufficiently broad class of surfaces. Our construction applies to piecewise-smooth surfaces, although in our implementation we focus on piecewise linear.

Requirement 2 (*Barrier*). $\psi_\epsilon(x; f)$ increases to infinity for f if the distance to contact $d_c(x, f)$ goes to zero. Combined with incremental potential time-stepping and CCD, this can be used to guarantee that all configurations remain contact-free.

Requirement 3 (*No spurious forces*). For the undeformed configuration $f = \mathrm{Id}$, $\psi_\epsilon(x; \mathrm{Id})$ is zero everywhere. This is necessary to ensure that the potential does not create artificial forces that would cause motion/deformations even if no external forces are acting on the object.

Requirement 4 (*Localization*). $\psi_\epsilon(x; f)$ has a locality parameter $\epsilon > 0$ ($\hat{d}$ in the notation of Li et al. [2020]), with no restrictions on its magnitude, on which it depends at least continuously, and the potential vanishes if $d_c(x, f) > \epsilon$. If we solve a sequence of problems with decreasing ϵ, we approach a solution of the standard inequality-constrained formulation of contact problems.

Requirement 5 (*Differentiability*). If f is defined by a finite number of parameters (in the simplest case, vertex positions of a mesh) then $\psi_\epsilon(x; f)$ depends differentially, and piecewise twice differentiably on the parameters of f. Then the potential leads to a force with
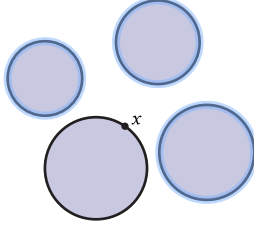
Fig. 3. In the rigid case, for a point x, the contact candidate set (highlighted in blue) is the set of points in all connected components of Ω that do not contain x.
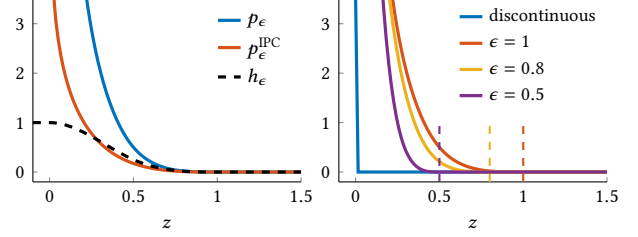

Fig. 5. **Barriers.** Left: plot of the cubic spline $h_\epsilon(z)$ and barrier function $p_\epsilon(z) = h_\epsilon(z)/z^{n-1}$ for $\epsilon = 1$. We also include a plot of the log barrier $p_\epsilon^{\text{IPC}}(z)$ of Li et al. [2020] for comparison. Right: our barrier improves approximation to the discontinuous function as ϵ goes to 0.

piecewise continuous Jacobian, allowing for second-order methods for implicit time-stepping.

Importantly, these properties are independent of whether we consider a smooth problem or its discretization: below, we require that both a smooth potential and its discretization satisfy these properties.

In the discrete case, the total potential is a sum of discrete potentials defined on pairs of discrete primitives. We note that it is important that *all requirements above are satisfied by the discretized total contact potential directly*, not only in its continuous formulation.

4 FORMULATION

Our approach to defining contact potentials is conveniently formulated in terms of *contact candidate sets* $C(x; f)$. These sets are defined for every point x as a subset of Ω, and the potential we construct is restricted to these.

To be able to use them in the definition of a contact potential, we choose the contact sets $C(x)$ with the following properties

- For f in contact, $C(x)$ contains the contact point of x;
- If is not in contact, the (Euclidean) distance from $f(x)$ to $C(x)$ is positive, and it decreases to zero if we interpolate between an f in contact and $\tilde{f}$ in contact.

We define the *distance to contact* $d_C(x, f)$ of a point x under deformation f as the distance to the contact candidate set of x:

$$d_C(x, f) = \min_{y \in C(x)} \|f(x) - f(y)\|.$$

If a potential $\psi(x; f)$ is obtained by integrating a positive function $g(x, y)$ supported on the contact set, depends on the distance between x and y, and increases to infinity if the distance is zero, then we can observe that using such candidate sets allows us to satisfy Requirements 2 to 4. In particular, the Requirement 4 (*Localization*) is satisfied if $g(x, y)$ is chosen to vanish for a small $d_C(x, f)$, and Requirement 3 (*No spurious forces*) is satisfied if it vanishes for $d_C(x, Id)$.

4.1 Rigid objects

To make the contact set definition less abstract, we briefly consider the case of rigid objects, for which f restricted to each connected component of Ω is rigid motion, as a guiding example for defining a contact candidate set and a corresponding potential satisfying the contact potential requirements (Section 3).

In this case, the choice is simple and intuitive: for x, the contact candidate set $C(x)$ is the set of points in all connected components of Ω that do not contain x (Figure 3). Then the distance $d_c(x)$ is simply the distance to the closest object. The requirements above for the contact set are satisfied, and one can define the potential as

$$\psi^R(x; f) = \int_{y \in C(x)} p_{\epsilon(x)}(\|f(x) - f(y)\|) \, dy \qquad (1)$$

where $p_{\epsilon(x)}(z) \to \infty$ as $z \to 0$, fast enough for the integral to be unbounded if x is in contact and there are points in $C(x)$ arbitrarily close to x (Figure 5). We choose

$$p_{\epsilon(x)}(z) = \frac{h_\epsilon(z)}{z^{n-1}}$$

for $n = 2, 3$,

$$h_\epsilon(z) = \frac{3}{2} B^3 (2z/\epsilon)$$

where $B^3 \in C^2(\mathbb{R})$ is a cubic spline (Figure 4)

$$B^3(v) = \begin{cases} \frac{2}{3} - v^2 + \frac{1}{2}|v|^3 & |v| < 1 \\ \frac{1}{6}(2 - |v|)^3 & 1 \leq |v| < 2 \\ 0 & 2 \leq |v| \end{cases}.$$

Therefore, $h_\epsilon \in C^2(\mathbb{R})$ satisfies

$$h_\epsilon(\epsilon) = h_\epsilon'(0) = h_\epsilon''(\epsilon) = 0.$$

We note that somewhat counter-intuitively, it is not strictly necessary for the potential $\psi(z)$ to be infinite to guarantee that no contact happens (see Appendix A). The choice of the power for the potential is important: on the one hand, e.g., the logarithmic potential used in IPC, in the continuum case may result in a finite integral, i.e., requirement 2 not satisfied; on the other hand, if the power

Fig. 4. Plot of $B^3(z)$.

is too high, that the pointwise potential may lead to divergent integrals for piecewise-smooth surfaces folding along sharp features (see Section A).

Then h_ϵ vanishes if $d_C(x; f)$ exceeds ϵ, and as a consequence, ψ^R is zero. By construction, $\psi^R(x; f)$ becomes infinite if $d_C(x; f) \to 0$, i.e., this potential satisfies Requirements 2, 4, and 5, as long as the

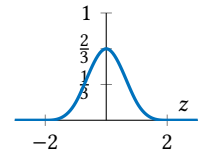potential grows fast enough. To satisfy Requirement 3 (*No spurious forces*), we choose a variable $\epsilon(x)$ for each point x to be equal to

$$\min(d_C(x, f_0)/2, \epsilon_{\text{trg}}),$$

where ϵ_{trg} is the global parameter determining the maximal distance at which the potential may be nonzero.

Remark. Note that $\epsilon(x)$ is not necessarily smooth or even continuous as a function of x. This may affect the convergence of the outer integral in computing Ψ; however, it does not affect the required properties, in particular, Requirement 5 (*Differentiability*) of the total potential's dependence on the shape parameters of f.

4.2 Deformable smooth surfaces

Contact candidate sets for self-contact. For an arbitrary deformation f, self-contact needs to be taken into account since points on the same objects cannot be excluded. We need a different approach for defining $C(x)$.

Our key observation is: *for a fixed point x and $y \neq x$ in contact, if we consider $d(y) = \|f(x) - f(y)\|$ with y varying over the surface, $d(y)$ has a local minimum at y* (Figure 8).

Clearly, $d(y)$ has also a local minimum at x which we exclude, but not at any point in a neighborhood of x. This leads to the idea of including in $C(x)$ points that are close to being local minima of the distance function $\|f(x) - f(y)\|$ in the candidate set, and, in addition, the vector $f(x) - f(y)$ points towards the exterior of the surface at $f(x)$ and towards the interior at $f(y)$. The latter is critical for handling thin objects.

We first consider smooth surfaces, to describe the main ideas without considering many cases needed for piecewise smooth surfaces, and then extend the formulation to a general case.

In this case, the local minima of the distance function from $f(x)$ to $f(y)$ for a contact-free f satisfy the condition

$$\Phi(x, y)^2 = \|(f(y)-f(x))_+ \times n(x)\|^2 = 1-((f(y)-f(x))_+ \cdot n(x))^2 = 0,$$

where $(\cdot)_+$ denotes normalization to a unit vector. The resulting distance to contact and potentials are related to one of the commonly considered *gap function* types in the FEM literature [Otaduy et al. 2009; Wriggers 1995] and the potential to tangent-point kernels [Buck and Orloff 1995; Strzelecki and von der Mosel 2013; Yu et al. 2021a,b].

Additionally, define

$$\Phi^n(x, y) = n(x) \cdot (f(y) - f(x))_+$$

Then $f(y)$ is on the side of the outward pointing normal from x, if $\Phi^n(x, y) > 0$.

This suggests the following definition of a contact candidate set for a smooth surface:

$$C(x) = \{y \mid \Phi^n(x, y) \geq -\alpha, \Phi(y, x) \leq \alpha\} \tag{2}$$

(Figure 6). Observe that $C(x)$ for smooth surfaces is isolated from x: as y approaches x, $f(y) - f(x)$ becomes tangent to the surface and, as a consequence, perpendicular to the normal $n(y)$, so the distance from $C(x)$ to x is positive for any $\alpha < 1$.

In the case of smooth surfaces $\Phi(x, y)^2 + \Phi^n(x, y)^2 = 1$, so restricting $\Phi(x, y)^2$ to be close to zero is equivalent to restricting $\Phi^n(x, y)^2$
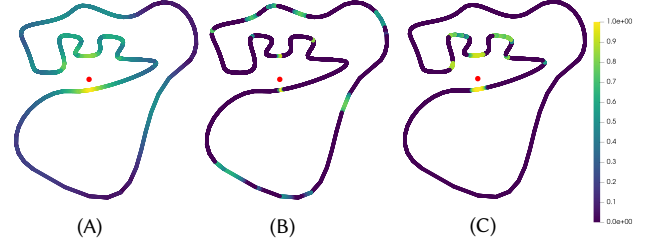

Fig. 6. Potential distribution on the surface with respect to the red dot. From left to right are potential distributions. (A) Without constraint on surface tangent, the potential distributes spherically around the red dot regardless of the surface shape. (B) With tangent constraints, only surfaces that are close to the local minima of distance from the red dot have high values. (C) With both tangent and normal constraints, only surfaces on the closer side of the volumetric object have high values.
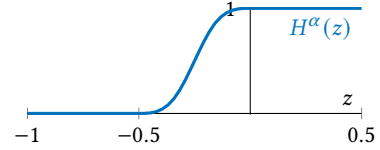

Fig. 7. Smoothed Heaviside function $H^\alpha(z)$ with $\alpha = \frac{1}{2}$.

to be 1. However, the relationship is more complex for piecewise smooth surfaces, so we treat these separately.

At the same time, as f approaches contact at x, the set will eventually contain the contact point y.

In this case, the property 1, i.e., the fact that contact potential is well-defined for a contact-free f: while $\psi(x)$ may be defined for any point x, it may have no upper bound, so the integral over x is not defined. However, there is a global positive lower bound on the separation distance for non-empty $C(x)$: the set of such points is closed, and the bound is positive for all points. This ensures that Requirement 1 (*Generality*) is satisfied for smooth surfaces.

Contact potential. Using a C^2 smooth function

$$\delta_\alpha(z) = \frac{2}{\alpha} B^3(2z/\alpha),$$

which

(1) has a maximal value at zero,
(2) decays to zero for $|z| = \alpha$ (i.e., on the boundary of the contact candidate set), and
(3) integrates to 1,

we also use the smoothed Heaviside step function $H^\alpha(z) = H(z/\alpha) \in C^2(\mathbb{R})$, with

$$H(z) = \begin{cases} 0 & z < -3 \\ \frac{1}{6}(3 + z)^3 & -3 \leq z < -2 \\ \frac{1}{6}(3 - 9z - 9z^2 - 2z^3) & -2 \leq z < -1 \\ 1 + \frac{1}{6}z^3 & -1 \leq z < 0 \\ 1 & 0 \leq z \end{cases} ;$$

Then we define the *angle factor* as

$$\gamma^S(x,y) = \delta_\alpha(\Phi(x,y))H^\alpha(\Phi^n(x,y)) \tag{3}$$

which ensures that the potential is supported within the contact set $C(x)$.

We can define a potential for deformable contact as

$$\Psi(f) = \int_{x\in\Omega}\int_{y\in\Omega}\gamma^S(x,y)\gamma^S(y,x)p_{\epsilon(x)}(\|f(x)-f(y)\|). \tag{4}$$

Here, similar to the rigid motion case, we choose

$$\epsilon(x) = \min(d_C(x;f_0)/2, \epsilon_{\mathrm{trg}}),$$

where ϵ_{trg} is the maximal distance at which the potential is non-zero so that in the undeformed state, the potential vanishes everywhere.

Relation to gap functions and tangent kernels. One of the gap functions described in FEM literature (e.g., [Wriggers 1995]) is obtained by shooting a ray along the normal direction $n(x)$ to obtain a point $y(x)$, defining in this way the correspondence between surfaces, and defining the gap function as $n(x) \cdot (f(y(x)) - f(x))$. This is similar to applying $\Phi(x,y)$ for a specific choice of y. We note that in general geometry, $y(x)$ may not exist, or move discontinuously. Similarly, repulsive surfaces [Strzelecki and von der Mosel 2013; Yu et al. 2021a] use potentials of the form $(n(x) \cdot (f(y) - f(x)))^p/\|x-y\|^{2p}$. Our potential can be considered a localization of repulsive potentials of this type in space and angle.

4.3 Piecewise smooth contact

We now extend the potential to piecewise smooth surfaces. There are two reasons for this extension: (1) piecewise smooth surfaces are ubiquitous in graphics and scientific computing as they can represent shapes with sharp features, and (2) smooth surfaces are often approximated by piecewise-linear for which efficient robust continuous collision detection is available.

We consider surfaces consisting of patches Ω_i, which form a curved manifold mesh satisfying the standard definition [Do Carmo 2016]. We assume that each patch has continuously varying normals defined everywhere including boundary (i.e. no cones are allowed) and the edge curves meeting at a vertex have distinct tangents. We refer to the curves separating patches as *edge curves*, or simply edges, and points shared by more than 2 patches as vertices.

In this case, there are six possibilities for a contact point, corresponding to the possible pairs of contacts between any two element types: Face-Face, Face-Edge, Face-Vertex, etc.

The smooth-surface definition can be applied in cases involving a face, as there is a well-defined normal on one side, and the definition of the potential requires only the normal at y. However, cases that do *not* involve a face (for example, Edge-Edge) require special handling and a generalization of our definition of $C(x)$ and distance to contact. Observe that we *cannot* assume that, e.g., nearby face-vertex contacts are sufficient: two edge points may be close, but all points may be far from contact in the sense of $d_C(x;f)$ defined above.

Contact candidate sets for piecewise-smooth surfaces. To generalize $C(x)$ to piecewise smooth surfaces, we take the following approach.
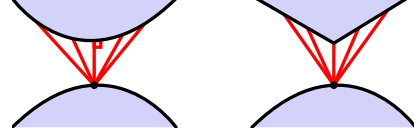

Fig. 8. Contact point as a local minimum of distance.

Consider the distance $\|f(y) - f(x)\|^2$ as a function of y. Then for the gradient w.r.t. y, we obtain

$$\nabla_y f(y)\,(f(y) - f(x)) = 0$$

at the closest point. The columns of the matrix $\nabla_y f$ are two tangents at y, i.e., this condition is equivalent to the condition of $n(y)$ being parallel to $f(y) - f(x)$ in the definition of C.

To generalize the minimum condition to arbitrary points of piecewise smooth surfaces, consider the set of faces Ω_i containing a point y (one face for face points, two for edge points, and any number for vertex points). For each Ω_i and each parametric direction at y, there is a well-defined tangent t.

The condition for a local minimum is that for any Ω_i, and any tangent t, the distance increases along the tangent, i.e.,

$$\forall t,\ t^\top \nabla_y f(y)\,(f(y) - f(x)) \geq 0,$$

where the gradient ∇_y is computed with respect to a parametrization of Ω_i.

If $i > 1$, for any face Ω_i, there are two unit tangent vectors $t_i^1(y)$ and $t_i^2(y)$ along two edge curves of Ω_i meeting at y (Figure 9). Any tangent vector t at y for Ω_i is a linear combination of t_i^1 and t_i^2 unless these are collinear; however, in general, one cannot infer that f increases along all directions t between t_i^1 and t_i^2 from increase along these two directions, unless the angle between them is strictly less than π, i.e., not for points on straight edges, or curved surfaces if the patch Ω_i boundary has a concave corner.

To treat all cases uniformly, we add a halfway vector t_i^3; corresponding to the direction halfway between t_i^1 and t_i^2. Note that the issue described above doesn't happen to piecewise linear surfaces, so t_i^3 is not used in this case.

Then every unit tangent vector t within Ω_i can be represented as a linear combination of t_i^1 and t_i^3, or t_i^3 and t_i^2, with *non-negative* coefficients, and consequently, positivity along these three directions is sufficient for all directions.

Denote $g(y) = \nabla_y f(y)\,(f(y) - f(x))$. Then if $t_i^k \cdot g(y) \geq 0$ for $k = 1, 2, 3$, the same holds for any $t = a_\ell t_i^\ell + a_m t_i^m$, $(\ell, m) = 1, 2$ or $2, 3$.

The only special case is the case of face points, which do not have any edge vectors; in this case, we can pick t_i^k as arbitrary tangent vectors spaced at equal 120-degree angles in the parametric domain. Note that this doesn't happen to piecewise linear surfaces.

The analog of $\Phi^n(x,y)$ requires a criterion for determining that a unit vector v at a point x of a piecewise surface points outside. Let P_i be the tangent plane to Ω_i, and let n_i be the normals. Suppose for at least one face, the projection of the vector v to P_i is inside the face. Then the planar sector defined by the rays along v and the closest projection v', to the plane P_i which has minimal value of $|v \cdot n_i|$, lies entirely on one side of the surface locally at $f(x)$. Thus
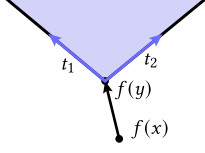
Fig. 9. 2D version of notation for Definition 1.

if $n_i \cdot v$ which is signed distance from the tip of v to P_i is positive, it is pointing to the outside of the surface. Define $d_i^+ = \min_i v \cdot n_i$, over i for which $v \cdot n_i > 0$, and v projects inside the face i, setting it to ∞ if this set is empty. Similarly define $d_i^- = \min_i -v \cdot n_i$ for $v \cdot n_i < 0$, and similarly v projects inside the face i, setting it to ∞ if the set is empty (both cannot be empty). Then v points outwards if $d_i^- - d_i^+ > 0$.

This leads to the following definition of a contact candidate set:

Definition 1 (Piecewise smooth contact set). For each point y of Ω, let I be the set of indices of patches Ω_i containing y, and for each patch, define $t_i^k(y)$, $k = 1, 2, 3$ as above. Then $y \in C(x)$ for a point y iff

$$\Phi_{ik}(x, y) := t_i^k(y)^\top \nabla_y^i f(y)(f(y) - f(x))_+ \geq \alpha \quad (5)$$

for $i \in I$, $k = 1, 2, 3$, with $\nabla_y^i f$ computed with respect to Ω_i, and α is a positive number, *and* in addition, $d_i^- - d_i^+ \geq -\alpha$

Contact potential for p.w. smooth surfaces. Using the new definition of contact candidate sets $C(x)$ we can define the contact potential in a similar way to smooth surfaces, except we use a new angle factor, defined as

$$\gamma^{PS}(x, y) = \prod_{i=1}^{m_y} \prod_{k=1}^{3} H^{\alpha,b}(\Phi_{ik}(x, y)) \quad (6)$$

where m_y is the number of patches meeting at a vertex, $m_y = 1$ for a face point y, $m_y = 2$ for an edge point, and $m_y > 2$ for a vertex point. The function $H^{\alpha,b}(z)$ is a smoothed Heaviside function (Figure 7) with shift, which is 1 for positive arguments $z \geq b$ and 0 below $-\alpha$:

$$H^{\alpha,b}(z) := H\left(\frac{3}{\alpha + b}(z - b)\right),$$

satisfying

$$H^{\alpha,b}(b) = 1, \ H^{\alpha,b}(-\alpha) = 0, \text{and} \ H^{\alpha,b} \in C^2(\mathbb{R}),$$

Additionally, we define

$$\Phi_i^n(x, y) = n_i(y)^\top \nabla_y^i f(y)(f(y) - f(x))_+$$

where $n_i(y)$ is the face normal of patch Ω_i at y, which is unique per patch on a piece-wise linear surface. To eliminate the contact between opposite sides of the object (Figure 6), the following should hold

$$\Phi_i^n(x, y) < 0 \quad \text{for some } i \in 1, \ldots, m_y.$$

Note that not all normals satisfy the constraint above when there is a collision, so the constraint is on only one normal direction. While this condition for determining the correct side is precise, its mollification is relatively complex. For this reason, we relax this to

a simpler, weaker condition, following a similar way of smoothing the constraints on Φ_{ik} and multiply $\gamma^{PS}(x, y)$ by

$$H^{1,0}\left(-1 + \sum_{j=1}^{m_y} H^{\alpha,b}(-\Phi_i^n(x, y))\right),$$

which reaches maximum 1 when there exists i such that

$$\Phi_i^n(x, y) \geq b,$$

and vanishes if for all i

$$\Phi_i^n(x, y) \leq -\alpha.$$

Since Φ_i^n, Φ_{ik} are essentially dot products between unit vectors, α acts as a threshold of contact potential on the cosine value of the angle between contact pairs. One can pick α based on the desired angle threshold.

Recall that contacts between vertices and faces may not necessarily lead to non-zero potentials for faces. As a consequence, the interaction potential in these configurations will vanish outside of a measure-zero set (edge curves and/or vertices), and direct integration of $\psi(x)$ will lead to a zero potential.

For this reason, we treat low-dimensional elements separately. Informally, we can think about the interaction of lower-dimensional elements, as expanding them into areas: e.g., an area of width L along curves, and an area of size L^2 assigned to vertices, with the potential constant along the direction in the area orthogonal to the curve, or on the whole area assigned to a vertex. More precisely, this corresponds to adding line integrals for edge curves weighted by L, and points sums for vertices, weighted by L^2. The potentials we integrate for a pair of elements of G and H, possibly of different dimensions, can be written in the form

$$P(x, y; G, H) = \gamma^{PS}(x, y; G)\gamma^{PS}(y, x; H)p_{\epsilon(x)}(x, y)$$

where $\gamma^{PS}(x, y; G)$ is the angle factor defined in Equation (6), *for x considered as a point on element G*. E.g., if x is a vertex point on a boundary of a face G, then it is considered as a face point, with a single face used. If G however is an edge, both faces are used for x.

Then the total potential can be written as

$$\sum_{(g,h)} L^{2-\dim g - \dim h} \sum_{i \in I_g, j \in I_h, G_i \neq H_j} \int_{G_i} \int_{H_j} P(x, y; G_i, H_j) \quad (7)$$

where g, h are one of V, E, F, the sum is over 6 unordered pairs I_g and I_h are sets of indices of elements of types g and h, and the integrals are area, line, or 0-dimensional, i.e., simple evaluations for vertices.

If we define the distance to contact as

$$\min_{y \in C(x)} d(x, y),$$

then the potential Equation (7) satisfies all required conditions.

Parameters of the potential. In principle, whether the potential satisfies contact potential requirements does not depend on the choice of L, ϵ_{trg}, and α, as long as these are positive.

- α determines the size of the contact candidate sets, and as a consequence how smoothly the potential depends on f: the closer α is to zero, the less smooth the dependence is.

- ϵ_{trg} determines the upper bound on how far the potential extends ($\hat{d}$ in IPC) but especially at concave corners, ϵ_{trg} as well as f_0, affect this.
- L determines the strength of potential for low-dimensional contact.

4.4 Discretization

To discretize these integrals, and still satisfy the contact potential conditions, we cannot directly apply an arbitrary quadrature: if we evaluate the potentials at arbitrary points on elements we may miss a contact point, and the potential will not satisfy the barrier property. To address this problem we choose the closest points as our quadrature points, possibly reducing integration precision, but guaranteeing that the potential remains a barrier.

In this case, the lowest-order discretization of the potential integral is:

$$\sum_{(g,h)} \sum L^{2-\dim g - \dim h} \sum_{(i,j)} P(x_i, y_j; G_i, H_j) A(G_i) A(H_j)$$

where (x_i, y_j) is a pair of closest points on elements G_i and H_j, respectively, and (i, j) are pairs of non-adjacent elements, and $A(G_i)$ and $A(H_j)$) are measures of elements (area for faces, length for edges and 1 for vertices). Here the adjacent elements are dropped from summation.

For the discrete potential to depend differentiably on the shape parameters (vertex positions in the piecewise-linear case), we need an additional modification. First, observe that the distance between two elements is already C^1 with the only exception being the distance between two parallel edges, which can be mollified as explained in [Li et al. 2020], so the term $p_{\epsilon(x)}$ in P does not require modifications. However, the terms $\gamma(x, y)$ depend on the direction between closest points, which may change non-smoothly. We use a mollification for the closest point position to obtain a smooth variation. Importantly, it does not affect the barrier property or the contact set properties.

Distance function mollification. For an edge $[p_0, p_1]$, and a point q, we define the mollified closest point as follows. Let $0 \le s < 1$ be the coordinate of q along the edge, and $0 \le c < 1/2$ be the mollification parameter (we use 10^{-3}). Our mollification is based on

$$h(z) := \begin{cases} 0 & z \le 0 \\ z(2 - z) & 0 < z < 1 \\ 1 & 1 \le z \end{cases} \qquad h_c(s) := h\left(\frac{s}{c}\right)$$

Define three basis functions $m_0(s) = h_c(c - s)$, $m_1(s) = h_c(s - (1 - c))$ and $m_2(s) = 1 - m_0(s) - m_1(s)$, and mollified closest point as $m_0(s_c)p_0 + m_1(s_c)p_1 + m_2(s_c)p_c$, where p_c is the closest point, and s_c its coordinate along the edge. Then the γ is computed using the mollified point.

Similarly for a triangle, we use barycentric coordinates (s_1, s_2, s_3) to define a mollified closest point, with 7 basis functions, one for each vertex and edge and the central one:

$$m_i^e(s_i) = m_0(s_i)(1 - m_0(s_j))(1 - m_0(s_k))$$
$$m_i^v(s_i) = (1 - m_0(s_i))m_0(s_j)m_0(s_k)$$

where i, j, k is a permutation of $(1, 2, 3)$ and m_i^e and m_i^v are the edge and vertex basis functions. The mollified closest point combines the vertex, edge, and face closest points with these weights.

Note that with the amplification above, the edge-edge mollification in IPC is not needed anymore. The edge-edge mollification was introduced to avoid non-smooth distance with almost parallel edges. In our case, not only the parallel edges case but also the case where edge-vertex distance reduces to vertex-vertex distance is mollified.

Adaptive ϵ_{trg}. We introduce adaptive ϵ_{trg} for our formulation. We first specify a fixed $\hat{d}_{\max}$ for the simulation, and collect contact pairs in the rest configuration within $\hat{d}_{\max}$ and with nonzero potential values. We pick ϵ_{trg} for every primitive (vertex/edge/face) so that none of the contact pairs are active. Then the ϵ_{trg} for every edge (face) is chosen to be the minimum among its neighboring vertices (edges). Benefited from our formulation with angle constraints, for a cube mesh there's no contact for a wide range of α, so the adaptive ϵ_{trg} can be any positive number without causing spurious contact forces between its primitives.

Adjacent elements can be excluded. Specifically, if adjacent elements are in contact, other, nonadjacent elements will be in contact. Adjacent elements such that one does not contain another can be vertex-adjacent face and edge, and vertex- and face-adjacent faces. Suppose an edge and a face are in contact, i.e., an edge is in the plane of the face, and they share a vertex and other points. Then the other edge endpoint is either inside the face, or the edge intersects the boundary of the face, crossing an edge which it does not share a vertex with. In the first case, we have vertex-face contact, in the second case, we have non-adjacent edge-edge contact. If it is a face-face contact, the argument reduces to the first, if applied to one of the face edges that is not a common edge with the other face.

5 EVALUATION

We implement our algorithm in C++ by extending the IPC Toolkit [Ferguson et al. 2020] (details provided in the supplemental document). We use Eigen [Guennebaud et al. 2010] for linear-algebra operations, Pardiso [Alappat et al. 2020; Bollhöfer et al. 2019, 2020] for solving linear systems, auto-differentiation [aut 2012] for the derivatives of our potential, and PolyFEM [Schneider et al. 2019] as the finite element simulation framework. All experiments are run on a system with an AMD Ryzen Threadripper PRO 3995WX 64-Cores (limited to 16 threads) and 440 GiB of memory. For comparisons against IPC [Li et al. 2020] and "Convergent IPC" [Li et al. 2023], we use the open-source implementations provided in the IPC Toolkit (together with PolyFEM). Please see our supplemental video for animations. All simulation parameters and statistics are summarized in Table 1. Our reference implementation, used to generate all results, will be released as an open-source project.

5.1 Resolving Spurious Forces

In this section, we consider a variety of unit tests where the original IPC formulation introduces spurious forces.

Using larger ϵ_{trg}. Suppose the input mesh is relatively fine, either globally (a dense mesh) or locally (adaptively refined e.g. to resolve fine features). Having a large ϵ_{trg} allows one to solve the implicit
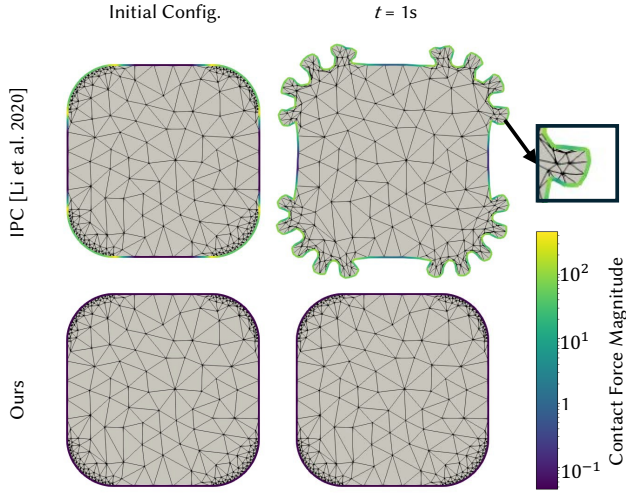
Fig. 10. **Spurious stresses: nonuniform squircle mesh.** Our method properly handles nonuniform meshes. In this case, we consider a 2D rounded block (of size 1 m x 1 m) with maximal edge length 0.21 m in its straight sections and minimal edge length 0.01 m at its corners. Top row: with $\hat{d} = 0.1$ m, IPC introduces spurious contact forces in the refined corners, resulting in a deformation. Bottom row: our method avoids this by calculating $\epsilon(x)$ such that the potential is zero in the rest configuration.
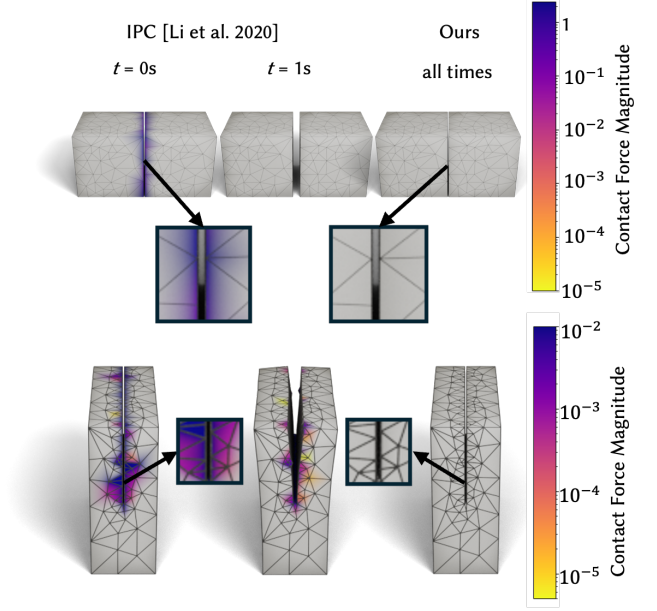

Fig. 11. **Spurious stresses: rest configuration 3D.** Top row: two cubes initially separated by less than $\hat{d} = 0.025$ m. IPC artificially repels the two blocks while ours does not. Bottom row: a block with a slit of width less than $\hat{d} = 0.0125$ m. Our method does not introduce spurious forces across the slit.
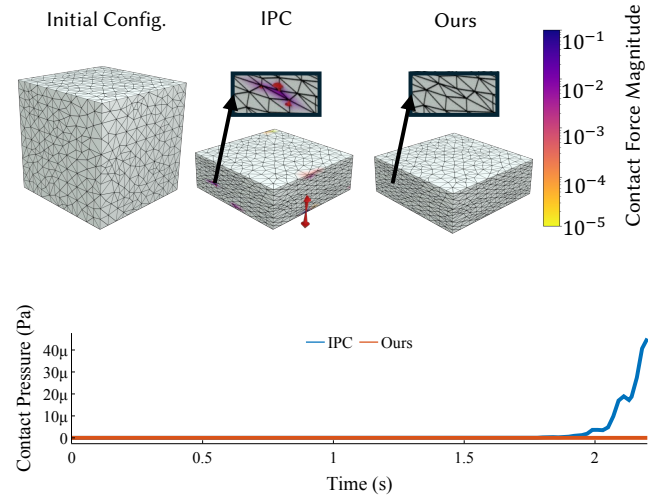
time-stepping problem faster because the barrier is less numerically stiff, and CCD needs to do less work in the line search if the objects are kept further apart by the potential. An even stronger reason, when for large deformations, the elements may shrink a lot, is given below Figure 12. We provide an example of this scenario in Figure 10, where it is natural to have small edges around the rounded corners of a square and long edges along the sides. However, doing so restricts the range of usable $\hat{d}$ for IPC. Using a value ($\hat{d} = 0.1$ m in this case) larger than the minimum edge length ($h_{\min} = 0.01$ m) results in spurious forces along the corners and artifacts upon simulation. In contrast, our utilization of an adaptive ϵ allows us to choose a starting ϵ_{trg} that results in zero initial contact force, while still being large for the non-refined regions.

Spurious stress in the rest configuration. Even if $\hat{d} < h_{\min}$ is satisfied, the IPC formulation may still have spurious stress in the rest configuration. We show two such scenarios in Figure 11.

In the first case (Figure 11 Top), two blocks in the initial configuration sit on the plane with the initial distance between blocks less than $\hat{d}$. With IPC, they incorrectly start sliding apart without external force applied where the blocks should stay still at all times. In the second case (Figure 11 Bottom), IPC causes the slit to expand at the top without external forces and stress appears at the bottom.

Our approach avoids spurious forces between close objects by using an adaptive ϵ_{trg}. For the contact between primitives on the same object, due to our constraints on the angle of contact pairs, we can use ϵ_{trg} larger than the edge lengths without activating contact at rest configuration (with a wide range of α).

Fig. 12. **Spurious stresses: compression.** An initially valid choice for $\hat{d}$ can lead to spurious contact (arrows and color) when using IPC. Left: cube at rest with a minimum edge length of 0.03 m. Center: upon compression with $\hat{d} = 0.0125$ m, IPC introduces spurious contacts. Right: our method avoids this by considering the angle when finding candidate contacts. Bottom: IPC introduces artificial contact pressure.
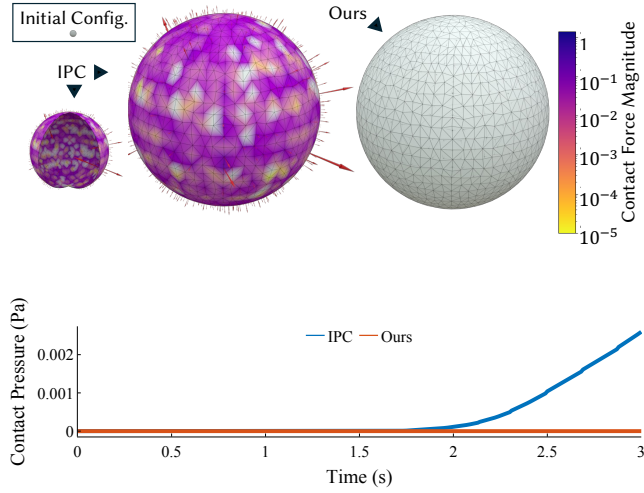
Fig. 13. **Spurious stresses: expansion.** Upper left: an initially deflated balloon with rest thickness of 0.1 m is expanded by an outward force of 5 kN. Center: with $\hat{d} = 0.02$, the balloon eventually becomes thin enough that spurious contact forces (represented as red arrows) between its inner and outer layers appear. Lower left: a cross-section showing spurious forces on both the inner and outer layer of the balloon. Right: our method can inflate the balloon until it is arbitrarily thin without introducing artificial contact forces. Bottom: artificial contact pressure is introduced by IPC.

Spurious stress under deformation. Even for a well-chosen $\hat{d}$ parameter and a benign rest configuration, IPC can still add spurious forces upon deformation.

In Figure 12, we compress a cube mesh to 33% of its original height. We assign a Poisson ratio of 0 to the cube to show an example where no bulging and/or folding of the surface occurs. The value for $\hat{d}$ is initially chosen such that no points are in contact, but upon compression distances shrink and IPC introduces spurious contact forces on the sides of the cube (despite them being flat). Our formulation does not have spurious forces because we use the angle of contact to build our collision candidate set.

In Figure 13, we inflate a balloon modeled as a volumetric membrane of thickness 0.1 m.[1] As the balloon inflates, its walls get thinner and eventually they become thinner than the initially chosen $\hat{d} = 0.02$ m. At this point, IPC treats the inner and outer sides as in contact, introducing forces between the two sides. Our collision candidate set for one side does not include the other as it accounts for the angle of contact, avoiding this issue.

In Figure 14, we extrude the extended χ-shaped structure from [Joodaky 2020] to 3D and simulate its compression. When the shape buckles it forms *cusp* contacts at the corners. We show that Convergent IPC [Li et al. 2023] exhibits large contact forces in these regions, IPC [Li et al. 2020] exhibits less forces, while our method has much less contact forces than both methods. Our method reduces the contact forces around the *cusp* because of the constraints on the angle

---

[1]While one could model the balloon using codimensional shell elements to avoid this issue, modeling the thickness may be important for analysis or design. For example, with constant outward pressure, the balloon oscillates in thickness, which a shell model would not capture.
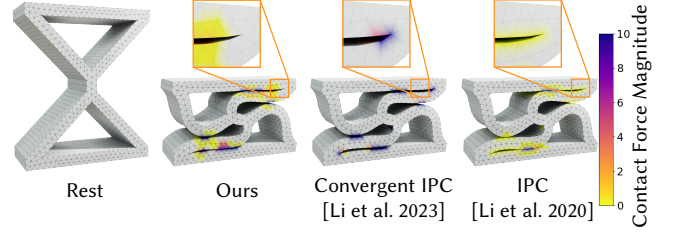
Fig. 14. **Cusp compression.** Compression of a χ-shaped structure. For comparison, we use the adaptive barrier stiffness for IPC, κ equals Young's modulus both for convergent IPC and our method. In the last time step, IPC took 257 iterations, convergent IPC 58, and ours 72.
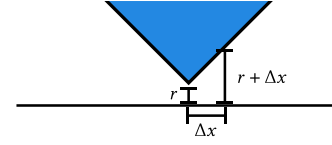

Fig. 15. Setup for corner hitting a plane discussed in Section 5.2.

between faces – only when edges are close enough to parallel, i.e. the angle of the *cusp* becomes small enough, the contact is activated. Without this constraint, large κ results in spurious stresses as in Convergent IPC, while small κ results in convergence issues in IPC – the number of iterations is much higher due to tiny distance of 10^{-11} m between contact pairs.

5.2 Infinite Potential

Consider a corner hitting a plane, and suppose they are both refined (Figure 16 Right). In this case, the integral of the Convergent IPC [Li et al. 2023] potential over the plane will be finite, even if the apex of the corner is directly on the plane (violation of Requirement 2). Although each descritization will be infinite, the implication is that the total potential will decrease as we refine, so the actual distance will decrease to zero.

Figure 15 show the configuration: for a gap of r and a distance from the closest point on the plane Δx, points along the square will be at a distance of $\Delta x + r$. Assume the potential for points on the corner is just $\ln\left(\frac{\Delta x + r}{\hat{d}}\right)$,[2] then integrating over x we get

$$(r + L) \ln\left(\frac{r + L}{\hat{d}}\right) - \ln\left(\frac{r}{\hat{d}}\right) r - L,$$

and for $r = 0$ (i.e. contact), a finite limit $\ln\left(\frac{L}{\hat{d}}\right) L - L$.

The force (i.e., the potential derivative) at $r = 0$ is $\sim \ln(r)$ though (i.e. is infinite), so for a finite force the static equilibrium problem will always have a valid solution. This is not the case however for a dynamic problem: if the kinetic energy exceeds the finite potential, then the barrier cannot prevent contact. We show this case in Figure 16. One can see the Convergent IPC model results in ever decreasing distances as the mesh is refined (without refining $\hat{d}$). Our method in comparison exhibits the same trajectory for all

---

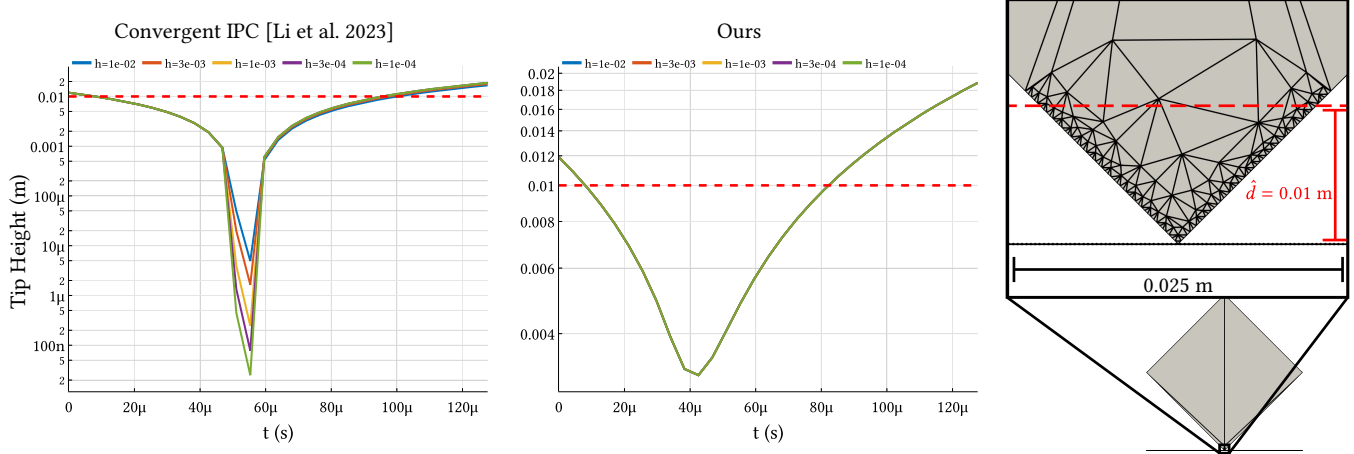[2]The extra quadratic term does not affect the conclusions.

Fig. 16. **Finite Potential for Zero Distance.** A square impacts a plane at the corner of the square with a velocity of 234 m/s. We refine both meshes (only at the corner of the square of efficiency). Here we plot the height of the squares tip over time. Left: for Convergent IPC [Li et al. 2023], we see that the minimum distance shrinks with every refinement of the mesh. This is a consequence of the continuous model having a finite potential. Middle: due to our choice of barrier function, our method exhibits the same trajectory (all plots overlap) under refinement.

levels of refinement. This is because we use a barrier function whose integral over the surface is not finite as the distance goes to zero, and because the potential converges under mesh refinement.

Given that the minimum distance doesn't depend on the mesh resolution, one can pick ϵ_{trg}, independent of the edge lengths, based on the deired accuracy on the contact handling.

5.3 Parameter Ablation and Convergence Study

To study the effects of our simulation parameters we conduct ablation studies for α and ϵ_{trg}.

Lower values of α (Figure 17) make the solution more accurate, preventing spurious forces from appearing but requiring more iterations. This effect is more pronounced in scenes with cusps: for simpler scenes the difference in iteration count is less noticeable (Figure 18).

We plot in Figure 18 the effect of ϵ_{trg} on the number of iterations. As the support of the contact potential increases the problem becomes softer and the number of iteration decreases accordingly.

We also perform a convergence study for three different scenarios show in Figure 19. We see convergence under mesh refinement in all of these scenes. Importantly, we use a fixed ϵ_{trg} (unlike Convergent IPC [Li et al. 2023] which requires co-refinement of $\hat{d}$).

We note that when our discretization is applied to PL surface, treating it as a piecewise smooth surface, the integrals for edge potentials and summations for vertex potentials are introduced, with the relative scale of potentials determined by the constant L. As the surface is refined, the scale of the constant also needs to be adjusted, for the refined mesh potential to approximate the smooth surface potential. We leave a rigorous study of convergence of the discrete potential to the potential of the limit smooth surface as future work.
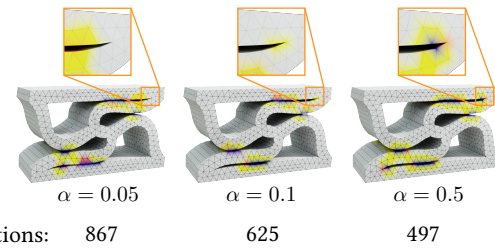

Fig. 17. **Parameter-study:** α. Simulation of Figure 14 with different α including the total number of solver iterations. As α increases, the nonlinear problem becomes softer, hence fewer iterations. However, large α introduces artifacts and spurious stresses in some cases.
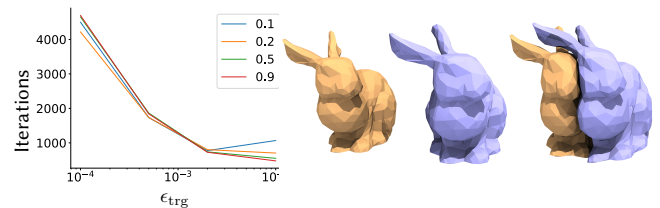

Fig. 18. **Parameter-study:** ϵ_{trg}. Simulation of two bunnies colliding with varying α and ϵ_{trg}. We show the plot (left) of the number of iterations over ϵ_{trg}, with a different α for each curve, as well as the initial frame (middle) and the colliding frame (right) of the simulation. For simple scenes, for a large range of α, the number of iterations is similar; as ϵ_{trg} increases, it takes fewer iterations to converge.

5.4 3D Examples

Lastly, we reproduce challenging 3D simulation examples of [Li et al. 2020]. First, we validate our method on the 3D unit tests proposed by [Erleben 2018] in Figure 20. We see similar results to those shown
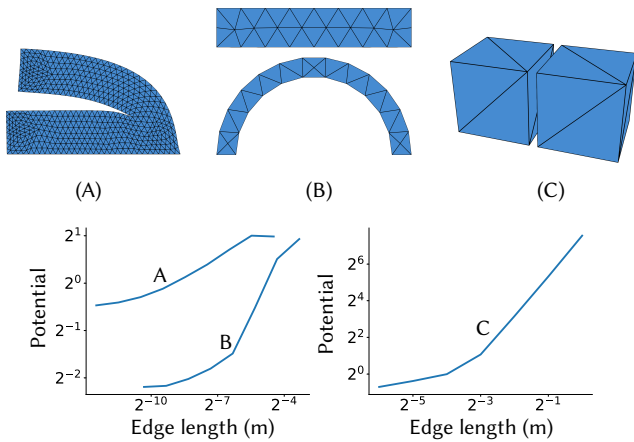
Fig. 19. Convergence of potential under mesh refinement in fixed scenes with fixed ϵ_{trg}. We pick 2 configurations in 2D and 1 in 3D, start with the coarse mesh, and compute the potential under mesh refinement. In 2D, the potentials of (A) and (B) converge with an order of 0.82 and 1.57. Due to the nearly singular cusp, (A) converges slower than (B). In 3D, although we cannot further refine the surface mesh due to memory restriction (the number of contact candidates grows quadratically under refinement), it shows a trend of converging.
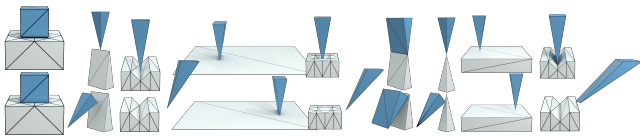

Fig. 20. **Erleben tests.** We reproduce the test-cases of Erleben [2018]. Top: initial conditions involving challenging exact point-point, point-edge, and edge-edge collisions. Bottom: as in [Li et al. 2020], our approach robustly passes all the tests.

in [Li et al. 2020], but we highlight one improved result in Figure 21 where we see reduced spurious tangential movement compared to [Li et al. 2020]. Second, we reproduce the dolphin funnel (Figure 22), trash-compactor (Figure 23), and mat twist (Figure 24) examples. Each of these examples features large deformations and complex contacts. Just as in [Li et al. 2020], we robustly handle these scenarios and prevent intersections and inversions at every step.

6 CONCLUSION

Incremental potential contact [Li et al. 2020] enabled a qualitative improvement in robustness of accurate simulation of deformable objects with complex contact. In this paper, we revisit the IPC contact formulation to relate it to a family of potentials defined for a broad class of surfaces, that satisfy a set of natural requirements for barrier potentials for contact. We demonstrate that applying these principles leads to a new formulation that is similar in efficiency to IPC but alleviates some of its shortcomings.

Future work. While we have not considered friction in this paper, the extension to friction is straightforward: the IPC formulation
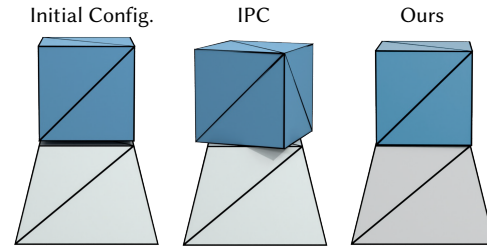

Fig. 21. **Erleben test: cliff edge.** We reproduce the cliff edge test-case of Erleben [2018]. Middle: [Li et al. 2020] passes the test but introduces spurious horizontal forces, causing the top block to rotate. Right: our method significantly reduces extra sliding due to the restriction in normal forces.
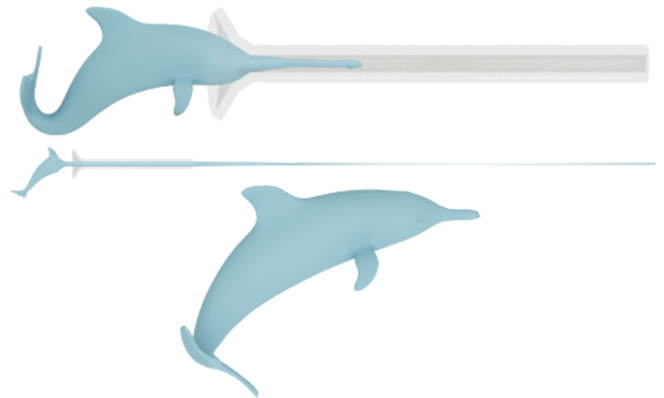

Fig. 22. **Dolphin in a funnel.** We reproduce the funnel test from [Li et al. 2020] using our method. Top: an elastic dolphin is pulled through a small tube. Middle: this causes extreme deformations. Bottom: the dolphin squeezes through without artifact and recovers its original shape.
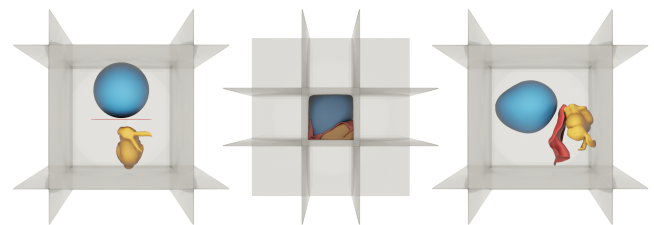

Fig. 23. **Trash-Compactor.** We reproduce the trash-compactor example from [Li et al. 2020] using our method. Left: three objects are placed in a compactor. Middle: the objects are compressed. Right: the compactor releases and the shapes return to their original shape without intersections.

applies essentially with no changes. Our derivation assumes surfaces without boundaries at several steps, which needs modifications for co-dimensional objects; a possible direction for future work is handling surfaces with boundaries and co-dimensional surfaces. Additionally, it may be beneficial to use a high-order quadrature to better integrate the continuous formulation (in addition to including the closest-point quadrature).

Table 1. **Simulation statistics.** For each simulation we report geometry, (minimum, average, and maximum edge length h), time step Δt, material (Young's modulus E, Poisson ratio ν, and density ρ), ϵ_{trg}, maximum memory used, as well as average timing and number of Newton iterations. A value of 0 for iterations indicates the optimization did not run because the initial configuration was at a force equilibrium (i.e., no spurious forces at rest).

| Example	# nodes, # cells	h (m) (min, avg, max)	Δt (s)	E (Pa), ν, ρ (kg/m^3)	ϵ_{trg} (m)	α	memory (MB)	timing (s), iterations (per timestep)
Slit block (2D) (Supp. Video)	199, 325	0.04, 0.18, 0.34	0.0047	4000, 0.2, 100	0.1	0.1	19	0.07, 0
Slit block (3D) (Fig 11)	1058, 3871	0.01, 0.04, 0.13	0.0047	100, 0.2, 10	0.0125	0.1	93	0.14, 0
Fillet block (2D) (Fig 10)	445, 672	0.01, 0.04, 0.21	0.0047	4000, 0.2, 100	0.1	0.1	26	0.09, 0
Fillet block (3D) (Fig 1)	1003, 3824	0.02, 0.19, 0.42	0.0047	100, 0.2, 10	0.07	0.1	92	0.14, 0
Compressed block (2D) (Supp. Video)	356, 630	0.09, 0.12, 0.22	0.0047	10000, 0, 100	0.04	0.1	63	0.17, 2.3
Compressed block (3D) (Fig 12)	2169, 9799	0.03, 0.10, 0.18	0.0047	10000, 0, 100	0.0125	0.1	644	0.78, 2.5
Two blocks (2D) (Supp. Video)	708, 1260	0.09, 0.12, 0.22	0.0047	4000, 0.2, 100	0.1	0.1	24	0.08, 0
Two blocks (3D) (Fig 11)	580, 2079	0.04, 0.21, 0.55	0.0047	10000, 0, 100	0.025	0.1	61	0.14, 0
Balloon (3D) (Fig 13)	421, 701	0.04, 0.11, 0.18	0.0047	4000, 0, 100	0.02	0.1	722	0.67, 1
Dolphin funnel (Fig 22)	4074,10511	0.0017,0.020,0.081	0.025	10000,0.4,1000	0.001	1	1403	15.4, 58.1
Mat twist (Fig 24)	45000, 133206	0.0067, 0.0088, 0.0126	0.04	20000,0.4,1000	0.002	0.4	17683	32.4, 195
Cusp compression (Fig 14)	3761, 12777	0.024, 0.047, 0.089	0.01	1000000, 0.3, 100	0.015	0.05	501	5.02, 15.8
Trash compactor (Fig 23)	6611, 21696	0.00033,0.049,0.36	0.002	10000, 0.4, 1000	0.002	1	2833	119.8, 125
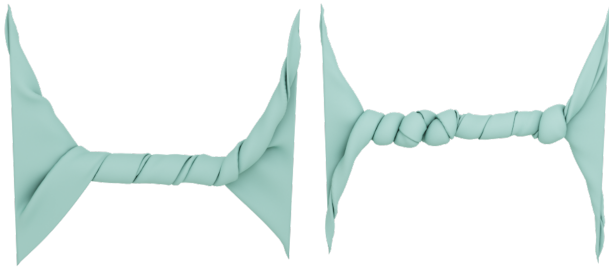
Fig. 24. **Mat-Twist.** We reproduce the mat-twist example from [Li et al. 2020], demonstrating, overall, our method works similarly to the original IPC. Left: our simulation at 10 s after 2 rounds of twisting at both ends. Right: at 27.5 s after 5.5 rounds of twisting.

REFERENCES

2012. Mitsuba Renderer. https://www.mitsuba-renderer.org/misc.html Accessed: 2023-10-10.

Christie Alappat, Achim Basermann, Alan R. Bishop, Holger Fehske, Georg Hager, Olaf Schenk, Jonas Thies, and Gerhard Wellein. 2020. A Recursive Algebraic Coloring Technique for Hardware-Efficient Symmetric Sparse Matrix-Vector Multiplication. *ACM Transactions on Parallel Computing* 7, 3, Article 19 (June 2020), 37 pages.

Francisco Armero and Eva Petőcz. 1998. Formulation and analysis of conserving algorithms for frictionless dynamic contact/impact problems. *Computer methods in applied mechanics and engineering* 158, 3-4 (1998), 269–300.

F.B. Belgacem, P. Hild, and P. Laborde. 1998. The mortar finite element method for contact problems. *Mathematical and Computer Modelling* 28, 4 (1998), 263–271. Recent Advances in Contact Mechanics.

Ted Belytschko, Wing Kam Liu, and Brian Moran. 2000. *Nonlinear Finite Elements for Continua and Structures*. Vol. 26. John Wiley & Sons, Ltd. 38–3926–38–3926 pages.

Ted Belytschko and Mark O Neal. 1991. Contact-impact by the pinball algorithm with penalty and Lagrangian methods. *Internat. J. Numer. Methods Engrg.* 31, 3 (1991), 547–572.

Matthias Bollhöfer, Aryan Eftekhari, Simon Scheidegger, and Olaf Schenk. 2019. Large-scale Sparse Inverse Covariance Matrix Estimation. *SIAM Journal on Scientific Computing* 41, 1 (2019), 380–401.

Matthias Bollhöfer, Olaf Schenk, Radim Janalik, Steve Hamm, and Kiran Gullapalli. 2020. State-of-the-Art Sparse Direct Solvers. *Parallel Algorithms in Computational Science and Engineering* (2020), 3–33.

Gregory Buck and Jeremey Orloff. 1995. A simple energy function for knots. *Topology and its Applications* 61, 3 (1995), 205–214.

Nicholas J Carpenter, Robert L Taylor, and Michael G Katona. 1991. Lagrange constraints for transient finite element surface contact. *International journal for numerical methods in engineering* 32, 1 (1991), 103–128.

PW Christensen, Anders Klarbring, Jong-Shi Pang, and Niclas Strömberg. 1998. Formulation and comparison of algorithms for frictional contact problems. *Internat. J.*

Numer. Methods Engrg. 42, 1 (1998), 145–173.

Peter Deuflhard, Rolf Krause, and Susanne Ertel. 2008. A contact-stabilized Newmark method for dynamical contact problems. *Internat. J. Numer. Methods Engrg.* 73, 9 (2008), 1274–1290.

Manfredo P Do Carmo. 2016. *Differential geometry of curves and surfaces: revised and updated second edition*. Courier Dover Publications.

Yinwei Du, Yue Li, Stelian Coros, and Bernhard Thomaszewski. 2023. No Free Slide: Spurious Contact Forces in Incremental Potential Contact. arXiv:2308.01696 [cs.GR]

Thang Duong and Roger Sauer. 2019. A concise frictional contact formulation based on surface potentials and isogeometric discretization. *Computational Mechanics* 64, 4 (2019), 951–970.

Kenny Erleben. 2018. Methodology for Assessing Mesh-Based Contact Point Methods. *ACM Transactions on Graphics* 37, 3, Article 39 (July 2018), 30 pages.

Zachary Ferguson et al. 2020. IPC Toolkit. https://ipc-sim.github.io/ipc-toolkit/. https://ipc-sim.github.io/ipc-toolkit/

Zachary Ferguson, Pranav Jain, Denis Zorin, Teseo Schneider, and Daniele Panozzo. 2023. High-Order Incremental Potential Contact for Elastodynamic Simulation on Curved Meshes. In *ACM SIGGRAPH 2023 Conference Proceedings* (Los Angeles, CA, USA) *(SIGGRAPH '23)*. Association for Computing Machinery, New York, NY, USA, 11 pages.

Zachary Ferguson, Minchen Li, Teseo Schneider, Francisca Gil-Ureta, Timothy Langlois, Chenfanfu Jiang, Denis Zorin, Danny M. Kaufman, and Daniele Panozzo. 2021. Intersection-Free Rigid Body Dynamics. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 40, 4, Article 183 (July 2021), 16 pages.

Gaël Guennebaud, Benoît Jacob, et al. 2010. Eigen v3. http://eigen.tuxfamily.org. http://eigen.tuxfamily.org

David Harmon, Daniele Panozzo, Olga Sorkine, and Denis Zorin. 2011. Interference-Aware Geometric Modeling. *ACM Transactions on Graphics* 30, 6 (Dec. 2011), 1–10.

David Harmon, Etienne Vouga, Breannan Smith, Rasmus Tamstorf, and Eitan Grinspun. 2009. Asynchronous contact mechanics. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 28, 3 (2009), 1–12.

David Harmon, Etienne Vouga, Rasmus Tamstorf, and Eitan Grinspun. 2008. Robust Treatment of Simultaneous Collisions. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 27, 3 (Aug. 2008), 1–4.

Patrice Hauret and Patrick Le Tallec. 2006. Energy-controlling time integration methods for nonlinear elastodynamics and low-velocity impact. *Computer methods in applied mechanics and engineering* 195, 37-40 (2006), 4890–4916.

Michael Hiermeier, Wolfgang A Wall, and Alexander Popp. 2018. A truly variationally consistent and symmetric mortar-based contact formulation for finite deformation solid mechanics. *Computer Methods in Applied Mechanics and Engineering* 342 (2018), 532–560.

S. Hüeber and B.I. Wohlmuth. 2006. *Mortar methods for contact problems*. Springer Berlin Heidelberg, Berlin, Heidelberg, 39–47.

Amin Joodaky. 2020. *Mechanics and Design of Polymeric Metamaterial Structures for Shock Absorption Applications*. Ph. D. Dissertation. Purdue University Graduate School.

Jonathan M. Kaldor, Doug L. James, and Steve Marschner. 2008. Simulating knitted cloth at the yarn level. In *ACM SIGGRAPH 2008 Papers* (Los Angeles, California) *(SIGGRAPH '08)*. Association for Computing Machinery, New York, NY, USA, Article 65, 9 pages.

Couro Kane, Eduardo A Repetto, Michael Ortiz, and Jerrold E Marsden. 1999. Finite element analysis of nonsmooth contact. *Computer Methods in Applied Mechanics and Engineering* 180, 1-2 (1999), 1–26.

G Kloosterman, Rudolf MJ van Damme, Antonius H van den Boogaard, and J Huetink. 2001. A geometrical-based contact algorithm using a barrier method. *Internat. J. Numer. Methods Engrg.* 51, 7 (2001), 865–882.

Rolf. Krause and Patrick. Zulian. 2016. A Parallel Approach to the Variational Transfer of Discrete Fields between Arbitrarily Distributed Unstructured Finite Element Meshes. *SIAM Journal on Scientific Computing* 38, 3 (2016).

Lei Lan, Danny M. Kaufman, Minchen Li, Chenfanfu Jiang, and Yin Yang. 2022. Affine Body Dynamics: Fast, Stable and Intersection-Free Simulation of Stiff Materials. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 41, 4, Article 67 (July 2022), 14 pages.

Lei Lan, Yin Yang, Danny Kaufman, Junfeng Yao, Minchen Li, and Chenfanfu Jiang. 2021. Medial IPC: Accelerated Incremental Potential Contact with Medial Elastics. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 40, 4, Article 158 (July 2021), 16 pages.

Egor Larionov, Ye Fan, and Dinesh K. Pai. 2021. Frictional Contact on Smooth Elastic Solids. *ACM Trans. Graph.* 40, 2, Article 15 (apr 2021), 17 pages.

TA Laursen and GR Love. 2002. Improved implicit integrators for transient impact problems—geometric admissibility within the conserving framework. *Internat. J. Numer. Methods Engrg.* 53, 2 (2002), 245–274.

Minchen Li, Zachary Ferguson, Teseo Schneider, Timothy Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M. Kaufman. 2020. Incremental Potential Contact: Intersection- and Inversion-free Large Deformation Dynamics. *ACM Transactions on Graphics* 39, 4, Article 49 (July 2020), 20 pages.

Minchen Li, Zachary Ferguson, Teseo Schneider, Timothy Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M. Kaufman. 2023. Convergent Incremental Potential Contact. arXiv:2307.15908 [math.NA]

Minchen Li, Danny M. Kaufman, and Chenfanfu Jiang. 2021. Codimensional Incremental Potential Contact. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 40, 4, Article 170 (2021).

Steve A. Maas, Benjamin J. Ellis, Gerard A. Ateshian, and Jeffrey A. Weiss. 2012. FEBio: Finite Elements for Biomechanics. *Journal of Biomechanical Engineering* 134, 1 (Feb. 2012).

Miguel Otaduy, Rasmus Tamstorf, Denis Steinemann, and Markus Gross. 2009. Implicit Contact Handling for Deformable Objects. *Computer Graphics Forum* 28 (April 2009), 559–568.

Laetitia Paoli and Michelle Schatzman. 2002a. A numerical scheme for impact problems I: The one-dimensional case. *SIAM J. Numer. Anal.* 40, 2 (2002), 702–733.

Laetitia Paoli and Michelle Schatzman. 2002b. A numerical scheme for impact problems II: The multidimensional case. *SIAM journal on numerical analysis* 40, 2 (2002), 734–768.

Michael A. Puso and Tod A. Laursen. 2004. A mortar segment-to-segment contact method for large deformation solid mechanics. *Computer Methods in Applied Mechanics and Engineering* 193, 6 (2004), 601–629.

Roger A Sauer and Laura De Lorenzis. 2013. A computational contact formulation based on surface potentials. *Computer Methods in Applied Mechanics and Engineering* 253 (2013), 369–395.

Teseo Schneider, Jérémie Dumas, Xifeng Gao, Denis Zorin, and Daniele Panozzo. 2019. PolyFEM. https://polyfem.github.io/.

Paweł Strzelecki and Heiko von der Mosel. 2013. Tangent-Point Repulsive Potentials for a Class of Non-smooth m-dimensional Sets in $\mathbb{R}^n$. Part I: Smoothing and Self-avoidance Effects. *Journal of Geometric Analysis* 23, 3 (2013), 1085–1139.

Robert L Taylor and Panayiotis Papadopoulos. 1993. On a finite element method for dynamic contact/impact problems. *International journal for numerical methods in engineering* 36, 12 (1993), 2123–2140.

Mickeal Verschoor and Andrei C Jalba. 2019. Efficient and accurate collision response for elastically deformable models. *ACM Transactions on Graphics* 38, 2, Article 17 (March 2019), 20 pages.

Didier Vola, Elaine Pratt, Michel Jean, and Michel Raous. 1998. Consistent time discretization for a dynamical frictional contact problem and complementarity techniques. *Revue Européenne des éléments finis* 7, 1-3 (1998), 149–162.

Etienne Vouga, David Harmon, Rasmus Tamstorf, and Eitan Grinspun. 2011. Asynchronous Variational Contact Mechanics. https://github.com/evouga/collisiondetection. *Computer Methods in Applied Mechanics and Engineering* 200, 25 (2011), 2181–2194. https://github.com/evouga/collisiondetection https://github.com/evouga/collisiondetection.

Peter Wriggers. 1995. Finite Element Algorithms for Contact Problems. *Archives of Computational Methods in Engineering* 2 (Dec. 1995), 1–49.

Peter Wriggers and Tod A Laursen. 2006. *Computational contact mechanics*. Vol. 2. Springer.

Tianyi Xie, Minchen Li, Yin Yang, and Chenfanfu Jiang. 2023. A Contact Proxy Splitting Method for Lagrangian Solid-Fluid Coupling. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 42, 4 (2023), 1–14.

Jonathan Youett, Oliver Sander, and Ralf Kornhuber. 2019. A globally convergent filter-trust-region method for large deformation contact problems. *SIAM Journal on Scientific Computing* 41, 1 (2019), B114–B138.
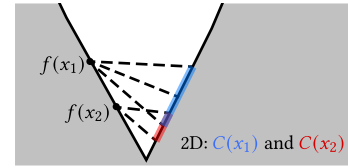
Fig. 25. The contact set for two point x_1 and x_2 on a concave corner may be arbitrarily close, but the contact area for such points becomes increasingly small.

Chris Yu, Caleb Brakensiek, Henrik Schumacher, and Keenan Crane. 2021a. Repulsive Surfaces. *ACM Trans. Graph.* 40, 6 (2021).

Chris Yu, Henrik Schumacher, and Keenan Crane. 2021b. Repulsive Curves. *ACM Trans. Graph.* 40, 2 (2021).

A VERIFICATION OF PROPERTIES

The properties *No spurious forces* and *Localization* are accomplished by construction. *Differentiability* is also attained by construction, as all functions used in the definition of the potential are C^1. One subtle point is that the distance functions used in the *discrete* version of the potential are also C^1, except the edge-to-edge distance that require mollification. However, the closest point on an element to a given point $f(x)$ is *not* a C^1 function of $f(x)$, and for this reason we introduce mollification whenever these points are used.

Our arguments are informal and we leave a complete proof with all the necessary technical assumptions as future work.

Generality. In the case of smooth surfaces, the potential is clearly well-defined as it involves integrating a function bounded from above, as there is a uniform bound on the distance between a point x and a point in its contact set y.

For piecewise smooth surfaces, including PL surfaces, the situation is more subtle, as points may be arbitrarily close to their contact set, if points are located on faces sharing an edge (Figure 25). However, the contact area for such points becomes increasingly small, so if the point-point potential does not grow very quickly, the integrated potential is well-defined.

We assume that the angle γ between the normals of two faces sharing an edge, is never zero. Then for a point at distance r from the edge the distance to the contact set is $d \sim r \sin \gamma$, and the area of the contact set, in 3D, $\sim r^2 \sin \alpha$. Thus, if potential power does not exceed 2, then the integral of the potential over the contact set for a point at distance r to the edge, remains bounded, as $d^{-n} r^2$ is bounded. Similarly, in $2D$, the potential needs to grow no faster than $1/r$.

Barrier. The interplay between the discretization and the continuum formulation is quite subtle in this case. In the continuum case, clearly, it is not sufficient to choose a point-to-point potential function that becomes infinite for coinciding points: the integral of such function may be finite, which would mean that the barrier is not preventing the contact.

Consider first the integral of the potential over one surface S_y with respect to the variable y, with x fixed. Further, assume an isolated contact point with the distance between surfaces growing

as r^n with the distance to the contact point. For a smooth contact $n = 2$, for a point contact (e.g. with a cone), $n = 1$.

Consider a model problem with two surfaces in planar contact on a disk of unit radius, at distance h, and assume $\hat{d}$ and α to be large. Then the total contact potential is approximated by

$$2\pi \int_{r=0}^{R} \int_{q=0}^{R} \int_{\theta=0}^{2\pi} ((r\cos\theta - q)^2 + r^2 \sin^2\theta + h^2)^p \, d\theta \, dr \, dq$$

where R is the contact zone radius, can be computed explicitly, and for the point-to-point potential with $p = 2$, can be shown to be $4\pi^2(\ln(2) - \ln(h))^2$, i.e., proportional to $\ln(h)^2$. If we further assume that one of the surfaces has a cone point, i.e., the distance between the surface and the plane increases linearly w.r.t. the distance along the plane to the contact point, with h replaced with $(h+r)$ above, then the rate of growth reduces to $\ln(h)$, but still results in a repulsive force.

For a potential with a lower rate of growth, the limit interaction potential can be finite, hence it does not result in a barrier in the continuum case.

We note however, that for our construction of the discrete potential this does not necessarily cause a problem for any fixed resolution: as the quadrature for the potential is defined using the minimal distance between contact pairs, for *any* barrier potential (e.g., $\ln(d)$ used in [Li et al. 2020]) one still has a barrier property *for any discrete approximation* but the strength of the barrier decreases with refinement, and the minimal distance between objects decreases.

Supplement to Orientation-aware Incremental Potential Contact

ZIZHOU HUANG, New York University, USA

MAX PAIK, New York University, USA

ZACHARY FERGUSON, Massachusetts Institute of Technology, USA

DANIELE PANOZZO, New York University, USA

DENIS ZORIN, New York University, USA

1 IMPLEMENTATION DETAILS

Our implementation makes use of most parts of the IPC algorithm, except for the potential formulation. We first utilize the existing broad-phase algorithm to collect contact candidates (edge-vertex in 2D, edge-edge and face-vertex in 3D) within $\hat{d}$.

Then in 2D, we split edge-vertex candidates into edge-vertex and vertex-vertex pairs, where the former only contain pairs whose minimum distance is not reached at the endpoints of the edge (otherwise the potential vanishes due to mollification). Similarly, in 3D, we split face-vertex candidates into face-vertex, edge-vertex, and vertex-vertex pairs. Face-vertex pairs only contain those whose minimum distance is reached in the interior of the triangle; edge-vertex pairs only contain those whose minimum distance is reached in the interior of the edge. Note that we don't need to further split an edge-edge candidate, since if it reduces to an edge-vertex pair, it's already included in the face-vertex candidates. In total, we have 2 types of collision pairs in 2D and 4 types in 3D. The reason why we have more types than IPC is that the restriction in the distance direction and edge/face orientation eliminates some pairs necessary to avoid intersections.

Since our formulation allows for a $\hat{d}$ larger than the mesh edge length, simply filtering the candidates by distance is not enough for efficiency, since every vertex has interaction with its 1-ring neighbors. To avoid a significant increase in cost, we first compute the distance types of pairs and evaluate the potential with double precision. For piecewise functions with large ranges of constant values h_ϵ, $H^{\alpha,b}$, δ_α, we cache in which range the variable belongs, to filter pairs with zero potential, and to avoid computation of gradient and hessian at trivial values (e.g. when $H^{\alpha,b}(z) = 1$). Note that even if $\hat{d}$ is larger than the edge length, it's undesirable to have contact everywhere on the surface (otherwise spurious stresses may appear), so the number of pairs with positive potential values is still sparse on the surface in most cases. We then use auto-differentiation to compute the gradient and hessian for every collision pair. We observe that in our examples, the FEM assembly and linear solve are far more expensive than our computation of potential derivatives, so we don't further optimize the efficiency of our algorithm.

2 CONVERGENT IPC LIMITATIONS

Here we evaluate the "Convergent IPC" [?] formulation and compare it to our own.

Authors' addresses: Zizhou Huang, zizhou@nyu.edu, New York University, USA; Max Paik, mp6569@nyu.edu, New York University, USA; Zachary Ferguson, zfergus@nyu.edu, Massachusetts Institute of Technology, USA; Daniele Panozzo, panozzo@nyu.edu, New York University, USA; Denis Zorin, dzorin@cs.nyu.edu, New York University, USA.

? define the continuum form of the IPC potential as

$$\psi^{IPC}(x;f) = \max_{y,\ \|x-y\|>r} p_\epsilon^{\text{IPC}}(\|f(x) - f(y)\|)$$

where the barrier function

$$p_\epsilon^{\text{IPC}}(z) = \begin{cases} -\kappa \left(\frac{z}{\epsilon} - 1\right)^2 \ln\left(\frac{z}{\epsilon}\right) & 0 < z < \epsilon \\ 0 & \epsilon \le z \end{cases}$$

and r (with $r \to 0$) is a small radius in which self-contact is ignored. From the above expression, one can make several observations.

First, instead of the integral ψ is computed as max; this is not smooth w.r.t. f. This is recognized by ? with a couple of options proposed (e.g., L_p-norm or LogSumExp) but not implemented. Instead, a smoother approximation is done in the discrete case. Therefore the method uses a smoother approximations to a non-smooth limit potential. As refinement progresses the approximation becomes less smooth.

A specific example for which the max is not smooth is when there is a switch between two parametrically distant closest points. Suppose for some x there are two points $f(y_0)$ and $f(y_1)$ equidistant from $f(x)$, and with the distance less than ϵ. If f is changing, i.e., we consider a family f_t, with a scalar parameter t, then at a point y, the velocity $v(y) = \frac{\mathrm{d}}{\mathrm{d}t} f_t(y)$ is defined. If these velocities are different at y_0 and y_1 and the closest point switches from $f(y_0)$ to $f(y_1)$, then the derivative of $\psi(x; f_t)$ has a discontinuity. This is shown in detail in Figure 1.

Second, the choice of r that would ensure that no contacts are missed, requires an estimate on the curvature of the curve f. Otherwise, no matter how small r is, the surface can fold onto itself, so that there is a contact at a point with $|x - y| < r$. Furthermore, if we would like to use a large ϵ, then r has to be equally large, excluding ever larger parts of Ω, i.e., potentially missing contact.

One smoothed version of $\psi^{IPC}(x; f)$ proposed in the paper uses L_p norm, i.e.,

$$\psi_{L_p}^{IPC}(x;f) = \left(\int_{y,\ \|x-y\|>r} p_\epsilon^{\text{IPC}}(\|f(x) - f(y)\|)^p\right)^{1/p}$$

which is similar to our potential for $p = 1$. However, the candidate set here would be all points excluding a small part of Ω near x determined by r.

The "smooth" discretization proposed for the IPC potential in 2D for a piecewise linear mesh is defined as

$$\psi^{IPC,d}(x;f) = \sum_{\text{edges } i} p_\epsilon^{\text{IPC}}(\|f(x) - f(y_i(x))\|) \\ - \sum_{\text{vertices } j} p_\epsilon^{\text{IPC}}(\|f(x) - f(x_j)\|) \tag{1}$$
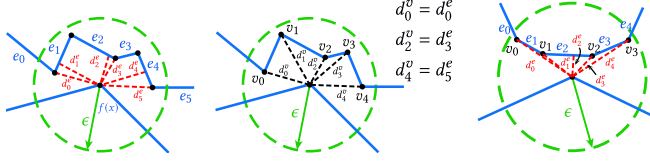
$$d_0^v = d_0^e$$
$$d_2^v = d_3^e$$
$$d_4^v = d_5^e$$

Fig. 2. Discretized convergent IPC geometry. Left: $d_i^e = \|f(x) - f(y_i(x))\|$ shown in red are the distances to the edges that are within ϵ ($\hat{d}$ in IPC notation) of the point $f(x)$, which is assumed to be a vertex. Middle: $d_j^v = \|f(x) - f(x_j)\|$ are distances to vertices within the same radius. Right: the "convex" case that motivates the convergent IPC discretization definition. In this case, the distance from the point $f(x)$ to the other candidate points (i.e., points not on the adjacent edges) that are within the ϵ ball, is increasing as one moves away from the closest point at the distance d_2^e. The distances to all edges other than the closest one are the distances to one of their vertices, which is canceled in the formula, leaving d_2^e only.
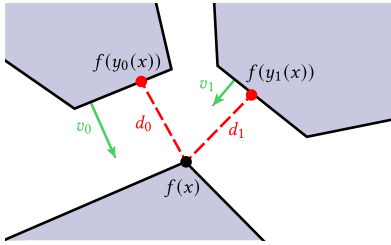

Fig. 1. Here we shows the situation for which the dependence of the minimum distance on f is discontinuous for the "Convergent IPC" formulation [?]. Suppose for $t < t_0$, $d_0 > d_1$ (i.e., d_1 is minimal) and it changes at the constant rate v_1 on some interval $t_s \le t_0 < t_e$. Similarly, d_0 is changing at a higher velocity v_0, but initially is further away, becoming the closest point for $t > t_0$. Although in f_t the dependence of f on t is smooth at every point, the minimal distance does not have a derivative w.r.t. t, at this point. If we think in discrete terms, and the endpoints of two segments are q_{ij}, $i, j = 0, 1$, then we can write $d_{\min}(t) = \min(d_0(t), d_1(t))$ as a function $d_{\min}(q_{00} + v_0 t, q_{01} + v_0 t, q_{10} + v_1 t, q_{11} + v_1 t)$, the derivative w.r.t. t is the directional derivative w.r.t. the vector of degrees of freedom q_{ij}, in the direction $[v_{00}, v_{01}, v_{10}, v_{11}]$, and it has a discontinuity.

where the summation is over all edges not containing x and all vertices excluding x if x is a vertex, and $y_i(x)$ is the closest point to x on the edge i. An illustration of this discretization is given in Figure 2.

Here one can observe that:

- As $y_i(x)$ may vary nonsmoothly with f, the expression may still be non-smooth w.r.t. f.
- The implicit assumption is that ϵ in p_ϵ^{IPC} is less than any edge length, otherwise, the potential will create an artificial repulsion between adjacent points.
- The proof of the discrete potential's positivity (which is not guaranteed by construction) uses the acceptable $\hat{d}$ assumption (ϵ in our notation), which requires either updating ϵ separately at each vertex as the mesh evolves (or setting the most conservative ϵ globally, also updating it as the mesh changes).
- For either the continuum version of the potential or for the discrete version, there is no guarantee that there is no repulsion in the undeformed shape.